

**The Semantics and Pragmatics of Polysemy:
A Relevance-Theoretic Account**

Ingrid Lossius Falkum

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy

University College London

January 2011

I, Ingrid Lossius Falkum, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Sign.

Abstract

This thesis investigates the phenomenon of *polysemy*: a single lexical form with two or multiple related senses (e.g. catch the *rabbit*/order the *rabbit*; lose a wallet/lose a relative; a *handsome* man/a *handsome* gift).

I develop a pragmatic account of polysemy within the framework of Sperber and Wilson's relevance theory, where new senses for a word are constructed during on-line comprehension by means of a single process of *ad hoc* concept construction, which adjusts the meanings of individual words in different directions.

While polysemy is largely unproblematic from the perspective of communication, it poses a range of theoretical and descriptive problems. This is sometimes termed the *polysemy paradox*. A widely held view in lexical semantics is that word meanings must consist of complex representations in order to capture the sense relations involved in polysemy. Contrary to this view, I argue that a conceptual atomist approach, which treats word meanings as unstructured atoms and thereby avoids the range of problems associated with decompositional theories of word meaning, may be at least as able to account for polysemy when paired with an adequate pragmatic theory.

My proposed solution to the polysemy paradox is to treat polysemy as a fundamentally communicative phenomenon, which arises as a result of encoded lexical concepts being massively underdetermining of speaker-intended concepts, and is grounded in our pragmatic inferential ability. According to this approach, the role of the linguistic system in giving rise to polysemy is to provide a minimal input, or clue, which the pragmatic system uses as evidence to yield hypotheses about occasion-specific, speaker-intended meanings. I further show how this pragmatic approach can account for cases of 'systematic polysemy', usually seen as prime candidates for an analysis in terms of lexical rule application. Finally, I develop an account of metonymy within the overall framework of relevance-theory.

Contents

Acknowledgements.....	7
<i>Chapter 1</i> Introduction.....	9
1.1 What is polysemy?	9
1.2 Background.....	10
1.3 The polysemy paradox	13
1.4 The problem of definition.....	16
1.4.1 Polysemy and homonymy.....	16
1.4.2 Ambiguity and vagueness.....	19
1.4.3 Polysemy and lexical pragmatics.....	23
1.5 Outline of the thesis.....	26
<i>Chapter 2</i> Lexical semantics and polysemy: the problem of representation	28
2.1 Introduction.....	28
2.2 Lexical decomposition approaches.....	29
2.2.1 The classical theory	30
2.2.2 Prototypes, cognitive semantics and polysemy	36
2.2.3 The generative lexicon	42
2.2.4 Neo-classical theories.....	47
2.3 Conceptual atomism.....	56
2.3.1 Radical nativism?	58
2.3.2 Conceptual atomism and polysemy.....	59
2.4 Psycholinguistic perspectives on polysemy.....	62
2.4.1 Core meanings or separate entries?	63
2.5 Conclusion	72
<i>Chapter 3</i> Language, theory of mind, and inferential communication: solving the problem of polysemy motivation?	75
3.1 Introduction.....	75
3.2 The code model of communication	77
3.2.1 The generative lexicon	79
3.3 Linguistic underdeterminacy	82

3.3.1	The (in)effability of thought	86
3.4	The inferential model of communication	91
3.4.1	Relevance theory	91
3.4.2	Polysemy and utterance comprehension	94
3.5	An evolutionary story about theory of mind and language	98
3.6	Implications of the pragmatic approach.....	106
<i>Chapter 4</i>	Relevance-theoretic lexical pragmatics: the <i>how</i> of polysemy	110
4.1	Introduction.....	110
4.2	The pragmatic adjustment of lexical meaning.....	112
4.2.1	A single process of <i>ad hoc</i> concept formation	116
4.2.2	Implications of the relevance-theoretic account.....	124
4.3	A lexical-pragmatic account of polysemy.....	127
4.3.1	A re-analysis of the Pustejovsky (1995a) cases.....	128
4.3.2	Polysemy, sense relations and lexical semantic change	132
4.3.3	Prepositional polysemy.....	137
4.4	Conclusion	146
<i>Chapter 5</i>	Systematic polysemy and the count-mass distinction.....	148
5.1	Introduction.....	148
5.2	Computational semantic accounts: The case of ‘grinding’	152
5.3	A relevance-theoretic account.....	160
5.3.1	The count-mass distinction and systematic polysemy.....	161
5.3.2	A (mainly) pragmatic account of systematic polysemy	168
5.3.3	The <i>book/window</i> cases.....	182
5.4	Conclusion	198
<i>Chapter 6</i>	Metonymy	200
6.1	Introduction.....	200
6.2	Previous approaches to metonymy	205
6.2.1	Metonymy as pragmatic function application	205
6.2.2	Cognitive linguistic approaches to metonymy.....	211
6.3	A relevance-theoretic approach.....	223
6.3.1	Lexical pragmatic processes: associative or inferential?.....	223
6.3.2	Papafragou (1996)	226

6.3.3	Metonymy as <i>ad hoc</i> concept construction?.....	236
6.3.4	Metonymy as ‘naming’ revisited	252
<i>Chapter 7</i>	Conclusion: a pragmatic account of polysemy.....	260
References	269

Acknowledgements

I am profoundly grateful to my supervisor, Robyn Carston, whose combination of deep insight, unfailing support, encouragement and patience has been greatly inspiring to me. I feel truly privileged to have had her as my supervisor.

I also owe a great intellectual debt to Deirdre Wilson, who introduced me to her work with Dan Sperber on relevance theory during my MA studies at UCL. I thank her for formative discussions and for always being supportive of my work.

At UCL I have had the privilege of working with highly competent linguists and philosophers, including Alison Hall, Kate Scott, Thiago Galery, Richard Breheny, Akis Kechagias, Nathan Klinedinst, Jonny McIntosh, Anna Pollard, Nausicaa Pouscoulous, Matthew Reeve, Paula Rubio-Fernández, Hanna Stöver, Hiroyuki Uchida, Catherine Wearing and Beata Zacharska. I have very much enjoyed being part of this stimulating environment. I also thank James Webb at UCL Psychology for discussions and for updating me on the psycholinguistic research on polysemy.

I would like to express my sincere thanks to the Research Council of Norway for funding the first three years of my PhD studies at UCL, and to UNIFOR for the financial support I received from *Programredaktør Andor Birkeland og hustru Halinas legat* during the final year of writing.

I am also grateful to everyone at the Centre for the Study of Mind in Nature at the University of Oslo, for taking me on as an affiliate during my PhD studies, and welcoming me into the environment of the centre during my research stays in Oslo. Thanks to Georg Kjøll, Trine Antonsen, Heine Alexander Holmen, Nick Allott, Endre Begby, Herman Cappelen, Einar Duenger Bøhn, Timothy Chan, Julian Fink, Christel Fricke, Jan Terje Faarlund, Olav Gjelsvik, Eline Busch Gundersen, Carsten Hansen, Torfinn Huvenes, Anders Nes, Mathias Sagdahl, Lalaine Siruno, Rachel Sterken, Anders Strand, Bjørn Ramberg and Kari Refsdal for stimulating discussions, great workshops and divine dinners.

I owe a great thanks to Ernie Lepore at Rutgers University in New Jersey, USA, who invited me to be his co-lecturer on the course 'On Words' at the CSMN in Oslo, and for inviting me to come and present my work at Rutgers in 2009. These events

made a substantial contribution to my PhD work. Thanks also to Kjell Johan Sæbø at the University of Oslo for useful comments on several draft chapters, and to Thorstein Fretheim at the Norwegian University of Science and Technology for his support and interest in my pragmatics studies at UCL.

Thanks to my beloved family; my father Erik, who is a constant source of inspiration to me, and whose coaching skills proved indispensable in the final stages of the writing up of this thesis; my mother Kristin and my stepfather Kjell for their presence and enduring support, and for always being interested in my work despite coming from very different academic backgrounds; my brothers Jon and Øyvind for making me laugh; my sister-in-law Irene and niece Ronja for being so lovely; my family-in-law the Engen Marsdals for their encouragement and warm presence. A profound thanks also to Katrine Ziesler for being such a good friend and for helping me in proofreading several chapters.

A great thanks goes to my friends and PhD colleagues, Katerini Storeng, Kjerstin Aukrust, Mia Eikeset, Guro Fløgstad, Frida Gullestad and Kathrine Asla Østby for great conversations and support during a both wonderful and challenging period of our lives. In particular, Katerini's company made long days of hard work at the British Library seem like a breeze. Thanks also to my other good friends for their support and welcome distractions from thesis writing, in particular Elisabeth Lackner, Lisa Vesterås and Beate Reme.

Finally, I thank my husband, Magnus, for his constant love, support and companionship, for always being available for discussions of the intricacies of polysemy with me, and for tolerating the company of this thesis on our honeymoon in Malaysia.

Chapter 1

INTRODUCTION

1.1 What is polysemy?

That a single word form can be associated with several different meanings is a well-known fact about language. Take the word *run*. Its meaning in the verb phrase *run a half marathon* is clearly different from the one it has in *run some water*, or, for that matter, in *run on gasoline*, *run on empty*, *run a shop*, *run late*, *run away from responsibilities*, *run in the family*, *run for President*, and so on. This phenomenon is described as *polysemy*, and it proliferates in natural languages. This is confirmed by the range of different senses (and/or uses) that any dictionary will have listed under a considerable number of its entries.¹ On top of this comes the range of non-established senses that lexical items may be used to express on different occasions of use, which are contextually derived on the spot.

Linguists, philosophers of language and psychologists have long been interested in the polysemy phenomenon due to the challenging issues it raises for theories of semantic representation, semantic compositionality, language processing and communication. Traditional approaches tend to regard polysemy as a matter of different senses being listed under a single lexical entry, with the comprehension of a polysemous word involving selection of the contextually appropriate sense from among the list of senses (so-called *sense enumeration lexicons*). Another traditional line of approach regards polysemy as being represented in terms of a single, maximally general meaning, from which the contextually appropriate senses are derived (so-called *core meaning approaches*). In more modern approaches, it is generally acknowledged that polysemy is the result of the interaction of several factors, some of which are linguistic, some cognitive, and some communicative, and the debate is more about which of these factors is the most important: Does polysemy have a primarily *linguistic* basis (an assumption held by most scholars working within computational semantic

¹ For instance, according to (Byrd et al. 1987) almost 40 per cent of the entries in *Webster's Seventh Dictionary* have two or more senses.

frameworks), is it essentially *cognitive* (as is claimed by scholars working within the cognitive grammar tradition), or is it a fundamentally *communicative* phenomenon?

1.2 Background

The fact that a word may be associated with several meanings was addressed at least as early as the writings of Aristotle (Barnes 1984). In the opening of *Categories*, Aristotle distinguishes between *synonymy* ('univocity') and *homonymy* ('multivocity', 'being spoken of in many ways'). Two things, *a* and *b*, are synonymous or univocal if they are both called by the same name *F*, and the definition of *F* is the same for both of them, whereas *a* and *b* are homonymous if they are called by the same name *F*, but the definition of *F* for *a* does not completely overlap with the definition of *F* for *b* (Shields 2009).² An example of synonymy is the occurrences of *human* in 'Socrates is a *human*' and 'Plato is a *human*', where the things named by the word *human* are the same in both cases.³ An example of homonymy is the occurrences of *bank* in 'John went to the *bank* to open a savings account' and 'Plato and Socrates had a picnic on the *bank*', where the things that are named by *bank* ('financial institution', 'riverbank') have distinct definitions.⁴ Furthermore, as the definition above states, homonymy also includes those instances of things called by the same name that have partially overlapping definitions. Consider the occurrences of *healthy* below:

- (1) a. Socrates is *healthy*.
- b. Socrates' exercise regimen is *healthy*.
- c. Socrates' complexion is *healthy*.

These three predications of *healthy* are non-univocal. In (1)a., the meaning of the predicate is roughly 'is physically fit', in (1)b. it means 'promotes health', and in (1)c., it means 'is indicative of health'. As Aristotle observed, the meanings of *healthy* in (1)b.

² In the classical approach of Aristotle, definitions describe necessary and sufficient application conditions. I return to this approach in section 1.4.2 below, as well as in detail in Chapter 2 in connection with the discussion of polysemy representation.

³ Aristotle's notion of synonymy is thus different from its contemporary usage where it describes different words with the same meaning.

⁴ It should be noted that Aristotle's theory is concerned with metaphysics, more specifically with the essences of things, and so his terms 'synonymy' and 'homonymy' apply in the first instance to things, and not to word meanings.

and (1)c. are both dependent on the meaning of *healthy* in (1)a, which is contained as part of their definitions. This is referred to as a kind of *core-dependent homonymy* (Shields 1999; Owen 1960 termed this 'focal meaning'), that is, an intermediate case between univocity and full homonymy.

Until relatively recently, almost all theories of linguistic semantics were based on these classical ideas about necessary and sufficient application conditions for concepts, as manifested most notably by Katz's semantic theory (1972; Katz and Fodor 1963; Katz and Postal 1964). As I will discuss in more detail in the next chapter, Katz's approach takes word meanings to be constituted by definitions, and lists as many meanings for a word as there are sets of necessary and sufficient conditions for its application. In contemporary linguistic theory this view is still held by a number of scholars, most notably by those working within Anna Wierzbicka's (1996) theory of Natural Semantic Metalanguage (NSM).

Another early mention of the issue of lexical meaning variation in the history of Western philosophy is Locke's (1975 [1689]) discussion of the English connective *but* and Leibniz's (1996 [1765]) subsequent critique of it (cf. Fieke Van der Gucht and De Cuyper 2007). Locke saw *but* as being associated with a number of different meanings (e.g. it may express opposition, coordination, etc.), and expressed doubts about the possibility that they could all be instantiations of a single more abstract meaning. Leibniz, on the other hand, took issue with Locke's claim that *but* has several different meanings and argued that instead, we should try to reduce all the uses of a word to "a determinate number of significations" (Leibniz 1996 [1765]: III, §4), by searching for a 'paraphrase' that is able to cover as much of the semantic variation of the word as possible.

Interestingly, this short discussion between Locke and Leibniz sums up the broad lines of the traditional debate over polysemy. As I mentioned above, theories of polysemy representation are often divided into 'sense enumeration lexicons' and 'core meaning approaches'. Sense enumeration lexicons, which take the representation of a polysemous lexical item to consist in a listing of all its uses, bear a clear resemblance to Locke's position above, while core meaning approaches, which see polysemous lexical items as being represented in terms of a highly abstract 'core meaning', which remains constant across all its uses, strongly resembles Leibniz's approach.

In general linguistics, Bréal (1924 [1897]) was the first to introduce the term 'polysemy' ('polysémie') to describe single word forms with several different meanings (cf. Nerlich 2003). For Bréal, polysemy was primarily a diachronic phenomenon, arising as a consequence of semantic change. Words acquire new meanings through use, but these do not automatically eliminate the old ones. Polysemy, then, is the result of the parallel existence of new and old meanings in the language; it is the 'synchronic side' of lexical semantic change. However, Bréal also observed that, at the synchronic level, polysemy is not really an issue, since the context of discourse determines the sense of a polysemous word and eliminates its other possible meanings (Bréal 1924 [1897]: 157). These early insights of Bréal also underlie much contemporary research in lexical semantics and pragmatics.

Following the advent of transformational-generative grammar in the late 1950s, with its main focus on syntax, polysemy received little attention for several years (some exceptions are Weinreich 1964, 1966; Anderson and Ortony 1975; Apresjan 1974; Caramazza and Grober 1976). However, with the development of cognitive grammar during the 1980s, polysemy reappeared on the research agenda as a central topic in lexical semantics, in particular as a result of the pioneering studies of prepositional polysemy conducted by Brugman (1988; Brugman and Lakoff 1988) and Lakoff (1987). A central claim of these studies was that polysemy is not so much a linguistic phenomenon as a cognitive one, resulting from the way in which our conceptual categories are structured.

Today, there are broadly two main trends in the research on polysemy. One is the well of polysemy research conducted within the cognitive linguistics framework, which has grown out of the work by Brugman and Lakoff, as mentioned above, as well as Langacker's (1987) foundational work in cognitive grammar (e.g. Geeraerts 1993; Tuggy 1993; Cuyckens and Zawada 1997; Dunbar 2001; Nerlich et al. 2003; Tyler and Evans 2003). The other is the polysemy research conducted within computational semantic frameworks, which includes most notably the generative lexicon account maintained by Pustejovsky (1995a), as well as several others (e.g. Copestake and Briscoe 1996; Kilgarriff 1992, 1995; Kilgarriff and Gazdar 1995; Asher and Lascarides 2003; Asher forthcoming). In contrast to the cognitive linguistic approaches, such

computational approaches see polysemy as primarily a linguistic phenomenon, arising from lexicon-internal computational processes.

However, as we are going to see in the coming chapters, in addition to these two main trends in the research on polysemy, much of the work conducted within the relatively new field of ‘lexical pragmatics’ has a direct bearing on the issue of polysemy (e.g. Recanati 1995, 2004; Carston 1997, 2002b, 2010; Blutner 1998, 2004; Wilson and Carston 2006, 2007). The main concern of these approaches is the communicative aspect of polysemy, that is, the interaction between linguistically-encoded content and contextual information in the derivation of speaker-intended lexical meanings.

Recently, there have also been several psycholinguistic studies investigating the representation of polysemy (e.g. Klein and Murphy 2001, 2002; Klepousniotou 2002; Klepousniotou, Titone, and Romero 2008; Beretta, Fiorentino, and Poeppel 2005; Pylkkänen, Llinás, and Murphy 2006). I will consider some of these studies in more detail in the next chapter.

1.3 The polysemy paradox

Already Bréal had noted that when talking to each other we rarely get confused by the multiplicity of meanings that a word can have. Sometimes, of course, speakers may exploit the polysemous potential of a lexical item to create confusion or a humorous effect, for instance a pun, but generally, polysemy is unproblematic from the point of view of communication. In normal circumstances, speakers can trust hearers to use contextual cues to quickly and reliably figure out the meaning they intend to communicate when using a word that could take on a different meaning in a different context. So the potential ambiguity that polysemy creates (and the process of disambiguation that would follow such ambiguity) does not seem to be much of an issue in human communication; it is something that we handle effortlessly and unconsciously, for the vast majority of the time.

In contrast, polysemy raises a host of challenging issues in semantic theory and semantic applications, such as lexicography, computational models of natural language processing and translation. This has led some scholars to talk of a ‘paradox’ in connection with polysemy, referring to the discrepancy between the relative ease with which it is used and understood in communication, and the range of theoretical and

descriptive problems it is associated with (Ravin and Leacock 2000; Taylor 2003). I outline the main issues in this section.

A first issue concerns how polysemy should be defined, and (if desirable) distinguished from homonymy (accidental multiple encoding), on the one hand, and contextual modulation of lexical meaning, on the other hand. I will refer to this as the *problem of definition*. In lexicography, the problem of definition translates into a methodological issue of distinguishing between those cases of multiple encodings that should be listed as distinct entries in the dictionary (homonymy), and those that should be listed under a single entry (polysemy), and, in the latter case, of determining which senses are established (hence should be listed) and which are not (hence should not be listed). I return to discussion of the problem of definition in more detail in Section 1.4.

A second issue is how polysemous lexical items are represented in the mental lexicon. The question is whether all the different senses of a polysemous lexical item are stored (as claimed by sense enumeration lexicons), or only a single representation is stored and all the other senses contextually derived, (as claimed by core meaning approaches), or only some senses are stored and the others contextually derived. I will refer to this as the *problem of representation*. A problem for fully encoding sense enumeration lexicons is that the mental lexicon would have to store indefinitely many semantic distinctions for each lexical item. A challenge for core meaning approaches is to decide on the appropriate level of abstraction for defining the core meaning of a lexical item, in order to capture what is common to all its possible, sometimes radically different, uses (cf. *run* above). A problem for more moderate approaches, which postulate only a limited number of senses of a word, is to distinguish between those meanings that are stored in the lexicon and those that are constructed on-line in language use.

Traditionally, polysemy is described as the association of two or more *related* senses with a single word form. Thus, a further issue that an account of polysemy representation must address is how the relations between the polysemous senses are represented in the mental lexicon, that is, if they are represented at all. For instance, on the classical approach, where word meanings are represented as necessary and sufficient features, sense relations are accounted for in terms of relations between the features contained in the lexical representations (e.g. the meanings of *healthy* in (1) above are all

related by virtue of containing the core concept ‘physically fit’ as part of their definitions). In the cognitive linguistics tradition, where a common view is that polysemous lexical items are represented as a network of senses centred around a primary, prototypical sense (e.g. Lakoff 1987), sense relations concern, in the first instance, adjacent members of the network, while members that are only indirectly connected in the network may be very different in semantic content.

The problem of polysemy representation is intimately connected to the broader question of what word meanings are, and, specifically, what kind of mental representation is encoded by a lexical form: Are lexical meanings complex, that is, do they consist of smaller units of meaning, and if so, what are these (necessary and sufficient features, conjunctions of basic concepts, lexical networks, or something else?), or are they more appropriately seen as simple, unstructured meanings (atomic concepts)? How one answers this question has significant consequences for an account of polysemy representation. In Chapter 2, I address this issue in detail.

A third issue that the proliferation of polysemy in natural languages raises is the fundamental question of why it exists. What is it about our language systems, specifically their lexical component, that makes them so susceptible to polysemy? Why do we rather use the same word to describe a set of different things or properties than have a distinct word for each sense, and what are the mental abilities that allow us to do this? This issue, which I refer to as the *problem of polysemy motivation*, has received little attention outside the cognitive linguistics paradigm (where, as I mentioned above, polysemy is seen as a result of how cognitive categories are structured more generally). In particular, the role of our communicative abilities in motivating polysemy has, in my view, not received the focus it deserves. While polysemy (in fact, ambiguity quite generally) represents an almost insurmountable challenge for computational models of natural language processing and translation, it is, as I mentioned at the beginning of this section, largely unproblematic from the point of view of human communication. A pragmatic theory that is able to explain how and why this is so, is also likely to shed light on the issue of what motivates the proliferation of polysemy in natural languages in the first place. I return to consideration of this issue in Chapter 3.

A final issue is *how* lexical meanings get extended into several different meanings. If we grant that some (possibly many) senses of polysemous lexical items are

derived or constructed during on-line processing, what are the processes or mechanisms involved, and how do they operate? This is a topic that has been widely studied, in particular by computational semanticists, cognitive linguists, and by pragmatists and philosophers of language interested in lexical meaning. Again, the main controversy is whether the mechanisms involved in meaning extension are primarily linguistic, cognitive or communicative (i.e. pragmatic). In Chapters 4, 5 and 6, I take a pragmatic approach to the question of how lexical meanings are extended during utterance processing.

1.4 The problem of definition

A standard textbook definition of polysemy is “the association of two or more related senses with a single linguistic form” (Taylor 1989/2003: 144). In this section, I consider three issues which have a bearing on the definition and delimitation of polysemy, including (i) the distinction between polysemy and homonymy, (ii) the distinction between ambiguity and vagueness, and (iii) the connection between polysemy and contextual modulation of lexical meaning.

1.4.1 Polysemy and homonymy

Traditionally, polysemy is distinguished from homonymy.⁵ In polysemy, the different senses of a single lexical item are seen as being related in some non-trivial way, whereas in homonymy, the multiple encoding is a matter of historical accident. An example of homonymy is the lexical form *coach*, which encodes the entirely unrelated meanings ‘bus’ and ‘sports instructor’. It is standard to see these as being represented as two different lexemes (COACH₁ and COACH₂) in the mental lexicon.

However, drawing the distinction between related and unrelated senses of a lexical form is often far from a straightforward matter. The question is: What does it mean for two senses to be related? The criteria that have been suggested include etymology and speaker intuitions about unrelatedness vs. relatedness of meaning (Lyons 1977b). According to the etymological criterion, two senses are homonymous if they are

⁵ The distinction between polysemy and homonymy corresponds to Weinreich’s (1964) distinction between *complementary* and *contrastive ambiguity*.

historically unrelated. On this approach, the linguistic form *file* would be an instance of homonymy, as the sense ‘folder or box for holding loose papers’ originates from the French word *fil*, and the sense ‘tool with roughened surface(s)’ comes from the old English word *féol*. That these two senses came to be associated with the same lexical form in contemporary English is thus a matter of historical accident. According to the same criterion, the noun *position*, which has the senses ‘a particular way in which someone or something is placed or arranged’ and ‘a person’s particular view or attitude toward something’, would be polysemous as a result of the shared etymological origin of its senses.

However, this way of distinguishing between polysemy and homonymy is problematic if we are concerned with characterising the linguistic knowledge of speakers and hearers. To illustrate, consider the English word *cardinal*. This word encodes the meanings ‘leader of the Roman Catholic Church’ or ‘North American songbird of the bunting family’. The two senses are historically related; the male cardinals are mostly red in colour and so this bird was named *cardinal* by virtue of its resemblance in colour to the red cassocks worn by cardinals. According to the etymological criterion, then, *cardinal* would be polysemous. However, many speakers of English may not be aware of this historical connection, and to them the two senses may seem entirely unrelated (i.e. homonymous). So, distinguishing between polysemy and homonymy on the basis of etymology does not, in many cases, capture differences in speakers’ intuitions of semantic relatedness, and, although such etymological considerations are no doubt useful to lexicographers in the making of dictionaries, it is doubtful whether they are relevant to a synchronic analysis of polysemy.

