

Part I

Introduction

The four chapters that make up Part I lay the foundations for the development of LCCM Theory in the remainder of the book. Chapter 1 addresses the inherent variation in word meaning in situated contexts of use, the central problem addressed in the book. Also reviewed—and rejected—is the standard account of meaning in linguistic semantics, referred to as *literalism*. Chapter 2 introduces the theoretical starting points and assumptions upon which LCCM Theory rests. Chapter 3 introduces the perspective provided by cognitive linguistics, and shows how this informs the development of LCCM Theory. Chapter 4 provides an informal introduction to the account of word meaning provided by LCCM Theory.

1

Words and meaning

[M]eaning is the “holy grail” not only of linguistics, but also of philosophy, psychology, and neuroscience... Understanding how we mean and how we think is a vital issue for our intuitive sense of ourselves as human beings. For most people, meaning is intuitively the central issue in the study of language—far more important than understanding details of word order or morphology.

Ray Jackendoff, *Foundations of Language* (2002: 267)

Providing an account of the nature of meaning and meaning construction processes is, as observed in the quotation above, the Holy Grail of linguistics as well as a range of related disciplines in the humanities and the social and cognitive sciences. In this book I am concerned with word meaning, and the role of words in meaning construction: how words mean. This is fundamental to an account of the role of language in giving rise to meaning. Nevertheless, accounting for the role of words in meaning construction has proved to be both controversial and problematic for much of the relatively short history of linguistics as a discipline, as well as for research on language within philosophy, and, indeed, for work more generally in cognitive science.

The specific problem that I address in this book is this: how do we account for the inherent variation of word meaning in language use? That is, the meaning associated with any given word form appears to vary each time it is used, in terms of the conceptualization that it, in part, gives rise to. To illustrate, consider the following examples focusing on the form *France*:

- (1) a. France is a country of outstanding natural beauty
- b. France is one of the leading nations in the European Union
- c. France beat New Zealand in the 2007 Rugby World Cup
- d. France voted against the EU constitution in the 2005 referendum

In these examples the meaning associated with *France* varies across each instance of use. In the first example, *France* relates to a specific geographical landmass coincident with the borders of mainland France. In the second example, *France* relates to the political nation state, encompassing its political infrastructure, political and economic influence, and its citizens, including those in French overseas territories. In the example in (1c) *France* relates to the

team of fifteen rugby players, drawn from the pool of rugby players of French citizenship, who represented the French nation in the 2007 Rugby World Cup. In the final example, *France* relates to the French electorate, and specifically that part of the electorate which voted against proceeding with ratification of a proposed EU constitution in a national referendum in 2005.

These examples illustrate that a word form such as *France* appears to be protean in nature: its meaning is flexible, in part dependent upon the **context** of its use. This notion of context must include, at the very least, all of the following, discussed in more detail later in the chapter: (i) the other words that make up the utterance itself, (ii) the background knowledge shared by the speaker and hearer, (iii) the physical venue and temporal setting of the utterance, and (iv) the communicative intention of the speaker, as recognized and interpreted by the hearer, in service of facilitating the interactional goal(s).

My task in this book is to provide a theoretical account of the flexibility associated with word meaning in language use. To do so, we will need to examine and develop an account of a number of issues. Firstly, I will develop an account of **semantic structure**, which is to say, the nature of much of the linguistic knowledge associated with words.¹ This must include an account of the knowledge of **usage patterns** associated with words, including what counts as an appropriate **context of use**, given the notion of the components of context just sketched, and elaborated on below. Secondly, I will develop an account of **conceptual structure**. This relates to the non-linguistic knowledge representations that words tap into and can draw upon in situated language use. Together, an account of semantic structure and conceptual structure constitutes an account of what I refer to as **semantic representation**. Thirdly, I develop an account of the linguistic processes that facilitate **composition**, giving rise to distinct conceptualizations associated with a word such as *France* as illustrated in the examples above. Finally, I attempt to do all this while bearing in mind that meaning construction constitutes a form of **joint action** (Clark 1996), in service of situated communicative goals. Hence, the approach I take to lexical and compositional semantics must be thoroughly grounded in a **usage-based** perspective (Langacker 2000). The tack I take, in presenting an account of the issues just outlined, is to develop and introduce a new—or at least a differently nuanced—theory of lexical representation and meaning construction. This is termed the **Theory of Lexical Concepts and Cognitive Models (LCCM Theory)**. I begin the presentation of this new approach in the next chapter.

However, we must first examine the received view of word meaning that has emerged in contemporary linguistics, and consider problems that arise for it. This will allow us to move towards a new account of lexical representation, and **compositionality**—how words are composed in service of situated meaning construction. This is our task in the present chapter.

¹ I will specify the nature of semantic structure assumed by LCCM Theory in Part II of the book.

The received view of word meaning

The standard account of word meaning, at least in the dominant Anglo-American tradition, I refer to as **literalism**; in this I am following Recanati (2004). In fact, literalism is less an account of word meaning, being more an account of the nature of linguistic semantics in general, of which word meaning is clearly a central aspect. Literalism is also less an account associated with any individual scholar. Nevertheless, it is probably fair to claim, as Recanati does, that it represents the dominant position in modern linguistics with respect to the nature of word meaning, **sentence meaning**, and **speaker meaning**. In particular, literalism takes as axiomatic the principled division of labour between **semantics**—the context-independent aspects of meaning—and **pragmatics**—the context-dependent aspects of meaning. In this section I first present the perspective provided by literalism, before going on to argue, in subsequent sections, why a new perspective on word meaning, and the role of words in meaning construction, is required.

Literalism views sentence meaning as a consequence of adding or composing smaller units of meaning, together with the grammatical configurations in which they appear. In other words, accounting for linguistic meaning, from this perspective, assumes that the “ingredients” of language are words and rules, with rules serving to conjoin “atomic” meaning elements encoded by words. On this view, a descriptively adequate account of linguistic semantics should provide an observationally accurate account of these “elements of meaning” (associated with words or a single word), and the “rules of combination” (resulting in a sentence).

Identification of the elements of meaning is often referred to as **componential analysis**. This approach seeks to work out how to represent the meanings of words, or more precisely, what are termed **lexemes**—the meaning that is held to underlie a series of related forms, for example, *sing*, *sang*, *sung*, *singing*, and so forth, which are assumed to all have the same meaning, SING. The essential insight of this approach is that word meanings are made up of atomic elements or components. Typically, lexical items are thought of as being tagged with syntactic, morphological, and semantic features.

