

# 1 Introduction

## 1.1 Semantics and semantic theory

In its broadest sense, semantics is the study of meaning and linguistic semantics is the study of meaning as expressed by the words, phrases and sentences of human languages. It is, however, more usual within linguistics to interpret the term more narrowly, as concerning the study of those aspects of meaning encoded in linguistic expressions that are independent of their use on particular occasions by particular individuals within a particular speech community. In other words, semantics is the study of meaning abstracted away from those aspects that are derived from the intentions of speakers, their psychological states and the socio-cultural aspects of the context in which their utterances are made. A further narrowing of the term is also commonly made in separating the study of semantics from that of pragmatics. Unfortunately, the nature of the object of inquiry of the discipline (what constitutes semantic meaning, as opposed to pragmatic meaning) and the domain of the inquiry (what aspects of meaning should be addressed by the discipline) remain difficult and controversial questions. There are, however, three central aspects of the meaning of linguistic expressions that are currently accepted by most semanticists as forming the core concern of linguistic semantics. These central concerns of semantic theory, adapted from Kempson (1977:4), are stated in (1) and may be adopted as criteria for ascertaining the adequacy of semantic theories which apply in addition to the general conditions on scientific theories of falsifiability and rigour.

(1) **A semantic theory must:**

1. capture for any language the nature of the meaning of words, phrases and sentences and explain the nature of the relation between them;
2. be able to predict the ambiguities in the expressions of a language;
3. characterise and explain the systematic meaning relations between the words, the phrases and the sentences of a language.

One may add to these the condition that a semantic theory should provide an account of the relation between linguistic expressions and what may be called 'things in the world'. In other words, it is a primary concern of a semantic theory to explain how human beings can use their language to convey information about the external world. We may thus require a semantic theory to conform also to the criterion of adequacy in (2).

- (2) A semantic theory must provide an account of the relation between linguistic expressions and the things that they can be used to talk about.

There are many other aspects of meaning that can be included in the domain of linguistic semantics, but a theory conforming to the four criteria in (1) and (2) will cover the main ground of the discipline and provide a firm basis for further research. In this book, we will be looking at a particular theory of semantics that goes a long way towards satisfying these criteria and that has been very influential in linguistic

## 1 Introduction

semantics over the last two decades. This theory is a formal theory of semantics and is distinguished from **general linguistic semantics** by its greater use of mathematical techniques and reliance on logical precision. This is not to say that **formal semantics** and general linguistic semantics are completely separate disciplines. It sometimes appears that these two approaches to the semantics of natural languages are mutually incompatible, but this is not obviously true. The former draws heavily on the long tradition of research in the latter which in turn benefits from the greater precision of the former. Both approaches enable us to understand more about meaning and greater integration between them would doubtless bring greater benefits to the discipline.

Formal semantics itself was devised as a means of providing a precise interpretation for **formal languages**, i.e. the logical and mathematical languages that are opposed to **natural languages** that are spoken or written as the native languages of human beings. Many logicians considered it to be impossible to apply the same rigour to the semantics of human languages, because of their supposedly inexact syntax, their vagueness and their ambiguity. In the late nineteen-sixties, however, the philosopher Richard Montague asserted that it was possible to use the same techniques in analysing the meanings of sentences in English. In three articles, *English as a formal language*, *Universal grammar* and *The proper treatment of quantification in English*, all published or presented in 1970, Montague gave arguments for his hypothesis that:

There is in my opinion no important theoretical difference between natural languages and the artificial languages of logicians; indeed, I consider it possible to comprehend the syntax and semantics of both kinds of language within a single, natural and mathematically precise theory.  
 Montague (1974: 222)

Throughout the nineteen-seventies, after his tragic death in 1971, Montague's work had a radical effect on the study of semantics in linguistics. Indeed, his ideas on the semantics of human languages have become central to the understanding of many of the questions and theories being discussed in linguistic semantics today. Owing to the relatively recent application of the tools of formal semantics to the analysis of natural languages, however, there are many topics in linguistic semantics that have not yet been formally analysed, but it is hoped that ultimately a good deal of linguistic meaning will be amenable to the sort of rigorous treatment envisaged by Montague. It is the exposition of Montague's theory in its now classical form that constitutes the subject matter of this book, but, before the main points of his semantic theory are introduced, the four criteria of adequacy in (1) and (2) above will be discussed in more detail in order to provide a clearer idea of the fundamental issues that underlie the development of the theory in later chapters.