Another criterion that has been suggested as a way to distinguish between polysemy and homonymy is speaker intuitions about related and unrelated senses. According to this criterion, two senses are polysemous if they are judged by native speakers to be related, and homonymous if they are judged to be unrelated (or at least their meanings are considered to be further apart than polysemous senses as in, e.g., *cardinal*). Distinguishing polysemy from homonymy would thus depend on a sort of ‘folk etymology’. A problem that arises in connection with this criterion is that sense relatedness appears to be a matter of degree, and, moreover, judgements about the relatedness of the senses of a given word are likely to be subjective (Lyons 1977b). For

instance, compare the different senses of the nouns *right* ('morally correct', 'the right-hand part, side or direction'), *letter* ('symbol of the alphabet', 'written communication'), and *position* ('a particular way in which someone or something is placed or arranged' and 'a person's particular view or attitude toward something'). Intuitively, we feel that the senses of *position* are more closely related than those of *right*, whose semantic content is arguably quite distinct, and than those of *letter*, whose senses might nevertheless be considered to be more closely related than those of *right*. The senses of these words thus seem to be related to different degrees. Furthermore, it is possible that some speakers may claim to see a relation between the senses of a word form where others do not. An example is the different meanings of the noun *ear*, 'organ of hearing' and 'seed-bearing head or spike of a cereal plant'. For some native English speakers, these senses may seem related, while others may not see any relation between them at all. Such facts make the prospect of drawing a clear-cut distinction between polysemy and homonymy on the basis of speaker intuitions slim.⁶

Most importantly, however, it is not clear that speakers' intuitions about relatedness and unrelatedness of senses have any bearing on the way in which individuals use and understand words (Lyons 1977b: 552), quite unlike, for instance, intuitions about grammaticality, which have been considered the basic data to be explained within generative grammar. This is because it seems that many of our intuitions about sense relations might be reflective (i.e. arrived at by thinking about language) and thus not a direct reflex of the way in which word meanings are represented in our linguistic systems.

It is possible to argue that the distinction between polysemy and homonymy is of no theoretical significance, and to see them both as instances of a more general phenomenon of linguistic ambiguity. This is the position taken by Kempson (1977: 82), who treats the lexical form *run* as being represented as several distinct entries in the mental lexicon, one for each of its meanings, on a par with clearly homonymous forms such as *file*, *coach* or *bank*. Kempson sees this as the appropriate way to account for the fact that *run* creates an ambiguous sentence when more than one of its interpretations are possible. To this kind of approach it has been objected (in particular by scholars

⁶ That is, unless the distinction is allowed to vary across different I-languages, which it clearly should do.

working within the cognitive linguistics paradigm, cf. Taylor 1989/2003: 107) that it reduces polysemy to an arbitrary, unmotivated phenomenon, and thus makes it impossible to explain patterns of polysemy that are observed repeatedly across languages (e.g. Sweetser 1990). Although this is a reasonable objection, and a semantic (or pragmatic) theory that can explain these facts is clearly desirable, Kempson's position may still be correct with regard to the representation of *conventional* polysemy, that is, the cases where the related senses of a lexical form have become established in the lexicon. What would be required, then, is an account of how polysemous senses are derived, and a means to distinguish between conventional and contextually derived instances of polysemy.

1.4.2 Ambiguity and vagueness

A much-discussed distinction in semantic theory is that between ambiguity and vagueness (e.g. Kempson 1977; Cruse 1986; Atlas 1989; Geeraerts 1993; Tuggy 1993; Williamson 1994). Traditionally, lexical ambiguity is seen as involving two or multiple lexemes with distinct senses (and may, as we saw above, include both homonymy and polysemy), and vagueness a single lexeme with a non-specific meaning (monosemy), which is contextually specified. Different kinds of vagueness are discussed in the literature, including 'indeterminacy of meaning', that is, cases in which the meaning of the lexical item appears to be quite intangible, and 'lack of specification', that is, cases in which the meaning of the lexical item is in principle quite clear but is very general (Kempson 1977: 125). The adjective *good* has been mentioned as an example of indeterminacy of meaning, due to the range of different senses it may express in describing different things (e.g. *good knife/football player/student/weather*, etc.), as well as in describing the same thing (e.g. a *good job* could be one that's well paid, offers interesting tasks, has an inclusive social environment, gives a certain social status, etc.). Examples of vagueness due to lack of specification are terms such as *teacher*, *cousin*, *neighbour*, etc., all of which are unspecified with regard to gender.

Several tests have been proposed for distinguishing between ambiguity and vagueness. These can be divided into the following types: (i) logical tests; (ii) linguistic tests, and (iii) definitional tests (Geeraerts 1993, 1994).

An example of a logical test is the one proposed by Quine (1960: 129), according to which a lexical item is ambiguous if it can at once be clearly true and clearly false of the same thing. For instance, an assertion of ‘Anne is wearing a *light* skirt’ would be true of a situation where Anne is wearing a black skirt made of a thin fabric of little weight, if *light* is taken to denote the property ‘of little weight’, and false if *light* is taken to denote the property ‘pale’. Thus, the adjective *light* is ambiguous according to the logical test.

Linguistic tests involve semantic restrictions on sentences that contain two occurrences of the lexical item under consideration. If a grammatical construction requires semantic identity between the two occurrences, ambiguous expressions will give rise to several readings for the construction. For instance, Kempson’s (1977: 129) anaphora-based test involves the use of the expression *do so too* (or *so did X, X did/has/will/is too*), which demands identity of meaning of two verb phrases.⁷ A verb phrase, then, is two-ways ambiguous if conjoining a *do so too* phrase to it renders the whole sentence two-ways ambiguous, as in (2) and (3):

- (2) Anne wore a *light* skirt and Jane did so too.
- (3) Tom went to the *bank* and his mother did so too.

Both sentences above have two readings because the same senses of *light* and *bank* must be selected in each of the conjuncts, indicating that the expressions are ambiguous. Compare with (4) and (5), where the expressions *neighbour* and *good* may be distinctly specified in the two conjuncts (that is, so-called ‘crossed’ readings are allowable):

- (4) John is my *neighbour* and Sue is too.
- (5) The book is *good* and the cake is too.

According to Kempson’s test, then, *neighbour* and *good* are not ambiguous but vague, and should thus be represented as single lexemes.

Another type of linguistic test is ‘co-predication’, which is taken as a diagnostic of ambiguity if it gives rise to ‘zeugma’, that is, the oddity that results when two or more

⁷ A similar kind of linguistic test for ambiguity is the so-called identity test proposed by Zwicky and Sadock (1975).

terms are inappropriately linked (Cruse 1986, 2000; Pustejovsky 1995b; Copestake and Briscoe 1996). Consider (6) and (7):

- (6) ? Peter went to the *bank* to deposit money and to watch the riverboats.
(7) ? Susan's grandfather and driver's license *expired* last week.

According to this test, the fact that we perceive both these sentences as odd is indicative of the ambiguity of the expressions *bank* and *expire*.

Finally, according to the definitional test, a word is ambiguous if more than a single definition is needed to account for its meaning. This type of test was originally proposed by Aristotle in *Topics* (Barnes 1984), and presupposes the classical view of word meanings, where these are represented in terms of necessary and sufficient features. For instance, while two distinct definitions are required to account for the two meanings of *bank* ('financial institution', 'riverside'), and for the meanings of *healthy* above, there may be only one definition needed to account for the meaning of *neighbour*. As mentioned above, in contemporary linguistic theory this view is espoused most notably by proponents of the 'natural semantic metalanguage' (NSM) method of semantic analysis (e.g. Wierzbicka 1996). Goddard (2000: 132) describes the NSM method for establishing polysemy as follows:

One assumes to begin with that there is but a single meaning, and attempts to state it in a clear and predictive fashion in the form of a translatable reductive paraphrase. Only if persistent efforts to do this fail is polysemy posited. The next hypothesis is that there are two distinct meanings, and attempts are made to state both in a clear and predictive fashion, and so the process goes, until the full range of application of the word can be captured within the specified range of senses.

On this approach, then, a lexical form has no more meanings than the number of maximally general definitions that are required to account for its range of applications.⁸

A problem with these three types of tests for ambiguity, as Geeraerts (1993) has demonstrated in detail, is that different tests may not always agree with each other, and that by manipulating the context, they can be made to yield inconsistent results. One of

⁸ A problem with the definitional test for ambiguity is that it rests on a view of lexical semantics that has been shown to be largely untenable. As I will return to the classical theory (as well as the problems associated with it) in more detail in the next chapter, I say no more about it here.

the examples Geeraerts discusses is the word *newspaper*, which is standardly seen as being ambiguous (polysemous) between the senses ‘printed object’ and ‘the management of the organisation that publishes the newspaper’. This word would be ambiguous according to the logical and definitional tests but not so according to the linguistic test:

(8) The *newspaper* decided to change its print.

The grammatical construction in (8) requires identity of reference between the subject of the VP ‘decided to change its print’, and the subject (PRO) of the infinitival complement ‘to change its print’. However, although in (8) *newspaper* is first used in its ‘management of organisation’ sense and then in its ‘printed material’ sense, this does not affect the acceptability of the sentence and so no ambiguity would be predicted (Geeraerts 1993).

Several authors have discussed examples of contextual influence on ambiguity tests. Consider the following instances of co-predication (Norrick 1981: 115):

- (9) a. ? Judy’s *dissertation* is thought provoking and yellowed with age.
b. Judy’s *dissertation* is still thought provoking though yellowed with age.

The sentence in (9)a., is zeugmatic, due to the use of *dissertation* to refer to informational content in the first conjunct, and to a material object in the second conjunct. However, when the sentence is slightly altered, as in (9)b., no zeugmatic effect occurs. Cruse (1986: 65-66) discusses a similar case, involving uses of the noun *door*:

- (10) a. The *door* was smashed in so many times it had to be bricked up.
b. ? Billy took the *door* off its hinges and walked through it.

Given that *door* is used to refer to different objects in the two conjuncts (‘door-panel’, ‘doorway’), the sentence in (10)a. should be zeugmatic; yet no such effect is obtained. In contrast, the sentence in (10)b., which appears to involve the same alternation in

meaning for *door* as in (10)a., is clearly zeugmatic. Thus, co-predication does not seem to be a reliable test for the ambiguity of an expression.⁹

Geeraerts's (1993) examination of the three types of ambiguity tests described above leads him to reach the following conclusion: Instead of presupposing, as the use of these tests does, that there must be a 'true' criterion for distinguishing between ambiguity (polysemy) and vagueness, it might be that the distinction itself is unstable; that is, it might not be possible to draw a clear-cut distinction between ambiguity and vagueness. If this is so, he argues, it should have consequences for a theory of lexical meaning:

The tremendous flexibility that we observe in lexical semantics suggests a procedural ... rather than a reified conception of meaning; instead of meanings as things, meaning as a process of sense creation would seem to become our primary focus of attention. (Geeraerts 1993: 260)

Geeraerts's conclusion is in accord with the approach to polysemy that will be put forward in this thesis, insofar as it acknowledges the central role of processes of sense creation in giving rise to polysemy. In the next section, I consider briefly this connection between polysemy and the pragmatic adjustment of lexical meaning ('lexical modulation', 'sense creation'), and sketch the broad outlines of my approach.

1.4.3 Polysemy and lexical pragmatics

So far, we have looked at the traditional distinctions between polysemy (related senses) and homonymy (unrelated senses), on the one hand, and ambiguity (two or more stored senses, including homonymy and conventional polysemy) and vagueness (a single sense subject to contextual specification), on the other hand. With regard to the first distinction, we saw that neither etymological considerations nor speaker intuitions provide satisfactory criteria for distinguishing between polysemy and homonymy. As to the second distinction, we saw that the number of tests proposed to distinguish between

⁹ However, Cruse himself is more optimistic about using co-predication as a diagnostic of ambiguity. He explains the difference between (10)a. and (10)b. in the following way: In (10)a., *door* is used in a 'global door' sense, which includes both the doorway and the door-panel, and this is what prevents the appearance of zeugma in this case. However, the presence of zeugma in (10)b. shows that the meanings of *door* are in actual fact distinct. Therefore, Cruse concludes, in applying the co-predication test we should try to avoid contexts such as (10)a., which may lead us to draw a wrong conclusion as to whether or not a lexical item is ambiguous.

ambiguity (homonymy and conventional polysemy) and vagueness did not yield clear-cut results. An important reason for this was the contextual influence on these tests, reflecting the fact that any lexical item, whether monosemous, homonymous or polysemous, may have its meaning modulated on a particular occasion of use by the linguistic or non-linguistic context in which it occurs. Consider (11)-(16):

- (11) Mary *opened* the shoebox. ('performed the kind of action that causes a shoebox to become open')
- (12) I didn't get enough *units*. ('university credit modules')
- (13) The steak is *raw*. ('grossly undercooked')
- (14) My ex-wife is a *ghost* to me now. ('distant memory')
- (15) John just phoned *the bank*. She is willing to support him. ('John's mother')
- (16) *The violin* is in a bad mood today. ('The violin player')

In (11), the verb *open* expresses the specific action one performs when opening a shoebox, which is different from the kind of action involved, for instance, in opening a pair of curtains or a dishwasher. It seems unlikely that the lexicon should encode such fine-grained distinctions, when they can easily be inferred from the context; thus, the use of *open* above appears to be a case of contextual modification of lexical meaning. In (12) the noun *units* is used to communicate 'university credit modules', an interpretation clearly not derivable from the linguistic context. In (13), where the adjective *raw* is used to convey that the steak is grossly undercooked, we seem to have a case of (extra-linguistic) contextual modulation of the conventional 'uncooked' meaning of the adjective. The examples in (14) and (15) are both instances of metaphorical uses of a noun; in (14) *ghost* is used to communicate something along the lines of a 'distant memory', while (15) contains a highly creative (metaphoric) use of the noun phrase *the bank* to refer to John's mother, derived on the basis of the encoded 'financial institution' sense of this word. The use of *the violin* to refer to the player of the violin in (16) is a case of a metonymic use of a word, where the metonymic meaning is derived on the basis of a contextually salient relation holding between the conventional and the metonymic denotations.

There is no doubt that such pragmatic adjustment of lexical meanings plays a central role in giving rise to polysemy. In fact, leaving the metaphoric and metonymic uses above aside, the kind of contextual modulation of encoded lexical meanings involved in the uses of *open*, *unit* and *raw* above is often described as a kind of polysemy. For instance, Sperber and Wilson (1998) discuss the range of meanings that the verb *open* may be used to express as an instance of polysemy. Other scholars may see *open* as an instance of vagueness, or indeterminacy of meaning, as described above, and restrict the term ‘polysemy’ to cases where the related senses of a lexical form are (hypothesised to be) encoded.

In this thesis, my aim is to investigate the close connection that exists between the cases commonly discussed under the heading ‘polysemy’ in the literature, and cases of on-line contextual modulation of meanings due to pragmatic processes operating at the level of individual words. More specifically, I will develop a pragmatic account of polysemy within the framework of relevance theory (Sperber and Wilson 1986/1995; Carston 2002b). The argument presented in this thesis is that polysemy is a fundamentally communicative phenomenon, which arises (mainly) as a result of encoded lexical concepts being massively underdetermining of speaker intended concepts, as illustrated by the examples in (11)-(16), and grounded in our pragmatic inferential capacity. As mentioned above, while most accounts of polysemy either take it to be a linguistically driven phenomenon or, as in cognitive linguistic frameworks, a result of how cognitive categories are structured quite generally, this thesis provides a novel perspective on the phenomenon by highlighting its communicative aspect, which is claimed to be essential to its development and proliferation in language. The specific pragmatic processes responsible for the construction of occasion-specific lexical meanings have recently been described in much detail in the relevance-theoretic framework (Carston 1997, 2002b; Wilson and Sperber 2002; Wilson and Carston 2006, 2007). This thesis endorses the account of lexical pragmatic processes proposed by relevance theory, and shows how it can be used to account for the construction (and the subsequent establishment) of polysemous senses for a lexical item. The thesis extends the relevance-theoretic analysis by developing an account of instances of systematic polysemy and metonymically derived polysemy, previously unexplored within the relevance-theoretic approach to lexical pragmatics.

Throughout this thesis, I will operate with a broad notion of polysemy, using it as a cover term for conventional polysemy (which nonetheless, arguably, has a pragmatic basis) and cases of clearly context-dependent, hence pragmatically derived, lexical meanings. Thus, this will include the cases of vagueness described as ‘indeterminacy of meaning’ (e.g. *open*, *good*), but exclude those cases of vagueness that involve a ‘lack of specification’ (e.g. *cousin*, *neighbour*, *teacher*, etc.), in which the contextual specification does not seem to involve an actual modulation of the encoded meaning. In cases where it is relevant, I will operate with an intuitive distinction between conventional (i.e. ‘encoded’, ‘semantic’) polysemy and contextually-derived polysemy, acknowledging that there may be no clear-cut way of drawing this distinction.

1.5 Outline of the thesis

In Chapter 2, I address the problem of polysemy representation, as described above. I consider several influential approaches to word meaning and assess their implications for polysemy representation. I also discuss some relatively recent psycholinguistic studies that apply experimental methods in order to glean evidence from on-line processing that may bear on the question of whether polysemous lexical items are represented as multiple or single entries in the lexicon. In Chapter 3, I consider the problem of polysemy motivation. I take a pragmatic approach to polysemy and argue that it is, at the deepest level, a consequence of how communication works; it results from our capacity to infer speaker meanings on the basis of not fully determining linguistic evidence. In Chapter 4, I address the nature of the pragmatic processes that are involved in the construction of polysemy. I outline the relevance-theoretic approach to lexical pragmatics (e.g. Carston 2002b; Wilson and Carston 2007), according to which there is a single pragmatic process of *ad hoc* concept construction that can adjust the meanings of individual words in different directions and show how it can account for the construction of polysemy during utterance comprehension. Chapter 5 is devoted to a discussion of so-called systematic polysemy (e.g. count-mass alternations such as ‘The *chicken* pecked the ground’/‘We had *chicken* in bean sauce for dinner’), which is usually taken as a prime case of linguistically generated polysemy. My main claim with regard to such cases is that, although they clearly have a linguistic component, the

contribution of the linguistic system to their generation and interpretation is less central than is often thought. In the final Chapter 6, I consider the process of lexical meaning extension referred to as metonymy, and discuss two possible directions for a relevance-theoretic account of this phenomenon.

Chapter 2

LEXICAL SEMANTICS AND POLYSEMY: THE PROBLEM OF REPRESENTATION

2.1 Introduction

In the previous chapter, we saw that one component of the so-called polysemy paradox was the issue of how polysemous lexical items are represented in the mental lexicon, which we called the *problem of representation*. This will be the topic of the present chapter.

There are broadly two ways in which one might approach this problem. The question of how polysemy is mentally represented is intrinsically connected to the more general question of what word meanings are, and, specifically, what kind of mental representation is encoded by a lexical form. A fundamental debate in lexical semantics is whether word meanings decompose into smaller units of meaning, and, if so, how (lexical decomposition approaches), or whether word meanings are unstructured atoms (non-decompositional approaches). Thus, one way to approach the problem of polysemy representation is to situate it within the larger context of this debate between decompositionalists and non-decompositionalists. A widely held view among decompositionalists is that in order to capture meaning relations between words (e.g. synonymy, analyticity, entailment) word meanings must consist of complex representations, and some scholars further think that this is also required in order to handle polysemy (Jackendoff 1992b, 2002; Murphy 2002; Vicente and Martínez Manrique 2008, 2010). As the psychologist Gregory Murphy writes:

The fact that people use words like *paper* to refer to a substance, a copy of a daily publication, and an editorial policy means that the word's representation must be complex. (Murphy 2002: 413).

Another way to approach the problem of polysemy representation is from the perspective of the overall structure of the mental lexicon. As I mentioned in the previous chapter, it is widely agreed that the meanings of a homonymous word form (e.g. *bank*) are represented as separate entries in the lexicon. Some scholars, most notably within the field of psycholinguistics, have therefore asked whether this is also the case for

polysemous lexical items, or if they should rather be seen as being represented as single entries in the mental lexicon (and the different meanings derived on the basis of lexical rules or pragmatic procedures). Based on these two options, lexical semantic theories of polysemy representation are often divided into two types (as we saw in the previous chapter): (i) *sense enumeration lexicons* (e.g. Katz and Fodor 1963; Katz 1972; Weinreich 1966; Lakoff 1987; Brugman 1988; Taylor 1989/2003), and (ii) *core meaning approaches* (Caramazza and Grober 1976; Allerton 1979; Ruhl 1989; Pustejovsky 1995a). As we are going to see, this is a very crude distinction, but it nevertheless singles out a question that arises largely independently of the debate concerning the internal representation of lexical meanings (or senses), and, in principle, whose answer might be compatible with both decompositional and non-decompositional approaches to lexical semantic representation.

The chapter will be organised in the following way. In the first part I discuss some influential decompositional approaches to word meaning, and assess their implications for polysemy representation (section 2.2). In the second part I discuss Fodor's (1998) non-decompositional approach, and consider the possibilities for accounting for polysemy within this framework (section 2.3). On the basis of the accounts reviewed in this chapter, I argue that a conceptual atomist approach, which treats word meanings as unstructured atoms, may be at least as able as decompositional approaches to account for polysemy when paired with an adequate pragmatic theory. In the final part of the chapter I discuss some relatively recent psycholinguistic studies that address the question of whether polysemous words are represented as multiple or single entries in the mental lexicon (section 2.4).

2.2 Lexical decomposition approaches

By far the majority view in linguistics is that most word meanings decompose into smaller units of meaning. There is, however, considerable variation as to how the internal structure of word meanings is conceived, whether it is taken to consist in, e.g., a definition, a prototype structure or semantic network, an underspecified template or a partial definition. In this section, I consider some of these approaches in more detail, as well as the implications of each of them for polysemy representation.

2.2.1 The classical theory

The classical theory of concepts, traceable back to at least the time of Socrates,¹⁰ takes lexical semantic representations to be constituted by definitions. The definition of a complex concept *C* specifies the necessary and sufficient conditions for something in the world to be a *C*. This kind of approach divides the lexicon into definable (complex) terms, on the one hand, and a set of primitives, on the other hand, typically acquired through the senses.¹¹ The primitives form the inventory from which complex concepts are constructed, on the basis of a set of combinatorial processes (conjunction, negation, etc.). The paradigmatic classical concept is encoded by the word *bachelor*, and can be defined as follows:

(1) [¬MARRIED & MALE & HUMAN & ADULT]

Each of the conjuncts in (1) is a necessary condition for a thing in the world to be a bachelor, and jointly they are sufficient, so that anything that meets all of these conditions is indeed a bachelor and anything that does not is not a bachelor.

A great advantage of the classical theory is that it appears to have considerable explanatory power. For one thing, it provides a natural explanation of categorisation, reference determination and concept acquisition, all of which are explained in terms of lexical semantic representations having definitional structure: Categorisation is a matter of checking whether the set of necessary and sufficient conditions specified by the definition of a concept is indeed satisfied by an object, a concept refers to the things that satisfy its definition, and acquiring a concept involves learning its definition. For another, it captures our intuitions about semantic relations, including synonymy (e.g. of *bachelor* and *unmarried man*), ‘analytic’ inferences (i.e. inferences drawn solely on the

¹⁰ Cf., for instance, Socrates’s search for the definition of *mud* (which he ends up deciding is ‘earth mixed with water’) in Plato’s (1987) *Theaetetus*. Cf. also the works of Aristotle (Barnes 1984), mentioned in the previous chapter.

¹¹ On strictly empiricist versions of the classical theory, the concepts occurring in a definition should ultimately decompose into basic, perceptual or sensory concepts, following Leibniz’s (1996 [1765]) ‘Empiricist Maxim’, stating that there is nothing in the mind which was not first in the senses. More recent manifestations of this view are found in Wierzbicka (1996) and Prinz (2002).

basis of the meanings of the words used, e.g. *John is a bachelor* → *John is unmarried*),¹² and entailment (e.g., on the assumption that the definition of *kill* is CAUSE TO DIE, it explains informally valid arguments of the type *John killed Harry* → *Harry died*).

There are two main options for describing polysemy within a framework based on the classical assumptions about word meanings: (i) the meaning of a polysemous lexical item is represented as a single (abstract) definition, and the different senses are contextually determined (the 'monosemy view', cf. Allerton 1979; Ruhl 1989); and (ii) each meaning of a polysemous word is associated with its own set of necessary and sufficient application conditions (Katz and Fodor 1963; Katz 1972; Wierzbicka 1996; Goddard 2000). The most influential of these has undoubtedly been (ii), largely due to the pioneering work in lexical semantics by Jerrold Katz (1972; Katz and Fodor 1963; Katz and Postal 1964), which I consider below.

