13

Cognitive semantics in context

In this chapter, which concludes Part II of the book, we compare and contrast cognitive semantics with two other modern approaches to linguistic meaning: **truth-conditional** (or **formal**) **semantics** and **Relevance Theory**. As noted at various points in this book, cognitive semantics emerged and developed as a reaction against formal semantics. For this reason, we look in more detail at the truth-conditional approach to sentence meaning in this chapter and present some explicit points of comparison between the formal and cognitive approaches (section 13.1). We also provide an introduction to Relevance Theory, a modern approach that attempts to account for the pragmatic aspects of linguistic communication within a broader cognitive framework (section 13.2). Although this model explicitly adopts the formal view of language by assuming a modular theory of mind as well as a truth-conditional model of semantic meaning, it rejects some of the received distinctions assumed within formal approaches to linguistic meaning, such as a clear-cut division of labour between semantics and pragmatics. In this, Relevance Theory represents a formally oriented model that is in certain respects consonant with cognitive semantics. By drawing some explicit comparisons between cognitive semantics and these two models, we set the cognitive linguistics enterprise within a broader theoretical context. However, because cognitive semantics represents a collection of distinct theories, some of which examine quite distinct phenomena, this comparison will be limited to the areas that truth-conditional semantics and Relevance Theory are concerned with: while truth-conditional semantics is primarily concerned with meaning construction (or sentence meaning), Relevance Theory addresses word meaning, sentence meaning, pragmatic meanings and figurative language such as metaphor and irony.

13.1 Truth-conditional semantics

In this section we briefly present some of the ideas developed within the discipline called 'philosophy of language' that go under the name of **truth-conditional semantics**. As we saw briefly in Chapters 5 and 11, these ideas relate to meaning, truth and reality, and how meaning can be represented according to a formal metalanguage developed from logic. These ideas came to be highly influential in formal linguistics following the pioneering work of philosopher and logician Richard Montague in the 1960s and early 1970s. Montague argued that many of the ideas from the philosophy of language could be systematically applied to natural language. The tradition that grew up in linguistics following Montague's theory came to be known as truth-conditional or formal semantics.

13.1.1 Meaning, truth and reality

The philosophical interest in the relationship between meaning, truth and reality has a long and venerable tradition dating back to the ideas of the ancient Greek philosophers over 2,000 years ago. Since Aristotle, philosophers who have attempted to understand the concept of truth have equated this notion with reality as a guarantor of truth. This approach is called the **correspondence theory** and holds that a **truth bearer** (for example, a natural language sentence) is true if it corresponds to a state of affairs holding in the world. From this perspective, truth is a property of sentences that correspond to a reality they describe. The twentieth-century philosopher Alfred Tarski was influential in arguing that meaning could be equated with truth defined in terms of its correspondence with the world: if a sentence is true by virtue of its correspondence with some state of affairs, then this truth condition constitutes its meaning. Consider the following excerpt from Tarski's classic paper first published in 1944:

> *Semantics* is a discipline which . . . *deals with certain relations between expressions of a language and the objects* (or 'states of affairs') *'referred to' by those expressions*. (Tarski [1944] 2004: 119; original emphasis)

From this perspective, linguistic meaning is truth defined in terms of correspondence to reality. Meaning can therefore be defined in terms of the conditions that hold for a sentence to be true.

13.1.2 Object language versus metalanguage

Tarski argued that truth can only be defined for those languages whose semantic structure has been exactly defined and that it is not possible to define

the semantic structure of a language that is self-defining. For example, in a natural language, words are defined using other words in the language: if we 'define' *bachelor* as 'an unmarried adult male', we are using other words from the same language to define the word. According to Tarski, this fails to provide an objective definition, because it relies on words from the same language to understand other words. Tarski describes languages that are self-defining as **closed** because they fail to provide an objective definition of a particular term or expression. Therefore he argues that in order to establish the meaning of a sentence from a given natural language, we need to be able to translate the sentence from that **object language** into a **metalanguage**, a language that can be precisely and objectively defined. Tarski argues that **predicate calculus**, which was pioneered by the philosopher Gottlob Frege in his work on logic, provides a logic-based metalanguage for capturing the 'invariant' (semantic or context-independent) aspects of meaning. According to this view, predicate calculus, or a similar 'logical' language, provides a means of capturing meaning in a way that is objective, precisely stated, free from ambiguity and universal in the sense that it can be applied to any natural language.

13.1.3 The inconsistency of natural language

It is important to note that Tarski was concerned with the study of semantics (meaning in general) rather than specifically linguistic semantics. While Tarski thought that the truth-conditions for formal languages like logic could be precisely specified, he argued that the meaning of natural languages could not be precisely specified in terms of truth conditions. Tarski expresses this view in the following way:

> *The problem of the definition of truth obtains a precise meaning and can be solved in a rigorous way only for those languages whose structure has been exactly specified*. For other languages – thus, for all natural 'spoken' languages – the meaning of the problem is more or less vague, and its solution can have only an approximate character. (Tarski [1944] 2004: 121; original emphasis)

A particularly clear illustration of the way in which natural language resists precise definition in terms of truth conditions emerged from J. L. Austin's work on **speech acts**. This theory was developed in Austin's 1955 lectures, which were published posthumously in 1962. Austin observed that only certain types of sentence relate to 'states of affairs in the world'. This sentence type, which Austin called **constative**, is illustrated in examples (1) to (4).

(1) It is raining. [constatives]
(2) My cat is black and white.
(3) Tony Blair is Prime Minister.
(4) She doesn't feel very well today.

Compare examples (1)–(4) with what Austin called **performative** sentences, illustrated in examples (5)–(11).

(5) I bet you £10 it will rain tomorrow. [performatives]
(6) I hereby name this ship the HMS *Sussex*.
(7) I declare war on the citizens of Mars.
(8) I apologise.
(9) I dub thee Sir Walter.
(10) I hereby pronounce you man and wife.