An early such componential-style analysis was that developed by Katz and colleagues (Katz and Fodor 1963; Katz and Postal 1964; Katz 1972). In this account, word meanings consist of semantic markers and distinguishers. Semantic markers comprise the information shared by words, while distinguishers constitute the idiosyncratic information specific to a given word meaning. For instance, based on Katz and Postal (1964), the polysemous senses for the word *bachelor* can be represented as in (2), where the semantic markers are given in parentheses and the semantic distinguishers are given in square brackets.

- (2) a. (human) (male) [who has never married]
 b. (human) (male) [young knight serving under the colours of another]

- c. (human) [recipient of the lowest academic degree]
- d. (non-human) (male) [young fur seal without a mate]

More recent and more sophisticated componential analyses of word meaning are provided by Anna Wierzbicka (e.g., 1996) in her **Natural Semantic Metalanguage** (NSM) account of word meaning, and Ray Jackendoff (1983, 1990) in his theory of **Conceptual Semantics**. Nevertheless, it is important to point out that neither Wierzbicka nor Jackendoff endorse all aspects of literalism. In particular, they do not take the view that compositional (i.e., sentence level) semantics patterns after reference, nor that sentence meaning should be truth evaluable (see the discussion below).

However, the hallmark of componential accounts, and the view of word meaning adopted under literalism, is that word meanings are assumed to be relatively fixed and stable. Put another way, the semantic primitives which make up a given word meaning can be identified independently of context.

Once identified, word meanings are integrated, by applying the rules of the grammar, in order to provide sentence meaning. Literalism, then, assumes that the contribution of language to meaning construction is essentially additive in nature, positing grammatical principles which ensure that the semantic units which result are unable to change or delete the meanings of the units which are conjoined to form a larger semantic unit or expression. This restriction serves to make a larger expression, for instance a sentence, **monotonic** with respect to its component parts, where the term “monotonic” has to do with the view that the component parts retain their original meanings in the larger expression (e.g., Cann 1993). Thus, the individual word meanings do not alter their meaning in the larger semantic units of which they form part.

Once composition has occurred, this gives rise to sentence meaning. Under literalism, sentence meaning, technically known as a **proposition**, is **truth evaluable**—although this issue is potentially problematic.² That is, a sentence—a well-formed grammatical string of words—is held to “carry” a meaning which patterns after reference: the conventional assignment of a worldly entity and state of affairs to the complex linguistic expression resulting from composition of the individual elements in forming a sentence. The meaning associated with the sentence constitutes the proposition, that is, the sentence meaning. Thus, in the following example sentence:

² A number of scholars working in the Pragmatic tradition (e.g., Bach 1997; Carston 2002; Recanati 2004) have observed that it is often (or usually) the case that the linguistic form uttered by an interlocutor underdetermines the sentence meaning. That is, utterances are often not propositional, but have to be completed by what has been termed **pragmatic intrusion**, such that inferential processes are required in order to render the utterance propositional and hence truth evaluable. Carston, for instance, refers to the notion that linguistic meaning underdetermines sentence meaning (i.e., the proposition expressed) as the **Underdeterminacy thesis**. For instance, while the following example from Carston (2002: 17): *On the top shelf*, relates to a specific location, as Carston notes, “[W]hat is meant by a speaker . . . is something sentence-shaped (propositional), presumably quite obvious in the context [for example, ‘the item you are looking for is on the top shelf’].”

(3) Brighton is 50 miles south of London

the proposition “carried” by the sentence can be evaluated as being true or false with respect to the state of affairs which holds in the world. In this case, the proposition expressed by (3): that Brighton is 50 miles from London, is true.

Thus far, we have been addressing the first half of literalism: the study of semantics. According to literalism, word meanings and the resulting sentence meaning, is context-independent. However, the full meaning of a sentence, what is referred to as speaker meaning, may also depend on context. This aspect of meaning falls under the purview of the sub-branch of linguistics known as pragmatics.

The distinction between sentence meaning and speaker meaning was introduced by the British philosopher Paul Grice (e.g., 1989). Grice distinguished between what a sentence means, its literal meaning, and what a sentence implicates, by virtue of the context in which it is deployed, and the speaker’s communicative intention in deploying it in the particular context of use. The latter sort of meaning is what Grice referred to as speaker meaning. According to literalism then, there is a principled distinction between semantics, which is concerned with literal or sentence meaning, and pragmatics, which is concerned with context-based speaker meaning: what is implicated.

To illustrate, let’s reconsider the sentence in (3). The literal meaning of this sentence relates to a state of affairs in the world referenced by the proposition expressed by this sentence. However, the proposition expressed is independent of any given context of use. To illustrate, now consider (3) as part of an exchange between two interlocutors in (4) who are driving to Brighton, are just north of London, and whose petrol gauge is hovering just above empty.

- (4) A: Do you think we can make it to Brighton without filling up?
B: Brighton is 50 miles south of London

According to literalism, the sentence expressed by B means what it does: Brighton is 50 miles south of London, which is truth evaluable independent of any given context because it can be assessed by virtue of a context-independent state of affairs: in the world, Brighton really is 50 miles south of London.

However, in the context associated with the exchange in (4), it means more than this. This is because the use of this sentence in this context implicates something in addition to the literal meaning expressed by the sentence. The implicature associated with the sentence uttered by B is that the travellers cannot reach Brighton unless they first obtain more petrol for their car. Thus, the speaker meaning is a consequence of interpreting the communicative intention of the speaker in deploying the sentence meaning in a given context. A somewhat simplified overview of the main elements of literalism are presented in Figure 1.1.

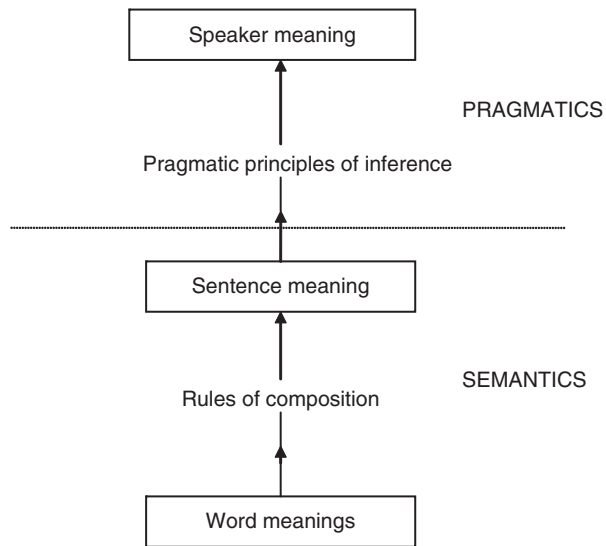


FIGURE 1.1. An overview of literalism

In sum, and from the perspective of literalism, word meanings involve relatively fixed and context-independent atoms of meaning. These atoms are concatenated, given the rules of the grammar, and then interpreted, by virtue of principles of language use. The context-independent atoms of meaning associated with words contribute to sentence meaning, and speaker meaning relates to the use to which sentences are put (including the context-independent word meanings which constitute them), which speaker meaning builds upon.