Katz's semantic theory

Katz's semantic theory (Katz and Fodor 1963; Katz and Postal 1964; Katz 1972) is a modern linguistic embodiment of the classical theory. Katz's primary aim was to provide a theory of natural language semantics that was able to explain semantic relations and contrasts between word meanings (e.g. synonymy, antonymy, contradiction, analyticity, entailment, etc.), and the relation between word meanings and sentence meanings, in a description of the 'semantic component' of the grammar. In his view, the only way to achieve this was by means of lexical decomposition of the sort provided by the classical theory.

On Katz's (1972: 36) account, the semantic component of the grammar is seen as containing (i) a *dictionary*, comprising a list of the meanings of the words in the language; and (ii) a set of *projection rules*, reflecting the speaker's ability to construct sentence meanings from word meanings. The dictionary lists under a single lexical entry the different senses of a word (which together constitute the *meaning* of that word),

¹² In this case, principles of valid inference are taken to apply to the semantic representation and not to the (surface) form of the sentence, and so this argument has the underlying form *John is a man and unmarried*. So *John is unmarried*, and is a case of conjunction elimination ($p \ \& \ q \rightarrow q$).

each of which can be broken down into *semantic markers* (or primitives), as in Katz and Fodor's (1963: 198) suggestion for the dictionary entry for the noun *ball*:¹³

(2) *ball*

1. *Ball* → Noun concrete → (Social activity) → (Large) → (Assembly) → [For the purpose of social dancing]
2. *Ball* → Noun concrete → (Physical Object) → [Having globular shape]
3. *Ball* → Noun concrete → (Physical object) → [Solid missile for projection by engine of war]

On Katz's theory, the noun *ball* is (at least) three ways ambiguous. When put into linguistic context, the projection rules operate to disambiguate its meaning, by amalgamating the features of the different senses of the word with those of other words in the surrounding linguistic context to check which features are compatible and which are not. Thus, an NP such as *large ball* would retain the three-way ambiguity, *metal ball* would exclude the reading of *ball* as a social activity, while *sponge ball* would be fully disambiguating, only picking out the second reading according to which *ball* refers to a physical object having globular shape.

Katz's theory is a prime example of a *sense enumeration lexicon*, where different senses of a lexical item are listed under a single dictionary entry. As shown by the dictionary entry for *ball* in (2) above, a single entry may contain both polysemous and homonymous readings of a lexical item. While the second and third senses of *ball* are polysemous (the 'globular shaped object' and 'type of missile' senses), the first sense (where *ball* means 'social activity') is homonymous with both of them. Katz (1972: 69) suggested that the distinction between homonymy and polysemy could be drawn on the basis of the notion of 'semantic similarity'. According to his definition (1972: 48), two constituents are similar (on a sense) if they have a semantic marker in common. Applying this to the dictionary entry for *ball*, we see that the polysemous senses under 2. and 3. share the semantic marker (Physical Object), and are thus similar according to the definition. Moreover, neither of them shares any semantic markers with the

¹³ The semantic markers are given in parentheses. The square brackets indicate what Katz and Fodor (1963: 187) called 'distinguishers', which were supposed to reflect idiosyncratic aspects of word meanings.

homonymous sense in 1. However, although it might work to distinguish polysemy from homonymy in this particular case, semantic similarity does not seem to be an adequate diagnostic of polysemy, as there would arguably be a number of instances of homonymy which exhibit semantic similarity as well (e.g. the homonymous senses of the noun *file*, ‘folder or box’ and ‘tool’ would also share the semantic marker (Physical Object)), and there are polysemous senses which do not share any semantic markers at all (as shown by Katz’s dictionary entry for the adjective *handsome* (ibid. 43)). So, unlike the other cases of semantic relations (for instance, synonymy, antonymy, analyticity and entailment) that can be ‘read off’ elegantly from the theory by means of sameness, overlap or incompatibility of semantic markers contained in the semantic representations of words, it appears that polysemy (and ultimately the distinction between polysemy and homonymy) is not so easily accounted for this way. Thus, it appears that, at least in the particular version of lexical decomposition that Katz was offering, there would be little support for the claim that lexical decomposition is a requirement for a proper treatment of polysemy.¹⁴

In addition to the more specific problems concerning polysemy and homonymy representation in Katz’s theory, philosophers have pointed out a number of more general problems related to the view that words meanings can be given in terms of necessary and sufficient application conditions, which, taken together, have made it nearly impossible to maintain the classical theory as an account of lexical semantic representation. Some of them are considered below.¹⁵

Problems with the classical theory

Perhaps the most basic problem for the classical theory is that there are in fact very few words or concepts that have proper definitions, referred to as *Plato’s Problem* (Laurence and Margolis 1999: 14). Two and a half millennia of philosophical investigations into the meanings of various important concepts such as KNOWLEDGE, JUSTICE, GOODNESS,

¹⁴ A further problematic aspect of Katz’s dictionary entries is that they list homonymous senses under single lexical entries (cf. the dictionary entry for *ball*), suggesting that this is also how such senses are psychologically represented (although, given Katz’s (1981) Platonist view of language, it is not clear that he is committed to this view). This would run counter to the widely accepted view that homonymous senses are psychologically represented as distinct lexical entries.

¹⁵ Much of what I will have to say in the next section is based on the excellent discussion of the classical theory in Laurence and Margolis (1999).

TRUTH and BEAUTY have shown that it is notoriously difficult to come by definitions that capture all the necessary and sufficient conditions for the application of a concept (not to mention the increase in difficulty if these have to be given in a vocabulary of sensory or perceptual terms), and so *bachelor* and a few other clearly definable concepts may be exceptional rather than typical. For most concepts, it appears that it is only possible to specify some necessary conditions, a point which is particularly clear in the case of natural kind terms. As pointed out by J. D. Fodor et al. (1975), for a colour term like *red*, the property COLOUR is a necessary condition for something to be red, but there is no property (other than the property RED itself) with which the property COLOUR can combine in order to provide a definition of *red*.¹⁶ For other concepts, there may not even be any clearly necessary conditions, as shown by Wittgenstein's (1953, sections 65-78) discussion of *game* (e.g. board-games, card-games, ball-games, etc.): There is probably no property that the range of phenomena falling under the concept GAME have in common.¹⁷

A related argument that has been advanced against definitional accounts of word meaning is that, even for apparently definable terms, definitions appear to be psychologically irrelevant. For instance, according to the definitional approach, a sentence such as 'John is a *bachelor*' should be representationally more complex than the corresponding sentence 'John is *unmarried*', since the definition of *unmarried* is a

¹⁶ J. A. Fodor and colleagues (Fodor, Fodor, and Garrett 1975; Fodor et al. 1980; Fodor 1970, 1981) have provided a range of arguments against definitional approaches to lexical semantic representation (and, more recently, against decompositional approaches more generally, cf. Fodor 1998; Fodor and Lepore 2002)). In particular, their arguments were directed at definitional analyses of causative verbs (McCawley 1968; Lakoff 1970, 1976; Faarlund 1978), which, in the early days of generative (transformational) grammar were treated as being represented by two events at the deep structure level (e.g. *kill* was represented as CAUSE TO DIE, *paint* as CAUSE TO BECOME COVERED IN PAINT, etc.). This made it possible to account for adverb scope ambiguities of the type 'John *almost* killed Harry' in terms of the adverb modifying both events (long scope), giving rise to the interpretation that John almost brought about Harry's death, or only the second event (short scope), giving rise to the interpretation that John brought it about that Harry almost died. What Fodor and colleagues convincingly showed was that the causal subcomponents that are clearly entailed by such causative verbs are only necessary application conditions, i.e. one-way entailments, rather than the two-way entailment relations required by definitions. In other words, every case of killing is indeed a case of causing to die, but not every case of causing to die is a case of killing (Fodor 1970).

¹⁷ Rather, Wittgenstein notes, the various activities subsumed under this label are "*related* to one another in many different ways" (1953: section 67). He introduces the term 'family resemblances' to characterise this situation. Various types of games form a family, and so the word *game* can be seen as giving access to a 'family of meanings'. This idea became central to the later developments of 'prototype theory' (Rosch and Mervis 1975; Rosch 1999 [1978]), which I return to in section 2.2.2.

proper part of the definition of *bachelor*. However, this prediction is not supported by the available experimental evidence, which suggests that definitions do not affect processing (Kintsch 1974; Fodor, Fodor, and Garrett 1975; Fodor et al. 1980; de Almeida 1999).¹⁸

A further problem for the classical theory is the vagueness of many concepts, which may pose difficulties in categorisation. For instance, does a carpet belong to the category of furniture or to some other category? The classical theory does not allow for such indeterminacy in category membership, nor does it allow for indeterminacy in our knowledge about category membership (that is, that we may be uncertain as to whether something belongs to a category or not). A related problem is that it is possible to have a concept in spite of being ignorant or mistaken about the properties we take its instances to have. For instance, although there are many people who think that whales are a kind of fish, the fact that they are mistaken about this does not mean that they do not have the concept WHALE.

The appeal of the classical theory has been largely owing to its ability to account for semantic relations, and analytic inferences in particular. However, Fodor et al. (1980) have shown that there are several cases of informally valid arguments whose validity cannot be reconstructed by reference to definitions. For instance, on the assumption that *kill* can be defined as CAUSE TO DIE, there is no rule of standard logic that guarantees the validity of the inference from *John caused Harry to die* to *Harry died*. In fact, it appears to rely on a further (non-standard) inference rule that allows for the elimination of the causal element in the definition (CAUSE P \rightarrow P).¹⁹ Furthermore, the very idea that there are statements that are analytic, that is, statements that are true *a*

¹⁸ There is another psychological fact that goes against the classical theory as concerns its account of lexical acquisition, called the 'basic-level advantage' (Rosch et al. 1976). Children have been shown to acquire the words for basic, middle-level concepts (e.g. DOG) before the words for more general, superordinate concepts (e.g. ANIMAL) and those for more specific, subordinate concepts (e.g. POODLE) (Brown 1958; Rosch et al. 1976; Horton and Markman 1980; Bloom 2000). This fact does not mesh well with the idea that learning a concept involves learning its definition: For instance, in order to acquire a (basic-level) concept such as DOG, the child has to already have the (superordinate) concept ANIMAL, as this concept would be part of its definition.

¹⁹ Fodor's suggestion, then, was that instead of treating our intuitions about analytic inferences as resting on underlying definitions, they could be accounted for in terms of inference rules (or meaning postulates) attached to concepts (Fodor et al. 1980; Fodor 1981). On this view, the inference from *John killed Harry* to *Harry died* would rely on stipulated inference rules attached to the concept KILL (KILL \rightarrow CAUSE TO DIE) and CAUSE (CAUSE P \rightarrow P). I return to the issue of meaning postulates in section 2.3 in the discussion of conceptual atomism.

priori, has come under serious attack, in particular by Quine (1999 [1953]). On Quine's view, it is impossible to determine in advance of empirical investigation the conditions under which a particular statement would be confirmed, nor is any statement immune to revision. Thus, it is impossible to draw a principled distinction between those inferences that are content-constitutive (analytic, *a priori*) and those that are merely contingent (synthetic, *a posteriori*). If Quine is right about there being no such thing as an analytic statement, this would undermine the whole basis for the classical theory, which is built on the assumption of the existence of analytic relations between concepts.

Finally, the psychological evidence for so-called typicality effects seems hard to reconcile with the classical theory. There is a robust body of experimental results showing that people have little difficulty ranking items with respect to how 'typical' members of a category they are (Rosch 1973, 1999 [1978]; Rips, Shoben, and Smith 1973; Rosch and Mervis 1975). For instance, a robin is considered to be a more typical member of the category BIRD than a chicken, or an ostrich, is. Such typicality effects have been found to correlate with speed of categorisation (typical members are categorised faster than non-typical members) and error rates (the more typical a member of a category, the less errors people make in categorising it). However, if category membership is a matter of satisfying a set of necessary and sufficient conditions, there is no reason to assume that all members should not be on equal footing, or that certain members should be faster, or easier to categorise, than others. Thus, typicality effects raise serious explanatory problems for the classical theory.

2.2.2 Prototypes, cognitive semantics and polysemy

The label 'prototype theory' describes a set of theories that build on the assumption that categories exhibit typicality effects (Rosch 1973, 1975a, 1975b; Rosch and Mervis 1975; Smith, Shoben, and Rips 1974; Osherson and Smith 1981; Smith and Medin 1981). On this view, rather than being represented as necessary and sufficient conditions, most concepts are structured in terms of 'prototypes'; based on a statistical analysis of the properties (features) that the members of a category tend to have.²⁰ Something falls

²⁰ The notion of a 'prototype' is used in (at least) two ways; as describing, for a given category (e.g. BIRD), (i) sets of features (e.g. has wings, flies, lays eggs, sings, is of such and such a size, colour, shape, etc.) (cf. Rosch 1975a), or (ii) exemplars (e.g. robins as best examples of birds) (cf. Smith and Medin 1981).

under a concept if it satisfies a sufficient number of features, where some features may be weighed more significantly than others. For instance, a robin is a member of the category BIRD by satisfying a large number of the features determined by this category, e.g. 'has wings', 'flies', 'nests in trees', 'sings', etc. (and hence is a typical instance of the category). A peacock, although satisfying less of them, still satisfies a sufficient number of features for it to belong to the category BIRD (and hence is a less typical instance of the category). In this way, prototype theory avoids the problems that the classical theory faced concerning the lack of definitions and the analytic/synthetic distinction. Typicality effects can be explained in terms of a member's similarity to the prototype of a category (Rosch and Mervis 1975).²¹ On this approach, category membership is a matter of degree; there are more or less typical members, better and poorer examples of a category, and the boundaries between categories are not necessarily clear-cut.

Prototype semantics

In linguistics, the prototype approach to categorisation has formed the basis for many accounts of word meaning and of polysemy, in particular within the cognitive linguistics paradigm (e.g. Fillmore 1982; Lakoff 1987; Brugman 1988; Brugman and Lakoff 1988; Taylor 1989/2003; Geeraerts 1997). In an early account, Fillmore (1982) suggested treating word meanings as being represented as 'semantic prototypes', which could be realised in several different ways. For instance, he took the meaning of the verb *climb* to be composed of the disjunction of two mutually compatible conditions, 'clambering' and 'ascending'. In its prototypical uses, both conditions are satisfied (e.g. a person climbing up a rock wall), while in less central uses only one of the conditions may be present (e.g. a gecko climbing up a wall (no clambering), or a cat climbing down a tree (no ascending)).²² In other cases, there is a disjunction of two conditions, but one of them has a more privileged status, and is satisfied in prototypical uses (e.g. the adjective

²¹ It has been noted that typicality effects do not in themselves provide evidence for concepts being structured in terms of prototypes, as even apparently well-defined categories may exhibit typicality effects. For instance, in a series of experiments, Armstrong, Gleitman and Gleitman (1999 [1983]) found that subjects judged the number 7 to be a better example of the category ODD NUMBER than, e.g., the number 57, and that a housewife was a better example of the category FEMALE than were both waitress and comedienne. They concluded that there cannot be any incompatibility between the definability of a concepts and its having a prototype structure.

²² See Coleman and Kay (1981) for a similar analysis of the verb *lie*.

long, where the spatial sense is privileged over the temporal sense). Thus, the polysemy of *long* is captured in the semantic representation of this word. In yet other cases, the meaning of a word consists in a fixed set of conditions (e.g. *bird*), and the most prototypical examples are those that satisfy the most conditions; the ones that come close to an idealisation of the category (e.g. robins) (cf. Rosch 1975a).

A widely discussed objection to prototype semantics as it is manifested in Fillmore's early approach, as well as against prototype theory as a theory of concepts more generally, is its inability to account for the composition of complex concepts (Fodor 1998; Fodor and Lepore 2002). The issue is not whether typicality effects are psychologically real or not (there are few scholars who would dispute that) but what role (if any) they should play in a semantic theory. The compositionality problem for prototype theory is that the meaning of a complex concept is generally not a complex prototype built up from the prototypes of its constituents. For instance, while a guppy may be a prototypical instance of the category PET FISH, it is not a prototypical instance of either the category PET or the category FISH. Several sophisticated suggestions have been made for ways to account for conceptual combination within prototype theory (Osherson and Smith 1981; Smith and Osherson 1984; Smith et al. 1999 [1988]; Huttenlocher and Hedges 1994; Kamp and Partee 1995), most of which, however, seem to be able to handle only a very limited range of data. Moreover, a range of further problems have been pointed out with prototype theory as a theory of conceptual representation, including the lack of prototypes for many concepts (e.g. THE NORWEGIAN PRESIDENT, NEW SPECIES, BELIEF), the possibility of having a concept without knowing its prototype (e.g. someone could have the concept THE WORKS OF VIRGINIA WOOLF without knowing any of the properties associated with it), and the difficulty of providing an adequate account of reference determination, in that concepts with prototype structure may fail to cover atypical instances (e.g. a zebra without stripes) and incorrectly include non-instances (e.g. dingoes and wolves could fall under the category DOG).

In view of these problems, it seems clear that an adequate theory of lexical semantics cannot be based on prototype structure alone. Indeed, some scholars have proposed a 'dual theory' of concepts, where concepts are composed of a classical core (i.e. a definition) plus an optional identification function with prototype structure (e.g. Armstrong, Gleitman, and Gleitman 1999 [1983]). Lakoff (1999 [1987]), on the other

hand, has claimed that the very idea that categories are structured in terms of prototypes is a wide-spread misunderstanding of prototype theory.²³ Instead, he argues, typicality effects are surface phenomena, and result from the organisation of knowledge in terms of so-called ‘idealised cognitive models’. Lakoff’s theory of knowledge representation, with its account of word meaning and polysemy, has been immensely influential within the cognitive linguistics framework, and we will spend the rest of this section looking at this approach.

Radial categories and polysemy

In Lakoff’s (1987) framework, idealised cognitive models (ICMs) are relatively stable mental structures that represent theories about the world with respect to a particular domain, and which guide categorisation and reasoning. They are ‘idealised’ in the sense that they are abstractions across a range of experiences, and may give rise to typicality effects in several different ways. The simplest type of prototype effect is illustrated by the concept BACHELOR (Lakoff 1987: 70-71). For instance, the ICM with respect to which a bachelor is defined includes information about the institution of marriage, a typical marriageable age, etc. The concept BACHELOR itself is not a graded category (either a man is unmarried or he isn’t), but against the background conditions specified by the ICM, some unmarried men may be considered typical bachelors, some less typical bachelors (e.g. the Pope). A second source of typicality effects is so-called ‘cluster concepts’, consisting of a combination of number of individual ICMs, illustrated by the concept MOTHER (Lakoff 1987: 74). A mother could be the person who gives birth (THE BIRTH MODEL), the female adult who nurtures and raises a child (THE NURTURANCE MODEL), the wife of the father (THE MARITAL MODEL), etc. More than one of these conditions may characterise a mother, and any one of them can be absent from such a characterisation. Typicality effects arise when one of the submodels is viewed as primary.²⁴ In addition, Lakoff claims, the concept MOTHER is an instance of a ‘radial category’ where there is a *central* subcategory (the cluster concept, including the birth

²³ In fact, Rosch (1999 [1978]: 200) herself had warned against this misunderstanding of the notion of a prototype for a category: “to speak of a *prototype* at all is simply a convenient grammatical fiction; what is really referred to are judgements of degree of prototypicality. ... Prototypes do not constitute a theory of representation of categories.”

²⁴ Lakoff (1987: Chapter 4) further identifies ‘metonymic models’ as a way in which typicality effects may arise on the basis of ICMs. I return to this issue in Chapter 6.

model, the nurturance model, etc.), combined with a set of non-central extensions, or variants, of the central category (e.g. adoptive mother, stepmother, surrogate mother, etc.) (ibid. 83-84). Rather than being generated by the central subcategory, the variants are extended by convention from the central category (i.e. they must be learned), on the basis of general principles of extension.

The notion of radial categories forms the basis for Lakoff's account of polysemy (Lakoff 1987; Brugman 1988; Brugman and Lakoff 1988), which has inspired a host of studies of polysemy within the strand of linguistics known as 'cognitive semantics'. On this approach, which takes linguistic categories to be no different from other kinds of conceptual categories, most word meanings are seen as a type of radial category in which the different senses of a word are organised with respect to a prototypical sense. The paradigmatic example is the preposition *over*, first discussed by Brugman (1988):

- (3) a. The bird flew *over* the house. ('above and across')
- b. The painting is *over* the couch. ('above')
- c. The truck ran *over* the rabbit. ('across')
- d. Sarah lives *over* the hill. ('on the other side')
- e. Mary nailed a board *over* the hole in the ceiling. ('covering')
- f. I will read the papers *over* the weekend. ('temporal')
- g. John has a strange power *over* Mary. ('control')

The idea is that *over* constitutes a radial category composed of a range of distinct but related senses, organised around the prototypical, or central, sense (which Brugman (1988) and Lakoff (1987) take to be the 'above and across' sense in (3)a.), in a lexical network structure.²⁵ The different senses of *over* exhibit typicality effects; more typical senses are located 'closer' to the prototypical sense in the network, while less typical senses are located in its periphery. Such peripheral senses are derived from more typical senses by a set of cognitive principles for meaning extension (e.g. 'conceptual metaphors', cf. Lakoff and Johnson 1980), giving rise to meaning chains (e.g. sense *A* is related to sense *B* in virtue of some shared attribute(s), sense *B* is related to sense *C*,

²⁵ A slightly different, albeit similar manifestation of the network model of polysemy representation is given by Langacker (1988). For reasons of space, I cannot get into the details of his approach here.

which is related to sense *D*, and so on). For instance, the ‘control’ sense in (3)g. is seen as being derived from the ‘above’ sense in (3)b. on the basis of the metaphorical schema CONTROL IS UP, LACK OF CONTROL IS DOWN (Lakoff 1987). Sense relations, then, concern, in the first instance, adjacent members of the category, while members that are only indirectly connected in the semantic network may be very different in semantic content.²⁶

A central aspect of Lakoff and Brugman’s cognitive semantic approach is that radial categories such as that associated with *over* are stored in the long-term semantic memory of speakers. This means that the different senses associated with a polysemous lexical form are taken to be conventional, rather than generated on the basis of a linguistic or pragmatic mechanism. In this respect, the radial category account of polysemy is a radical version of the type of approach we referred to above as *sense enumeration lexicons*, in that the full range of senses are taken to be stored as part of a semantic network (hence, it is sometimes referred to as the ‘full-specification approach’ to lexical semantics, cf. Evans and Green 2006). A common criticism of the full-specification approach is that it entails a (potentially) indefinite proliferation of mentally stored senses in order to cover the range of uses of lexical forms (for instance, Brugman (1988) identifies nearly a hundred different uses of *over*). In addition to the consequence of an enormous demand on the storage capacity of the language user, this account fails to distinguish between those aspects of meaning that are part of the word meaning proper and those that result from its interaction with the context, a problem that is sometimes referred to as the ‘polysemy fallacy’ (Sandra 1998). Indeed, the polysemy account of Lakoff and Brugman appears to ignore the contribution of pragmatics to the derivation of word senses altogether. A further criticism against the sense network accounts is their lack of methodological constraints in postulating different senses of a word. Sandra (1998: 370-371) writes:

When looking at such [semantic network] analyses one has the strong impression that many of the distinctions are unnecessary – at least the authors typically do not provide proof, linguistic or other, of the necessity for making

²⁶ Thus, Wittgenstein’s (1953) metaphor ‘family resemblance’ (cf. Rosch and Mervis 1975) is often used to describe such polysemous categories within the cognitive linguistics framework (Taylor 1989/2003).

their distinctions. One can only conclude that the analysis reflects the particular approach of the individual linguist and that different linguists would probably come up with different analyses at different points in time.

More recently, several scholars working within the cognitive linguistics paradigm have taken a more moderate approach to polysemy, acknowledging the context dependency of word meanings (Tyler and Evans 2001, 2003; Allwood 2003; Zlatev 2003; Evans 2005, 2009; Taylor 2006). In particular, Tyler and Evans (2001, 2003) have developed an account of polysemy that, while espousing the Lakoff-Brugman idea that polysemous senses are represented in terms of sense networks centred around a prototypical sense, proposes a set of criteria that makes it possible (i) to determine whether a particular sense of a word counts as a distinct sense; and (ii) to establish the central sense of a polysemous lexical item. This account, termed 'the Principled Polysemy approach' by its authors, seeks to avoid the polysemy fallacy by distinguishing between those senses that are stored in semantic memory and those that are pragmatically constructed during 'on-line' processing. I return to consideration of the Principled Polysemy approach in Chapter 4, in a further discussion of prepositional polysemy.

2.2.3 The generative lexicon

The generative lexicon (Pustejovsky 1991, 1995a, 1998a, 1998b) stands out from the other theories covered in this chapter by being designed with the sole purpose of accounting for polysemy. Pustejovsky sought to provide a more explanatory account of polysemy than that given by sense enumeration lexicons. In his view, accounts that describe polysemy in terms of multiple listings of senses in the lexicon are inadequate, primarily because they are unable to explain how words may take on an infinite number of meanings in novel contexts. Not only is it impossible for such accounts to list all the possible meanings of a lexical item, they also miss the generalisations that can be made on the basis of what appear to be regular patterns of sense alternations, and fail to capture how polysemous senses may partially overlap and be logically related to one another (Pustejovsky 1995a: Chapter 4). A more promising approach, he argues, which is able to meet these explanatory requirements, is a lexicon where items are decomposed into templates (rather than sets of features), combined with a generative framework for the composition of lexical meanings (*ibid.* 58).