### 1.1.1 Compositionality

A fundamental property that any semantic theory must have is the ability to pair the syntactic expressions of a language with their meanings. In the first condition of adequacy in (1.a), above, this property is characterised as a requirement that a semantic theory account for the nature of the meaning of linguistic expressions and be able to pair every expression in a language (words, phrases and sentences) with an appropriate meaning. As already mentioned, the characterisation of meaning is a

*Semantics and semantic theory*

controversial matter, but whatever meanings are taken to be within a theory, it is obvious that there must be some way of associating them with appropriate linguistic expressions. This is not a trivial matter, however, and there are a number of important points that need to be discussed with respect to this property.

In the first place, let us consider more closely what it is that is to be assigned meanings by a semantic theory. Condition (1.a) refers to words and sentences as the carriers of meaning. The term **sentence** here is being used in its abstract sense, common in linguistics, as the largest unit of syntactic description, independently of its realisation in spoken or written texts. Like the term sentence, **word** is also ambiguous in everyday English. Within semantics, the notion of word that is most useful is that of the **lexeme** which is an abstract grammatical construct that underlies a set of **word forms** which are recognised as representatives of 'the same word' in different syntactic environments. For example, the word forms *sing*, *sings*, *singing*, *sang* and *sung* are particular inflectional variants of a lexeme which we may represent for the time being as *SING*. It is to lexemes and not to word forms that meanings should be assigned, because while the inflectional properties of the verb *SING* may vary in different syntactic environments, the sort of action described by the verb remains the same. For this reason, the meaning of words is referred to in this book as **lexical meaning**, rather than word meaning.

Although it is possible for the meanings assigned to lexemes and sentences to be very different from each other, it is reasonable to expect the meanings of sentences to be related to the meanings of the lexemes underlying the word-forms they contain. It is intuitively implausible for there to be a language where the relation between the meaning of a sentence and the meanings of its component lexemes is entirely random. While languages do contain idiomatic phrases and sentences where lexical and sentential meaning are not transparently related (e.g. *kick the bucket* meaning the same as *DIE* in English), this is never the general situation. If there were no direct relation between lexical and sentential meaning, of course, the meaning of each sentence in a language would have to be listed. Since the number of sentences that make up a language is infinite, this would mean that no human being would be able to determine the meanings of all the sentences of any language owing to the finite resources of the brain. This is absurd, of course, and just as sentences are defined recursively by syntactic rules, taking words (or morphemes) as their basis, so their meanings should also be defined recursively from the meanings ascribed to the lexemes they contain.

Thus, in addition to associating each expression in a language with a meaning, an adequate semantic theory must also be able to explain how the meanings of smaller expressions contribute to the meanings of larger ones that contain them. A theory that derives the meaning of larger expressions from those of smaller ones is said to be **compositional**. The **Principle of Compositionality**, given an initial definition in (3), is generally attributed to the German philosopher Gottlob Frege, and is thus sometimes referred to as the **Fregean Principle of Compositionality** or just the **Fregean Principle**, although it is unlikely that he ever stated the principle in precisely this way.

- (3) The meaning of an expression is a function of the meaning of its parts.

The notion of a **function** will be discussed in more detail in Chapter 4, but essentially it is an operation that derives a single result given a specified input. Thus,

### 1 Introduction

the principle (3) minimally requires that the meaning of a larger expression be uniquely determined from the meanings of its component parts. This cannot be all there is to compositionality, however, since, otherwise, we would expect that sentences containing the same words mean the same thing. This is, of course, not true. The sentence *Jo kicked Chester* does not mean the same as *Chester kicked Jo*. It must also be the case, therefore, that the syntactic structure of an expression is relevant to the derivation of its meaning. Indeed, we may strengthen the principle of compositionality so that, in deriving the meaning of a composite expression, the meaning of its component expressions are combined in some way that refers to the way they are combined by the syntax. This implies that wherever meanings are combined in a particular way to derive the meaning of a composite expression, all other composite expressions of the same sort have their meanings determined in the same way. In other words, the construction of meanings is **rule-governed**, in the same way that the construction of the well-formed syntactic expressions of a language is rule-governed. For example, whatever rule derives the meaning of the sentence *Jo sang* from the subject *Jo* and intransitive verb *sang* applies to all declarative sentences derived by combining a subject noun phrase with the appropriate form of an intransitive verb.