Problems with the received view

Literalism as an approach to meaning construction suffers from a fatal problem: the principled separation between context-independent (sentence) meaning and context-dependent (speaker) meaning. Put another way, the difficulty at the heart of literalism is the principled division of labour that it posits between semantics and pragmatics. In terms of the approach to word meaning adopted by literalism, words are assumed, apart from a number of notable exceptions such as **indexicals** (for instance *he*, or *here*), to have meanings tied to them which are context-independent. This follows as word meaning falls under the purview of semantics (rather than pragmatics).

However, a by now large number of scholars have argued that the principled separation of context-independent and context-dependent meaning

(the semantics/pragmatics distinction) is illusory.³ From this it follows that the position that word meanings are context-independent is potentially problematic. For instance, in the Pragmatics tradition, researchers have shown that the meaning of a given word, and hence the truth conditions of the sentence to which the word contributes, is typically (perhaps always) a function of context/background knowledge (see in particular Carston 2002; Searle e.g., 1983; Recanati 2004).

By way of illustration, consider the following examples of *open* based on those discussed by Searle (1983):

- (5) a. John opened the window
 b. John opened his mouth
 c. John opened the book
 d. John opened his briefcase
 e. John opened the curtains
 f. The carpenter opened the wall
 g. The surgeon opened the wound
 h. The sapper opened the dam

As Searle observes, in examples such as these the meaning of *open* is a function of what he refers to as the “background”, which is to say our knowledge of the sorts of ways in which entities and objects of different kinds are opened. Crucially, the different ways in which we can open things is a function of our **encyclopaedic knowledge**, which is to say knowing about and experience with the very different sorts of operations involved. For instance, opening a wound involves, for instance, the skilled use of a scalpel on flesh, to create an aperture of a certain size and shape for a particular purpose, such as to clean the wound and/or remove potentially damaged or diseased tissue. The opening of a wall involves different sorts of tools, typically carpentry tools of a particular kind, which are applied to a wall, made typically of wood, and resulting in an aperture of a certain size and shape for a very different sort of purpose: for instance to create or insert a doorway. Both of these operations differ from opening a mouth, which involves muscle gestures on a pre-existing aperture, or opening curtains, which doesn’t involve an aperture at all, both of which serve very different functions. Finally, opening a dam by a sapper involves knowledge relating to warfare—a sapper is a military explosives expert—and destroying the dam in question as part of a military action. Thus, understanding what *open* means in (5h) involves knowledge of a very different sort of event, agents, and purposes.

³ For a flavour of the range and nature of the problems that have been raised for a principled separation between context-independent and context-dependent dimensions of meaning, see, for example, the approaches to language and situated communication highlighted by the following: Carston 2002; Clark 1996; Coulson 2000; Croft 2000; Evans 2004a; Fauconnier 1987; Lakoff 1987; Langacker 1987; Recanati 2004; Sperber and Wilson 1995; Sweetser 1999; Tyler and Evans 2003).

In addition, in each of these examples the sort of encyclopaedic knowledge involved is a function of the utterance context in which the word is embedded. Thus, not only is the meaning of the word a function of quite distinct sorts of encyclopaedic knowledge, the sort of encyclopaedic knowledge to which the word provides access is a function of the context in which the word is embedded. That is, the linguistic context in part serves to narrow the sort of encyclopaedic knowledge to which *open* relates in each example. Thus, and as Searle observes, the semantic contribution that *open* makes to the truth conditions of sentences, such as these, varies, being a function of the sentential context in which it is embedded.

While the examples above relate to literal sentences, the context dependence of *open* is even more marked if we consider uses that are, intuitively, more figurative in nature. Consider the following indicative set of examples:

- (6) a. The discussant opened the conversation
 b. John opened a bank account
 c. John opened the meeting
 d. John opened a dialogue
 e. The Germans opened hostilities against the Allies in 1940
 f. The skies opened
 g. He opened his mind to a new way of thinking
 h. He finally opened up to her

The meaning of *open* in each of these examples relates to distinct sorts of actions, events, and situations. In the first example, opening a meeting requires a designated authority: a meeting “chair”, who, in declaring the meeting open, performs a specific speech act, thus facilitating the meeting process. In opening a dialogue, two (or more) interlocutors begin and continue a conversation that can take place face-to-face, electronically via email, on the telephone, or via the exchange of letters. To open such an exchange relates to the initiation of the exchange. To open a bank account involves completing certain formalities such as an interview with a bank official, financial checks, and the filling in of paperwork. In contrast, to open hostilities, as in the example in (6e), concerns the initial actions involved in warfare. Thus, each of these uses of *open* relates to very different forms of initiations, involving different sorts of events, procedures, and agents. In contrast, in the example in (6f), the usage of *open* relates to a sudden and heavy downpour of rain, while the last two examples relate to flexibility of thinking and emotional responses and/or being more expansive in terms of spoken, physical, or emotional interactions.

What examples such as those in (5) and (6) illustrate is the following. Firstly, a word such as *open* provides access to an impressively diverse array of encyclopaedic knowledge involving distinct scenarios, actions, events, and agents. As we have just seen, things that can be “opened” include an array of

different sorts of physical entities and abstract events—which is related to Searle’s notion of “background”.⁴ Understanding the examples in (5) and (6) involves complex and detailed knowledge about the sorts of scenarios that *open* relates to in each example and, thus, the specific way in which *open* applies in each case. After all, opening a mouth involves a very different form of opening than when a carpenter opens a wall, or when a sapper opens, and thus destroys, a dam. Hence, the meaning of *open* in each example is, in part, a function of tapping into the encyclopaedic knowledge, in order to determine the specific meaning of *open* in each example. Put another way, it is the scenario that *open* relates to that, in part, determines the nature of the meaning associated with *open* in each case.

Secondly, in each case it appears to be the sentential context, which is to say the other words in the sentence, which serve to direct the sort of encyclopaedic knowledge that *open* provides access to. That is, while *open* has a large body of knowledge, in the sense of a sophisticated range of scenarios and events that it can be applied to, what I will refer to as its **semantic potential**, the sentential context serves to guide and narrow the specific sorts of knowledge that a given instance of *open* actually relates to. In sum, the meaning of *open* appears to be a function of (i) (sentential) context which guides the (ii) encyclopaedic knowledge to which *open* relates in a given instance of use.