Only sentences of the kind in (1) to (4) can be said to have truth conditions because they can be verified against the corresponding state of affairs that they describe. In contrast, it makes little sense to think of the sentences in (5) to (11) as 'describing' states of affairs because these sentences are performing verbal acts rather than describing situations. Observe that performatives license the adverb *hereby*, and are restricted to the first person present tense. If these sentences are changed to the third person and/or to the past tense, they become descriptions of states of affairs rather than performatives (11a), and do not license *hereby* (11b). Furthermore, only certain verbs function as performatives (11c).

(11) a. He sentenced you to ten years of hard labour yesterday.
b. He hereby sentenced you to ten years of hard labour yesterday.
c. I hereby love you.

As these examples illustrate, only a subset of sentence types can be understood in terms of their correspondence with 'states of affairs' or situations that they describe. Furthermore, this observation is not limited to the distinction between the types of examples illustrated here. For example, interrogative sentences like *Do you want a cup of tea?* and imperative sentences like *Shut the door!* cannot be described as 'true' or 'false' with respect to a given state of affairs in the world.

13.1.4 Sentences and propositions

Before exploring how truth-conditional semantics was developed into the basis of a formal approach to linguistic meaning, we first need to introduce the important distinction between **sentence** and **proposition**. A sentence is a

linguistic object, a well-formed grammatical string of words that can be described according to its grammatical properties. The meaning 'carried' by a sentence is a proposition. Crucially, there is no one-to-one correspondence between sentence and proposition because the same sentence can carry different propositions (e.g. *I love you* expresses a different proposition depending on who *I* and *you* refer to), and the same proposition can be expressed by different sentences. This is illustrated by example (12), in which both the active sentence (12a) and the passive sentence (12b) describe the same state of affairs and thus represent the same proposition. This means that these two sentences have the same truth conditions.

(12) a. Shakespeare wrote *Romeo and Juliet*.
b. *Romeo and Juliet* was written by Shakespeare.

In truth-conditional semantics, it is the meaningful proposition that is the truth-bearer. In other words, truth conditions relate to the proposition expressed by a sentence rather than directly to the sentence itself.

13.1.5 Truth-conditional semantics and the generative enterprise

Despite reservations expressed by philosophers of language like Tarski and 'natural language philosophers' like Austin, the philosopher and logician Richard Montague (e.g. 1970, 1973) argued that natural language semantics could be modelled in terms of truth conditions. According to this perspective, a crucial aspect of natural language semantics relates to logical properties and relations so that natural language can be 'translated' into the metalanguage of predicate calculus, exposing its meaning to rigorous scrutiny and definition. In this section, we present an overview of this tradition.

Montague's ideas have appealed to formal linguists because of the precision offered by the application of truth-conditional semantics to natural language. In particular, this approach has appealed to scholars who have sought to integrate the field of linguistic semantics with the generative grammar model developed by Chomsky. As we have seen in earlier chapters, language is viewed as a modular system in the tradition pioneered by Chomsky (see Figure 13.1). Within this model, each module represents an encapsulated system of linguistic knowledge that contains principles operating over primitives of a specific kind. For example, while the syntax module operates over grammatical categories like noun, verb, tense and so on, the phonology module operates over speech sounds representing bundles of articulatory features. Many semanticists influenced by the generative enterprise sought to develop an approach to natural language semantics that could provide a semantic representation for the grammatical representation generated by the syntax module: the sentence.

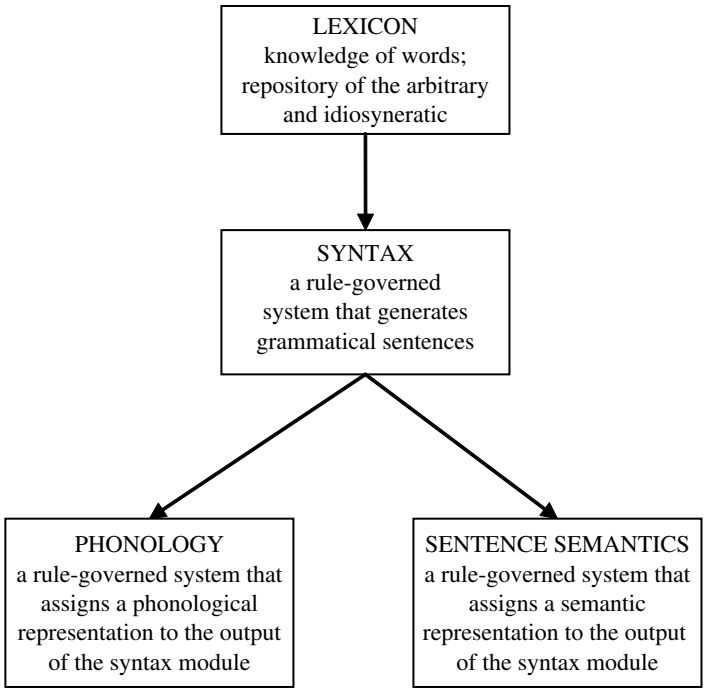

Figure 13.1 The generative model

13.1.6 Compositionality of meaning

Formal semanticists adopt the **Principle of Compositionality**. This principle states that the meaning of a complex expression is determined by the meanings of its parts, affected only by the grammatical structure in which these parts coexist. The fact that grammatical structure plays a role in linguistic meaning is illustrated by examples (13a) and (13b). These examples contain the same words, but express different propositions precisely because those parts are arranged differently within the syntactic configuration.

(13) a. Joe gave Sally a lift.
b. Sally gave Joe a lift.

The fact that syntax can affect the semantic interpretation of a sentence explains why, in the generative model, there is a semantic component that assigns a semantic representation to the output of the syntax module. While the lexicon accounts for a speaker's knowledge of word meaning, this model also requires a module that accounts for the meaning of a complex expression in which those words have been combined into a particular grammatical structure.