Furthermore, it is a general property of human languages that all the sub-expressions of a grammatically well-formed phrase have a role to play in the interpretation of a sentence, even if, on occasions, this role is predictably redundant (as, for example, in double negative constructions in certain dialects of English like *I never did nothing* where the second negative expression merely reinforces the idea of negation introduced by the first). Semantic rules should, therefore, not be allowed to delete meanings during the derivation of the meaning of a composite expression. The effect of this restriction is to make the creation of the meanings of larger expressions **monotonic** with respect to their component parts where a derivation is said to be monotonic if all properties of previous parts of a derivation are maintained throughout. In other words, once information is introduced into a monotonic derivation, it is not lost thereafter. The initial definition of compositionality in (3) may thus be strengthened to give the statement in (4).

- (4) **The principle of compositionality:** The meaning of an expression is a monotonic function of the meaning of its parts and the way they are put together.

The implications of this interpretation of the principle of compositionality is that meanings should be ascribed not only to lexemes and sentences but also to other syntactic constituents. It is thus generally assumed that meanings should be assigned to all the well-formed constituents of a language, not just to its words (lexemes) and sentences. Indeed, the concept of syntax as a bridge between phonology and semantics, current in many grammatical theories, would seem to require that all constituents be assigned a meaning by the semantics and, furthermore, that (surface) syntactic structure should directly determine how the meanings of sentences are derived. It is common to assume that semantic constituency parallels syntactic constituency and hence that an adequate semantic theory must be able to ascribe appropriate meanings to noun phrases like *the old cat*, *Jo's mother*, *Chester* and verb phrases like *sang*, *kicked the cat*, *ran slowly*, etc, according to their syntactic structure.

One way in which this may be achieved is to adopt the hypothesis that for each syntactic rule of the grammar (or syntactic structure admitted by the grammar) there is a corresponding semantic rule that derives the meaning of the resultant expression (or structure). For example, assuming that there is a rule that defines a sentence in English as consisting of a noun phrase plus a verb phrase, then the adoption of the **rule-to-rule hypothesis** in (5), together with the principle of compositionality in (4), requires that there be a corresponding semantic rule deriving the meaning of the sentence from the meanings of its immediate constituents, NP and VP.

- (5) **Rule-to-rule hypothesis:** For each syntactic rule there is a corresponding semantic rule.

The principle of compositionality in (4) is assumed to be a constraint on semantic theories and, indeed, will be seen to be the primary motivator behind much of the discussion in later chapters. The rule-to-rule hypothesis, on the other hand, is not a necessary requirement of a semantic theory, but a means of achieving compositionality. We will have cause to question the validity of this hypothesis in the later chapters of this book, but it is used in the earlier chapters to maintain a transparent relation between syntactic structures and semantic representations.

### 1.1.2 *Meaning relations*

Another aspect of meaning that must be accounted for by any semantic theory is the systematic relations that hold between different expressions in a language. According to the condition of adequacy in (1.b), a theory must satisfactorily analyse the intuitions speakers of English have about the semantic relations between lexemes and between sentences. This assumes that expressions in a language which may not be syntactically related may be related semantically and, indeed, such is the case. Consider the sentences in (6). Assuming that the reference of the name *Jo* and the discourse context are held constant for all the sentences in (6), then the sentences in (6.b) to (6.i) are semantically related to that in (6.a), even though it is not always the case that there is a direct syntactic relation between them.

- (6) a. Jo stroked a cat.  
 b. A cat was stroked.  
 c. There was a cat.  
 d. No-one stroked a cat.  
 e. There are no such things as cats.  
 f. A cat was stroked by Jo.  
 g. It was Jo who stroked the cat.  
 h. Jo touched a cat.  
 i. Jo stroked an animal.

The relation between the sentences in (6.a) and those in (6.b) and (6.c) is one of **entailment**, as it is intuitively impossible for it to be true that Jo stroked a cat on some occasion without it also being true that a cat was stroked or that there existed a cat to be stroked on that occasion. We may thus define a sentence  $S_1$  as **entailing** a sentence  $S_2$  if the latter truly describes a situation whenever the former also does. The

## 1 Introduction

negation of an entailment always derives a **contradiction** and a sentence  $S_1$  may be said to **contradict** a sentence  $S_2$  if the former must be false when the latter is true (or vice versa). For example, the assertion that Jo stroked a cat is contradicted by the non-existence of cats, making (6.e) a contradiction of (6.a). The relation between (6.a) and the sentences in (6.f) and (6.g) is also primarily semantic, although most syntactic theories recognise a syntactic relation between the sentences as well. Using a common term in a technical way, we may say these sentences are **paraphrases** of each other, since they all have the same core meaning. Another way of putting this is to say that they mutually entail each other. Hence, we may say that a sentence  $S_1$  is a paraphrase of a sentence  $S_2$  if  $S_1$  entails  $S_2$  and  $S_2$  entails  $S_1$ . An adequate theory of semantics must, therefore, provide an account of entailment, contradiction and paraphrase that allows one to identify which sentences are entailed by, or contradict or paraphrase, another in a language. Indeed, this concern, along with compositionality, is a major motivation for the theoretical programme developed in this book.