While the general problem in literalism is the strict separation between context-independent meaning (semantics), and context-dependent meaning (pragmatics), this gives rise to two problems for the resulting view of word meaning. Under literalism, word meaning falls under the purview of semantics. We saw in the previous section that under literalism word meanings are held to be: (i) stable and relatively circumscribed knowledge units, and (ii) context-independent. Hence, word meanings, which while susceptible to contextual interpretation (at least if meaning is understood in referential terms as in a possible world semantics), are held to constitute circumscribed knowledge units which are stored and can be deployed independently of other sorts of knowledge. Words meanings are thus separable from other kinds of knowledge such as the kind of representation(s) I have referred to as encyclopaedic knowledge. They are conceived as constituting fixed and relatively stable bundles of semantic elements, additionally tagged with syntactic and morphological features.

As we have just seen with our discussion of *open*, word meanings do appear to relate to and draw upon a potentially large body of knowledge, which following other scholars (e.g., Haiman 1980; Langacker 1987) I have been

⁴ While encyclopaedic knowledge, in the sense that I use it here, and as developed in cognitive linguistics (see the discussion in Chapter 3), is arguably related to Searle’s notion of “background” it is not quite the same. For Searle, background has to do with what we might think of as knowledge which constitutes entrenched, non-representational practice. What I take from Searle is the idea that word meaning is always contextualized with respect to knowledge which, in (large) part, determines the linguistic meaning.

referring to as encyclopaedic in nature. Moreover, the meaning of *open* only ever appears in given contexts of use, even when these are the minimal contexts of use deployed by the linguist: a numbered “linguistic example” set off and embedded in the running text of technical articles published in academic journals.⁵ In other words, word meaning emerges from a large semantic potential which is narrowed by the sentential (and extra-linguistic) context in which it is embedded. As such, word meaning appears to be guided by and a function of context: words, I suggest, do not mean independently of context. Thus, the fundamental problem with literalism is that it attempts to artificially divorce (word) meaning from (situated meaning in) context of use. More precisely, literalism lives in something of a fool’s paradise. It holds that language users retain an idealized, timeless meaning for *open* which they neatly keep apart from the situated meanings of *open* which arise from its use in examples such as in (5) and (6). The mistake that literalism makes, then, is in being reductionist and simplistic about meaning.

An additional challenge: figurative language

As we have just seen in our discussion of *open*, the protean nature of word meaning relates both to literal and figurative uses. A challenge for any theory of **lexical representation**—which is to say, the mental representations associated with words, consonant with the protean nature of word meaning discussed in this chapter—is to provide an account of literal and figurative language. Under literalism, these are treated as radically different sorts of language. It is often assumed, from this perspective, that figurative language involves the “defective” use of literal language, as argued, for instance, by Searle ([1979] 1993). On this view, the use of figurative language arises from the context-dependent interpretation of literal language, and thus involves principles of pragmatic inference being applied once the context-independent sentence meaning has been derived. Put another way, figurative language is a function of language use, and thus falls under the purview of pragmatics, rather than semantics proper.

The difficulty for what we might refer to as the **literalism perspective on figurative language**, is as follows. This perspective predicts that understanding a literal sentence should be faster than understanding a figurative expression: we must first understand what the sentence means before we can interpret what the speaker intends us to infer by using the sentence in a non-literal way. However, as has been shown, based on investigations of psycholinguistic processing, language users often appear to be equally as efficient in computing the meaning of figurative language utterances as they are non-figurative ones (Gibbs 1994; Glucksberg 2001, 2003; see also Giora 1997, 2003).

⁵ See similar arguments made by scholars including Clark (1983); Coulson (2000); Evans (2006); Fauconnier (1997); Langacker (1987); Sweetser (1999); Tyler and Evans (2003).

The challenge, then, that awaits an account of lexical representation and the role of words in meaning construction is to work out the difference, if any, between the role and function of literal and figurative word use in meaning-construction processes. To illustrate the nature of the challenge, let's consider the following example:

(7) John's boss is a pussycat

Presumably this utterance doesn't mean that John's boss is a pussycat, in the sense of a four-legged organism, with a tail and pointy ears that utters "miaow." Rather, the meanings associated with the phrases *John's boss* and *pussycat* have to be integrated with the **predicate nominative construction**, which ordinarily carries a class-inclusion meaning.⁶ Informally, this construction has the following syntax: "SUBJECT is an NP," and means, again informally: "The subject is a type of the entity specified." To illustrate, consider the following:

(8) John's boss is a pianist

The meaning that a language user would ordinarily derive, for an example such as this, would be that John's boss is included in the category of those who play the piano and thus constitutes a pianist, and that this situation persists through time. But, the same construction does not provide a class-inclusion reading for the previous example in (7). The challenge then, for our account of the variation in word meaning, is to be able to provide an explanation as to why (7) means something other than what it literally says, while (8) means what it does literally appear to say.

The nature of context

I suggested above that the fundamental problem with literalism is that it attempts to artificially divorce meaning from context of use. Before proceeding with an attempt to identify the ingredients of a theory of word meaning and meaning construction, we must first get an initial sense of the different sorts of context which serve to narrow the meaning of a word. Accordingly, we will begin to see that the notion of context is a complex and multifaceted phenomenon crucial for language use and language understanding. Accordingly, the notion of context is fundamental to the development of LCCM Theory that I begin to sketch in the next chapter, and develop in detail in the

⁶ The nominative predicative construction involves the copular or "linking" verb *be* which combines with a nominal, e.g., "a pianist." The nominal functions as the essential part of the clausal predicate: "is a pianist." Langacker (1991a) in his analysis of the nominative predicate construction argues that *be* encodes the "continuation through time of a stable situation characterized only as a stative relation" (*ibid.* 65).

rest of the book.⁷ Hence, the account of word meaning provided is diametrically opposed to that offered by literalism.

As the approach I take is usage-based, I use the term **utterance**, rather than sentence, in discussing word meaning. This reflects my assumption that it is only by taking account of language in use that we can hope to fully understand the nature of word meaning. It also follows from the position that sentences, as understood in linguistic theory, are artificial theoretical constructs, abstracted from actual **usage events**, which is to say, utterances. I will have more to say about the distinction between sentences and utterances in Chapter 4.

Utterance context

As we saw with the examples relating to *France* and *open* above, the utterance elements which occur in a given utterance contribute, in part, to determining the meaning of the word. That is, and as suggested above, the utterance provides a context which assists in narrowing the meaning of the word in question. To illustrate, consider the following examples:

- (9) a. On May 1st my grandfather expired
b. On May 1st my driving licence expired

The meaning of *expired* in each example is a function of the utterance in which it is embedded. In the first example, *expired* relates to an event involving death, while in the second, *expired* relates to expiry of the term for which an individual's right to drive on the public highway was sanctioned or "licensed."