13.1.7 Translating natural language into a metalanguage

Predicate calculus, the logical metalanguage into which formal semanticists translate natural languages like English, contains a range of expressions. These expressions represent the meaning expressed by units of language like nouns, verbs and adjectives by means of **terms**. There are two kinds of terms: **individual constants** and **predicates**. Constants are expressions that relate to specific entities (like *James Bond* or *the spy*) and are represented by lower-case letters of the alphabet like *a, b, c* and so on. Predicates are expressions that represent processes (expressed by verbs like *eat*), properties (expressed by adjectives like *funny*), roles (expressed by nouns like *a top British spy*) and relations (expressed by prepositions like *under*). One-place predicates like *funny, die* or *a top British spy* only require a single participant to complete their meaning (e.g. *James Bond is funny; James Bond died; James Bond is a top British spy*), while two-place predicates like *appreciate* or *under* require two participants (e.g. *James Bond appreciates Miss Moneypenny; James Bond is under the desk*). Predicates are represented by upper-case letters of the alphabet, like A, B, C and so on. When constants and predicates are combined, this results in a **formula**. For example, the sentence in (14a) can be expressed by the formula in (14b), where upper-case S represents the predicate *sings* and lower-case f represents the constant *Fred*. By convention, the predicate occurs first in the predicate calculus formula, so the 'translation' does not reflect the word order of English.

(14) a. Fred sings.
b. S(f)

Example (15) illustrates a formula in which a two-place predicate combines with two constants. The relative order of the constants is important, because this reflects the difference in meaning contributed by the syntactic structure: like the natural language sentence in (15a), the formula in (15b) says that Jane loves Tom, not that Tom loves Jane.

(15) a. Jane loves Tom.
b. L(j, t)

In sentences like *Jane loves Tom and Tom loves Jane*, which consist of two or more conjoined clauses and thus express two or more propositions, the clauses are connected by natural language **connectives** like *and*, *or*, *but* and so on. In sentences like *Jane does not love Tom* or *Jane loves Tom but not Bill*, the negation word *not* is an **operator**, an expression that takes **scope** over some part of the sentence and affects its meaning. Natural language expressions like *all, every*

Table 13.1 Connectives and operators in predicate calculus

Connective	Syntax	English
$\wedge$	$x \wedge y$	X and y
$\vee$	$x \vee y$	X and/or y
$\vee_e$	$x \vee_e y$	X or y but not both
$\rightarrow$	$x \rightarrow y$	If x, then y
$\equiv$	$x \equiv y$	X if and only if y
Operator	**Syntax**	**English**
$\neg$	$\neg x$	not x
$\forall$	$\forall x$	every/all x
$\exists$	$\exists x$	some x

and *some* are also operators. These are **quantifiers** and take scope over some part of the sentence by quantifying it (for example, the sentences *Every policeman witnessed some crimes* and *Some policemen witnessed every crime* each express a different proposition due to the positions of the quantifiers, despite the fact that they contain the same predicates and constants). Connectives and operators are represented by the logical symbols in Table 13.1, where the column 'syntax' shows how these symbols can be combined with other units.

Example (16) shows how the sentence in (16a) is translated into a predicate calculus formula (16b). The expression in (16c) shows how the predicate calculus can be 'read'. In this example, *x* represents a **variable**. This is an expression that, like a constant, relates to an entity or group of entities (hence the lower-case symbol); unlike a constant, a variable does not indicate a specific entity. The lower case letters x, y and z are reserved for variables.

(16) a. Every pupil sat an exam
b. $\forall x\,(P(x) \rightarrow S(x, e))$
c. For every entity x, if x is a pupil, then x sat an exam

13.1.8 Semantic interpretation and matching

Of course, the translation of a sentence from object language to metalanguage does not in itself tell us anything about what the sentence means. To accomplish this, the symbols in the metalanguage must be assigned a semantic interpretation or **value**, at which point the formula, which represents the proposition expressed by the original natural language sentence, must be **matched** with the state of affairs it describes. The process of assigning values

and matching the proposition to the state of affairs it describes can be divided into four steps.

Assigning values

The first step is to assign the symbols of predicate calculus a semantic interpretation. This idea was implicit in the previous section, where we assigned the symbols a semantic value. For example, predicates expressed by *eat* and *love* are represented by E, L and so on, and constants expressed by proper nouns like *Jane* and *Tom* are represented by j, t and so on. Because natural language connectives and operators are closed-class expressions, these correspond to fixed logical symbols. In contrast, predicates and constants can be expressed by upper- or lower-case letters of the alphabet, with the exception of *x*, *y* and *z*, which by convention are reserved for variables.

Establishing a model of the world

The second step is the establishment of some model of the world against which the symbols in the metalanguage can be matched. Within formal semantics, models are typically represented in terms of **set theory**. For example, in a model of the world in which all women love chocolate, the sentence *All women love chocolate* would be true. However, in a model in which only a subset of women love chocolate, a further subset love chips and an intersection of these two subsets love both, the sentence all women love chocolate would be false, whereas the sentences *Some women love chocolate, Some women love chips, Some women love chocolate and chips* and *Not all women love chocolate* would be true. It is because the symbols are matched with a model of the world that this type of approach is also known as **model-theoretic semantics**. This idea is illustrated by Figure 13.2.

Matching formula with model

The third step is a matching operation in which the symbols are matched with appropriate entities in the model. This is called **denotation**: expressions in the metalanguage denote or represent elements in the model, and the meaning of the sentence is equivalent to its **denotatum**, or the sum of what it corresponds to in the model. Matching of predicates establishes the **extension** of individuals over which the predicate holds, which is represented in terms of sets. For example, in the sentence *All women love chocolate*, the predicate *love* represents a relation between the set of all entities described as *women* and the set of all entities described as *chocolate*. Once this matching operation has taken place, then the truth value of the sentence can be calculated.