Other sorts of **implication** between sentences are also recognised in general linguistic semantics. Some of these derive from lexical meaning as in (6.h) and (6.i) which are related to (6.a) by virtue of the meanings of the lexemes *STROKE* and *CAT*, respectively. It is part of the meaning of the lexeme *STROKE* that an action of stroking also involves an action of touching, so that (6.a) implies (6.h). Furthermore, it is part of the meaning of *CAT* that anything that is a cat is also an animal and thus (6.a) implies (6.i). The meaning relations that hold between the lexemes of a language (or between lexemes and larger expressions) are called **sense relations** and include **hyponymy**, which holds if the sense of one lexeme includes that of another (e.g. between *CAT* and *ANIMAL*); **synonymy**, where two lexemes have the same sense (e.g. between *MERCURY* and *QUICKSILVER*); and **oppositeness**, where two lexemes have opposing senses (e.g. *BIG* and *SMALL* and *DEAD* and *ALIVE*). Hence, an adequate theory of semantics must give some account of lexical meaning and, in particular, of the sense relations that hold between lexemes in a particular language.

Other implicational meaning relations hold between sentences in addition to those that have been noted above. One of these is so-called **conversational implicature**, which is derived not from the conventional meanings of words or sentences, but from principles of discourse and context. For example, in the exchange in (7), the apparent irrelevance of Ethel's reply in (7.b) to Bertie's question in (7.a) leads the latter to infer (7.c). The reason behind inference has to do with Bertie's expectations about Ethel's co-operativeness in providing him with the information he needs. The fact that she has not given a straight answer leads Bertie, not to assume that she is being deliberately unhelpful, but to look for some piece of information that is relevant to his question that is indirectly implied by Ethel's response. This example is discussed in more detail in Section 1.3.1 below.

- (7) a. Bertie: Is Fiona a good lecturer?  
 b. Ethel: She has a good line in sweaters.  
 c. Bertie (thinks): Fiona is not a good lecturer.

Another sort of implication between sentences is called **presupposition**. A sentence is said to presuppose another if its truth and that of its negation both imply that the presupposed sentence is also true. In other words, presupposition deals with aspects of meaning that are assumed to hold of a situation when a sentence is uttered

*Semantics and semantic theory*

to describe that situation. For example, the use of the definite article in a noun phrase is said to presuppose the existence of something that has the property described by the common noun in the same NP. This is illustrated in (8) where the truth of the sentences in (8.c) and (8.d) is presupposed by that of (8.a) and its negation in (8.b), because of the use of *the* in the subject noun phrase. (8.c) is not implied by (8.e) which replaces *the* in the subject NP by *every* (as shown by the bracketed causal clause which denies the truth of (8.c), apparently without contradiction) and, while (8.f) implies (indeed, according to the discussion above, entails) (8.c), its negation in (8.g) does not.

- (8) a. The Duchess of Muckhart terrorised the village.  
 b. The Duchess of Muckhart didn't terrorise the village.  
 c. There is a Duchess of Muckhart.  
 d. There is a village.  
 e. Every Duchess of Muckhart terrorised the village (because there is no Duchess of Muckhart).  
 f. A Duchess of Muckhart terrorised the village.  
 g. A Duchess of Muckhart didn't terrorise the village.

It is usually assumed that implicatures such as that in (7) result from principles of conversation and thus form part of the domain of pragmatics rather than semantics (see Chapter 7 for some further discussion). More controversial, however, is the status of presupposition. Whether it should be included in semantic or pragmatic theory is an extremely vexed question, as indeed is the definition and status of the phenomenon itself. As the inclusion of this topic would require considerable discussion, it is omitted from consideration in this book. This is for convenience only and should not be taken to reflect on the importance of the topic, only on the controversiality of its analysis. The reader is referred for further information on this vexed topic to the books and articles mentioned at the end of this chapter.