Now consider another example involving a verb. This involves the following well-known context-dependent alternation associated with the verb *bake*:

- (10) a. Fred baked the potato
b. Fred baked the cake

While the example in (10a) relates to a change-of-state reading, the example in (10b) relates to a creation reading. That is, in (10b) the meaning of *bake* can be paraphrased by "made" or "created", while the meaning of *bake* in (10a) cannot be paraphrased in this way. The shift in meaning associated with *bake* appears to be a function of the object associated with *bake*: potato versus cake, and thus the specific consequence(s) that baking has for particular entities designated. While a potato is rendered edible by virtue of baking, as its interior becomes soft and it is thus easier to consume, an "uncooked cake" is not in fact normally thought of as a cake, but as a "potential cake." While the process of baking does not affect the existential status of a potato, but

⁷ By incorporating the notion of context into the theory, the approach I take is fundamentally concerned with language in use, and thus, as already observed, is usage-based in nature.

rather affects its state, a cake only in fact exists once it has been baked, as baking is one of the requisite stages involved in making a cake.

My final example of the role of utterance context in contributing to the meaning of a given word relates to what Schmid (2000) terms “shell nouns.” According to Schmid, “Shell nouns make up an open-ended functionally-defined class of abstract nouns that have, to varying degrees, the potential for being used as conceptual shells for complex, proposition-like, pieces of information” (*ibid.* 4). Common examples of shell nouns include: *case, chance, fact, idea, news, point, problem, position, reason, report, situation, thing*. The significance of shell nouns for the present discussion is that the semantic value of the shell noun is normally determined by the utterance context. Moreover, the shell noun itself serves to characterize and encapsulate the idea whose meaning it simultaneously takes on. Thus, the meaning associated with the shell noun is, paradoxically, both a function of and a contributor to the utterance context in which it is embedded. To illustrate, consider the following example drawn from Schmid (2000):

- (11) **The Government’s aim is** to make GPs more financially accountable, in charge of their own budgets, as well as to extend the choice of the patient

In the example in (11) the shell noun is in bold. The idea the shell noun relates to is underlined. The shell noun, the noun phrase in which it occurs, and the idea it relates to, which here is mediated by the copula *is*, are collectively termed the “shell-content-complex.”

According to Schmid, the meaning of the shell-content-complex in examples such as this are a function of the specific combination of the shell noun and the idea it relates to. That is, the shell-like function of the shell noun is not an inalienable property of the noun itself, but rather derives from the way it is used. In this example, the speaker presents a particular idea (“to make GPs more financially accountable, in charge of their own budgets, as well as to extend the choice of the patient”) as an “aim”. This provides a particular characterization for the idea. Moreover, by providing this characterization, the shell noun also serves to **encapsulate** the various components and complex ideas contained in the idea as a single, relatively stable, albeit temporary, concept. It does so by casting “this complex piece of information into one single noun phrase” (*ibid.* 7). Evidence for this unity comes from the next sentence presented in (12):

- (12) **The Government’s aim is** to make GPs more financially accountable, in charge of their own budgets, as well as to extend the choice of the patient. **Under this new scheme**, family doctors are required to produce annual reports for their patients . . .

Here we see that once the complex idea has been encapsulated, it can be glossed with a different characterization as signalled by the shell noun phrase *this new scheme*, marked in (12) in bold. In essence, the content associated with shell nouns comes from the ideas, that is, the utterance context, they relate to. Yet, the ideas receive their characterization, and even their construal as a single unified idea, from their participation in a shell-content-complex.

Manner of utterance

The manner of the utterance can provide a context which serves, in part, to determine the meaning of a particular word. For instance, whether a particular word receives stress or emphasis of some kind can contribute to the meaning of the word. Consider the following by way of illustration.

- (13) a. Look at that *blàckbird*
b. Look at that black *bìrd*

The compound *blackbird* receives primary stress on the adjective *black*. In contrast, a bird that happens to be black, but is not a blackbird, receives primary stress on *bird*, as in the second example. Here, stress serves as a type of **contextualization cue**, serving to determine, in part, the semantic contribution of *black* to the utterance.⁸

Extra-linguistic context

The time, venue, or medium (e.g., spoken or written), or the genre of the medium (e.g., newspaper report versus spoken lecture) of an utterance can contribute to the meaning of given words, and thus provide a context. In this case, the context is extra-linguistic as it constitutes the “location,” broadly construed, in which the utterance occurs. To illustrate, consider the following utterance:

- (14) “I watched the young lady approach the bar.”

The meaning of *bar* in this utterance is determined, in part, by the kind of venue to which the utterance relates. For instance, if uttered in a court of law, the notion of bar would refer to the raised platform at which the judge sits. If said in a public house, it would refer to the area at which alcohol is ordered and purchased.

Consider another example of extra-linguistic context, this time employing the word *safe* in the context of a child playing on the beach. The examples are based on Sweetser (1999):

⁸ The term “contextualization cue” was coined by Gumperz (1982). In borrowing the term here, I am using it in a slightly different way from that of Gumperz who applied it in the context of his work on code-switching.

- (15) a. The child is safe
b. The beach is safe
c. The shovel is safe

In this context, the meaning of (15a) is that the child will not come to any harm. However, given the extra-linguistic context, (15b) does not mean that the beach will not come to harm. Instead, it means that the beach is an environment in which the risk of the child coming to harm is minimized. Similarly, (15c) does not mean that the shovel will not come to harm, but that it will not cause harm to the child using it to dig in the sand. These examples illustrate that there is no single fixed property that *safe* assigns to the words *child*, *beach*, and *shovel*. In order to understand the utterances we must interpret them, in part, with respect to a specific extra-linguistic context, a scenario, which holds. In this scenario, there is a child on a beach, employing a spade to dig in the sand. In order to successfully interpret these utterances we must also draw upon our encyclopaedic knowledge relating to children, beaches, and shovels, and the potential harm that shovels can cause if mis-used, for instance.

Encyclopaedic knowledge

Earlier in this chapter I noted that the utterance context serves to narrow that part of the encyclopaedic knowledge to which a word potentially provides access. What I have in mind by encyclopaedic knowledge has been referred to by a range of terms in the linguistics and cognitive science literature. These include the following: **background knowledge**, **common-sense knowledge**, **sociocultural knowledge**, and **real-world knowledge**. By encyclopaedic knowledge I have in mind the highly detailed, extensive, and structured knowledge we as humans appear to have access to in order to categorize the situations, events, and entities we encounter in our everyday lives and in the world, and the knowledge we draw upon in order to perform a range of other higher cognitive operations including conceptualization, inference, reason, choice, and the knowledge which language appears to rely upon. This kind of knowledge is primarily non-linguistic, or conceptual in nature, and appears to constitute a vast structured body of relational information which psychologists sometimes refer to as **frames** (e.g., Barsalou 1992, 1999; Barsalou *et al.* 1993). Although I will revise the notion of encyclopaedic knowledge as the book proceeds, the notion of encyclopaedic knowledge will be central to the theory of word meaning and compositional semantics developed in this book.