It is clear that when faced with Pustejovsky's requirements for an adequate theory of lexical semantics, a sense enumeration lexicon such as Katz's (1972) would be unsuccessful on all counts. For network theories, the central worries (as discussed in the previous section) would be the issue of indefinite proliferation of senses and the failure to account for the interaction of word meaning with the context, as the main force of these theories is precisely their ability to provide generalisations over patterns of meaning extension and to account for relations between different senses of a word. However, it is a genuine question whether we want our *semantic* theory to handle all this, that is, whether we should expect it to account for the range of possible (context-dependent) meanings of a word (as well as the number of ways in which these meanings may be related), *without the intrusion of pragmatics*, which is in effect what Pustejovsky is suggesting. Before I discuss this issue in more detail, I would like to look in some detail at the central aspects of his generative lexicon theory.

On Pustejovsky's account, the semantics of a lexical item is viewed as a structure consisting of four components, described as *argument structure* (number and type of arguments), *event structure* (event type of a lexical item, including subevents), *qualia structure* (a structured differentiation of the meaning of a lexical item), and *lexical inheritance structure* (the ways in which a lexical item is related to other lexical items in the lexicon) (Pustejovsky 1995a: 61). The qualia structure of a lexical item is the hallmark of Pustejovsky's theory, and consists in a specification of four different roles: The *constitutive* role captures the relation between an object and its constituents, or proper parts; the *formal* role specifies what distinguishes the object within a larger domain; the *telic* role defines the purpose and function of the object (if there is one); and the *agentive* role describes the factors involved in the origin or coming into existence of the object.²⁷ In (4), we have a simplified description of the qualia structure for the noun *novel*, including values for each of the qualia roles (adapted from Pustejovsky 1991:

²⁷ Although the notion of 'qualia structure' seems to apply mainly to the meanings of nouns (and to nouns denoting concrete objects in particular), Pustejovsky maintains that a qualia structure can be defined for all types of lexical items, but that not all lexical items need carry a value for each role.

427), specifying that a novel is a narrative, it has the form of a book, its purpose is to be read and it comes into being by a process of writing:²⁸

(4) *novel*

QUALIA STRUCTURE:

CONSTITUTIVE ROLE = narrative

FORMAL ROLE = book

TELIC ROLE = read

AGENTIVE ROLE = write

Complementing these underspecified lexical entries is a set of generative mechanisms, described as *type coercion*, *co-composition* and *selective binding*, which operate to yield compositional interpretations. The process of ‘type coercion’ is defined as a “semantic operation that converts an argument to the type that is expected by a function, where it would otherwise result in a type error” (Pustejovsky 1995: 111). Consider (5):

- (5) a. Mary *began* reading a novel.
b. Mary *began* a novel.

In (5)a. the VP complement of *begin* is of the semantic type ‘event’, and in (5)b. its NP complement specifies an object. Instead of there being different lexical entries for *begin*, that is, one for each complement type the verb may select for, it is treated as having a single lexical entry, its argument being specified as an event (Pustejovsky 1995a: 116). In cases where this requirement is not directly satisfied by the surface syntactic structure, as in (5)b., coercion applies in order to change the type of the complement NP into an event, consistent with event information contained in the qualia structure for the noun. Thus, the denotation of the NP *a novel* is ‘coerced’ into an event denotation compatible with either of the interpretations ‘Mary began to write a novel’ and ‘Mary began to read a novel’ (cf. the qualia structure for *novel* in (4) above).

The generative operation referred to as ‘co-composition’ enables the information carried by a complement to act on the governing verb, by taking the verb as argument

²⁸ For instance, this very schematic representation does not indicate that *novel* belongs to a group of nouns that are seen as being polysemous between a ‘physical object’ sense and an ‘information type’ sense. These cases will be discussed in Chapter 5.

and shifting its event type (Pustejovsky 1995a: 123). Consider the uses of *bake* in (6) below:

- (6) a. John *baked* a potato.
b. John *baked* a cake.

Rather than treating the different senses of *bake* as separate lexical entries, Pustejovsky sees the change-of-state meaning in (6)a. as basic, and the creative meaning in (6)b. as being derived by co-composition, applying at the level of the entire verb phrase. When combined with the complement *potato*, whose agentive role is specified as ‘natural kind’, the change-of-state interpretation of *bake* remains unchanged. The semantics of the complement *cake* however, shifts the change-of-state meaning of *bake* to a creation sense, a result of the agentive role of *cake* being specified as an artefact (i.e. originating out of the act of baking).

Finally, ‘selective binding’ is described as a generative process whereby an adjective “is able to make available a selective interpretation of an event expression contained in the qualia for the head noun” (Pustejovsky 1995: 128). Consider the uses of *good* below:

- (7) a. Every chef needs a *good* knife.
b. ‘American Pastoral’ is a *good* novel.

In (7), *good* selectively modifies the event description given by the telic roles of the nouns; in the case of *knife* the action of cutting (giving rise to the interpretation ‘a knife that cuts well’) and in the case of *novel* the action of reading (giving rise to the interpretation ‘a good read’).

Problems with the generative lexicon

There is little doubt that Pustejovsky’s generative lexicon provides a considerably more explanatory account of polysemy than sense enumeration lexicons of the kind advocated by Katz (1972). The problems associated with this approach, however, are numerous (Fodor and Lepore 1998, 2002; Blutner 2002; de Almeida 2004; de Almeida and Dwivedi 2008; Willems 2006). First, it cannot avoid making a range of wrong predictions concerning the construction of compositional interpretations. For instance,

it predicts that *Mary baked the pizza* should have a creative reading, while it seems that a non-creative reading (on which Mary is heating up a frozen pizza) is just as likely to be the preferred one. Similarly, there are several other (context-dependent) ways in which a knife could be good that do not involve ‘cutting’ (e.g. ‘good for stabbing people with’, ‘good for threatening people with’, ‘good for holding’, etc.). To solve this problem, Pustejovsky could allow for default interpretations to be defeasible, so that specific (linguistic) contexts may suggest other interpretations not inherent to the qualia of the lexical item, as suggested by Copestake and Briscoe (1996). However, incorporating defeasibility into the semantic system would not help with the cases of clearly infelicitous interpretations generated by it. It predicts, for instance, that a VP such as *begin a car* should be interpreted as ‘begin to drive a car’ and that *a fast cake* should mean ‘a cake that is fast to eat’, based on the telic roles for *car* and *cake*. Moreover, there are cases of uses of evaluative adjectives for which the generative system makes no interpretive predictions at all, due to the lack of a telic role for the adjective to selectively modify, as in *good weather* and *good children*. Yet there is little doubt that in these cases the compositional process proceeds as normal, giving rise to a different sense of *good* in each case.

Another criticism that has been levelled against the generative lexicon theory is that it lacks a distinction between linguistic knowledge and general world knowledge. The claim is that Pustejovsky’s lexical entries contain a considerable amount of information that is more likely to belong to the conceptual system than to the level of linguistic representation. For instance, as Fodor and Lepore (1998) argue, although trolley cars and knives are both artefacts, the VPs *bake a trolley car* and *bake a knife* resist a creative reading just as much as *bake a potato* does, and the reason they do this must lie in our real world knowledge about these objects. Further, the lack of a principled distinction between linguistic and conceptual knowledge paves the way for a range of *ad hoc* solutions to interpretive problems of the kind discussed above, as there is nothing in the theory that prevents us from building in all sorts of information in the lexical entry for an object. In principle, a preference for a change-of-state interpretation of *bake the pizza* could be explained in terms of the lexical representation of *pizza*

containing a distinction between home-made and pre-fabricated exemplars (and then linguistic cues could be used to choose between them), which would be clearly absurd.²⁹

A final theoretical consideration is the significant amount of work that the generative lexicon leaves for the pragmatic interpretive system to do, despite its information-rich lexical entries and elaborate generative-semantic mechanisms. In principle, as Žegarac (2006) has argued, a pragmatic theory that can serve this purpose (i.e. one that is capable of overriding default interpretations in the absence of linguistic cues and correcting the clearly wrong interpretations generated by the system) should also be capable of handling that part of the interpretive work that the generative lexicon does adequately. The question, then, is whether this would not in fact make the generative lexicon redundant. It seems unnecessary to expect the semantic system to do all this interpretive work, if we have a pragmatic system that can perform all the same tasks, only better and more accurately. I will discuss this issue further in Chapters 3, 4 and 5.

2.2.4 Neo-classical theories

The label ‘neo-classical theories’ describes a rather heterogeneous group of approaches that take word meanings to be constituted by partial definitions (Laurence and Margolis 1999). In this section I discuss two linguistic manifestations of the neo-classical approach, proposed by Pinker (1984, 1989, 2007) and Jackendoff (1983, 1992b, 1997, 2002, 2010).

Pinker’s theory of verbal decomposition

Pinker (1984, 1989, 2007) espouses a decompositional account of verb meaning, where verbs are partially defined in terms of a small set of universal, recurring, grammatically relevant elements. One of his main arguments for this approach is that it provides a

²⁹ A further problem, related to the previous discussion of the classical theory, is the inability of Pustejovsky’s lexical-semantic representations to pick out correct denotations. Despite being rich in conceptual content, in some cases they do not distinguish between what are obviously distinct denotations (for instance, the lexical-semantic representation for the noun *cake* (Pustejovsky 1995a: 123) would cover a range of similar objects as well, e.g. *pizza*, *bread*, *muffin*, etc.). In other cases, they predict a narrower denotation than is in fact the case (e.g. *good knife*), and sometimes, as we have seen, they predict a wrong denotation (e.g. *begin a car*). Moreover, Plato’s Problem equally applies to Pustejovsky’s theory; for many concepts, it will be difficult to give any precise values for the qualia roles (try, for instance, to provide qualia structures for abstract nouns such as *semantics*, *freedom*, *justice*, etc.), or indeed for any of the other levels of linguistic representation proposed by the theory.

solution to ‘Baker’s Paradox’ (1989: 7-9); the problem of explaining how children acquire syntactic properties of verbs (e.g. how they learn which verbs are dativisable, passivisable, and so on, and which are not), faced with the lack of parental correction and the (apparent) arbitrariness of such syntactic properties (e.g. near-synonyms may have different argument structures, e.g. *give/donate*), indicating that there is no simple semantic criterion guiding them.

In his proposal for a resolution to the paradox, Pinker suggests that verb alternations are in fact *not* based on arbitrary rule applications, but that the child is using semantic criteria to constrain the application of alternation rules (1989: 30). Based on the work by Beth Levin (1985; 1993), he suggests treating verbs as being divided into ‘lexical subclasses’, “defined by a distinctive, grammatically relevant subset of the semantic structures that constitute the meaning of a verb” (Pinker 1989: 103). A verb’s syntactic behaviour with respect to a given alternation is predictable on the basis of membership in a lexical subclass.

Pinker’s point of departure is simple, transitive action verbs, in which X does something to Y. On the basis of the various alternations that they enter into, he identifies a set of grammatically relevant bits of meaning that he takes to be components of the semantic representations of the verbs. For instance, he shows that some verbs enter quite naturally into the ‘conative alternation’ (which conveys attempting), e.g. *cut* (*Mary cut the rope/Mary cut **at the rope***) and *hit* (*John hit the wall/John hit **at the wall***) while other verbs that convey actions which, on cognitive grounds, should also be eligible for entering into this alternation, do not, e.g. *touch* (*Susan touched the cat/*Susan touched **at the cat***) and *broke* (*Bill broke his bicycle/*Bill broke **at his bicycle***). He concludes that the kind of verbs compatible with this construction must express a motion resulting in a kind of contact. This includes verbs of cutting (*cut, slash, chop, hack, chip*, etc.) and verbs of hitting (*hit, beat, elbow, kick, punch, poke, rap, slap, strike*, etc.), but excludes verbs of touching (*touch, kiss, hug, stroke, contact*, etc.) and verbs of breaking (*break, shatter, crack, split, crumble*, etc.) (Pinker 1989: 104).

In the ‘part-possessor ascension’ (involving a shift from construing the object as a ‘possessor of body part’ to ‘person as body part’), only the subclass of verbs that denote physical contact may alternate, including verbs of cutting (*Sam cut Brian’s arm/Sam cut Brian **on the arm***), hitting (*Mary hit the dog’s leg/Mary hit the dog **on the leg***), and verbs

of touching (*Richard touched Jane's cheek/Richard touched Jane **on the cheek***). Verbs of breaking, however, do not alternate (**Jim broke Mary **on the arm***) (ibid.: 105). As regards the 'middle alternation' (specifying the ease with which an action can be carried out on a patient), only the verbs that denote an effect may alternate (ibid.: 106). Thus, the alternation applies to the breaking and cutting sub-classes (*I broke the glass/The glass breaks **easily**, I cut the bread/This bread cuts **easily***) while not to the verbs of hitting and touching (*He hit the wall/*This wall hits **easily**, The touched the wire/*This wire touches **easily***) Based on these, and a set of other alternations involving transitive verbs (see Pinker 2007: 103-107 for a summary), Pinker sorts verbs in to classes, depending on the type of concepts they express:

- (8) *hit*: MOTION, CONTACT
cut: MOTION, CONTACT, EFFECT
break: EFFECT
touch: CONTACT

(8) suggests that the reason why *break* and *touch* do not enter into the conative alternation is that their lexical representations do not contain both the concepts MOTION and CONTACT. Similarly, *break*, failing to express the concept CONTACT, does not enter into the part-possessor ascension, while *touch*, and *hit*, which do not express EFFECT may not enter into the middle alternation.³⁰

On Pinker's approach, the features that sort verbs into classes (and hence alternations) in (8) are "not arbitrary markers" (2007: 107), but part of their necessary application conditions. Furthermore, membership in one of the lexical subclasses is seen as a sufficient condition for a verb to alternate (1989: 103).³¹ However, verbal polysemy appears to pose a difficulty for this view. Compare the uses of *cut* in (9):

³⁰ The status of these elements included in Pinker's semantic vocabulary is not entirely clear. For instance, do concepts such as MOTION, CONTACT and EFFECT decompose into smaller units of meaning, or are they primitive? Do they have lexical counterparts, or are they abstract semantic features, as, for instance, in Jackendoff's framework (1983, 1990, 1992b, 2002)?

³¹ Pinker denies the possibility that inferential relations of the sort *hit* → MOTION could be part of our general reasoning system (in the form of meaning postulates) instead of part of the lexical-semantic representation of the verbs. If they determine the syntactic constructions that the verb may enter into, they must be "part of the language engine proper" (Pinker 2007: 107). On the face of it, this seems to provide an argument in favour of lexical decomposition, at least in terms of partial semantic

- (9) a. John *cut* the rope./John *cut at* the rope.
 b. The bank *cut* its interest rates/*The bank *cut at* its interest rates.

As we have seen, the theory predicts that *cut* should enter into the conative alternation as a result of its semantics specifying ‘a motion resulting in a kind of contact’, as in (9)a. above. As to (9)b., Pinker might say that in this particular sentence *cut* does not express ‘a motion resulting in a kind of contact’, hence does not alternate. This would leave him with two possibilities concerning the meaning of the linguistic form *cut*. The first is that it encodes more than one meaning, so that the *cut* in (9)b. differs from the *cut* in (9)a. by having its own set of application conditions (whatever these might be). Only the *cut* in (9)a. would be a member of the group of ‘cutting’ verbs, and enter into the conative alternation. This solution, although possible, misses the intuition that the two uses of *cut* are closely related in meaning – they seem very much like different uses of the same word. Alternatively, Pinker could modify his claim about the semantics of *cut*, and say that the *cuts* in (9)a. and (9)b. are both instances of a single lexical item, and that the (linguistic) context determines whether or not *cut* in fact expresses ‘a motion resulting in a kind of contact’ (for instance, when its object denotes an abstract entity it follows that *cut* cannot express this meaning). This would amount to a modification of the claim that the features MOTION and CONTACT are part of the necessary application conditions for the verb *cut* (i.e. they are only necessary for the verb to enter into the alternation).

The polysemy of *hit* in (10) poses a perhaps more serious problem for Pinker’s account, as it seems to challenge the claim that membership in a lexical subclass is in fact a sufficient condition for a verb to alternate:

- (10) a. The woman *hit* the mugger with her umbrella/The woman *hit at* the mugger with her umbrella.
 b. The driver *hit* a pedestrian/*The driver *hit at* a pedestrian.
 c. The sniper *hit* a policeman./*The sniper *hit at* a policeman.

representations, and against an account that treats such inferential relations as part of the conceptual system, where there is no prediction that they should have any syntactic reflexes. However, as I will argue shortly, there is at least reason to question whether this link between the verbal semantic properties he postulates and the syntactic alternations they enter into is as neat as Pinker claims.

The verb *hit* has slightly different, but clearly related meanings in these examples: in (10)a., it conveys ‘direct a blow at’, in (10)b., it expresses ‘come into contact with (something/someone) quickly and forcefully’, and in (10)c., it conveys ‘strike’. The problem is that in all three cases *hit* expresses a motion resulting in a kind of contact, yet it is only the use of *hit* in (10)a. that licences the conative alternation. (On cognitive grounds, it is not difficult to imagine a situation where a driver attempts to hit a pedestrian, or where a sniper attempts to hit a policeman). I believe this poses a real problem for Pinker’s account, since it seems that in order to get around it, he either has to propose a more fine-grained semantic requirement for a verb to enter into the conative alternation (i.e. one that might distinguish between different senses of a polysemous verb), or he has to modify the claim that membership in one of the lexical subclasses is a sufficient condition for a verb to alternate.

Jerry Fodor, who is an ardent critic of lexical decomposition (defending himself an atomist position on word meanings, cf. section 2.3), has claimed, in a discussion of Pinker’s account of verbal decomposition, that although there might be semantic facts that have syntactic reflexes, this is not in itself an argument for lexical decomposition (Fodor 1998: Chapter 3). His argument is that it is fully compatible with a *non-decompositional* (i.e. atomist) account of word meaning to say that a lexical entry may contain features which are capable of affecting its syntactic behaviour, but that this is not the same as claiming that there is a *semantic level* at which only semantic properties of expressions are specified.³²

However, further to this comment of Fodor’s, I believe that the challenge from polysemy weakens Pinker’s claim about the necessary connection between verbal semantic features and syntactic alternations, and, as a consequence, his claim about the necessity of lexical decomposition. At least, it seems that this cannot be the whole story with regard to these alternations, and that the role of context (linguistic and extra-linguistic) in licensing them needs to be further explored. So I agree with Fodor in that, although Pinker has undoubtedly provided a valuable account of a range of generalisations among aspects of verb meanings and their acquisition, he hasn’t necessarily shown that verb meanings need be decompositional. Moreover, as to the

³² Borg (forthcoming), who also defends an atomist approach to word meaning, makes a similar point.

claim that lexical decomposition is required for an adequate theory of polysemy, it seems clear at least that the kind of decomposition proposed by Pinker does not get us much closer to an explanation of the phenomenon.

Jackendoff's conceptual semantics

While Pinker is explicit that his goal is not to capture the totality of a verb's meaning, Jackendoff's theory of partial decomposition (Jackendoff 1983, 1990, 1992b, 2002, 2010) is intended as a full-fledged theory of concepts. Jackendoff's primary argument for a decompositional account of word meaning comes from the creativity of language (1983: Chapters 5 and 7; 1990, 1992b, 2002). In generative grammar, this refers to the infinite generative power of the human language capacity, enabling language users to produce and understand, in principle at least, an infinitely large number of sentences, most of which they have not heard before. Given that our brains are finite, this creativity is explained in terms of syntactic knowledge being represented as a finite set of primitives and combinatorial principles that together generate all the possible sentences of a language. Extending this argument to the conceptual domain, Jackendoff claims that, given the indefinitely large class of possible concepts, these must also be constructed from a set of primitives and principles of combination contained in an innate 'grammar of lexical concepts' (1992b: 25-26).

In Jackendoff's framework, the thoughts expressed by language are structured in terms of a level of representation called *conceptual structure* (CS), which is linked to the syntactic component through a set of correspondence rules. The primitives of CS are described as a set of 'major ontological categories', including features such as THING, PLACE, DIRECTION, ACTION, EVENT, MANNER, AMOUNT (Jackendoff 1983: 50), which need not have any direct linguistic realisation. In this respect, the conceptual primitives are similar to phonological features in that we have no conscious access to them.³³ CS

³³ This is, according to Jackendoff (2002: 336), what makes it so hard to expound on word meaning on the basis of raw intuition. On his view, the principles governing the combination of conceptual features into word meanings need not be the same as the principles determining the combination of word meanings into phrase meanings. So, even if some of the primitives that compose a word meaning could, in principle, be expressed as words, there would be no phrase composed of those words that could express what the original word expresses. This would make concepthood a property of conceptual constituents, and not of conceptual features. This approach to lexical decomposition distinguishes itself clearly from the definitional approaches that Fodor (1998) criticises, on which word meanings are composed out of

contains no formal distinction between semantic and pragmatics (Jackendoff 2010: 8). Instead, ‘semantics’ is viewed as providing “the part of conceptual structure of an utterance that is related directly to linguistic expressions” and ‘pragmatics’ as providing “the part that arises though inference, heuristics, world knowledge, and understanding of the context”.

Jackendoff acknowledges the range of arguments against classical decompositional accounts in terms of necessary and sufficient conditions (1983: 121), and agrees that word meanings decompose only into *necessary* conditions (for instance, COLOUR is part of the conceptual representation of *red*), but include other information as well. In addition to CS, Jackendoff (1992b, 2002) envisages a further level of representation, called *spatial structure* (SpS), at which the spatial understanding of the physical world is encoded. Jackendoff’s suggestion is that the lexical entry for a physical object word includes, in addition to its conceptual structure, a sophisticated 3D-model representation (in the sense of Marr 1982), which can be thought of as an imagistic representation of the prototypical instance of a category. On this view, *dog* may be represented in conceptual structure as a kind of animal, typically domestic, a member of the canine family, etc., while its spatial structure may include a 3D-representation specifying information about its shape, coat, how it moves, etc. Words that encode perceptual concepts can be encoded directly in SpS, so for instance, *red* is represented as KIND-OF-COLOUR in conceptual structure, while the perceptual distinction between it and all the other colours is encoded in SpS.³⁴ Furthermore, although this possibility is not explored by Jackendoff himself, the assumption that there are SpS representations to word meanings provides an account of ‘systematic polysemy’ of the kind in (11):

lexical concepts using a subset of the principles required for the composition of word meanings into phrases.

³⁴ This would appear to solve the ‘problem of completers’, associated with partial decompositional approaches (Laurence and Margolis 1999: 54): With only partial definitions they have no account of reference determination, but, if their definitions are ‘fleshed out’, all the problems of the classical theory return. Supplementing necessary conditions with SpS in the representation of word meanings makes it possible to determine which things fall under a concept without having to specify this in terms of actual primitives. In particular, it appears to provide a solution to Fodor et al.’s (1975) problem of ‘missing properties’ in the case of natural kind terms (*red* → COLOUR & X) . However, Laurence and Margolis (1999: 56) note that spatial structure does not in fact offer a satisfactory solution to the reference determination problem, because, just as is the case with prototypes, something can satisfy the spatial properties specified by a concept without falling under it (e.g. an Australian dingo looks very much like a dog) and something can fall under the concept without satisfying the spatial properties (e.g. a hairless dog).

- (11) a. John opened the *window*. (physical object)
b. Mary went through the *window*. (aperture)
c. Susan broke the *window*. (pane)

In particular, the uses of *window* in (11)b. and (11)c. would make reference to different aspects of the 3D-representation stored for this object, specifying spatial knowledge about the various parts that compose a window. I return to consideration of this analysis in Chapter 5.

The notion of ‘semantic field features’ is central to Jackendoff’s analysis of verbal and prepositional polysemy (1992b: 37). The evidence for the existence of semantic field features is taken to be provided by the following kinds of examples (where the polysemous words *go*, *change*, *be*, *keep*, *from*, and *to* appear in different semantic fields):

- (12) a. *Spatial location and motion*
i. The bird went from the ground to the tree.
ii. The bird is in the tree.
iii. Harry kept the bird in the cage.
- b. *Possession*
i. The inheritance went to Philip.
ii. The money is Philip’s.
iii. Susan kept the money.
- c. *Ascription of Properties*
i. The light went/changed from green to red.
Harry went from elated to depressed.
ii. The light is red.
Harry is depressed.
iii. Sam kept the crowd happy.
- d. *Scheduling of activities*
i. The meeting was changed from Tuesday to Monday.
ii. The meeting is on Monday.
iii. Let’s keep the trip on Saturday.