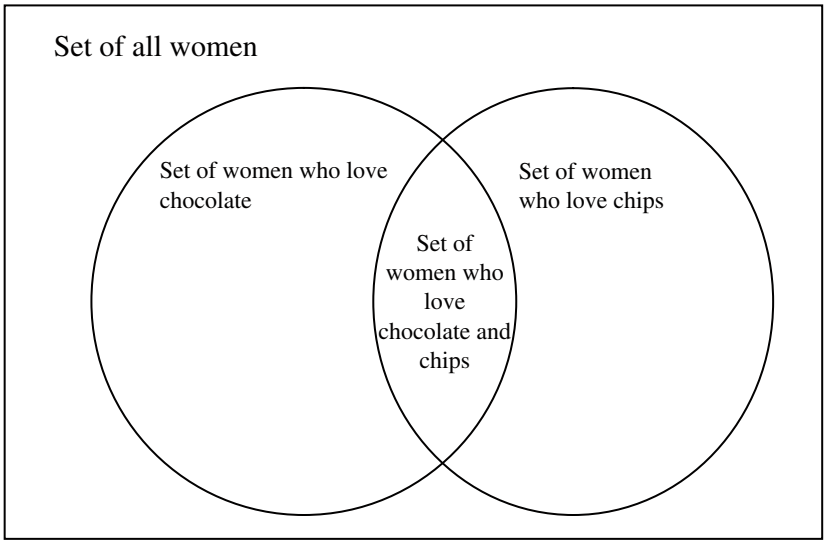

Figure 13.2 Set-theoretic model

Table 13.2 Steps for calculating truth conditions

Assigning values	This assigns a semantic value to the symbols in the formula. Upper case A, B, C correspond to predicates; lower case a, b, c to constants; and x, y, z to variables. Fixed symbols represent connectives and operators.
Establishing a model of the world	This set-theoretic model represents a 'state of affairs' against which the sentence is matched.
Matching formula with model	On the basis of correspondence theory, the denotatum of a sentence is its correspondence with the state of affairs represented by the model. The denotatum is composed of assignment of individual constant terms and the representation of predicates as a set (extension) of individuals over which the predicate holds.
Calculating truth values	Since meaning is defined in terms of truth and truth in terms of correspondence, the truth value is calculated on the basis of the correspondence between the sentence and the model.

Calculating truth values

The fourth step involves the calculation of truth values. If the formula matches the model, then the sentence is true. If it does not, then the sentence is false. These steps are summarised in Table 13.2. As this brief overview shows, in truth-conditional semantics the meaning of a sentence is equivalent to the conditions that hold for that sentence to be true, relative to a model of the world. Central to this approach is the correspondence theory of truth that

we considered earlier (section 13.1.1): meaning is defined in terms of the truth of a sentence, understood as conditions in the world (or a model of the world) to which the sentence corresponds.

We illustrate each of these steps with example (15), which is repeated here.

(15) a. Jane loves Tom.
b. L(j, t)

Once the sentence is translated into predicate calculus (15b), values are assigned to the logical symbols (e.g. j = Jane; t = Tom) and a model is established that identifies the entities corresponding to the linguistic expressions *Jane* and *Tom*. This model might represent the set of all people {Bill, Fred, Jane, Mary, Susan, Tom. . .}. Within this model is a domain or subset of entities who stand in the relation expressed by the predicate *love* (L). This is represented by (17), in which each ordered pair (inside angled brackets) stands in the relevant relation.

(17) L = {<Jane, Tom>, <Fred, Mary>, <Mary, Susan>}

Next, the formula is matched with the model so that constants and predicates are matched with entities and relations in the model. As (17) shows, this set contains an ordered pair, which means that Jane loves Tom. Finally, the truth condition of the proposition expressed by (15) is evaluated relative to this model. The rule for this evaluation process is shown in (18).

(18) [L(j, t) = 1 ≡ [<j, t>] ∈ [L]]

In this rule, the number '1' represents 'true' (as opposed to '0', which represents 'false'). This rule says '*Jane loves Tom* is true if and only if the ordered pair <Jane, Tom> is a member of the set L'. Since the set L contains the ordered pair <Jane, Tom> in the model, the sentence is true. Table 13.3 completes this brief overview of the truth-conditional approach to sentence meaning in formal semantics by summarising the properties that characterise this approach as it is conceived by generatively oriented semanticists.

13.1.9 Comparison with cognitive semantics

While the assumptions presented in Table 13.3 stand in direct opposition to those adopted within cognitive semantics, there are nevertheless some important similarities between the two approaches. Firstly, both approaches are concerned with explaining sentence meaning and with the nature of the relationships between the words in a sentence, as well as between the words

Table 13.3 Truth-conditional formal semantics

The nativist hypothesis is widely assumed.
The modularity hypothesis is widely assumed: linguistic knowledge emerges from an encapsulated cognitive system, and the language module itself has a modular structure.
Semantic (context-independent) knowledge is separable from pragmatic (context-dependent) and encyclopaedic (non-linguistic) knowledge.
A correspondence theory of truth is assumed, hence this approach is 'objectivist' in the sense that sentence meaning relies upon an objectively defined world or model of the world.
Sentence meaning can be modelled using a logical metalanguage.
The meaning of complex expressions is compositional. Figurative language is non-compositional and therefore exceptional.
In practice, this approach is focused upon the logical properties of a carefully selected set of declarative sentences.

and the grammatical structure in which they occur. Secondly, both formal semantics and cognitive semantics accept the existence of a real external world which bears upon the nature of linguistic meaning. For example, both theories distinguish between entities, properties, processes and relations. Thirdly, both approaches assume that humans have stable knowledge of the external world which is reflected in language, and attempt to model this knowledge. While the earliest truth-conditional models relied upon a direct link between language and external world (**referential** or **denotational** models), modern formal semantics attempts to model the system of human knowledge that mediates between linguistic symbols and external reality. Therefore, like cognitive semantics, formal semantics aims to construct a **representational** model.

Despite these important similarities, the differences remain significant. Beginning with fundamental assumptions, while formal semanticists assume an innate and modular system of specialised linguistic knowledge, cognitive semanticists reject this view in favour of a semantic system that provides 'prompts' to the rich conceptual system that it reflects. In adopting an objectivist approach to cognition, truth-conditional semanticists see human thought as 'disembodied' because linguistic meaning is conceived in terms of correspondence theory. In contrast, in adopting a broadly experientialist or empiricist approach to cognition, cognitive semanticists conceive meaning as the imaginative projection of bodily experience onto abstract cognitive models.