While speakers and hearers call upon encyclopaedic knowledge in using language, this knowledge thereby serves as a kind of context against which words receive and achieve meaning. For instance, the meaning of *France* in each of the examples in (1) above, draws upon a different body of knowledge. In the example in (1a) we draw upon our knowledge of the geographical

landmass associated with France, while in (1b) we draw upon our knowledge of France as a political entity, a nation state.

Interactional norms as context

A particular sort of encyclopaedic knowledge which provides a salient form of context relates to interactional or behavioural norms. This notion is sometimes referred to as a **cultural script** or a **cultural routine**, or simply as a **script**, particularly as developed in the computational literature associated with the work of Schank and Abelson (1977). For instance, the following restaurant script is adapted from Schank and Kass (1988: 190):

- (16)
1. Agent goes to restaurant
 2. Agent is seated
 3. Agent orders meal from waiter
 4. Waiter brings meal to agent
 5. Agent eats meal
 6. Agent gives money to restaurant
 7. Agent leaves restaurant

A cultural script such as this constitutes an interactional norm which provides the context against which words derive a particular meaning.⁹ For instance, the meaning of the word *restaurant* is, in part, informed by knowledge relating to the script captured in (16).

Interactional goals as context

Another form of context which serves, in part, to determine the meaning of a given word constitutes the interactional goals of the interlocutors. According to Clark (1996), linguistic communication is a form of joint action, in which interlocutors negotiate, establish, and attempt to achieve interactional goals.¹⁰ These goals, which can be explicitly signalled, or arise due to the extra-linguistic context or some aspect of encyclopaedic knowledge such as a cultural script, serve as the context against which the meaning of lexical items can be, in part, determined.

For instance, consider the following service encounter in a fast-food restaurant:

- (17) Customer: [Waits at serving counter]
 Server: [Appears after a short delay after fetching another customer's order] Hi!
 Customer: A double whopper meal please.

⁹ Fillmore's (e.g., 1982) notion of a **semantic frame**, discussed in the next chapter, provides a related construct to that of script.

¹⁰ These issues are discussed in more detail in Chapter 11.

Clark observes that in service encounters such as this one, interlocutors, through joint action, negotiate the accomplishment of communicative goals. In this example the server indicates their availability to take the customer's order by greeting the customer. That is, the customer takes the greeting, the utterance *Hi!*, as signalling an offer to receive the customer's order. Clearly, in order for *Hi!* to have this meaning, the server and customer must share an understanding as to the nature of the interaction and its objectives: the server is there to receive a food order (which is achieved by the greeting) and the customer wishes to place an order.

Discourse topic as context

The final kind of context I will mention relates to the notion of discourse topic. In general terms, interlocutors often appear to derive word meaning from what they take the discourse topic to be. For instance, consider the following utterance:

(18) That hike is killing me

In the context of a conversation on a recent central bank base-rate increase, this mention of *hike* might relate to the financial pain involved in an increase in mortgage repayments. However, in the context of a discussion of a recent cross-country walk, the pain might be more physical in nature.

The point, then, of this discussion has been the following. Context is a complex and multifaceted phenomenon. Moreover, the meaning associated with a word in any given instance of use is, in part, a function of the particular sort of context, linguistic or otherwise, in which it is embedded, and of which it forms a part. Put another way, word meaning is protean, its semantic contribution sensitive to and dependent on the context which it, in part, gives rise to.

This bears on the discussion of the nature of word meaning under literalism in the following way: the precise semantic contribution of each word appears to be a function of the context in which it is embedded. Put another way, words do not have discrete, timeless (i.e., context-independent) meanings, contra the assumption under literalism.

A possible solution? Sense Enumerative Lexicons

If the fixed, componential view of word meaning offered by literalism fails, what then? A possible solution to the apparent variation in word meaning exhibited in language use might be to posit a vast number of distinct senses. For instance, rather than assuming that the range of meanings associated with, say, *open* in the examples above are somehow due to context and/or encyclopaedic knowledge, we might assume that *open* has exactly the same

number of distinct meanings, technically known as **senses**, as the number of different sentences in which it appears, and that each of these are stored in long-term semantic memory.

Pustejovsky (1995) in his pioneering work on lexical semantics refers to approaches which posit a large number of distinct senses for given lexical items as **Sense Enumerative Lexicons** (or SELs for short). However, as Pustejovsky observes, even such accounts cannot predict the creative use of words in novel contexts. That is, even lexicons which assume a high degree of granularity fail on the score of descriptive adequacy in the face of the linguistic facts. Thus, word meaning in language use cannot be predicted from knowledge of the conventional range of uses to which words are put, even when one assumes a highly granular lexicon: one that posits a large number of distinct senses. This follows as the number of distinct word senses required, even for a single word, would need to be infinite, a position that, given memory constraints, is untenable, even allowing for the significant capabilities that language users have in terms of semantic memory.

To illustrate the foregoing, consider the lexical item *fast*, discussed by Pustejovsky. It is commonly assumed that this word has a number of **conventional senses**—mentally stored semantic units—associated with it. These include the following:

- (19) a fast car [fast₁: to move quickly]
- (20) a fast typist [fast₂: to perform some act quickly]
- (21) a fast decision [fast₃: to require little time for completion]

However, the definitions provided do not fully capture the “type”-semantics that these examples of *fast* are instances of. For instance, *fast* illustrated in (19) relates to an entity capable of moving quickly, whilst the type illustrated in (20) relates to entities capable of performing actions quickly, and so on. That is, each putatively conventional sense of *fast* has associated with it selectional restrictions, what I will refer to as **selectional tendencies**. The “to move quickly” sense, for instance, selects for members of the class of movable entities.

However, now consider the following example:

- (22) a fast driver

This usage of *fast* concerns not the actions of the driver. That is, it is not the actions of the driver which are performed quickly. Nor would this utterance normally refer to such actions, even if they were performed quickly. Rather, this expression refers to the speed at which cars controlled by the driver in question ordinarily proceed relative to some norm, such as the established speed limit for a particular road. In other words, this is an instance of fast₁, rather than fast₂. Yet, *fast*, in this example, relates to the vehicle driven by the

driver, rather than, strictly, the driver. Thus, the combination of *fast*₁, with *driver*, produces a novel reading in which *fast* might be paraphrased as “to cause to move quickly”.