In (12), each of the sets contains a sentence involving the verb *go* or *change*, expressing a change of some sort, the verb *be*, describing the resulting states of the changes in question, and the verb *keep*, which, in all of the sentences in which it occurs, denotes ‘the causation of a state that endures over a period of time’ (ibid.). Thus, Jackendoff sees the sets in (12) as having parallel conceptual structures, as schematised below (ibid. 38):

- (13) a. $[_{\text{Event}} \text{GO} ([], \left(\begin{array}{c} \text{FROM} ([]) \\ \text{Path TO} ([]) \end{array} \right))]$
 b. $[_{\text{State}} \text{BE} ([], [_{\text{Place}}])]$
 c. $[_{\text{Event}} \text{STAY} ([], [_{\text{Place}}])]$

What distinguishes them is a semantic field feature that specifies the field in which the EVENT or STATE is defined, and associates it with a particular inference pattern. This is supposed to explain how different occurrences of a word can license different inferences, and at the same time be similar in meaning. Consider again the *keep* examples in (12) above. Although the verb *keep* arguably expresses slightly different kinds of processes in each of the sentences in (12), *keep* denotes the ‘causation of a state that endures over time’ in all of them. The different uses are distinguished only by the kind of semantic field feature that it carries in each case.

Fodor (1998: 49-56) objects to the analysis of the intuitive polysemy of verbs such as *keep* in terms of semantic field features, claiming that it leads to infinite regress. His argument goes as follows: If the univocal meaning of *keep* is CAUSE A STATE THAT ENDURES OVER TIME, then the concepts contained in the definition, ‘CAUSE’, ‘STATE’, ‘TIME’, ‘ENDURE’, must also be univocal across semantic fields. Fodor asks if it is reasonable to take ‘CAUSE’ to be univocal in ‘CAUSE THE MONEY TO BE IN SUSAN’S POCKET’ and ‘CAUSE THE CROWD TO BE HAPPY’, and claims that Jackendoff faces difficulties no matter what his response is. It seems clear that he would avoid the first option, that is, to say that ‘CAUSE’ means different things in the two cases, as this would lead to regress (in order to explain the intuitive univocality of *keep* we need to provide a definition ‘X’ of CAUSE that explains its univocality across fields, and then a definition ‘Y’ of X that explains the univocality of X, and so on). The second option, to hold that CAUSE is a primitive, Fodor argues, must be as unappealing to Jackendoff as the first one, since it

implies that definitions are not required to explain a word/concept's univocality; CAUSE would be a case to the contrary. Thus, there would be no reason to believe that we need a definition to explain the univocality of *keep* across semantic fields either. On Fodor's view, *keep* is univocal because it always means KEEP, *cause* is univocal because it always means CAUSE, and, furthermore, "all words are univocal across semantic fields because semantic fields don't affect meaning" (1998: 52, Fodor's italics). I return to consideration of Fodor's analysis of *keep* in the next section.

However, Laurence and Margolis (1999: 59) do not think this critique of Fodor's need worry Jackendoff too much. In their view, Jackendoff should choose the second option, and admit that partial definitions are not necessary for a word to retain its meaning across semantic fields; in other words, that polysemy does not require partial decomposition. This does not mean that there may not be an explanatory advantage to postulating partial definitions in some cases (for instance, a partial definition may be invoked to explain the polysemy of *keep* but not of CAUSE), but this remains for Jackendoff to demonstrate. In my view, what is needed is a further specification of the conditions under which a semantic field feature applies to affect the meaning of a word (i.e. the contribution of context), as well as a more precise specification of what counts as a semantic field. If a new field feature can be postulated for each distinct sense of a word, Jackendoff's theory of polysemy would, in principle, not be much different from a sense enumeration account.

2.3 Conceptual atomism

Despite their significant differences, all the theories that we have looked at so far are versions of the *containment model* of conceptual structure (Laurence and Margolis 1999: 5): they take concepts to have their components (definitions, prototype structures, qualia structure, partial definitions) as proper parts. Conceptual atomism, on the other hand, takes most lexical concepts to be primitive; that is, to have no proper parts. There are (at least) two versions of this view, which I will consider (briefly) in this section: (i) 'pure' conceptual atomism, as recently defended by Fodor (1998; 2004a, 2004b; 2008), and (ii) atomism with meaning postulates (Fodor 1975; Fodor, Fodor, and Garrett 1975; Fodor et al. 1980; Fodor 1981; Sperber and Wilson 1986/1995; de Almeida 1999; Horsey 2006). The latter is a version of what is sometimes called the *inferential model* of

conceptual structure (Laurence and Margolis 1999: 5), on which there may be content-constitutive inferences that capture the logical properties of concepts, but which do not arise from definitions. I will also discuss the implications of conceptual atomism for polysemy representation.

Conceptual atomism claims that lexical concepts have no internal structure. On this view, the content of a primitive concept is determined by a constitutive link between the mind and the world; for Fodor (1998), this link is a *nomic* (lawful) relation that reliably locks the concept to a mind-external entity. For instance, the word *cat* encodes the concept CAT because there is a lawful relation between the property of being a cat and the mental entity CAT. In his early writings, Fodor argued (on the basis of the failure of the definitional approach to provide adequate definitions for, *inter alia*, natural kind terms and causative verbs), that analyticities such as *cat* → ANIMAL, *red* → COLOUR, *kill* → CAUSE TO DIE, etc. could be captured in terms of ‘meaning postulates’, that is, principles of inference that constitute one-way entailment relations between lexical concepts (Fodor 1975; Fodor, Fodor, and Garrett 1975; Fodor et al. 1980; Fodor 1981). For instance, on this approach, ANIMAL could be constitutive of the meaning of *cat* by standing in an inferential relation to the concept CAT, but would not be *contained* as part of its semantic representation. This made it possible to maintain an atomist account of lexical concepts while capturing intuitions about analytic relations between concepts.

On Fodor’s recent account, however, the content of a concept is seen as being exhaustively determined by the property it denotes (e.g. the content of CAT is exhausted by the property of being a cat, or ‘cathood’), hence the term ‘pure’ conceptual atomism (Fodor 1998; Fodor 2004a, 2004b; Fodor 2008). This position denies that a concept has its content in virtue of its relation to other concepts; in particular, it denies any role at all for necessary, or ‘analytic’ connections between concepts (for instance, between being a cat and being an animal). Fodor’s (1998) reason for rejecting meaning postulates comes mainly from Quine’s (1999 [1953]) convincing arguments against the possibility of drawing a principled analytic/synthetic distinction between those inferences that are content-constitutive of a concept and those that are not. Fodor’s recent approach has a clear advantage in that it sidesteps many of the problems associated with decompositional theories (in addition to Plato’s Problem and the analyticity issue, the

problem of ignorance and error does not arise: so long as CAT is appropriately connected to ‘cathood’ someone’s beliefs about cats are irrelevant to the possession of the concept).

A problem with pure conceptual atomism, however, is that, barring the issue of whether or not there are such things as analytic truths, it lacks an account of people’s clear intuitions about analyticities (Rey 1993). Horsey (2006) uses this as an argument to reinstate meaning postulates in an atomistic theory of lexical concepts. On his view, the notions of content constitutivity and analyticity must be pulled apart, the reason being that “it is perfectly possible for an inference to be content constitutive for a subject if the subject regards the inference as valid” (ibid.: 25). But, he notes, this does not require that the inference *is* in fact valid (i.e. analytic).³⁵ Thus, Horsey argues in favour of a psychological conception of meaning postulates, where these are seen as mentally represented inference rules, which provides no guarantee of their analyticity.

2.3.1 Radical nativism?

However, with or without (mentally represented) meaning postulates, a standard objection to conceptual atomism is its (apparent) commitment to radical nativism (e.g. Jackendoff 2002; Pinker 2007). A widely held view in cognitive science, defended also by Fodor himself (1981; Fodor et al. 1980; Jackendoff 1992b, 2002; Pinker 2007) is that primitive concepts are innate. The argument is that when we learn something, we construct it from previously known parts, using previously known means of combination. Since primitives have no parts, it follows that they cannot be learned; hence they must be innate. Both Jackendoff (2002) and Pinker (2007) take this to be a strong argument in favour of (partial) decomposition; in their view, it shows that lexical learning must also be a case of constructing word meanings from previously known parts. However, the conceptual atomist claim that most lexical concepts are primitive thus amounts to a claim that most lexical concepts are innate, including such unlikely candidates as TELEVISION, KIMONO, and NEURON. A possible way out of this problem, however, is to claim that it is not the concepts themselves that are innate; rather, we have innately-determined concept-acquisition mechanisms which are triggered by

³⁵ For instance, for some people, the inference from WHALE to FISH may be content constitutive of their concept WHALE, but the inference is not valid.

experience and which construct (atomic) concepts that go well beyond the available evidence. I cannot go into the details of how the mind-world relations that constitute conceptual content may come to obtain here, but some interesting proposals are given by Fodor himself (1998: Chapters 6 and 7), Margolis (1998), Sperber (1994b) and, in particular, Horsey (2006).

Another criticism against conceptual atomism is its ‘explanatory impotence’. (Laurence and Margolis 1999: 64). If we treat concepts as atoms, how are we, for instance, to make sense of typicality effects in categorisation, or make generalisations among inferential properties of different lexical items (e.g. the element shared by causative verbs) (Jackendoff 1992b: 50)?³⁶ The aspect of this criticism that is most important to our purposes here is that conceptual atomism has, it appears, no account of polysemy. In the remainder of this section I will look at this apparent problem.

2.3.2 Conceptual atomism and polysemy

In *Concepts* (1998: 53) Fodor explicitly states, “there is no such thing as polysemy”. Recall the discussion of *keep* in the previous section, where Fodor denied the apparent polysemy exhibited by this verb (which, on Jackendoff’s account, involved the supposed core meaning of the verb, ‘causation of a state that endures over time’, carrying different ‘semantic field features’, e.g. spatial location and motion, possession, ascription of properties, scheduling of activities). The examples are repeated in (14):

- (14) a. Harry *kept* the bird in the cage.
b. Susan *kept* the money.
c. Sam *kept* the crowd happy.
d. Let’s *keep* the trip on Saturday.

On Fodor’s view, *keep* is univocal; it expresses nothing more than ‘keeping’ in all of its uses above. However, recognising the intuitive polysemy of *keep*, Fodor (1998: 54) suggests that the apparent differences in meaning between (14)a.-d. do not have to arise from its being polysemous. What distinguishes, for instance, *Susan kept the money* in

³⁶ In fact, conceptual atomism is fully compatible with the existence of prototypes, if these are seen as being represented as part of an individual’s general world knowledge about conceptual denotations, and not as part of their lexical-semantic representations.

(14)b. and *Sam kept the crowd happy* in (14)c. is that in (14)b., the ‘keeping’ relation holds between the NP and the money, and in (14)c. it holds between the NP and the crowd’s being happy. These are clearly different things; hence the difference in meaning between (14)b. and (14)c.

As a possible objection to this univocality view, Fodor considers the case where a language A has a single unambiguous word, which may be translated into either of two words in another language B, depending on the context, and asks: Shouldn’t this be taken as evidence that the language A word is polysemous? No, he argues. He gives the following example to back up this claim (Fodor 1998: 55):

Suppose English has two words, ‘spoiled’ and ‘addled,’ both of which mean *spoiled*, but one of which is used only of eggs. Suppose also that there is some other language which has a word ‘spoilissimoed’ which means *spoiled* and is used both of spoiled eggs and of other spoiled things. The right way to describe this situation is surely *not* that ‘spoiled’ [sic.] is *ipso facto* polysemous. Rather the thing to say is: ‘spoiled’ and ‘addled’ are synonymous and are (thus) *both* correctly translated ‘spoilissimoed’.³⁷

Although the situation Fodor describes is clearly possible (but would probably strike most speakers of more than one language as being rather exceptional), it does not take much by way of cross-linguistic investigation to see that the conclusion that the two language B words have to be synonymous does not hold. For instance, consider the Norwegian translations of the *keep* sentences in (14) below:

- (15) a. Harry *holdt* fuglen i buret. [Harry *kept* the bird in the cage.]
b. Susan *beholdt* pengene. [Susan *kept* the money.]
c. Sam *underholdt* publikum. [Sam *kept* the crowd happy.]
d. La oss *fastholde* turen på lørdag. [Let’s *keep* the trip on Saturday.]

In the Norwegian sentences, the processes expressed by *keep* in the English equivalent sentences are all expressed by different verbs. So this would be an actual instance of Fodor’s scenario above, where language A (English) has a single (‘unambiguous’) word, which translates into several different words in language B (Norwegian), in lack of a single language B word that covers the meaning of the language A word. However, on

³⁷ According to the scenario we are presented with it is the word *spoilissimoed* that could be argued to be polysemous, not *spoiled*. I take it that Fodor has made a slip here.

no account would the meanings of the Norwegian verbs *holde*, *beholde*, *underholde* and *fastholde* be claimed to be synonymous.³⁸ Rather, it would be correct to say that their meanings are related (in fact, they all denote the ‘causation of a state that endures over time’),³⁹ but there is clearly a different kind of process involved in each case.⁴⁰ This appears to give Jackendoff’s analysis of *keep*, where such cross-linguistic differences may be predicted on the basis of the different meanings this verb may be used to express, an advantage over Fodor’s, where the assumption that *keep* expresses nothing more than KEEP renders such differences unexplainable.⁴¹

Contrary to Fodor, I believe that there is such a thing as polysemy, if what is meant by this is that words can be used to express different meanings in different contexts. There are two ways in which this situation could be described in Fodor’s framework: (i) in the case where the senses are conventionalised (e.g. *paper*, meaning ‘sheet of paper’, ‘newspaper’), they could be represented as separate lexical entries with the same linguistic form (e.g. PAPER₁, PAPER₂), on a par with homonymous senses (cf.

³⁸ This includes Fodor’s own account, where the words *holde*, *beholde*, *underholde* and *fastholde* would encode the unstructured concepts HOLDE, BEHOLDE, UNDERHOLDE, FASTHOLDE.

³⁹ All the verbs contain the stem *holde*, which corresponds to English *hold/keep*.

⁴⁰ Note that Fodor could not reasonably respond to this that the differences in meaning between the Norwegian verbs in (15) are all governed by the object (bird in cage, money, crowd), that is, that all of them simply mean ‘causation of a state that endures over time’ (i.e. encode the concept HOLDE) with different usage constraints on each one (just as *addled* has the usage constraint that it can only be applied to eggs). The Norwegian verbs clearly encode different concepts and have a range of other uses than those in (15) above.

⁴¹ In fact, such cross-linguistic differences illustrate a further problematic aspect of conceptual atomism, which is seldom discussed: As we have seen, on Fodor’s account, *keep* means KEEP, and the Norwegian verbs *holde*, *beholde*, *underholde*, and *fastholde* would have to be treated in the same way; as giving access to the lexical concepts HOLDE, BEHOLDE, UNDERHOLDE, FASTHOLDE. Given conceptual atomism, the meaning of *keep* and its Norwegian cognates would bear no relation to each other (e.g. the English concept KEEP and the Norwegian concept HOLDE would be no more similar in meaning than, e.g., KEEP and LOSE are in English). It thus seems to follow from the theory that speakers of different languages would have entirely different conceptual repertoires, depending on the respective vocabularies of their languages. This would imply a linguistic determinism very much in opposition with Fodor’s own ‘Publicity Constraint’ on a theory of concepts, on which concepts are seen as *public* entities, which people share: “all sorts of concepts (...), are ones that all sorts of people, under all sorts of circumstances, have had and continue to have” (1998: 29). It is hard to see how this claim could be reconciled with speakers of different languages having different conceptual repertoires.

It may be, however, that Fodor’s Publicity Constraint is too strong; perhaps we need to allow for some linguistically-determined differences in conceptual repertoires, cf. Gentner and Goldin-Meadow (2003). In any case, it seems that what Fodor would need is to give an account of how we compare and match lexical meanings in moving between different languages (otherwise, second language acquisition would be an even more laborious affair on this account than it already is); and if this cannot be done on the basis of atomic lexical concepts (which, as far as I can tell, it can’t), it has to be described in terms of some other level of representation (e.g. in terms of real-world knowledge about the denotations of lexical concepts).

Fodor and Lepore 2002); (ii) in all other cases, there would be a single encoded meaning of the word, and the differences in meaning perceived between its uses in different contexts could be treated as instances of pragmatic adjustment of the encoded meaning. Thus, pairing the atomist account of lexical semantics with an adequate theory of pragmatics would help solve the problem posed by *keep* above. If the different meanings of *keep* are seen as instances of pragmatic adjustment of the concept KEEP, then Fodor would be right to say of *keep* that its *semantic encoding* amounts to nothing more than KEEP. The cross-linguistic differences between English and Norwegian observed in (15) could be explained as an instance where a set of meaning distinctions which are derived pragmatically (on the basis of *keep*) in English, are lexicalised in Norwegian. That there is this kind of asymmetry between the vocabularies of different languages is hardly controversial (cf. Falkum 2007).

2.4 Psycholinguistic perspectives on polysemy

In the previous sections, we have been looking at a range of theoretical approaches to the nature of word meaning and polysemy. The central issue has been whether word meanings are mentally represented in terms of smaller units of meaning (lexical decomposition), and if so, what the nature of these units is, or if they are represented in terms of unstructured concepts (conceptual atomism). A further related issue has been to what extent the various approaches to lexical semantics are able to account for polysemy, and whether the claim that lexical decomposition is required for a proper account of polysemy is warranted (which, I think we may conclude at this stage, it is probably not).

In this section, I will approach these issues from a different angle by discussing some recent psycholinguistic investigations into the representation of polysemy. The main focus of the psycholinguistic literature on this topic has been the overall structure of the mental lexicon; how it represents polysemy compared with homonymy (cf. Chapter 1). On the assumption that the meanings of a homonymous word, e.g. *bank* ('financial institution', 'riverside'), are represented as separate entries in the lexicon, psycholinguistic studies have sought to find out whether this is also true of polysemous words, or if only a single meaning is represented and the others derived (either lexicon-internally or pragmatically). This question is largely independent of whether or not

represented meanings are decompositional, but has a bearing on the adequacy of the accounts of polysemy considered thus far. This section will be devoted to a discussion of some of the studies addressing this issue.

2.4.1 Core meanings or separate entries?

Empirical studies of polysemy representation are few in number and have, to some extent, yielded conflicting results. While the main body of experimental evidence points to a differential representation of homonymy and polysemy (Frazier and Rayner 1990; Williams 1992; Pickering and Frisson 2001; Klepousniotou 2002, 2007; Klepousniotou and Baum 2007; Klepousniotou, Titone, and Romero 2008; Beretta, Fiorentino, and Poeppel 2005; Pylkkänen, Llinás, and Murphy 2006; Brown 2008), some studies have found no difference between the processing of homonymous and polysemous lexical items (Klein and Murphy 2001, 2002). In a set of recent publications, Klepousniotou and colleagues (2002; Klepousniotou and Baum 2007; Klepousniotou, Titone, and Romero 2008) have suggested that the diverging results could be due to the type of polysemous stimuli used, and further that degree of sense relatedness may have an effect on processing. I consider some of the studies that have given rise to this debate in more detail here.

Klein and Murphy (2001, 2002) ran a series of experiments designed to investigate whether polysemous words are represented in terms of a common core meaning, from which the various senses are derived in context (the 'core meaning' or 'monosemy' view, cf. Caramazza and Grober 1976; Ruhl 1989; Pustejovsky 1995a), or if they are represented in the same way as homonymy, involving separately stored senses. Klein and Murphy (2001) used memory performance and lexical decision as measures of the representation of polysemous senses. Their experiments involved two phases. In phase 1, subjects were presented with a set of phrases which they were told to study for a later memory test. The phrases included polysemous words that were biased in interpretation toward one of the senses (e.g. the phrases *daily paper* and *liberal paper* biased the 'newspaper' sense of *paper*, while *wrapping paper* and *shredded paper* biased the 'sheets of material' sense). In phase 2, subjects viewed similar phrases in which one of the two words was capitalised: In the memory task, they were asked to judge whether they had seen the word in the previous task; in the lexical decision task they were asked

to make a sense/nonsense judgement (a 'sensicality judgement'). There were three experimental conditions: (i) 'repeated items'; the capitalised polysemous word occurred in the same phrase as in phase 1 (e.g. [*daily paper*], *daily PAPER*); (ii) 'consistent sense phrases'; the capitalised word occurred in a different phrase but with the same sense (e.g. [*daily paper*], *liberal PAPER*); and (iii) 'inconsistent sense phrases'; the capitalised word occurred in a phrase biasing the alternative sense (e.g. [*daily paper*], *shredded PAPER*). Klein and Murphy (2001) found that in the memory task, the repeated items were the most accurately evaluated, and the consistent sense phrases were more accurately evaluated than inconsistent sense phrases. Similarly, in the lexical decision tasks, subjects were faster and more accurate in their sensicality judgements in the consistent condition than in the inconsistent condition. Klein and Murphy took these results to suggest that polysemous senses are stored separately in the lexicon. If polysemous senses had a common core meaning, there should be no difference in memory performance or in reaction times for the consistent and inconsistent senses.

Klepousniotou (2002) conducted an experiment that aimed at testing the specific predictions about the processing and representation of polysemy made by sense enumeration lexicons (SEs) (e.g. Katz and Fodor 1963; Katz 1972; Weinreich 1966) and generative lexicons (e.g. Pustejovsky 1995a). As we saw in section 2.2.1, SEs list each meaning of an ambiguous word under a single entry. They contain no distinction between homonymy and polysemy, and thus predict that all ambiguous words should be processed in the same way. Generative lexicons list homonyms separately, while the senses of a polysemous word are derived compositionally on the basis of a single semantic representation and the operation of generative mechanisms (cf. section 2.2.3). This approach predicts a differential processing depending on the type of ambiguity. The study used a cross-modal priming lexical decision task. Participants heard a sentence that primed one of the meanings of an ambiguous target word (not incorporated in the sentence primes), e.g. 'He cannot read small print' (*eye*, 'body part', 'hole in a needle'). At the offset of the sentence, either a real word (the ambiguous target or an unrelated control, e.g. *mop*) or a phonologically legal nonword (e.g. *scling*) was presented on the computer screen for lexical decision. The ambiguous words were divided into four conditions: (i) homonymy condition (e.g. *coach*); (ii) metaphor condition (e.g. *eye*); (iii) metonymy (or count-mass polysemy) condition (e.g. *turkey*,

'animal'/'meat');⁴² and (iv) name condition (e.g. the producer/product metonymic extension, such as *Dalí*, 'the artist'/'a painting by that artist').⁴³ Priming was measured as time to recognise the primed ambiguous word relative to the unrelated control words. The results showed a significant difference between reaction times for conditions (i) and (iii) – the latter were responded to significantly more quickly than the former – indicating a differential processing of homonymy and count-mass polysemy. No significant differences were found between the homonymy and metaphor conditions, or between the metaphor and the count-mass polysemy conditions.⁴⁴ The results, in particular the faster reaction times and greater priming effects that were observed for the count-mass polysemy than for homonymy, were interpreted as supporting the hypothesis that type of ambiguity has an effect on processing, and to rule out a SEL account of homonymy and polysemy (predicting that all types of ambiguity should be processed in the same way), and to be consistent with the generative lexicon approach (associating homonymy with separate entries and polysemy with single entries, and where extended senses are generated by means of lexical rules).⁴⁵ Klepousniotou's (2002) results were thus directly conflicting with those obtained by Klein and Murphy (2001), by suggesting a processing difference between homonymy and polysemy (of a certain kind). However, it is hard to see that any specific conclusion concerning the adequacy of the generative lexicon theory was warranted by them; they would be consistent with any theory that distinguishes between homonymy and polysemy representation.⁴⁶ Also, the

⁴² Klepousniotou (2002) treats such count-mass alternations as a form of 'metonymic polysemy'. I am not certain that this is the right way to characterise them. See my Chapter 5 for a detailed discussion of systematic polysemy and the count-mass distinction.

⁴³ The choice to distinguish between 'metaphorically ambiguous words' and 'metonymically ambiguous words' (a distinction originally drawn by Apresjan 1974) reflects the assumption that the kind of process motivating the polysemy has a bearing on processing. Although this may be correct at a general level, it is possible that within each of these categories, there may be differences between individual words with regard to how they are stored in the lexicon (and, consequently, with regard to how they are processed). For instance, while the metaphorically generated polysemy of *foot* ('body part', 'base or bottom of something') is conventional and likely to be stored in the lexicon of many speakers, so could the meanings of *chicken* ('animal', 'meat'). If this were the case, we would not predict any processing differences between them.

⁴⁴ As to the 'name' condition, it had the largest error rate, and did not show any priming effects compared to controls. Klepousniotou attributes this result to a different cognitive process being involved in the recognition of proper names compared to the one(s) involved in the recognition of regular nouns. This seems correct. I return to consideration of the metonymy of proper names in Chapter 6.

⁴⁵ However, it is not entirely clear that the generative lexicon account gets much support from the results, as it might take time for the lexical rules to operate. See footnote 60.

⁴⁶ For instance, a theory where polysemous senses are derived via pragmatic (inferential) processes.

lack of a significant difference between the homonymy and metaphor conditions could be taken to indicate that certain types of polysemy do, in fact, behave similarly to homonymy.