Turning to how each model views the nature of linguistic meaning, formal semanticists argue that one of the primary goals of a theory of linguistic meaning is to address the **informational significance** of language. From this perspective, language is used primarily to describe states of affairs in the 'world', which are thus central to the account of linguistic meaning, as we have seen. This idea is represented by Figure 13.3.

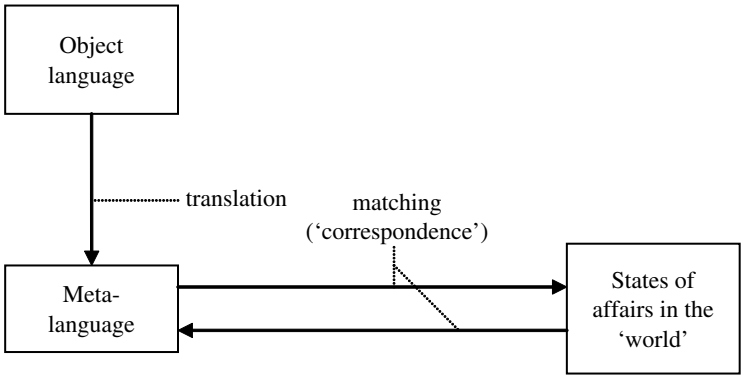

Figure 13.3 The construction of sentence-meaning in formal semantics

In Figure 13.3, the arrow from the object language to the metalanguage represents the translation process, which gives rise to a representation in the unambiguous and universally applicable language of predicate calculus. Meaning then derives from how well the values associated with the metalanguage correspond to a given state of affairs in the 'world', real or hypothetical.

In contrast, cognitive semanticists argue that the role of language is to prompt for conceptual representations (including simulations in the sense discussed in Chapter 7), so that meaning derives not from an objectively defined 'world' but from structured mental representations that reflect and model the world we experience as embodied human beings. According to the view in cognitive semantics, these mental representations are partly stable (stored) knowledge systems and partly dynamic (on-line) conceptualisations. It follows from this view that linguistic meaning resides not within a specialised system of linguistic knowledge but at the conceptual level itself. The cognitive view of the nature of linguistic meaning is represented by Figure 13.4.

Figure 13.4 represents the idea that two basic kinds of experience (sensory-perceptual experience of the external world and subjective experience from the introspective 'world') give rise to conceptual representations which can lead to simulations. Language prompts for these conceptual representations, serving as 'points of access' to relatively stable encyclopaedic knowledge (this is indicated by the arrow from 'language' to 'representation'). Conceptual representations are also subject to further processes of dynamic meaning construction. Meaning-construction can in turn have consequences for language, for example by giving rise to language change (this is indicated by the arrow from 'meaning-construction' to 'language'). For example, using the lexical item *mouse* to refer to a piece of computer hardware that 'resembles' a mouse is a consequence of single-scope blending; recall from the previous chapter that this involves the frame from one input space serving to organise the structure projected to the

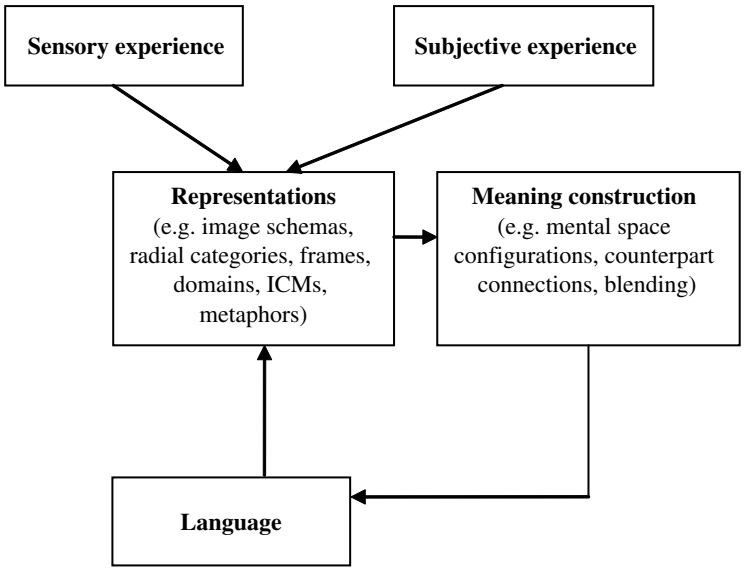

Figure 13.4 The nature of meaning construction in cognitive semantics

blended space. However, this blend has consequences for language: as a consequence of the perceived resemblance between a mouse and an item of computer hardware, the conceptual integration network that results affects conventional language use. Indeed, the conventional application of the lexical item *mouse* to the 'computer mouse' can be seen as testimony to the impact of blending on language. This illustrates the usage-based nature of the cognitive model, where language both gives rise to (= prompts for) conceptualisation (affecting our conceived 'reality') and in turn is modified and transformed by the resulting conceptual representations.

A further important difference relates to the nature of the relationship between semantics (context-independent meaning) and pragmatics (context-dependent meaning). As we have seen, cognitive semanticists adopt an encyclopaedic view of meaning together with a dynamic context-driven view of meaning construction, which entails that there is no principled distinction between semantic and pragmatic knowledge. In contrast, formal semanticists assume a sharp boundary between the two types of knowledge. According to this view, semantic knowledge is stable, conventionalised knowledge that is expressed by predictable form-meaning correspondences and is contained within the linguistic system. In contrast, pragmatic inferences cannot be predicted from linguistic form; pragmatic knowledge involves more generalised inferencing processes that do not relate specifically to language but operate over the output of the language system together with non-linguistic contextual factors. This is the issue that Relevance Theory addresses, to which we turn directly.