In addition to accounting for these semantic relations, a semantic theory may also be required to provide some account of **anomaly** in the meaning of expressions in some language. It should, therefore, be able to explain why certain expressions which are syntactically well-formed are unacceptable or deviant from the semantic point of view. For example, the sentence in (9.a) is syntactically well-formed and semantically coherent in English. Those in (9.b) and (9.c), however, are semantically anomalous despite the fact that they have the same syntactic structure as (9.a). Such sentences can, however, be given some sort of non-literal interpretation (although (9.c) is harder to find an interpretation for than (9.b)), unlike the completely ill-formed expression in (9.d) which is simply not English. This decline in acceptability from (9.a) to (9.c) and the incoherence of (9.d) should thus be explained by an adequate theory of semantics.

- (9) a. Green Wellington boots are very popular now.  
 b. Green ideas are very popular now.  
 c. Green corollaries are very popular now.  
 d. \*very Wellington are boots popular now green.

## 1 Introduction

### 1.1.3 Ambiguity

The third area of meaning that Kempson (1977) suggests must be explained by a semantic theory is **ambiguity**. A sentence is said to be ambiguous whenever it can be associated with two or more different meanings. Ambiguity can arise in a sentence for a number of reasons: through the ascription of multiple meanings to single words (e.g. (10.a)); through the assignment of different syntactic structures to a sentence (e.g. (10.b)); or through the use of certain expressions that may have different semantic **scope** (e.g. (10.c)).

- (10) a. Ethel's punch was impressive.  
 b. The strike was called by radical lecturers and students.  
 c. Every good politician loves a cause.

The first sort of ambiguity occurs where an expression is associated with two or more unrelated meanings, as in (10.a) where the word *punch* may be interpreted as a drink or as an action. Lexemes whose word forms have this property are called **homonyms** and can be subdivided into **homophones**, where the forms of the lexeme sound the same but may be written differently, e.g. *draft* and *draught* which can both be represented phonemically as /draʃt/, and **homographs**, e.g. *lead*, which are written the same, but which are pronounced differently. Some lexemes are both homophones and homographs, like *PUNCH*. Homonyms can be divided into full homonyms (like *BANK*, *PUNCH*), where all of the lexeme's associated word forms are phonetically or orthographically identical, and partial homonyms (like *FIND*, *FOUND*), where just some of its word forms are identical.

Homonymy is often contrasted with **polysemy**. A polysemous lexeme is one that is interpreted as having multiple senses that are not entirely distinct, as is the case in the standard examples of homonyms. The classic example of a polyseme in English is the lexeme *MOUTH* which has different interpretations depending on what sort of entity is described as having a mouth. There are, for example, human mouths, mouths of caves, mouths of bottles, mouths of rivers, and so on. In each of these cases, the properties of the entity described by *MOUTH* are different, but not absolutely different, as each one refers to an opening of some sort. The difference between homonymy and polysemy is one of degree, and precise definitions of these terms are difficult and controversial. As this book is not primarily concerned with lexical meaning, no attempt will be made to differentiate the two notions or to incorporate polysemy within the theory at all. As will be seen in Chapter 2, the approach to homonymy taken here is very simplistic: the senses of homonymous lexemes are simply differentiated formally by the use of superscripts, where necessary. Although an account of polysemy and a better approach to homonymy may be possible within the theory of formal semantics presented in later chapters, these matters are not central to the concerns of this book and an adequate discussion of the issues involved would only serve to increase the size of the book without serving any great purpose. The decision to exclude polysemy from consideration and to take a simplistic view of homonymy is taken on the grounds of expository convenience and readers are again referred to the further reading noted at the end of the chapter.

A more interesting source of ambiguity from the point of view of the formal



*Semantics and semantic theory*

semanticist is illustrated in (10.b). Here the ambiguity results from the possibility of assigning two or more syntactic structures to a single grammatical string of words. To ascertain the meaning of (10.b), for example, it is necessary to know whether the adjective *radical* modifies the nominal phrase, *lecturers and students*, in which case both the lecturers and the students who called the strike are all radical, or whether it modifies just the noun *lecturers*, in which case the lecturers who called the strike are said to be radical but the political attitude of the students who did so is not specified. These two readings are illustrated in (11) where the labelled bracketings of the agentive noun phrase in (11.b) and (11.d) correspond to the readings indicated in (11.a) and (11.c), respectively.

- (11) a. The strike was called by lecturers who are radical and by students.  
 b. [<sub>NP</sub> [<sub>N1</sub> [<sub>N1</sub> radical lecturers] and students]].  
 c. The strike was called by lecturers who are radical and by students who are radical.  
 d. [<sub>NP</sub> [<sub>N1</sub> radical [<sub>N1</sub> lecturers and students]]].