Klepousniotou, Titone, and Romero (2008) sought to uncover the reasons behind the conflicting results obtained by Klepousniotou (2002) and Klein and Murphy (2001, 2002). They hypothesised that the type of polysemous words studied could be a contributing factor, as Klein and Murphy (2001) had deliberately chosen stimulus words that had fairly distinct senses. Klepousniotou et al. (2008) employed the sense decision task of Klein and Murphy (2001). The experimental stimuli consisted of three groups of words: low overlap polysemes (e.g. *control panel/advisory panel*), moderate overlap polysemes (e.g. *laser beam/wood beam*), and high overlap polysemes (e.g. *marinated lamb/baby lamb*), divided into sub-groups of dominant and subordinate senses.⁴⁷ Sense dominance was manipulated as a between-subjects factor; one group of subjects were presented only with pairs of dominant senses, and the other group only with pairs of subordinate senses. Klepousniotou et al. (2008) found that for dominant senses, moderate and low overlap polysemes were responded to faster in a consistent context (e.g. *laser beam/shining beam*) than in an inconsistent context (e.g. *laser beam/balance beam*), but no such difference was observed for high overlap polysemes. In contrast, for subordinate senses, there were no significant differences found in reaction times between the polysemy categories; participants were significantly faster in the consistent contexts regardless of the degree of sense relatedness. Klepousniotou et al. (2008) took their results to suggest that degree of sense relatedness of polysemous words affects their processing, and thus, reveals differences in how they are mentally represented. More specifically, their results suggested that 'high overlap' words are processed differently from 'moderate-' and 'low overlap' words. This is an interesting finding, which indicates that not all polysemous words need be represented in the same way. As Klepousniotou et al. (2008) point out, it is compatible with their results that the senses of 'high overlap' words (e.g. *lamb, book, newspaper*) are represented as single lexical entries (e.g. in the form of a 'core' meaning) while the senses of 'moderate-' and 'low' overlap words (e.g.

⁴⁷ The stimuli were grouped on the basis of a pre-experimental norming study where subjects had rated polysemous senses in terms of degree of relatedness, and of dominance-subordinance.

panel, patient, beam, orange) are represented in a similar way as homonymy (i.e. separately), as originally suggested by Klein and Murphy (2001, 2002).

Beretta, Fiorentino, and Poeppel (2005) conducted an MEG⁴⁸ study to examine the neural correlates of homonymy and polysemy. They measured changes in MEG recordings during a visual lexical decision task, which included both homonymous and polysemous words. The aim of their study was to test the competing separate entries and single entry hypotheses about polysemy, using a more fine-grained on-line measure of lexical processing than that used in behavioural studies, which had given rise to the conflicting results of Klein and Murphy (2001) and Klepousniotou (2002). The 'M350', which is a response component produced in the left superior temporal cortex at 300-400 ms after the onset of a visual word, is hypothesised to be an index of initial lexical activation. Beretta et al. (2005) predicted that if the separate entries account was correct, polysemy and homonymy should behave in the same way at the initial stage of lexical access shown by the M350. However, if the single entry account for polysemy was correct, homonymy should yield distinct processing behaviour from polysemy.

Participants were presented with a word (a homonym, a word with multiple senses, few senses, or a single sense) on a screen and asked to make a lexical decision as quickly and accurately as possible. In this behavioural task, homonyms were responded to more slowly than monosemous words, while multiply polysemous words were responded to more quickly than words with few senses. Thus, homonymous words appeared to slow access (possibly due to competition between the senses) relative to words with one meaning, while words with many polysemous senses appeared to speed access relative to words with few senses. In this way, the behavioural results supported the single entry account of polysemy and the separate entries account for homonymy (thus contradicted the results obtained by Klein and Murphy (2001)).

The MEG results showed that homonyms elicited a later M350 peak latency than words with a single meaning, while polysemous words elicited an earlier peak than words with few senses. These results mirrored the behavioural results obtained at around 600-650 ms, and showed that the effect also held prior to the reaction time

⁴⁸ MEG (Magnetoencephalography) is an imaging technique used to measure the magnetic fields produced by electrical activity in the brain.

judgements, at around 300 milliseconds earlier. Thus, Beretta et al. (2005) took their study to support a single entry account of polysemy and a separate entries account of homonymy. The competition between homonyms could be what caused the delay in the M350 latency for these items.

Beretta et al.'s (2005) study had an additional finding: It showed an apparent processing *advantage* for multiply polysemous words compared with words with few senses. This finding replicated an earlier result obtained by Rodd, Gaskell and Marslen-Wilson (2002), who had found a significant advantage for many senses in both response time and accuracy compared with monosemous words. Rodd et al.'s study (2002) had been designed to test the so-called 'ambiguity advantage'; the finding that visual lexical decision is faster for semantically ambiguous words than for monosemous words (reported by Rubenstein, Lewis, and Rubenstein 1971; Jastrzembski and Stanners 1975; Jastrzembski 1981; Millis and Button 1989; Borowsky and Masson 1996, among others). In three lexical decision experiments, Rodd et al. (2002) showed that the ambiguity advantage was entirely due to polysemy, and that homonymous words in fact delayed access. They suggested as a possible explanation of the multiple senses advantage that words with many senses could be semantically richer than words with few senses. The additional semantic information may produce more stable representations, which in turn lead to faster recognition times. Another possible explanation is that the advantage is due to a so-called 'context availability effect': This has been proposed as an explanation of the faster processing of concrete compared with abstract words (Schwanenflugel, Harnishfeger, and Stowe 1988). The hypothesis is that comprehension is facilitated by the addition of contextual information (retrieved from the hearer's encyclopaedic knowledge or from the stimulus environment) to the materials that are to be understood. Since contextual information can usually be accessed more easily for concrete than abstract words, the former are processed with greater ease than the latter (also when controlled for effects of frequency and familiarity). Relating this to polysemy, it may be that words that have many senses, and are used in a wider range of contexts, have contextual information more accessible to them, and that this is what causes the

processing advantage. Note that this explanation does not imply that this information needs to be part of the semantic representation of the polysemous word.⁴⁹

Klepousniotou and Baum (2007) replicated Rodd et al. (2002) and Beretta et al.'s (2005) results, and found, in addition to a significant processing advantage for polysemous words compared with homonymous words, that count-mass polysemous words (e.g. *turkey*) were recognised significantly faster than metaphorically motivated polysemy. This result adds to the evidence suggesting that not all types of polysemy are necessarily processed (or represented) in the same way.

Pylkkänen, Llinás, and Murphy (2006) conducted a MEG study in a reinvestigation of the core meaning and the separate entries hypotheses of polysemy representation. They started out with the following hypotheses: Given that semantic relatedness is often facilitatory, and phonological relatedness inhibitory (similar-sounding words may compete in recognition), on the separate entries hypothesis, different senses of a polysemous word (*green book-interesting book*) should be expected to prime each other less (due to form-based inhibition) than purely semantic controls (*green novel-interesting book*). On the core meaning proposal, where polysemous senses should share a morphological root (cf. Halle and Marantz 1993) polysemous senses should not show the effects of sound and meaning relatedness that are found in competing words. Rather, they should show repetition priming (of the morphological root), and, potentially, sense competition. For instance, seeing the word *paper* used in two different senses should activate a single lexical entry, and thus reveal repetition effects. Pylkkänen et al. (2006) employed a priming paradigm, where polysemy was contrasted with homonymy and semantic relatedness. The stimuli had the following design (ibid. 99):

⁴⁹ A further, related explanation of the multiple sense advantage may be that it is a direct result of using a lexical decision task in which the words are presented in isolation. The great number of contexts in which a polysemous word may be adequately used may result in the development of a relatively context-independent representation of the word. As words with few senses may be used in a more restricted range of contexts, it may be that their access is more dependent on the word appearing in the appropriate context. Making a lexical decision would presumably be more difficult for these words than for the polysemous words when presented in the absence of a context.

	<i>Related Prime</i>	<i>Unrelated Prime</i>	<i>Target</i>
Homonymous	river bank	salty dish	savings bank
Polysemous	lined paper	military post	liberal paper
Semantic	lined paper	clock tick	monthly magazine

This design was similar to that used by Klein and Murphy (2001, 2002), who, as we saw, found that using a word in the same sense twice led to faster judgements than switching senses. This result was replicated by Pylkkänen et al. (2006) in a behavioural study using the above materials. Thus, prior to the MEG study, their behavioural evidence still pointed in favour of the separate entries hypothesis.

As already mentioned, the M350 component has been taken to reflect initial stages of lexical activation. It is also argued to be the first MEG component in response to visual words that is sensitive to repetition (Pylkkänen et al. 2000). Furthermore, it has been claimed that the M350 tracks morphological constituent frequency rather than whole-word frequency in compound processing, suggesting that it is an index of morphological root access (Fiorentino and Poeppel 2004). In view of this, Pylkkänen et al. (2006) expected polysemy, if involving distinct sense representations but morphological root sharing, to show shorter M350 peak latencies for targets preceded by sense-related primes, compared to unrelated controls. On the other hand, if polysemy is represented as homonymy, they expected the M350 effects obtained in the polysemy condition to be explainable as a combination of the effects elicited in the homonymy and semantic conditions.

The M350 results showed priming for the semantic targets, with M350 sources peaking earlier for related targets than for unrelated controls. Homonymous targets elicited a later M350 peak than unrelated controls, indicating competition. However, no such delay was found for the polysemous targets at M350, whose amount of priming was comparable to the priming observed for the semantic targets. Also, polysemous targets peaked earlier than unrelated targets. The results contradicted the previous behavioural data obtained by Klein and Murphy (2001, 2002) and pattern, like Beretta et al.'s (2005) findings, in favour of a single entry account of polysemy.

Taken together, the studies considered in this section indicate that type of lexical ambiguity affects processing, and, by inference, representation. Assuming that

homonymous words are represented separately, this does not seem to be the case for at least certain kinds of polysemy. As we have seen, the results are to some extent conflicting. In large part, this seems to be due to differences in the kind of polysemous stimuli used, as pointed out by Klepousniotou and her colleagues (Klepousniotou 2002; Klepousniotou and Baum 2007; Klepousniotou, Titone, and Romero 2008). For instance, Klein and Murphy (2001) used polysemous words that had fairly distinct senses (e.g. *liberal paper*, *shredded paper*), giving them a ‘homonymic character’, and found no differences in processing between homonymous and polysemous words. When dividing the polysemous stimuli into different categories, Klepousniotou et al. (2008) found a significant difference between the processing of homonymy and count-mass polysemy but no significant difference between homonymy and metaphorically-based polysemy. Overall, the results are consistent with the view that not all polysemous words are represented in the same way. Rather than storing only a single core meaning or a total list of all the possible distinct senses, the lexicon may store some reasonable number of senses. In some cases, the senses of a polysemous word may have become conventional to the extent that they are not obviously distinct from homonymy (as in many cases of metaphorical polysemy, e.g. *foot*). In other cases, only a single meaning may be stored and the other meanings contextually derived (either lexicon-internally or pragmatically). This view of polysemy representation is directly incompatible with a Katz-style sense enumeration lexicon that lists homonymous and polysemous senses of a linguistic form together under a single lexical entry. Nor does it mesh well with Lakoff (1987) and Brugman’s (1988) ‘full-specification’ approach, which takes all the senses associated with a polysemous lexical form to be represented as part of a semantic network. However, the experimental evidence is, in principle, compatible with several of the other lexical semantic accounts discussed in the first part of this chapter, including the generative lexicon theory (Pustejovsky 1995a), Jackendoff’s conceptual semantics (1992b, 2002), and conceptual atomism (despite Fodor’s (1998) claims to the contrary). Given the lack of more specific empirical evidence, a choice between these accounts has to be made largely on theoretical grounds.⁵⁰

⁵⁰ Recently there has been a string of publications investigating the processing of so-called type shifting constructions (e.g. *begin a memo*, *finish a play*, *enjoy an ice-cream*), discussed in connection with Pustejovsky’s (1995) generative lexicon account in section 2.2.3. These studies address the issue of

2.5 Conclusion

In this chapter, we have looked at the problem of polysemy representation. Two central questions were: (i) Does an adequate account of polysemy representation require lexical decomposition? (ii) Are polysemous lexical forms represented as separate lexical entries (sense enumeration) or as a single entry (core meaning)?

Starting with the second question, we have seen that the experimental evidence, although far from conclusive, indicates that not all instances of polysemy are necessarily represented in the same way. While conventionally polysemous lexical forms might be represented in a similar way to homonymous forms, less established cases of polysemy seem to be represented as a single entry with other senses pragmatically derived in on-line interpretation. This view of polysemy representation does not support the Katz-style sense enumeration lexicon or Lakoff and Brugman's 'full-specification' approach, but is, in principle, compatible with both the rich lexical entries approach of Pustejovsky's generative lexicon theory or Jackendoff's conceptual semantics, and with the much leaner approach of conceptual atomism.

With regard to the first question, I think we can conclude that lexical decomposition is not required for an adequate account of polysemy representation. Recall that the main reason for postulating complex lexical representations for polysemous forms is that it is seen as a way to capture relations between their senses. However, in the course of this chapter we have seen that it is not always clear that this is,

whether the processing of such constructions involves the application of a lexicon-internal generative mechanism to a semantically complex lexical entry (McElree et al. 2001; Traxler, Pickering, and McElree 2002; Pickering, McElree, and Traxler 2005; McElree, Frisson, and Pickering 2006), or the operation of a general pragmatic-inferential process to atomic concepts (de Almeida 2004; de Almeida and Dwivedi 2008). Some psycholinguistic studies are taken as providing support for so-called 'type shifting effects' (i.e. an extra processing load hypothesised to result from the operation of a type coercion mechanism in the processing of type shifting constructions). For instance, McElree et al. (2001) found, on the basis of a self-paced reading experiment, that type shifting constructions (e.g. 'The secretary began the memo before the annual sales conference') induced longer reading times (hence were associated with greater processing complexity) compared to preferred constructions (e.g. 'The secretary typed the memo before the annual sales conference') and non-preferred constructions (e.g. 'The secretary read the memo before the annual sales conference'). However, this result was not replicated by de Almeida (2004). In two self-paced reading experiments, type shifting effects were not obtained when no context was provided before the sentence, or with sentences embedded in contexts that specified the nature of the activity performed over the complement VP. In a recent survey of empirical studies of the processing of type shifting constructions, de Almeida and Dwivedi (2008) claim that the empirical results favouring a type coercion mechanism are, at best, weak, and that even if there exist such 'coercion' effects, there are other possible explanations for them (e.g. they could result from verbal gaps in the logical form of such constructions, which require contextual saturation). I return to this possibility in Chapter 4 (section 4.3.1).

in fact, achieved by all decompositional theories. For instance, in Katz's 'classical' theory of semantics, one of whose primary aims was to explain semantic relations and contrasts between word meanings (e.g. synonymy, antonymy, analyticity, entailment, etc.), it appears that the distinction between polysemy (related senses) and homonymy (unrelated senses) cannot be so easily drawn by reference to underlying semantic markers. It is also unclear how verbal polysemy can be handled in Pinker's theory of partial verbal decomposition. However, relations between polysemous senses are captured by the lexical network theory proposed by Lakoff and Brugman, and to a certain extent also by Pustejovsky's generative lexicon theory. A problem with both these approaches was that they appeared to ignore (or at least, drastically downplay) the contribution of pragmatics in the derivation of word senses.

Among the decompositional accounts discussed here, Jackendoff's theory of partial definitions seems, in my view, to be the one best suited to handle polysemy (e.g. the polysemy of physical object words such as *window* could be treated in terms of 3D-representations contained in their lexical representations; verbal and prepositional polysemy could be handled in terms of his theory of semantic field features). However, as Fodor has pointed out, this does not mean that Jackendoff has shown that polysemy requires lexical decomposition.

Another question is whether all of our intuitions about semantic relatedness need be captured at the level of *lexical semantic* representation. For instance, as discussed in the Introduction, the reason it is difficult to draw a clear-cut distinction between homonymy and polysemy is that it is not entirely clear what we mean when we talk of two senses being related (is it to do with etymology or speaker intuitions?). Furthermore, since it does not seem to be the case that our intuitions about relatedness and unrelatedness of senses have any bearing on the way in which we use and understand words, they can probably not be taken as a direct reflex of the way in which word meanings are represented in our linguistic systems.

To sum up, I do not think that the existence of polysemy provides a very strong argument in favour of a decompositional approach to word meaning. With regard to the possibilities of giving an account of polysemy representation, I cannot see why a conceptual atomist approach, if paired with an adequate theory of pragmatics, should not be at least as able as decompositional approaches to provide an account of

polysemy. In Chapter 4, I show how this can be achieved within the context of Sperber and Wilson's (1986/1995) relevance theory. Adding to this are the range of theoretical problems associated with decompositional accounts in general, some of which have been considered in this chapter (e.g. the incompleteness of many decompositions, the vagueness of many concepts, the fact that we may be ignorant or mistaken about the properties we take the instances of a concept to have, etc.). Conceptual atomism has a clear advantage in that it avoids many of the problems associated with decompositional approaches, although, as we saw in section 2.3, it is associated with some problems of its own.

An implicit assumption of the decompositional accounts that we have considered in this chapter is that a theory of lexical semantics should be able to explain much of what we would consider to be aspects of language meaning in use. However, I think it is an open question how much of the interpretive work should be attributed to the linguistic system itself, given that we have an independently motivated pragmatic interpretation system, automatically activated by verbal utterances, which is capable of rapidly generating new senses in contexts. This question will be addressed in detail in the next chapters.

Chapter 3

LANGUAGE, THEORY OF MIND, AND INFERENTIAL COMMUNICATION: SOLVING THE PROBLEM OF POLYSEMY MOTIVATION?

3.1 Introduction

In *The Big Book of Concepts*, Murphy (2002: 405) identifies (i) and (ii) as the main questions that a theory of polysemy has to answer:

- (i) How do we represent the meanings of a polysemous word so that we are able to understand its different uses?
- (ii) Why and how is it that word meanings get extended to have these different senses?

The first question, concerning the *representation* of polysemous words, was addressed in Chapter 2, and is, together with the definitional problem discussed in Chapter 1, the one that lexical semanticists, and psychologists to some degree, have been mainly concerned with. As to Murphy's second question about polysemy, the *why* question has to be distinguished from the *how* question. The latter has been widely studied, in particular within the cognitive linguistics tradition (e.g. Lakoff and Johnson 1980; Lakoff 1987; Brugman 1988; Taylor 1989/2003; Dirven and Pörings 2003; Tyler and Evans 2003), computational semantics (e.g. Kilgarriff 1992; Pustejovsky 1995a; Copestake and Briscoe 1996) and in the relatively new field of lexical pragmatics (e.g. Recanati 1995, 2004; Carston 1997, 2002b; Blutner 1998, 2002, 2004; Wilson and Carston 2006, 2007). I will address the question of *how* word meanings get extended in detail in the next chapter. However, the question as to *why* it is that word meanings get extended, or, more fundamentally, what it is about our language systems, specifically their lexical component, that makes them so susceptible to polysemy, has not received nearly as much focus in the literature.⁵¹

⁵¹ Lakoff (1987), as discussed in Chapter 2, sees polysemy as a consequence of our cognitive organisation, where word meanings are a subtype of 'radial category' in which different concepts are organised with respect to a prototype, and extended by convention on the basis of general principles (e.g. metaphor). Although, for reasons stated in that chapter, I do not think that his account of polysemy is tenable, I agree

This question, which we called the *problem of polysemy motivation* in Chapter 1, will be the topic of this chapter. What makes our languages exhibit polysemy? What is the motivation for it? Is it fundamentally linguistic, communicative or cognitive, or some combination of these? Why do we rather use the same word to describe a set of different things than have a distinct word for each sense? I will take the ‘polysemy paradox’ as a starting point. As I mentioned in the Introduction, polysemy (in fact, ambiguity quite generally) represents an almost insurmountable challenge for computational models of natural language processing and translation, but is largely unproblematic from the point of view of human communication. A pragmatic theory that is able to explain how and why this is so, is also likely to shed light on the issue of what motivates the proliferation of polysemy in natural languages in the first place. Most existing theories of polysemy (with the notable exception of the cognitive linguistic paradigm) either explicitly or implicitly take polysemy to be a fact about language, that is, a phenomenon requiring an analysis in terms of the functioning of the linguistic system, with little, or no attention devoted to how it arises in communication. In this chapter, I question this mainstream lexical semantic view, and argue that although there is, needless to say, an important linguistic aspect to polysemy, it is, at the deepest level, a consequence of how communication works; it results from our capacity to infer speaker meanings on the basis of not fully determining linguistic evidence.

The chapter is organised as follows. In section 3.2, I present a model of communication, *the code model*, which has been the basis for many linguistic approaches over the years. I show its implications for the interpretation of polysemy, exemplified by Pustejovsky’s generative theory, and argue that, faced with the challenge of *linguistic underdeterminacy* (the systematic gap between speaker meanings and linguistic meanings), it does not provide a workable account. In section 3.3, I consider the issue of linguistic underdeterminacy in more detail, discussing two possible views on the relation between linguistic meanings and the propositions that are communicated by them, both of which have important consequences for an account of polysemy. In

with Lakoff in that polysemy is dependent on the cognitive processes that enable us to construct, *inter alia*, metaphorical and metonymic meanings. While my focus in this chapter is the communicative aspect of polysemy, which I take to be fundamental, I will consider the specific processes of meaning extension responsible for the production of polysemy in Chapters 4, 5, and 6.

section 3.4, I present a viable alternative to the code model of communication, *the inferential model*, on which verbal comprehension involves the inferential attribution of speaker meanings on the basis of the linguistic evidence provided. I argue, on the basis of Sperber and Wilson's (1986/1995) relevance-theoretic framework, that only within this kind of approach can an appropriate account of polysemy comprehension be given. Finally, in section 3.5, I consider two hypotheses about the evolutionary basis for our ability to attribute speaker meanings, and suggest that the development and proliferation of polysemy in natural languages are best accounted for within an approach that takes public languages to have evolved against the background of an already existing cognitive capacity for attributing mental states to others (our 'theory of mind', or 'metarepresentational' capacity), specifically, communicative intentions, as argued by Sperber (2000) and Carston (2002b). The theme of 'language does not create polysemy, but polysemy is created by use of language' will recur throughout this chapter.

3.2 The code model of communication

Before Grice ([1967] 1989), virtually all theories of communication were based on the so-called *code-model*. The guiding assumption of this model is that communication is a matter of encoding and decoding of messages in the form of signals (Shannon and Weaver 1949; Peirce 1955; Saussure 1974; Vygotsky 1986, and many others). On this approach, linguistic communication proceeds by a speaker encoding a thought into a sentence of a language – where a language is seen as a code that pairs phonetic and semantic representations of sentences – and by the hearer decoding the uttered sentence into an identical thought. For instance, if Mary, a speaker of Norwegian, wants to communicate to John that their daughter Susanne is hungry, she looks up in her mental grammar of Norwegian to find the appropriate sound-meaning pairs, and produces an utterance of (1):

(1) Susanne er sulten.

John decodes this utterance by looking up the meanings of the sounds that Mary has produced in his mental grammar of Norwegian, applies the appropriate compositional rules to them, and ends up representing the meaning *SUSANNE IS HUNGRY*.

A well-known problem with the code model of communication is that linguistic utterances typically contain context-sensitive and/or linguistically ambiguous expressions. Consider (2):

(2) I lost my bat yesterday.

In order to figure out what the speaker has expressed by uttering (2), the hearer has to assign the appropriate referents to the indexical expressions *I*, *my* and *yesterday*, and disambiguate between the linguistically-encoded meanings of *bat* (e.g. does the speaker have a ‘wooden instrument’ or a ‘flying rodent’ in mind?). In addition, he has to determine whether the conventionally polysemous verb *lost* is used in its ‘deprived of’ sense or its ‘deprived of through death’ sense. In such cases, the hearer has to use information from the context (linguistic or extra-linguistic), information that it is reasonable to suppose the speaker intended him to use, to form a representation of ‘what is said’ by the utterance, as these context-sensitive elements make reference to the discourse situation in which they have been uttered.⁵² Thus, code-models of communication need some way to account for how the hearer fixes the referents of indexical expressions and resolves instances of lexical ambiguity, including conventional polysemy.

One possibility would be to supply the code model with a set of rules for pragmatic interpretation. For instance, language users could be equipped with rules telling them that every occurrence of *I* should be substituted by a reference to the speaker, and every occurrence of *yesterday* should be substituted by a reference to the day before the utterance. Similar rules would have to be postulated for every other indexical expression of the language.⁵³ Much less straightforward is the postulation of a

⁵² The notion of ‘what is said’ here refers to the proposition that is directly expressed by the utterance, i.e. something truth-evaluable.

⁵³ Within formal semantic approaches (e.g. Kaplan 1989 [1977]; Predelli 2005), which are not concerned with explaining how indexical reference is resolved but with providing truth-conditional contents to natural language sentences, it is common to treat indexicals as encoding not fully-fledged contents but rather functions from contexts to contents (Kaplan 1989 [1977]). This makes it possible to abstract away

rule for the resolution of lexical ambiguity (e.g. one that might account for how the hearer derives the contextually appropriate meanings of *bat* and *lost* in (2) above). When taken in isolation, natural language sentences often do not contain any linguistic indicators of the correct meaning of an ambiguous lexical form. Within computational semantic approaches (Asher and Lascarides 1995; Lascarides, Copestake, and Briscoe 1996), it has been suggested that lexical ambiguity is resolved by the hearer choosing the interpretation that is coherent with the discourse context (possible coherence relations being temporal succession, causal relation, elaboration, etc.). On this view, lexical ambiguity is resolved by means of pragmatic rules that operate on the linguistically expressed discourse context.⁵⁴

3.2.1 The generative lexicon

A modern, sophisticated version of the code-based approach to communication applied to the interpretation of polysemy is the generative lexicon theory (Pustejovsky 1991, 1995a), discussed in the previous chapter. On this approach, the polysemy interpretation involves a combination of decoding of the semantic content (qualia structure) of the lexical items in the utterance, and application of a set of *lexicon-internal* generative mechanisms. For instance, recall that the interpretation of so-called type shifting constructions ('*Mary began a book*'), was seen as a function of the interaction between the requirements of the argument structure of the verb (*begin* takes an event as its second argument) and the semantics of the complement noun (*book* contains the event 'read' as part of its qualia structure, thus yielding the interpretation 'Mary began reading a book'), a process referred to as 'type coercion'. The polysemy of

from the context-sensitivity of indexical expressions, and to capture the truth-conditional content of utterances in terms of their semantics (i.e. linguistically-encoded content).