13.2 Relevance Theory

Relevance Theory was developed by psychologist Dan Sperber and linguist Deirdre Wilson, and develops key insights from the well-known theory of pragmatics proposed by Paul Grice (1975). We base our discussion here on the 1995 edition of their landmark book, *Relevance: Communication and Cognition*, which was originally published in 1986. Relevance Theory represents a modern approach to pragmatics that adopts an explicitly generative view of language, and aims to provide a mentalist account of communication that can be integrated with the generative model of language. Despite its generative orientation, in its emphasis on linguistic communication within the context of general cognition, Relevance Theory is consonant with cognitive semantics in a number of respects. For example, Sperber and Wilson reject the semantic decomposition account of word meaning that characterises the standard formal view, and argue in favour of the incorporation of encyclopaedic meaning within the lexical representation of words. In this section, we focus on the Relevance Theoretic account of meaning construction, or sentence meaning.

13.2.1 Ostensive communication

Relevance Theory is a theoretical approach to communication in general, which views **verbal communication** as one instance of **ostensive-inferential communication**. According to Sperber and Wilson, the defining characteristic of communication is that it involves revealing or making manifest a particular **communicative intention**. In other words, the communicator's intention is revealed by some kind of ostensive behaviour. For example, in response to the question *How are you getting home?* you can perform a manual gesture representing a car's steering wheel. This is a form of ostensive behaviour signalling a specific communicative intention, namely that the 'addressee' should infer that you will be driving home. Equally, if you are at a party that you wish to leave, you can raise your arm and tap your watch to indicate to your partner that it's time to go. In both cases, the act would fail as an instance of communication if it were not ostensive. For example, if you were sitting in the bathroom by yourself, the act of tapping your watch would fail to achieve ostensive-inferential communication.

13.2.2 Mutual cognitive environment

Of course, for speaker and hearer to communicate successfully, particularly where inference is concerned, they must rely upon shared information. For example, the person in our earlier example who indicates that s/he will be driving home relies upon the assumption that his or her 'addressee' knows that cars have steering wheels and can recognise the gesture that represents this.

Sperber and Wilson describe this shared knowledge upon which inferences depend as the 'mutual cognitive environment'. Consider the following excerpts from Sperber and Wilson.

> The cognitive environment of an individual is a set of facts that are manifest to him . . . A fact is manifest to an individual at a given time iff [if and only if] he is capable at that time of representing it mentally and accepting its representation as true or probably true . . . an individual's total cognitive environment is the set of all the facts he can perceive or infer . . . a function of his physical environment and his cognitive abilities . . . The total shared environment of two people is the intersection of their two total cognitive environments, i.e. the set of all facts that are manifest to them both. (Sperber and Wilson 1995: 39–41)

As these excerpts make explicit, inference depends upon the speaker's knowledge, and the knowledge s/he can assume on the part of the hearer.

13.2.3 Relevance

According to Sperber and Wilson, human cognition is driven by **relevance** in the sense that information (whether sensory-perceptual or linguistic) is selectively processed on the basis of the search for **contextual effects**: information that will affect our existing knowledge in some useful way or will allow us to construct an inference. For example, imagine driving down the road in your car with the radio on. In this context, you are bombarded with sensory-perceptual stimuli including visual stimuli as well as linguistic and non-linguistic sounds. Suppose that you have been worried about your car lately. In this context, you might 'tune out' the linguistic sounds coming from the radio and focus your attention on the sounds coming from under the bonnet. Depending on whether these sounds are out of the ordinary or not, this information will interact with what you already know about your car and allow you to draw some conclusions. In this context, the car's sounds are more relevant than the radio's sounds. Now imagine that you are late for work and concerned about the time. You transfer your attention to the linguistic sounds coming from the radio and listen for the newsreader to announce the time. In this context, the radio's sounds are more relevant than the car's sounds. As this simple example illustrates, the human mind constantly searches for relevant information. This idea is captured by the '**Cognitive Principle of Relevance**', which states that 'Human cognition tends to be geared to the maximisation of relevance' (Sperber and Wilson 1995: 158).

Sperber and Wilson argue that ostensive-inferential communication is driven by the presumption of relevance. In other words, a hearer will assume that any act of (linguistic or non-linguistic) ostensive-inferential communication is

relevant, and moreover will search for the **optimally relevant** interpretation. It is this assumption that allows us to deduce or **infer** the communicative intention signalled by an act of ostensive communication. This idea is captured by the '**Communicative Principle of Relevance**', which states that 'Every act of ostensive communication communicates a presumption of its own optimal relevance' (Sperber and Wilson 1995: 260). 'Optimal relevance' is defined in the following way:

> Presumption of optimal relevance (Sperber and Wilson 1995: 158)
> 1. The set of assumptions **I** which the communicator intends to make manifest to the addressee is relevant enough to make it worth the addressee's while to process the ostensive stimulus.
> 2. The ostensive stimulus is the most relevant one the communicator could have used to communicate **I**.

Consider example (19) from Sperber and Wilson (1995: 189). Imagine that this utterance is made in a jeweller's shop in response to an enquiry from a customer about how long they might expect to wait for the watch to be repaired.

(19) It will take some time to repair your watch.

It is obvious that a watch repair must take 'some time' (as opposed to no time), so the customer assumes that the communicative intention behind the utterance cannot be to convey this uninformative and therefore irrelevant interpretation. Sperber and Wilson argue that our presumption of relevance in everyday communication guides us to a more appropriate interpretation of the utterance. If the customer knows that it usually takes about a week to get a watch repaired, then the most relevant reason for mentioning the time it will take is probably because the repair will take significantly longer than a week.

13.2.4 Explicature and implicature

Sperber and Wilson follow the formal view in distinguishing between what they call **explicature** and **implicature**. The term 'explicature' describes an assumption that is explicitly communicated. In relating to explicit or context-independent meaning, this term roughly corresponds to the traditional idea of semantic meaning. The term 'implicature', which is adopted from Grice (1975), relates to implicit or inferential (context-dependent) meaning, and corresponds to the traditional view of pragmatic meaning. Sperber and Wilson also follow the standard formal view in assuming that semantic 'decoding' takes place prior to the calculation of pragmatic inferences. However, they depart from the standard formal view in arguing that meaning construction relies to

considerable extent upon inference, even in the 'decoding' of explicatures. This idea is illustrated by example (20) from Sperber and Wilson (1995: 186).

(20) The child left the straw in the glass.