Now consider the following example:

(23) the fast lane (of the motorway)

Presumably this usage of *fast* also relates to *fast*₁. Yet, *the fast lane* is a venue for rapid locomotion rather than an entity capable of rapid locomotion. In other words, both the uses of *fast* in (22) and (23) while seemingly related to the meaning of *fast* in (21) have different semantic selectional tendencies, and somewhat novel meanings. We could posit that both (22) and (23) constitute distinct senses. However, we can continue finding novel uses of *fast*, for which we could produce a virtually infinite listing. Indeed, the same argument applies to sense 2 and 3 of *fast*.

In addition, a particular novel use can appear to feature nuances of different senses:

(24) We need a fast garage for our car, as we leave the day after tomorrow

As Pustejovsky (1995) notes, this use of *fast* appears to be a “blend” of both *fast*₂ and *fast*₃: a garage which carries out repairs quickly and takes little time to do so.

What this discussion of *fast* reveals, then, is that all the examples we have considered, and might wish to consider, upon close analysis predicate in a slightly different way. In other words, each unique instance has a distinct utterance context, and is associated with a slightly different semantic value. Thus, we can conclude from this that, in principle, every instance of use of a word such as *fast* has a different meaning. To take a “Sense Enumerative” approach to word meaning would be to sanction an infinite proliferation of word senses stored in memory by language users. Such a position is psychologically untenable.

Words as contextual expressions

The observation with which this book proceeds, then, is that words are never meaningful independent of the utterance in which they are embedded, and the encyclopaedic knowledge and extra-linguistic context which guide how words embedded in an utterance should be interpreted. Indeed, evidence from the perspectives of social psychology, cognitive psychology, interactional sociolinguistics, cognitive linguistics, corpus linguistics, and computational linguistics reveals that the view that words constitute fixed, context-independent structures, and that meaning construction is appropriately modelled

in terms of the straightforward approach to compositionality sketched above is untenable.

As observed by a large number of scholars, the meanings associated with words are flexible, open-ended, and highly sensitive to utterance context. Such scholars include, but are by no means limited to Allwood (2003), Carston (2002), Clark (1983, 1996), Coulson (2000), Croft (1993, 2000), Croft and Cruse (2004), Cruse (2002), Evans (2004*a*), Fauconnier (1997), Fauconnier and Turner (2002), Goffman (1981), Gumperz (1982), Harder (2009), Herskovits (1986), Lakoff (1987), Langacker (1987), Pustejovsky (1995), Sperber and Wilson (1995), Sweetser (1999), Dancygier and Sweetser (2005), and Tyler and Evans (2003). Indeed, as Croft (1993) observes, meaning construction appears to proceed by virtue of the meaning associated with a given word being interpreted once the meaning of the entire utterance has been established. That is, individual word meaning is determined by the encyclopaedic knowledge to which words provide access, as guided by context, rather than utterance meaning being a consequence of concatenating context-independent word meanings. As such I argue that words are **contextual expressions**. From this perspective, as utterance meaning is the result of assigning meaning to words in both linguistic and non-linguistic context, the end product is due to all three factors. Hence, meaning cannot be assigned unambiguously to words alone. Rather the semantic contribution associated with individual words emerges from the *mélange*: words *are* contextual expressions. From a usage-based perspective on language (e.g., Croft 2000; Langacker 2000; see Evans and Green 2006 for a review), this state of affairs is entirely natural, as I shall suggest in later chapters.

Accordingly, in this book I argue against the received view that words “carry” meaning. In point of fact, I will be arguing that meaning is not a property of words, or even language, *per se*. Rather my contention is that meaning arises as a function of the way in which words (and language) are deployed by language users in socioculturally, temporally, and physically contextualized **communicative events**, which is to say utterances, due to a complex battery of linguistic and non-linguistic processes, in service of the expression of situated communicative intentions.

Of course, to say that words do not “carry” meaning does not entail the claim that the semantic structure associated with linguistic units such as words is wholly indeterminate. This position, which may be associated with some usage-based approaches to language (e.g., Thompson 2002; Croft and Cruse 2004; see Harder 2009 for a description of the risk of ending up in the extreme position he calls “usage fundamentalism”), is hard to maintain. After all, as pointed out by Sweetser (1999), the very distinct readings typically derived from utterances of the following kind:

- (25) a. John ran up the stairs
b. John ran down the stairs

have to do with the fact that lexical forms are associated with relatively well-established—in the sense of conventionalized—semantic representations. For instance, the fact that (25a) means something quite different from (25b) is a consequence of switching the particle *up* for *down*. As we shall begin to see in the next chapter, my claim is not that words do not have stable semantic representations associated with them. I argue that they do, and refer to these as **lexical concepts**. Rather, my claim is that these lexical concepts provide access to encyclopaedic knowledge—a semantic potential—which is constrained and determined by context. Thus, the semantic structure (lexical concept) that a word is conventionally associated with does not in fact equate with the word's meaning. Word meaning, from this perspective, is always a function of a situated interpretation: the context in which any word is embedded and to which it contributes.

A further problem: compositionality

In the foregoing we have considered the nature of word meaning. I suggested that the problem to be accounted for, the inherent variation of word meaning in language use, is, in part, a function of words providing access to encyclopaedic knowledge. This in turn is narrowed by context, effectively delimiting which part of the encyclopaedic knowledge—the semantic potential—available to any given word is activated in any given utterance.

Yet, providing such an account is not enough if we are to fully get to grips with the contribution of words to meaning construction. To do so, we must, in addition, be able to account for how utterance (i.e., sentence) meaning arises. Utterance meaning involves several, often many, linguistic units, each of which individually exhibits great variability (Goldberg 2006; see also Kay and Michaelis forthcoming). That is, one must also be able to account for the integration of lexical and constructional meanings: we require an account of semantic compositionality, one that is coherent with the observable facts of language, and, of course, one which is cognitively plausible.

One of the most sobering realizations for any cognitive scientist attempting to grapple with the role of language in meaning construction is that despite the apparent ease with which we construct and interpret utterances in our everyday lives, the nature of semantic composition is a deceptively complex process. Moreover, the details of this process are far from being fully understood. For instance, the way in which the meaning of even a “simple” sentence is constructed is incredibly complex.

To illustrate, consider the example of: *The cat jumped over the wall*, discussed by Tyler and Evans (2003). This utterance describes a jump undertaken by a cat. Figure 1.2 presents some diagrams which present possible trajectories of the jump.