⁵⁴ Sperber and Wilson (1987, 1982; 1986/1995: Chapter 1, Section 3) have argued that, given that language contains context-sensitive elements such as indexicals and ambiguity, for the code model of communication to work, the context used by the hearer in understanding an utterance should always be identical to the one envisaged by the speaker. On their view, this 'mutual knowledge hypothesis' is intuitively incorrect, given our everyday experience with communication, and leads to infinite regress: In order to distinguish the assumptions they share and those they do not share, speakers and hearers must each make first-order assumptions about the assumptions they share (I know that *p*), then second-order assumptions about the first order assumptions they share (I know that B knows that *p*), then third-order assumptions about the second-order assumptions they share (I know that B knows that I know that *p*) and so *ad infinitum*. This 'mutual knowledge paradox' was first discussed by Lewis (2002 [1969]) and Schiffer (1972), and various solutions have been proposed to it, in particular by Bach and Harnish (1979), and Clark and colleagues (Clark 1992, 1996).

bake ('creative' vs. 'change-of-state') was generated as a function of the semantics of its complement, giving rise to a creative interpretation of *bake a cake* (artefact), a result of the process of 'co-composition', and a change-of-state interpretation of *bake a potato* (natural kind). The process of 'selective binding' generated different interpretations of evaluative adjectives (e.g. *good* in *good knife*, *good novel*), by allowing them to modify event descriptions contained in the semantic representation of the head nouns ('good for cutting', 'good read').

A problem that was pointed out in connection with the generative lexical semantic approach was that, by modelling polysemy interpretation entirely in terms of lexicon-internal processes, it is unable to account for its inherently flexible nature, and thus makes a range of wrong predictions. For instance, while it is clear that the tendency to interpret 'type shifting' constructions such as *begin a book* as 'begin reading a book' holds when the VP is considered in isolation (in a 'null context'), a more specific context may easily serve to point the hearer towards a different interpretation. Consider (3):

(3) *Mary, John and Sue work as book conservators at the British Museum. They are working on restoring a collection of medieval books, all of which are in a poor condition after having been stored on the shelves for many years. Since they are completely covered in dust, each book has to be carefully dusted before being rebound. One day, after hours of hard work, John asks if they should all take a break and go for coffee. Sue has just finished her pile and is ready to follow John to the coffee bar in the Great Court, when Mary utters:*

'Hang on a minute! I've just begun a huge old book.'

The most accessible interpretation of the last part of Mary's utterance above is clearly that she began *dusting* a book, not reading it, as would be the default interpretation predicted by the generative account, which would have to be overridden by context in this case. It does not take much imagination to think of other contexts in which the correct (speaker-intended and easily retrieved) interpretation of the VP *begin a book* is 'began binding a book', 'began mending a book', 'began ripping up a book', and so on. The problem for the generative, code-based account is that it is unable to predict such

interpretive flexibility, which would involve taking speaker intentions into account, even if the predictions it makes are in many cases correct.⁵⁵

The existence of clear interpretive tendencies for type shifting constructions in null contexts has been taken as evidence of a linguistic-semantic process (Copestake and Briscoe 1996). However, as we saw in Chapter 2 (section 2.2.3), the type coercion mechanism makes heavy use of information that, on most accounts, would count as general world knowledge (i.e. not linguistic knowledge). Therefore, it is possible that interpretive preferences (e.g. the preference for interpreting *begin a book* as ‘begin reading a book’) in the absence of further context may stem from highly accessible real-world knowledge about the denotations of the lexical concepts in the utterance (e.g. that books are typically read), and not from lexically stored information.

However, there are cases of polysemy for which there seems to be a *prima facie* reason to take the polysemy to be derived linguistically. One such case is the meaning alternations that appear to depend on whether a noun is used with count or mass syntax. An example is terms for animals. When they are used as count nouns, they designate the animal (e.g. The *chicken* pecked the ground), and when they are used as mass nouns, they tend to designate the meat (e.g. We ate *chicken* in bean sauce for dinner). Copestake and Briscoe (1996) suggest that this occurs due to a lexical rule of ‘meat grinding’, which creates meat-denoting mass nouns from animal-denoting count nouns. For instance, on the basis of this rule, we can predict that *chicken* will have a meat sense in (4) below:

(4) I love *chicken*.

In this case, an important indicator of the meat sense of *chicken* is that the noun occurs with mass morpho-syntax. But, unfortunately, syntax can only be part of the story of how this sense is derived. We can also get the meat sense when the word is used with count syntax, and the animal sense when it is used with mass syntax. Consider (5) and (6):

(5) John went to the store and came home with a *chicken*.

⁵⁵ For more detail, see the discussion of Pustejovsky’s account in Chapter 2.

- (6) To raise *chicken* in the backyard has become a new trend in many cities in the United States.

The use of *chicken* in (5) is compatible with both a ‘bird’ and a ‘meat’ sense. If the correct interpretation is the meat sense, which is arguably the most accessible one, it cannot be derived via a lexical rule of ‘meat grinding’, but has to be derived via some other (preferably pragmatic) mechanism that is capable of adjusting the count sense of the noun. As to the use of *chicken* in (6), it provides a counter-example to the rule of ‘meat-grinding’, as, in this case, the lexical inference from animal-denoting count nouns to meat-denoting mass nouns does not hold. So even in cases where, on the face of it, there is good reason to take an alternation in meaning to be linguistically-motivated, matters are rarely this straightforward.⁵⁶

3.3 Linguistic underdeterminacy

The challenge for code models of communication, even of the highly sophisticated kind proposed by Pustejovsky and others, is that what is conveyed by linguistic communication generally goes well beyond what can be coded, and does so in a highly flexible way. Since Grice’s influential work ([1967] 1989), the view that sentence meanings typically underdetermine the *implicit* content of a speaker’s meaning (‘what is implicated’), has been quite uncontroversial in philosophy of language and linguistics.⁵⁷ In recent years, there has been a growing consensus that the decoded sentence meaning may also underdetermine the *explicit* content of an utterance (‘what is said’), beyond containing indexical references and linguistic ambiguities (Searle 1978, 1983, 1992; Travis 1981, 1985, 1997; Sperber and Wilson 1986/1995; Recanati 1989, 1993, 1995,

⁵⁶ See Chapter 5 for extensive discussion of the interaction of polysemy and the count-mass distinction, including a critique of Copestake and Briscoe’s (1996) approach.

⁵⁷ One of Grice’s famous examples of the sentence meaning underdetermining the implicit content is B’s utterance in the following exchange:

A. I’m out of petrol.

B. There is a garage around the corner.

On the assumption that B is being relevant, A is warranted to infer from B’s utterance that he thinks the garage is open, that it has petrol to sell, and so on, even though these meanings are not explicitly indicated by the decoded sentence meaning (‘what is said’); rather, they are *conversational implicatures* of B’s utterance (Grice [1967] 1989: 32).

2004; Carston 1988, 2002b).⁵⁸ A number of other sources of linguistic underdetermination of the explicitly communicated meaning of an utterance have been discussed, including, following Carston (2002b: 22-28), (i) unarticulated constituents; (ii) underspecified scope of elements; (iii) underspecificity or weakness of encoded conceptual content; and (iv) overspecificity or narrowness of encoded conceptual content.⁵⁹ Examples are given below:

- (7) a. She's leaving.
b. She didn't butter the toast in the bathroom with the knife.
c. Susan is happy.
d. The steak is raw.

The utterance in (7)a. is fully sentential, but the location from where the person is leaving is unarticulated and must be supplied in order for the utterance to express a complete truth-evaluable proposition (e.g. 'Mona's leaving Oslo'). On the basis of this kind of example, several pragmatists have argued for the existence of a process of 'free' pragmatic enrichment that provides such unarticulated constituents to the proposition expressed when these are appropriately relevant (Sperber and Wilson 1986/1995; Carston 1988, 2002b; Bach 1994; Stainton 1994; Recanati 2002b; Hall 2008).⁶⁰ The utterance in (7)b. has several different interpretations depending on which constituent the negation takes scope over (e.g. 'the toast in the bathroom with a knife', 'in the bathroom', 'with a knife', etc.), each of which is associated with its own set of truth-conditions. The speaker-intended interpretation in a given context needs to be

⁵⁸ Grice's ([1967] 1989) treatment of indexical references and linguistic ambiguity as contributing to the proposition explicitly expressed by an utterance ('what is said') recognised their context-sensitivity but gave no account of how they were resolved. It is only in more recent years that it has been fully appreciated how these two components of context-sensitivity are not at all trivial, and that giving an account of how they are resolved must employ a fully pragmatic apparatus (e.g. of maxims, etc.).

⁵⁹ However, some scholars, most notably within the 'minimalist' camp of philosophy of language, make a distinction between what they call 'the proposition semantically expressed', which is a minimal proposition determined by linguistic meaning (and narrow context), including linguistically-mandated processes such as indexical reference resolution and disambiguation, and what the speaker said (which is very much pragmatically permeated) (e.g. Borg 2004; Cappelen and Lepore 2005).

⁶⁰ Recently, however, a group of authors have rejected this view, maintaining instead a more semantically oriented analysis, on which the interpretation of utterances such as (7)a. depends on the presence of covert indexicals in the logical forms of sentences, which are assigned values in context by a pragmatic process of saturation similar to that involved in assigning a referent to a pronoun (e.g. Stanley 2000, 2005; King and Stanley 2005). For instance, the logical form of the utterance in (7)a. would contain a location variable, which could be assigned the value 'Oslo'.

pragmatically inferred (at least in the absence of prosodic cues, as when the utterance occurs in written form) (Carston 2002b: Chapter 4). In (7)c., the adjective *happy* denotes a general concept that is usually made more specific in context (expressing, for instance, that Susan is ‘happy in a low-key, peaceful sort of a way’, ‘happy in an ecstatic, excited sort of a way’, etc.), while the encoded ‘uncooked’ meaning of *raw* in (7)d. could be relaxed in order to describe food that has had some, but grossly insufficient, cooking. This kind of data has been used to argue for the existence of a pragmatic process that adjusts the meaning of individual words (Recanati 1995, 2004; Carston 1997, 2002b; Wilson and Carston 2006, 2007). Carston (2002b: 28) notes that, although in cases such as (7)c. and (7)d. the operation of this pragmatic process is not necessary in order for the utterances to express full propositions, it “seems to be required if we are interested in finding that proposition which it is rational to assume that the speaker intended to express”. In Chapter 4 I argue that this pragmatic process, which constructs context-dependent lexical meanings on the basis of their encoded meanings, plays a central role in giving rise to polysemy.

There are at least two possible views one might take on the nature of linguistic underdeterminacy, described by Carston (2002b: 29) as the *convenient abbreviation view* and the *essentialist view*. On the convenient abbreviation view, underdeterminacy is a matter of effort-saving convenience for the speaker. Although sentence meaning more often than not underdetermines the proposition expressed by it, a sentence that fully encodes the speaker’s meaning could always be supplied. This is the view held by Quine (1960) and Katz (1972, 1976, 1981), who argued that every proposition expressed (or ‘statement made’) by a natural language sentence was describable in terms of a context-independent, ‘eternal’ sentence. Katz termed this the *principle of effability* (1972, 1976, 1981). Crucial to the transformation of a non-eternal sentence into an eternal sentence is the replacement of each indexical element with a description that picks out, for each indexical nominal expression, the intended referent in a unique context-independent way, and for each indexical tense indicator, a referentially unique time designation. Below is a suggestion for a fully encoding ‘eternal sentence’ counterpart for the utterance in (2) above, repeated as (8) for convenience:

(8) a. I lost my bat yesterday.

b. 'Ingrid Lossius Falkum caused Ingrid Lossius Falkum to be deprived of the wooden implement with a handle and a solid surface used for hitting the ball in baseball that belonged to Ingrid Lossius Falkum, between 2pm and 4pm on 22nd of August 2010 somewhere in the Majorstua area, Oslo, Norway'.

On the convenient abbreviation view of underdeterminacy, the idea would be that a speaker, in order to save herself the effort of having to express a long, complex sentence such as that in (8)b., can choose to use a sentence which does not fully encode her intended meaning, and rely on the hearer using his pragmatic inferential capacity to turn it into a fully propositional representation. On this approach, Carston (2002b: 35) points out, our pragmatic ability would be a 'useful add-on' to our language capacity, but its contribution would not be strictly essential in enabling us to express ourselves and communicate the way we do.

On the essentialist view, however, linguistic underdeterminacy is seen as an essential feature of the relation between sentence meanings and speaker meanings (Carston 2002b: 29). The view is that, given the complexity and fine-grainedness of the thoughts that speakers can entertain and communicate to each other, they generally do not lend themselves to a full encoding by natural language sentences. So although the sentence in (8)b. comes closer to encoding the speaker-intended meaning than (8)a., any attempt to be fully explicit (by providing uniquely denoting descriptions for entities and properties) is bound to fail. The essentialist view comes in a weak and a strong version. The weak version takes underdeterminacy to be the rule but a full encoding may be possible in a small number of cases (possible candidates are so-called analytic, i.e. context-free, sentences, e.g. 'Cats are animals'). The strong version takes underdeterminacy to be universal and denies that a full encoding is possible in any circumstances. On either of these versions, our ability to make pragmatic inferences about speaker-intended meanings would provide the essential foundation for our expressive and communicative abilities.

Both the convenient abbreviation view and the essentialist view of linguistic underdeterminacy provide a basis for polysemy. On the first view, polysemy could be motivated by a goal of 'economy of expression', representing an effort-saving convenience for the speaker. Instead of having to go through the laborious task of fully

encoding the sense she has in mind (which she could do, if she wanted to), the speaker can rely on the hearer's pragmatic abilities enabling him to arrive at the intended interpretation (for instance, the expression *mouth of the river*, containing the polysemous expression *mouth*, could be seen as a convenient shorthand for the fuller encoding *the place where the river enters the sea*). On the second essentialist view, however, where the linguistic codes that modern humans make use of in communication are not capable of fully encoding speaker meanings (and have to be supplied by pragmatic inference), polysemy would be a *necessity*. That the sentences of our languages can be used to express a number of different propositions depending on the context would, on this approach, be a necessary consequence of the nature of our linguistic codes. At the lexical level, if the vocabularies of our languages are not capable of encoding the range of concepts we can entertain and communicate, polysemy – understood as the ability of words to express different meanings in different contexts – would then follow as a necessary consequence.

3.3.1 The (in)effability of thought

The question as to whether there exist 'eternal sentences', which underlie Katz's effability principle, is central to the evaluation of the two positions on linguistic underdeterminacy presented above. In a recent survey paper, von Stechow and Matthewson (2008) examined the evidence for potential semantic universals on the basis of a range of cross-linguistic data. Comparing the expressive powers of different languages had been suggested by Katz (1972, 1976) as a way his effability principle could be tested:

[I]t follows from the effability thesis that each natural language is capable of expressing the same body of thoughts ... for any sentence of one natural language, there is at least one sentence in every other natural language that expresses the same proposition. Accordingly, if we find one sentence in one natural language that has no translation in some other natural language, the thesis is false.⁶¹ (Katz 1972: 20)

Von Stechow and Matthewson (2008) conducted investigations in three areas: (i) the lexicon (content morphemes); (ii) functional morphemes/categories and composition principles; and (iii) pragmatics. Acknowledging that cross-linguistic uniformity cannot be found at any kind of structural level (e.g. what is encoded by means of a simple expression in one language may be encoded by a complex expression in another language), they allowed for cross-linguistic equivalence to be understood as “what language X expresses [= encodes] is also expressible in language Y but at the price of some complexity” (von Stechow and Matthewson 2008: 144). Their results were as follows: They found no convincing examples of semantic universals at the level of the lexicon (e.g. nouns, verbs, adjectives); no individual lexical items are universally attested. Cross-linguistic differences were found both with regard to inventories of functional categories (e.g. some languages ‘lack’ certain quantifiers), and in the distinctions encoded by functional morphemes (e.g. apparently equivalent modals in two languages may differ with regard to their encoded content). While they found reason to agree with the view that there exists a small set of universally available composition principles, they point out that more work is needed to determine exactly what those principles consist of. As to the domain of pragmatics, considering the phenomenon of presupposition, von Stechow and Matthewson observe that although the literature devoted to presupposition concentrates almost exclusively on English, the existing empirical evidence indicates that languages may differ in their ability to express certain aspects of meaning as presuppositions (e.g. the existence of non-presuppositional *it*-clefts). On this basis, von

⁶¹ There are two ways in which this claim could be interpreted, depending on what is meant by ‘expressed’. Either it could mean that for any sentence of one natural language, there is at least one sentence in every other language that *encodes* its meaning, which is the interpretation likely to be intended by Katz. Or it could mean that for any sentence of one natural language, there is at least one sentence in every other language that *communicates* its meaning. In this case, one would be making a much weaker claim, allowing for non-equivalence at the level of encoded content, but assuming equivalence of content at the level of proposition/thought. Under this interpretation, the principle seems much more plausible.

Fintel and Matthewson conclude that Katz's effability principle (1972, 1976, 1981) cannot be maintained.⁶²

Sperber and Wilson (1986/1995: 191-192) and Carston (2002b: 30-42) argue against Katz's principle of effability and the existence of 'eternal' sentences on psychological and philosophical grounds. Sperber and Wilson maintain that the strong version of Katz's principle, on which it entails that every human *thought* corresponds to a fully encoding natural language sentence, is incompatible with the psychological, individualist notion of thought. Thoughts, in this sense, are likely to contain private references, to time, space, people, events, and so on, fixed in terms of a 'private logbook' and 'ego-centred map'. For instance, my mental representation of the man who is my husband contains a number of such private references, and is likely to differ from the representation his mother has of him, and the representation a person who has watched him give a talk at a political meeting may have of him. Sperber and Wilson's point is that thoughts that contain such private references do not seem amenable to a full encoding in a natural language.

Carston argues against a weaker version of the effability principle, on which it entails that only the *proposition expressed* by a natural language sentence (not thoughts in general) can be fully encoded by another natural language sentence,⁶³ drawing on arguments from Wettstein (1979) and Recanati (1987, 1994). She maintains that it is impossible to provide eternal sentences that correspond to propositions expressed by sentences containing indexical, referential, and predicate expressions. As observed by Wettstein (1979), the object referred to by an indexical expression can be picked out by a number of *non-synonymous* (i.e. truth-conditionally distinct) descriptions. For

⁶² However, von Fintel and Matthewson (2008: 191) are optimistic about the prospect of discovering 'purely semantic' universals, given their observation that "languages often express strikingly similar truth-conditions, in spite of more trivial differences in lexical semantics or syntax". Thus, they seem to espouse a view of effability on which there may be cross-linguistic differences in presuppositions (and implicatures, expressive content, etc.), but where, "at the level of core truth-conditional content, what one language can express any other can as well" (ibid. 7). I cannot assess this weakened effability claim here, or whether there is any psychological basis for postulating a level of 'core truth-conditional content' in this strictly semantic sense. What seems highly plausible, however, is that given our pragmatic abilities, for any given proposition expressed by use of a sentence in a language there is a sentence that can be used to express (= communicate) that proposition in every other language, in spite of the considerable differences between the linguistic codes used.

⁶³ Carston (2002b: 34) refers to this as the 'Third Effability Principle', originally considered (and rejected) by Recanati (1994: 157).

instance, the indexical expression *She* in an utterance of '*She* arrived late' can be uniquely described as 'The woman who danced with the Norwegian Crown Prince at t_1 ' 'The female teacher in the pink dress with a white waistband who was at Oslo Castle between t_2 and t_3 ', 'The brown-haired Labour Party representative who smoked three cigarettes outside Oslo Castle between t_4 and t_5 ', etc. It is not clear which of these expressions should be taken to be a constituent of the eternal sentence, i.e. that which is an actual formulation of the proposition that the speaker has asserted. Carston (2002b: 36) writes:

Indexical reference ... appears to be irrevocably context-bound; it is not reformulable in terms of a uniquely denoting expression, but depends on the addressee's capacity to identify the intended entity by some means which is non-linguistic, or at least, not wholly linguistic.

Furthermore, Recanati (1987) has claimed that it is not only indexical expressions that are referentially context-dependent; quite generally, there cannot be *reference* without a context (and thus no uniquely denoting descriptions). For instance, if proper names are treated as a variety of indexical (a position that both Recanati and Carston seem to endorse), it is, in view of Wettstein's argument above, unlikely that they can be given eternal descriptions. In the case of apparently complete definite descriptions (e.g. 'The Crown Prince of Norway'), Recanati (1987: 62) argues that their reference always depends on the 'domain of discourse', which is "that with respect to which the speaker presents his or her utterance as true".⁶⁴ Carston, following Travis (1985) and others, shows that this also includes the predication function of language. For instance, in cases where a communicated predicate is more specific than the concept it encodes (e.g. an utterance of 'Susan is *happy*' in a context where she has just learned the result of her exam) it would often be difficult to find a description that fully encodes this concept (e.g. Susan could be happy to the extent that she would consider continuing her studies,

⁶⁴ Recanati (1987: 63) discusses the following kind of example. Consider a situation where an old person, who is suffering from dementia, insists, in 2010, that the Crown Prince of Norway is still Harald (when he is, in fact, Håkon Magnus). It is conceivable that, for instance, his children could use the apparently complete definite description 'the Crown Prince of Norway' to refer to the previous crown prince, Harald, in a conversation with each other. In this case, the interpretation of this definite description would depend on the identification of the domain of discourse with respect to which it is intended to be evaluated (i.e. the belief world of their father).

but not happy enough to want to go out celebrating, etc.).⁶⁵ Moreover, cases where the concept communicated by a predicate depends on contextually-determined standards of precision (e.g. 'France is *hexagonal*', 'The steak is *raw*', etc.) are unlikely to have eternal sentence counterparts, as the approximate values of the predicates do not seem to be possible to express in words.⁶⁶

Taken together, these arguments point to the non-existence of eternal sentences in natural languages, on the basis of which Katz's effability principle must be rejected. It follows that, rather than being just an effort-saving convenience for the speaker, linguistic underdeterminacy is an essential feature of the relation between linguistic codes and the thoughts/propositions that can be expressed by them (Carston 2002b: 42). At the lexical level, this would make polysemy a necessary consequence of the fact that the concepts that are lexically encoded are far fewer than the range of concepts that can be communicated by use of them. In a sense, the meanings of polysemous lexical items are intrinsically ineffable in that, in most cases, it will be impossible to describe in natural language terms the full range of meanings a polysemous lexical item may be used to express in different contexts (e.g. *happy*, *good*, *over*, *run*, *open*, etc.). In addition, the sense that is communicated by use of a polysemous lexical item in a given context will arguably be subject to the more general issue of ineffability of reference and predication, as described above. In fact, the examples discussed by Carston (2002b: 40-41) of communicated concepts being more specific than the encoded concepts (e.g. *happy*), or those communicated concepts depending on a contextually relevant standard of precision (e.g. *hexagonal*, *raw*) are good candidates for polysemy.

In the next section, I present an alternative to the code model of communication, the *inferential model*, in which the essential nature of linguistic underdeterminacy can be accommodated. I claim that the inferential model, more specifically the relevance-theoretic pragmatic approach developed by Sperber and Wilson (1986/1995), also provides a natural framework for analysing polysemy comprehension.

⁶⁵ See Chapter 4 for a more detailed discussion of such cases.

⁶⁶ However, some philosophers of language, following Lewis (1979), will maintain that predicates such as *hexagonal* have the same content throughout, but come with standard of precision parameters that have to be fixed in context.

3.4 The inferential model of communication

Grice ([1957], [1967] 1989) was the first to hold the view that instead of being a matter of encoding and decoding of meanings, communication is *inferential*, and involves a kind of mindreading: A speaker provides evidence of her intention to communicate a certain content to the hearer, and the hearer recovers this intention by a rational maxim-guided inferential process using the evidence provided. Grice's claim was that recognising the intentions of others in this way is a form of inference to the best explanation, that is, to the explanation that best accounts for the available evidence. On this approach, communication can be achieved just by expression and recognition of intentions, even in the absence of a code (e.g. by pointing, gazing, eyebrow raising, etc.). However, using a language allows for a degree of precision and complexity that can usually not be achieved in non-verbal communication.