This sentence is straightforwardly interpreted to mean that a child left a 'drinking tube' in a glass drinking vessel. This meaning is the explicature expressed by the sentence. However, as Sperber and Wilson observe, even this straightforward sentence requires some inferential work, because the expression *straw* is lexically ambiguous: it could mean the child left a 'cereal stalk' in the glass. To derive the more likely or accessible 'drinking tube' interpretation, the hearer has to access encyclopaedic information relating to children and the typical scenarios involving a 'straw' and a 'glass'. The availability of the most salient interpretation might also depend on contextual information, such as whether the child in question was in a kitchen or a farmyard. As this example illustrates, many explicatures will rely upon inference on the part of the hearer in order to retrieve the intended meaning. Indeed, all explicatures containing referential expressions like *that man* or *him* rely upon inference for **reference assignment**: matching a referring expression with the 'right' entity. Sperber and Wilson's model therefore departs from the standard formal model in emphasising the role of inference in deriving explicit meaning. The exchange in example (21) illustrates how an implicature is derived (Sperber and Wilson 1995: 194).

(21) a. Peter: Would you drive a Mercedes?
b. Mary: I wouldn't drive ANY expensive car.

In this exchange, Mary fails to answer Peter's question directly (because Peter's utterance is a 'yes-no question' a straightforward 'yes' or 'no' would provide a direct answer). The presumption of relevance allows Peter to assume that Mary has answered the question in the most relevant way possible and to infer her intended meaning. Mary's utterance interacts with Peter's encyclopaedic knowledge and gives rise to the fact that a Mercedes is an expensive car. This fact interacts with Mary's assertion that she wouldn't drive ANY expensive car, and by a process of logical deduction gives rise to the explicature that Mary wouldn't drive a Mercedes. Mary's utterance counts as the optimally relevant way of answering Peter's question because it is maximally informative. Her utterance gives rise to a greater number of contextual effects than a direct 'no' response, because Peter now knows not only that Mary wouldn't drive a Mercedes, but also that she wouldn't drive a BMW, a Bentley, a Jaguar and so on. From this perspective, the extra effort or processing 'cost' involved in the retrieval of the implicature(s) is rewarded by the 'benefit' of a greater number of contextual effects.

13.2.5 Metaphor

Finally, we briefly consider Sperber and Wilson's account of figurative language, focusing on their discussion of metaphor. Sperber and Wilson argue that relevance and inference are also central to the interpretation of figurative language. Consider example (22) from Sperber and Wilson (1995: 236).

(22) This room is a pigsty.

According to the Relevance Theory account, the hearer is licensed to assume that the speaker is aiming for optimal relevance in uttering (22). Because the utterance is literally false (the room is not literally a pigsty), the literal interpretation is uninformative and therefore irrelevant. The hearer therefore assumes that the speaker intends some other interpretation and draws upon encyclopaedic knowledge and contextual knowledge in order to construct an inference. Encyclopaedic knowledge gives rise to the fact that a pigsty is associated with filth and untidiness. The resemblance between the encyclopaedic representation of a pigsty and the condition of the room (contextual information) allows the hearer to infer that the speaker intends to convey that the room is filthy and untidy. As Sperber and Wilson point out, the use of this metaphor carries additional contextual effects that could not be conveyed by the utterance *This room is filthy and dirty*. By comparing the room to a pigsty, the speaker provides a much richer representation of the condition of the room which might give rise to further implicatures (e.g. the filth and untidiness goes 'beyond the norm' for a room inhabited by humans rather than animals, the room smells bad, and so on). In this way, metaphor also rewards the hearer's extra processing cost with a richer set of contextual effects than a literal utterance: 'the wider the range of potential implicatures and the greater the hearer's responsibility for constructing them, the more poetic the effect, the more creative the metaphor' (Sperber and Wilson 1995: 236). Table 13.4 summarises the main assumptions of Relevance Theory.

13.2.6 Comparison with cognitive semantics

In many respects, the Relevance Theory view of meaning construction is similar to the view taken in cognitive approaches, including Mental Spaces Theory and Blending Theory. Both Relevance Theory and cognitive semantics are concerned with describing the mental processes involved in meaning construction. Like cognitive semantics, Relevance Theory focuses upon developing a psychologically plausible account of communication, and in emphasising inference, encyclopaedic knowledge and contextual knowledge, it relates to the processes that mental spaces and blending theorists refer to as projection, mapping, schema induction and integration. Furthermore, both Relevance Theory and cognitive

Table 13.4 Relevance Theory

Primarily concerned with accounting for ostensive-inferential communication; language is just one form of this.
Shared knowledge is the 'mutual cognitive environment'.
Cognition is driven by the search for relevance (Cognitive Principle of Relevance); relevance yields contextual effects.
Acts of ostensive communication (including utterances) presume their own optimal relevance.
Optimal relevance means that the information is worth retrieving and that the hearer has chosen the most relevant means of communicating.
While explicature and implicature roughly correspond to semantic and pragmatic meaning, respectively, both rely upon inference, which is relevance-driven.
Metaphors (and other forms of figurative language) are interpreted according to the same principles as literal utterances; they are relevance-driven in nature and provide a richer set of inferences than literal utterances.

semantics emphasise the idea that meaning construction is in large measure due to these mental processes rather than a simple matter of composing a sentence's meaning from its parts. Indeed, Sperber and Wilson explicitly reject what they call the 'code model' as a descriptively adequate account of communication. Furthermore, Sperber and Wilson claim that explicature, as well as implicature, require extensive inferencing (in processes such as disambiguation and reference assignment). In this respect, and in relying upon contextual and encyclopaedic information in these processes, Sperber and Wilson's view is consonant with the claim made by cognitive semanticists that words represent 'prompts' for meaning construction, and with the idea that a strict dividing line between semantics and pragmatics cannot be straightforwardly upheld. Finally, Sperber and Wilson argue that metaphor and other types of figurative language are unexceptional in the sense that they exploit the same cognitive processes by maximising relevance. In this respect, although the details of the Relevance Theoretic account of metaphor focus more on communication than on cognition, the integration of figurative and literal language is also consonant with the cognitive account.