In the above example, what is at issue is the scope of the adjective, *radical*. In (11.a), it modifies, and thus has scope over, the noun *lecturers*, while in (11.b) its scope is the nominal phrase *lecturers and students*. Scope is an important concept in semantics and a primary source of ambiguity which involves not only adjectives, but also conjunctions, like *and*, *or*, etc and quantifiers, like *every*, *all*, and *some* in English. **Structural ambiguity** of this sort thus has its source in the syntax of a language, but there are other scope ambiguities that do not directly depend on the syntactic structure of a sentence. Such ambiguity usually involves negation (*not*), quantification (*every*, *some*) and other elements like tense, which do not vary their syntactic position according to the reading of the sentence. For example, the two readings of the sentence in (10.c) can be made clear by those in (12). In (12.a), there is only one cause that every good politician loves, while in (12.b) each politician may love a different cause. The sentence in (10.c), however, is usually only assigned a single surface constituent structure, so that this ambiguity cannot be directly attributed to a syntactic source and is referred to as a **semantic scope ambiguity**.

- (12) a. Every politician loves a cause and that is their own career.  
 b. Every good politician loves a cause and each one loves a cause that everyone else loathes.

An adequate semantic theory must thus be able to predict where structural ambiguity is likely to arise in a language and provide a means of differentiating the interpretations of the different structures to an ambiguous sentence by the grammar, where this is relevant. It should also ensure that sentences that have two (or more) syntactic derivations, but only one semantic interpretation, are not assigned more than one meaning (see Chapter 3 for examples involving the conjunctions *and* and *or*). The theory should also provide an account of scope ambiguities where these are not directly reflected in syntactic derivations, and be able to differentiate the scopes of particular expressions independently of the syntax.

## 1 Introduction

### 1.1.4 Denotation

The final criterion of adequacy that is considered here is stated in (2), above, and is the most important for our purposes, since it forms the basis of the semantic theory to be proposed in the rest of this book. This criterion requires a semantic theory to give an account of the relation between linguistic expressions and what they can be used to talk about. Since language can be used to talk about what is outside the linguistic system, it is essential that a semantic theory should be able to associate linguistic expressions with extra-linguistic objects. Language is not used solely to talk about itself, but rather it is most commonly used to convey information about the situations in which human beings find themselves. Since a listener can in general understand the meaning of what is being said by a speaker, meanings must be publicly accessible in some sense. One way that this public accessibility must be realised is in the association of linguistic expressions with publicly identifiable entities and situations. For example, the utterance of a sentence like *The book is on the table* conveys information about two entities, one of which is conventionally called a book in English and one of which is conventionally called a table, and the relation between them. Someone who hears an utterance of this sentence associates it with the situation pictorially represented in (13). Although (13) is itself a representation of an actual (or possible) situation, it is nonetheless a non-linguistic representation and a theory of semantics should be capable of relating the meaning of the sentence to the picture and, indeed, to concrete, non-representational situations where there is a (single) book on the table.

The association between the sentence *The book is on the table* and the situation represented in (13) depends in part on there being, in the situation described, an instance of a thing that is conventionally called a book and one that is conventionally called a table in English. In other words, part of the meaning of the sentence depends on the sorts of extra-linguistic **entities** that can be referred to by the lexemes *BOOK* and *TABLE*. The aspect of the meaning of an expression that concerns its relation to such objects is called its **denotation** and an expression is said to **denote** particular sorts of extra-linguistic objects. Although this relation has often been called the **reference** of an expression, this book follows the usage of Lyons (1977) and reserves this latter term for the act of picking out a particular entity denoted by the expression through the utterance of that expression on some occasion. For example, in uttering the sentence *The book is on the table*, a speaker is said to be referring to two particular, contextually unique, entities. The entities being referred to by the use of the definite noun phrases, *the book* and *the table*, are single elements in the class of entities denoted by the lexemes *BOOK* and *TABLE*.

Thus, a speaker may use linguistic expressions to refer, but linguistic expressions themselves denote. No more will be said here about the act of reference, and the differences between denotation and reference, but for more details the reader is urged to consult the further reading at the end of this chapter.

Informally, we may think of the denotation as the relation between an expression and a class of various sorts of individuals, events, properties and relations that may be referred to by the use of the expression on some particular occasion. The lexeme *BOOK* may, therefore, be thought of as denoting the set of all books, *TABLE* as denoting the set of all tables, while the preposition *ON* may be thought of as denoting