While there are at least four possible trajectories associated with this utterance, the canonical interpretation is that the cat begins the jump on one side of




FIGURE 1.2. Possible trajectories for: *The cat jumped over the wall*

the wall, moves through an arc-like trajectory, and lands on the other side. Figure 1.2(d) best captures this interpretation. The issue to be accounted for is why it is that the reading typically derived relates to the trajectory diagrammed in 1.2(d) rather than one of the others. That is, what is it that excludes the trajectories represented in Figures 1.2(a–c)? After all, the utterance contains a number of words that have a range of interpretations. The behaviour described by *jump* has the potential to involve a variety of trajectory shapes. For instance, jumping from the ground to the table involves the trajectory represented in Figure 1.2(a). Jumping on a trampoline relates to the trajectory in 1.2(b). Bungee jumping involves the trajectory in 1.2(c). Finally, jumping over a puddle, hurdle, wall, etc., involves an arc-like trajectory as in 1.2(d). If the lexical item *jump* does not, in itself, specify an arc-like trajectory, but is vague with respect to its shape, then perhaps the preposition *over* is responsible.

Yet, *over* can also have several possible interpretations. It might be associated with an “across” interpretation: when we walk *over* a bridge (a horizontal trajectory). It can be associated with an “above” interpretation, as when an entity such as a hummingbird is *over* a flower (higher than but in close proximity to). Equally, *over* can have an “above” interpretation, as when a plane flies *over* a city: much higher and lacking close proximity. The point is that a word such as *over* can be used when different kinds or amounts of space are involved, and with a number of different trajectories/paths of motion. Hence, the received view that words are associated with fixed meanings, and that utterance meaning comes from concatenating the meanings of the individual words combined in a given utterance, underestimates the complexity involved in combining words, and the principles involved in their combination. An important aspect of the theory to be developed in this book relates to semantic composition, which is the subject of Part III.

Research issues to be addressed

The issues highlighted in this chapter relate to two issues central to my concerns in this book: the role of words in meaning construction, and the

nature of semantic composition. My first objective, and the subject of Part II of this book, is to provide an account of lexical representation. As already noted above, I advance the perhaps controversial claim that words do not in fact have meaning, although this position is not without precedent, particularly in the psychology literature (e.g., Barsalou *et al.* 1993; Murphy 1991). On my account, meaning is a function of an utterance, rather than a given lexical representation associated with a word, or other symbolic (i.e., linguistic) unit. I make the case for words, and symbolic units in general, being associated with the construct of the lexical concept, a unit of semantic structure. A lexical concept is a conceptual representation specialized for being encoded in and externalized by language. This idea is developed in more detail later in the book, beginning with discussion in the next chapter. Additionally, an account of lexical representation would be incomplete without considering the level of conceptual structure to which lexical concepts provide access. This level is populated by what I will refer to as **cognitive models**, for reasons that will become apparent in later chapters.

Having developed an account of lexical representation, my second concern is to provide an account of the meaning-construction processes which make use of the semantic and conceptual levels of representation in service of situated utterance meaning. This issue, which I refer to as **semantic, compositionality**, is the subject of Part III of the book. This involves an account of how lexical concepts are integrated in specific utterances: linguistically mediated usage events. The chapters in Part III address two key aspects of this process respectively, namely, the mechanisms of **lexical concept selection**, and **fusion**. Part IV of the book applies the theory of meaning construction developed to figurative language, arguing for a dynamic usage-based approach to figurative language understanding. One of the main claims to arise here is the position that language use often identified as constituting metaphor and metonymy arises from regular meaning-construction processes, which are, in principle, no different from those that give rise to non-figurative language. Thus, the present approach argues that figurative meaning derives from a meaning-construction process which marshals conventional linguistic resources (lexical concepts) together with the non-linguistic conceptual resources to which lexical concepts afford access. Thus, the position to be developed argues that there is continuity between literal and figurative language understanding. The treatment presented complements Conceptual Metaphor Theory (e.g., Lakoff and Johnson 1980, 1999), as I shall argue.

While on the face of it a new theory, LCCM Theory is, in fact, grounded in recent advances in the theoretical movement known as **cognitive linguistics**. As we shall see in Chapter 3, in certain respects it is not a new theory at all, but rather a synthesis of several extant approaches and theories that populate cognitive linguistics. However, the synthesis itself is genuinely novel, especially in so far as it serves to integrate cognitive linguistics approaches to grammatical

organization, lexical semantics, semantic composition, and figurative language. In so doing, it attempts to unify the complementary and sometimes competing theories and approaches that abound in cognitive linguistics. Moreover, while the role of language in semantic composition is crucial to cognitive grammarians (e.g., Langacker 1987, 2008; Goldberg 1995, 2006), it has tended to be ignored (or at least downplayed) in contemporary accounts of meaning construction in cognitive linguistics (e.g., Fauconnier and Turner 2002). LCCM Theory also serves to restore, or at least redress, the centrality of language to semantic compositionality, whilst also recognizing the importance of non-linguistic processes in meaning construction, pointed to by Fauconnier and Turner, and indeed others, not least in the work of George Lakoff (see, for example, Lakoff 1993, 1996, 2006; Lakoff and Johnson 1980, 1999).

As we shall see, LCCM Theory takes its name from the two central constructs upon which it is built, the lexical concept and the cognitive model. The purpose of the next three chapters then, is to begin to sketch an account of LCCM Theory. The rest of the book will work out the details.

Summary

This chapter has argued that the received view of meaning in linguistics, what I refer to as literalism, is flawed in a number of respects. The distinction it posits between sentence meaning and speaker meaning makes a principled distinction between context-independent meaning (semantics) and context-dependent meaning (pragmatics). The consequence of this for word meaning is that word meanings are assumed to be stable and relatively delimited “atoms of meaning,” which are context-independent. I have argued, on the contrary, that word meaning is inherently variable in language use. This is a function of both encyclopaedic knowledge and context of use. I have suggested that word meaning provides access to a sophisticated and structured body of non-linguistic encyclopaedic knowledge. This constitutes a word’s semantic potential. The precise part of this semantic potential which is relevant in any given utterance is a function of context, which serves to narrow or constrain the semantic potential. Thus, word meaning is always, in part, a function of and determined by context. I have also argued that the notion of context is a complex and multifaceted phenomenon which includes linguistic as well as non-linguistic aspects of the communicative event. In addition I have argued that a Sense Enumerative Lexicon approach to word meaning is unable to capture the rampant variation in meaning exhibited by words in language use. I have also pointed to the problem for any theory of compositionality that arises by acknowledging such variation in word meaning. This follows as the meaning of any utterance is a function, in part, of the word meanings which comprise it, and yet, each of these word meanings varies on each occasion of use.