Despite these areas of agreement, there are some fundamental differences between the two approaches. Most importantly, Relevance Theory assumes as its background a generative model of language; this model assumes the nativist hypothesis and the modularity hypothesis. In addition, Relevance Theory assumes a logical truth-conditional account of certain aspects of linguistic meaning. As a theory of communication, Relevance Theory provides an account of linguistic meaning with an emphasis on pragmatics, and sets out to account for the on-line process of meaning construction in more detail than it accounts for the stable knowledge systems that comprise knowledge of language or competence in the Chomskyan sense. In this respect, Relevance Theory accepts the distinction between linguistic knowledge and non-linguistic knowledge, and focuses on how the two interact to give rise to interpretation in communicative

contexts. This relatively broad focus explains why certain aspects of the model resonate with cognitive approaches, despite starting assumptions that stand in direct opposition to the cognitive view. A further difference relates to the fact that Relevance Theory places the emphasis on communication (the speaker's intentions and the hearer's assumptions in deriving inferences), while cognitive semantics emphasises the nature of the conceptual system and conceptual processes. For example, while Relevance Theory emphasises the communicative aspects of metaphor, conceptual metaphor theorists emphasise the structural dimensions of metaphor within the conceptual system. Finally, each approach focuses on a largely distinct range of phenomena. Relevance Theory, although it develops a new perspective, is nevertheless concerned with accounting for the phenomena that have traditionally been of concern within approaches to linguistic meaning, such as ambiguity, the nature of the relationships between word meaning and sentence meaning, between explicit and implicit meaning, and between literal and figurative language. In contrast, cognitive semantics addresses a wider range of phenomena, and is concerned not only with addressing long-standing concerns within approaches to linguistic meaning, but also with phenomena revealed by other related disciplines that cast light upon the nature of the conceptual system.

13.3 Summary

In this chapter we compared and contrasted cognitive semantics with two other modern approaches to linguistic meaning: **formal (truth-conditional) semantics** and **Relevance Theory**. As we observed, while the assumptions of truth-conditional semantics stand in direct opposition to the assumptions of cognitive semantics, certain claims made within Relevance Theory are more consonant with the cognitive approach. Truth-conditional semantics takes an **objectivist** approach to meaning, and is concerned with modelling sentences in terms of their **correspondence** to the 'world'. This is achieved by first translating natural language sentences into a logical **metalanguage**, and then by establishing how the logical form derived corresponds to a particular model of reality, represented in terms of **set theory**. Formal semanticists have been primarily concerned with **sentence meaning**. Relevance Theory, in contrast, is a theory of **communication**. The main architects of the theory, Sperber and Wilson, emphasise the role of **ostensive-inferential communication, relevance** and **inference**. They argue that both explicit and implicit meaning construction relies upon contextual and encyclopaedic knowledge in giving rise to inferences, and that metaphor relies upon the same communicative goals as literal language. Despite these similarities, Relevance Theory assumes a generative model of language and therefore accepts the distinction between linguistic and non-linguistic knowledge. In these respects, Relevance Theory is formally

oriented and rests upon guiding assumptions that stand in direct opposition to those of cognitive semantics.

Further reading

Readings in formal semantics

- **Bach (1989).** This is one of the most accessible book-length introductions to formal semantics.
- **Cann (1993).** This textbook is a challenging introduction for the novice, but is to be commended for attempting to introduce Montague's approach to natural language semantics without presupposing a particular theory of grammar.
- **Chierchia and McConnell-Ginet (2000).** A relatively accessible introduction to formal semantics.
- **Heim and Kratzer (1998).** This textbook explicitly attempts to relate formal semantics with grammatical phenomena from the perspective of Generative Grammar.
- **Portner (2005).** Another very accessible introduction to formal semantics.
- **Saeed (2003).** Saeed's excellent general introduction to semantics includes a chapter-length introduction to formal (truth-conditional) semantics. This is the most accessible chapter-length introduction around. Saeed also provides an overview of Jackendoff's Conceptual Semantics theory of linguistic meaning, which we briefly mentioned in Chapters 3 and 5. The reader is strongly encouraged to investigate Jackendoff's theory in order to gain insights into a non-truth-conditional formal model of linguistic meaning.

Relevance Theory

- **Carston (2002).** An extended application of Relevance Theory to a range of linguistic phenomena.
- **Sperber and Wilson (1995).** The seminal text by the architects of Relevance Theory, this book provides a remarkably accessible introduction.

Exercises

13.1 What's 'cognitive' about cognitive semantics?

In view of the discussion in Part II of the book, can you provide a rationale for the use of the term *cognitive* in cognitive semantics? In what respects can the

formal or generatively oriented models we have discussed at various points in Part II of the book, as well as in this chapter, also be described as 'cognitive'?

13.2 Comparison between approaches

Make an annotated table of the points of similarity and contrast between the approaches compared in this chapter.

13.3 Propositions versus construals

One of the key distinctions between formal and cognitive approaches relates to their different views about grammatical structure. As we saw in Chapter 6, cognitive approaches view grammatical structure as independently meaningful while formal approaches do not. An important idea that we will discuss in detail in Part III relates to the notion of *construal*: the idea that different grammatical forms, like different words, give rise to distinct construals or 'ways of seeing'. Consider the following examples.

(a) John kicked the ball.
(b) The ball was kicked by John.

From the perspective of truth-conditional semantics, these sentences both encode the same proposition and therefore express the same 'meaning'. From what you have learned in this part of the book, (i) say what the difference in meaning is, and (ii) explain how it is encoded linguistically. How might these differences be accounted for within the formal approach? Comment on what these examples reveal in terms of differing assumptions between cognitive semantics and formal semantics.

13.4 Metaphor

Consider the following sentence.

John is a block of ice.

Provide analyses of this example from the perspective of both Conceptual Metaphor Theory and Relevance Theory. In order to do so, you will need to be explicit about the context you are assuming. What do your analyses reveal about the similarities and differences between these two approaches?