3.4.1 Relevance theory

Sperber and Wilson (1986/1995) draw on Grice's central insight about communication as an inferential process that requires the attribution of speaker intentions, in their formulation of *relevance theory*. On their approach, inferential communication, or *ostensive-inferential* communication, involves two layers of intention: (i) an *informative intention*, which is the intention to inform an audience of something; and (ii) a *communicative intention*, which is the intention to inform an audience of one's informative intention (Wilson and Sperber 2004: 611). To communicate this way, the communicator has to use an *ostensive stimulus*, that is, a stimulus "designed to attract an audience's attention and focus it on the communicator's meaning", which provides evidence that she intends the audience to come to a certain conclusion (ibid.). Proper understanding depends on the fulfilment of the communicative intention, that is, on the hearer recognising the speaker's informative intention. In this way, utterance interpretation is seen as involving the attribution of several layers of intention; it is a sophisticated form of mindreading.

Relevance theory addresses a fundamental problem for human cognition. Given the vast amounts of information that our senses monitor, that are stored in our memories, and the possible conclusions that our inferential system can draw from all this information at any given time, the question is: how do we select which inputs from

the environment to attend to, and how do we, when processing the inputs, decide which contextual assumptions we should bring to bear, and what conclusions we should draw? Sperber and Wilson (1986/1995: 260) claim that there is a single answer to these questions, stated in their *Cognitive Principle of Relevance*: “Human cognition tends to be geared to the maximisation of relevance”. Relevance is defined as a potential property of all types of input to cognitive processes, and may be assessed in terms of the amount of effort (of perception, memory and inference) it takes to process the input, and the positive cognitive effects the individual may derive from it (where a ‘positive cognitive effect’ is described as a ‘worthwhile difference to the individual’s representation of the world’, and may include strengthening or elimination of existing assumptions, derivation of contextual implications). Other things being equal, the more cognitive effects an input yields to an individual and the less effort it takes to process it, the higher the degree of relevance of that input to that individual at that time.

Sperber and Wilson (1986/1995) further claim that ostensive stimuli (utterances) create expectations of relevance not raised by other types of stimuli, which are precise and predictable enough to guide the hearer toward the communicator’s meaning. This is stated in their *Communicative Principle of Relevance*: “Every act of ostensive communication communicates a presumption of its own optimal relevance” (ibid. 260). By requesting the addressee’s attention, the communicator conveys that her ostensive act is more relevant than alternative stimuli competing for his attention at the time. An utterance is optimally relevant if (a) it is at least relevant enough to be worth processing, and (b) it is the most relevant one compatible with the speaker’s abilities and preferences (ibid. 270).⁶⁷ To make her utterance optimally relevant, the speaker should achieve at least enough cognitive effects to make the utterance worth processing while avoiding causing the hearer any gratuitous effort in achieving those effects. The hearer’s goal in communication is to find an interpretation of the speaker’s meaning that meets the expectations of relevance raised by the ostensive stimulus itself.

⁶⁷ Expectations of relevance vary across speakers and times, depending on what alternative inputs are available. In general, to be worth processing, the relevance of the utterance must be higher than other inputs that the hearer could have been attending to at the time. Clause (b) takes into account that the speaker may be unable or unwilling to make her utterance more informative or economical (Sperber and Wilson 1986/1995: 266ff.).

In modern pragmatic theory, and relevance theory in particular, the capacity to infer speaker meanings on the basis of the evidence provided is taken to be closely tied to, and reliant on, the more general capacity to infer and attribute contentful mental states to others: our ‘theory of mind’, ‘mindreading’, or ‘metarepresentational’ capacity (Premack and Woodruff 1978; Wimmer and Perner 1983; Baron-Cohen, Leslie, and Frith 1985; Leslie 1987; Baron-Cohen 1995, and many others).⁶⁸ The claim is that the theory of mind capacity provides the foundation for the kind of ostensive-inferential communication that humans engage in; it is what enables us to bridge the gap between not fully encoding linguistic meanings and intended speaker meanings, including inferring speaker-intended concepts that differ from the concepts lexically encoded by the words used. Within this view, as Carston (2002b: 46) has pointed out, there would be a good reason for our having linguistic codes that do not fully encode the thoughts we communicate to each other: it would just not be necessary.⁶⁹

The role of the pragmatic inferential capacity in bridging the gap between lexically encoded concepts and communicated concepts is rarely explicitly considered in connection with polysemy (but cf. Tyler and Evans 2003, to be discussed in Chapter 4). This is, of course, understandable if the goal is to capture polysemy in terms of the semantic content of lexical items and linguistic operations over this content. However, on the view that linguistic underdeterminacy is a necessary aspect of the relation between linguistic meanings and speaker meanings, and where polysemy is seen as an instance of underdeterminacy applying at the level of individual lexical items, providing a full account of polysemy in terms of the workings of the linguistic system should, in principle, not be feasible. At least, there would be no need for polysemy to be entirely

⁶⁸ According to Sperber and Wilson (2002a), pragmatics is, in fact, a submodule of the mental system that is responsible for our ability to attribute mental states to each other.

⁶⁹ A real alternative to the inferential view of communication can be found in Millikan’s (1984, 1987, 2005) theory of utterance comprehension. Millikan rejects the Gricean (and relevance theory) assumption that utterance interpretation is a matter of inference about the speakers’ intentions, and argues that it should rather be modelled as a form of ‘direct perception of the world’, much like the way we perceive the world through vision. Her claim is that when we understand language, “it is the *world*, not meanings, and not speaker intentions, that is immediately perceived” (Millikan 2005: 206). Recanati (2002a, 2004, 2007) holds a similar view with regard to the derivation of the *explicit* content of an utterance (‘what is said’). Relevance-theorists, on the other hand, take the process of verbal comprehension to be entirely inferentially based, including the derivation of the proposition explicitly expressed. I cannot get into the details of the debate between Millikan/Recanati and the relevance-theorists here, but the question of whether explicit content is inferentially derived comes up again in Chapter 6, in the discussion of metonymy.

linguistically generated if communicators possess a powerful enough pragmatic inferential capacity. In the next section, I outline how the comprehension of polysemy, both of the conventional and non-conventional kind, can be naturally analysed within the relevance-theoretic framework presented above, as part of the overall inferential process of constructing hypotheses about speaker-intended meanings on the basis of the linguistic evidence provided.

3.4.2 Polysemy and utterance comprehension⁷⁰

Consider the following example. On 27 January 1967, the command module of the *Apollo 1* spacecraft was destroyed by a fire during a test and training exercise. The three astronauts aboard were killed. When one of the control centre employees called up another NASA astronaut to tell him about the tragedy that had taken place, he uttered (9):

(9) We *lost* our crew!

On his own account, the astronaut at the receiving end was at first unsure whether his colleague was telling him that the crew were missing, that they had been unable to find them and so couldn't carry out the testing, or if he was in fact telling him that they had all died.⁷¹ The difficulty, of course, concerned the conventional polysemy of the verb *lose*; whether his colleague in the control centre had intended to convey 'become unable to find' or 'be deprived of (someone) through death', both of which were compatible with the overall (linguistic and extralinguistic) context. However, when his colleague continued 'There has been a fire!', in an agitated voice, it started to dawn on the astronaut what was being communicated. Although the contextual information was still compatible with an interpretation on which his colleague was informing him that the testing had been unsuccessful due to two problems, (a) they couldn't locate the crew and

⁷⁰ The relevance-theoretic analyses outlined in this section do not specify the details of the pragmatic processes involved in the derivation of speaker meanings. Instead, they are intended to show that the comprehension of polysemy, even of the conventional kind, depends on a constructive pragmatic inferential process that forms hypotheses about speaker meanings on the basis of linguistic evidence and accessible contextual assumptions. I return in more detail to the specific pragmatic processes responsible for the construction of polysemy (that is, of the non-conventional kind) in Chapters 4, 5, and 6.

⁷¹ Attested example from the British documentary film *In the Shadow of the Moon* (2007) directed by David Singleton.

(b) there had been a fire, he inferred the causal-explanatory connection between the first and the second utterance and, taking the agitation in the speaker's voice into account, understood that what his colleague had intended to communicate by uttering (9) was that his three astronaut colleagues aboard the *Apollo 1* had all died.

On the relevance-theoretic account, the astronaut's interpretation of (9) proceeds according to a comprehension heuristic, which is applied automatically to verbal input (Wilson and Sperber 2004: 613). According to this heuristic, the hearer (i) takes the decoded linguistic meaning, follows a path of least effort in considering interpretive hypotheses, and (ii) stops when the interpretation he arrives at satisfies his expectations of relevance. The interpretation the astronaut arrived at was more relevant than the other possible interpretation. Given the available contextual assumptions (e.g. 'There had been a fire', 'The control centre employee was agitated', etc.) it was the most accessible interpretation (the least effort demanding), and it was the one that yielded the most cognitive effects (it provided a causal connection between the two utterances, and probably made the astronaut draw a great many contextual implications concerning the consequences for the crew's families, future Apollo missions, and so on). The proposition that was explicitly expressed by the utterance in (9) was something along the lines of (10):

(10) NASA was deprived, through death, of the crew of the *Apollo 1* spacecraft.

Arriving at this explicitly expressed proposition, including an interpretation of the polysemous verb *lose*, required a non-demonstrative pragmatic inference process geared to the recovery of the speaker's meaning.⁷²

In the example above, recovering the speaker-intended meaning of the conventionally polysemous verb *lose*, involved a pragmatic process of disambiguation.

⁷² It might be argued that the disambiguation of *lose* in (9) can also be accounted for within a computational approach that takes lexical ambiguity to be resolved by the hearer choosing the interpretation that is coherent with the discourse context (e.g. Asher and Lascarides 1995; Lascarides, Copestake, and Briscoe 1996). On this approach, the ambiguity could be resolved by the application of a pragmatic rule or constraint that tells the hearer to choose the interpretation that maintains discourse coherence (in this case, it would involve retrieving the causal connection between the two utterances). However, I think the relevance-theoretic analysis is to be preferred for reasons of psychological plausibility, in that it accounts for the fact that a variety of contextual information, not only that which was linguistically expressed in the discourse context, can be used in forming a hypothesis about the speaker-intended meaning.

However, in the examples containing the verb *open* (11) below, there seems to be only a single linguistically-encoded meaning of the verb, and the more specific interpretations are to be derived by means of a pragmatic process that constructs occasion-specific senses in accordance with the hearer's expectations of relevance:

- (11) a. Susan *opened* the door.
b. Mary *opened* a bottle of wine.
c. John *opened* a tin of beans.
d. Peter *opened* the curtains.

In these cases, the direct object of *open* indicates which concept the speaker intends the hearer to construct. Sometimes this requires little effort on the part of the hearer by way of pragmatic inference (e.g. in the 'normal' interpretation of *opened the door*, where Susan uses the door handle to open it), whereas in other cases, inferring the speaker's intended meaning may demand more of the hearer's pragmatic abilities (e.g. in a case where *opened the door* is used to convey that Susan picked its lock, smashed it open, etc.).

A similar analysis can be given to Pustejovsky's (1995a) examples of selective binding, where the meaning of an adjective is assumed to vary as a function of the semantics of the noun it modifies. These can be treated on the relevance-theoretic account as involving pragmatic specification of the concept linguistically-encoded by the adjective. For instance, the adjective *good*, discussed by Pustejovsky as subject to the process of selective binding, can be seen as encoding a very broad concept which, on most occasions of use, will have to be pragmatically adjusted in order for the hearer to arrive at the interpretation that was intended by the speaker. In this way, the concept that is communicated by *good* when it modifies, e.g. the noun *book*, would have, as a result of pragmatic specification, a narrower denotation than that which is encoded (e.g. 'good read'). However, unlike the generative lexicon, the relevance-theoretic approach predicts that *good book* may communicate different occasion-specific senses (e.g. a good book could be one that is easy to read, one that is entertaining, one of high academic or intellectual quality, one that is good to kill flies with, etc.), by treating lexical interpretation as a matter of adjusting the interpretation of individual words in

accordance with one's context-specific expectations of relevance.⁷³ I return to a more detailed consideration of this sort of pragmatic process in Chapter 4.

On the relevance-theoretic approach, the interpretation of type shifting constructions such as the VP *begin a book* in (3) above, would proceed in terms of a pragmatic process that recovers the contextually appropriate event associated with the VP, rather than in terms of a lexicon-internal generative mechanism (cf. Pustejovsky 1995a). In the case of (3), where the most accessible interpretation of the utterance 'I've just begun a huge old book' was that Mary had just begun *dusting* a huge old book, this would be the most relevant one the hearer could derive. It would be the interpretation that required the least processing effort, and gave rise to a set of cognitive effects of the expected sort (i.e. it provided an adequate explanation for Mary's previous utterance of 'Hang on a minute!'). Thus, it achieved relevance in the way that the hearer might have expected it to do.⁷⁴

Given the dependence of polysemy comprehension on this constructive pragmatic inference process, geared to the recovery of the speaker's meaning, I would like to suggest that pragmatics may play a more fundamental role in the development and motivation of polysemy in natural languages than has previously been recognised. In the next section, I discuss two hypotheses about the evolutionary basis for our ability to attribute speaker meanings, and suggest that the development and proliferation of polysemy in natural languages are best accounted for within an approach that takes language to have evolved against the background of an already existing cognitive capacity for attributing mental states to others. My suggestion is that the fundamentally pragmatic nature of polysemy has an evolutionary basis.

⁷³ It is worth mentioning that the flexibility lacking in Pustejovsky's account is allowed for in those formal semantic accounts that operate with a larger set of context-sensitive elements (i.e. extending beyond indexicals and linguistic ambiguity). For instance, Szabó (2001) postulates a variable in the semantic representation of the adjective *good*, requiring a pragmatic operation of saturation, which accounts for its different senses in, e.g. *good knife*, *good weather*, *good book*. However, any account that postulates such invisible indexical elements needs to be backed up by syntactic evidence that it is really there.

⁷⁴ It is compatible with this pragmatic account, however, that type shifting verbs such as *begin* come with some kind of semantic parameter indicating that their complement is an event or activity, in which case supplying the relevant event or activity to the proposition expressed by the utterance would be the result of a linguistically-mandated process of saturation. In any case, though, whether the process is a matter of saturation or a top-down, 'free' recovery of unarticulated constituents, the specific event or activity that is recovered in each case would be a matter for context and pragmatic inference. See Chapter 4 for a more detailed analysis.

3.5 An evolutionary story about theory of mind and language

It is widely thought that language and theory of mind are made possible by biologically evolved mental mechanisms, which are both domain-specific and distinctive to humans. The uniqueness of human language is attested by the fact that no linguistic system of equivalent power appears to exist elsewhere in the animal kingdom. Humans also appear to be unique in their ability to attribute beliefs, desires and intentions to others, although there is some experimental evidence suggesting that chimpanzees and other primates may have some rudimentary metarepresentational capacities (Call and Tomasello 2005). The interdependence of language and theory of mind in humans suggests that they have co-evolved. However, given the view that language and theory of mind are distinct mental mechanisms, we may hypothesise that one of them developed first to a degree sufficient to enable the development of the other (Sperber 2000). This gives rise to two different hypotheses about the origin of the co-evolutionary process. On the first, language evolved first and subsequently provided the basis for the development of a theory of mind ability. On the second, the theory of mind, or the metarepresentational capacity,⁷⁵ provided the foundation for the development of the language capacity (ibid.). In this section, I consider the implications of each of these two hypotheses about the evolution of language and theory of mind in humans for the development of polysemy. Drawing on arguments by Sperber (2000) and Carston (2002b), I will suggest that the second hypothesis (theory of mind-first) provides a plausible explanation of the development and motivation of polysemy in natural languages.

According to the language-first hypothesis, the capacity for language preceded the metarepresentational ability in human evolution (cf. Dennett 1991). On this hypothesis, a metarepresentational ability might have developed as a result of our ancestors, already linguistically equipped to assign content to linguistic utterances, gradually becoming aware of the representational character of the linguistic utterances in their environment. At first, the domain of such an ability would be restricted to

⁷⁵ Throughout his discussion of these two evolutionary hypotheses, Sperber (2000) uses the term 'metarepresentation', by which he understands the ability to form and entertain representations of representations quite generally, including the ability to attribute beliefs, desires and intentions to others (the theory of mind).

linguistic utterances, then it would extend to other kinds of representations, and to mental representations in particular. This hypothesis meshes well with the view of linguistic underdeterminacy as convenient abbreviation, as discussed in section 3.3 above. If language evolved first, it seems likely that there was a time in human evolution where our ancestors fully encoded their messages (and where their communication contained no implicatures), and explicit content was derived entirely via decoding.⁷⁶ The subsequent evolution of a metarepresentational capacity would have provided the communicators with a short-circuiting, effort-reducing tool, which saved them the laborious process of fully encoding and articulating their thoughts.

Sperber's (2000) main argument against the language-first hypothesis is that if ancestral language was entirely a matter of encoding-decoding, there must have been a radical change in the human linguistic mechanism at some point in its evolution. In order to consider this claim, it is useful to look at some characteristics of animal communication.

The code model of communication is probably correct for most animal communication. Well-known examples of coded animal communication include the dances of honeybees (von Frisch 1976) and vervet monkey vocalisations (Struhsaker 1967; Seyfarth, Cheney, and Marler 1980). While honey bees dance to convey information about the location and quality of food, vervet monkeys produce acoustically different alarm calls in response to different classes of predators (large cats, birds of prey, and snakes), which trigger specific escape responses among their conspecifics (leopard alarms make them run up into trees, eagle alarms cause them to look up, run into the bush, or both, and snake alarms make them stand on tiptoe, looking down at the ground). Early interpretations of vervet vocalisations took them as evidence that the vervet monkeys shared a semantic system of communication, by which they could refer to specific types of objects and events they encountered in their environment, and it was hypothesised that the alarm signals could be evolutionary precursors, or homologs, of human words (Seyfarth, Cheney, and Marler 1980; Cheney and Seyfarth 1990). This hypothesis has subsequently been questioned by many scholars, in view of the many

⁷⁶ But cf. Millikan (1984, 2005), who would dispute the claim that this is the only way in which communication may proceed in a non-metarepresentational environment. See footnote 69.

radically different properties of human language, which indicate the possibility of an independently evolved, distinctively human, mechanism (Bickerton 1990; Hauser, Chomsky, and Fitch 2002; Tomasello 2008; Fitch 2010). Unlike vervet vocalisations, which have an innate, genetic basis (all individuals have the same basic vocal repertoire, including those raised in isolation or with other monkey species (Tomasello 2008)), the human lexicon is learned - it is only the *ability to acquire* a lexicon that is innately determined. Furthermore, while vervet vocalisations are situation-specific (the vocal calls appear to be tightly fixed to the triggering situation), human words are non-situation-specific; they are rarely associated with specific functions, but can be linked to virtually any concept that humans can entertain (Hauser, Chomsky, and Fitch 2002; Jackendoff 2002). As Hauser, Chomsky and Fitch (2002: 1576) write, “[e]ven for the simplest words, there is typically no straightforward word-thing relationship, if ‘thing’ is to be understood in mind-independent terms”.

This difference between the properties of human and animal signals is crucial. First, it may be taken to illustrate how the co-evolution of language and theory of mind in humans has shaped our public codes, in particular their lexical component, into being characteristically underdetermining of speaker-intended content (thus, providing a strong basis for polysemy). Second, it illustrates the limitations of communication that takes place in the absence of an ability to recognise the mental states of others (although other factors obviously play a role as well). For a species like the vervet monkey, where there has been a clear evolutionary advantage to having vocal calls tightly fixed to the triggering situation (allowing for faster and more efficient escape reactions), it seems plausible that the progress of evolution may have involved a drift toward increasing differentiation and specialisation of signals (thus univocality), as indicated by Seyfarth and Cheney (2003) in a more recent review of research on animal vocalisations:

Natural selection has favoured callers who vocalise in ways that change the behaviour of listeners and ultimately return benefit to the caller. Simultaneously, selection has favoured listeners who detect the links between specific calls and particular events, thereby extracting whatever information may be relevant to them. In many group-living species, where callers and recipients have overlapping reproductive interests, selection has favoured callers who give acoustically different vocalisations in different circumstances, thus allowing listeners to extract more specific information. (Seyfarth and Cheney 2003: 168)

Thus, the properties of animal vocalisations seem, in evolutionary terms, to be a very long way from human-style communication. Sperber (2000: 123) argues that if our ancestors were such encoders-decoders, and had no mindreading abilities, “there is no sensible story of how the presence of utterances in their environment would have led them to discover their representational character, to metarepresent their content, and to use for this their own mental representations in a novel, opaque, manner”, and therefore, the language-first hypothesis must be rejected.

According to the theory of mind-first hypothesis, the metarepresentational ability (specifically, the ability to attribute contentful mental states to others) preceded the capacity for language in human evolution. An advantage of this hypothesis is that there are independent reasons for the evolution of a metarepresentational ability (having to do with competition, exploitation and co-operation), and the possession of this ability would lay the ground for other forms of communication, for instance, pointing and pantomiming, quite independently from any language. To illustrate this point, Sperber (2000: 123-126) invites us to imagine a situation involving two of our hominid ancestors, Peter and Mary, and five different scenarios it may give rise to, in increasing metarepresentational complexity.

In the first scenario, Mary is picking berries, and Peter happens to be watching her. Peter uses his metarepresentational ability and infers from Mary’s behaviour that if she’s picking berries, she must believe them to be edible. In this case, Peter would have a first-order metarepresentational belief:

- (12) Mary believes
that these berries are edible.

In the second scenario, Mary intends Peter to draw, from her behaviour, the inference that the berries are edible. Mary has a first-order informative intention:

- (13) That Peter should believe
that these berries are edible.

In the third scenario, Peter is aware that by picking berries, Mary intends him to come to the belief that they are edible. Peter has a second-order metarepresentational belief:

- (14) Mary intends
that he should believe
that these berries are edible.

In the fourth scenario, Mary intends Peter to become aware of the fact that she intends him to come to the belief that the berries are edible. Mary has a second-order informative intention (which is a third-order metarepresentation):

- (15) That Peter should believe
that Mary intends
that he should believe
that these berries are edible.

In the fifth scenario, Peter recognises that Mary intends him to become aware of the fact that she intends him to come to the belief that the berries are edible. Mary might achieve this by ostensibly ensuring that she has Peter's attention; e.g. by establishing eye contact with him, by picking the berries in a particular (perhaps stereotypical/exaggerated) way, etc. Peter comes to have a fourth-order metarepresentational belief:

- (16) Mary intends
that he should believe
that she intends
that he should believe
that these berries are edible.

It is at this level of metarepresentational sophistication that communication proper takes place, although in (16) it does not involve the use of a code. Sperber argues that this is where the substantial change occurs: Instead of being restricted to using direct

evidence (physical actions/behaviour) to fulfil her first-order informative intention, Mary is now in a position to use symbolic cues (e.g. to mime the action of eating the berries) to fulfil her second-order communicative intention (i.e. to make Peter recognise her intention to inform him that the berries are edible). If Peter and Mary had shared a common code, Mary could have produced an utterance to achieve the same goal. For instance, she could have produced an utterance of the coded symbol *good* while pointing at the berries, on the basis of which Peter would be warranted in inferring that Mary intended to communicate to him that the berries were edible.

This form of sophisticated metarepresentation is the normal situation in ostensive-inferential communication, as described in section 3.4.1, of which linguistic communication is the prime example. This ability, Sperber argues, is what made inferential communication possible, perhaps initially as a side effect, which then proved to be so highly beneficial that it created an environment conducive to the development of linguistic codes. From the development of a linguistic ability followed a co-evolutionary mutual enhancement of the two abilities.⁷⁷

This evolutionary scenario fits well with the essentialist view of underdeterminacy, discussed and supported in section 3.3. Indeed, it offers an explanation for why the linguistic codes that modern humans have are not capable of fully encoding speaker meanings, as observed by Carston:

⁷⁷ Some researchers claim that in addition to a developed theory of mind ability, human communication requires the ability to co-operate, that is, to participate in processes involving shared intentionality (e.g. Hurford 2007; Tomasello 2008). Sperber and Wilson (1986/1995) see ostensive-inferential communication as taking place within a 'mutual cognitive environment', that is, against the background of a set of assumptions that is manifest to a group of individuals (where an assumption *P* is manifest to an individual iff she is capable of representing *P* as true or probably true), and in which it is manifest to those individuals that they share that set of assumptions with each other. Every manifest assumption in a mutual cognitive environment is 'mutually manifest'. On this approach, communication involves the use of an ostensive stimulus in a mutual cognitive environment, where a communicator makes mutually manifest an intention to make some information manifest or more manifest to her audience, and where the communicative act itself provides evidence for the informative intention. Thus, within the relevance-theoretic framework, the claim is that this is sufficient for communication to take place.

[P]ublic language systems are intrinsically underdetermining of complete (semantically evaluable) thoughts because they evolved on the back, as it were, of an already well-developed cognitive capacity for forming hypotheses about the thoughts and intentions of others on the basis of their behaviour. (Carston 2002b: 30)

On this view, the development of polysemy could be seen as being ultimately motivated by the evolutionary dependence of language on a pre-existing mind-reading ability. I will end this section with an (inevitably speculative) suggestion about how polysemy might have evolved in the early manifestations of human language, against the background of an already evolved capacity for attributing mental states (specifically, communicative intentions) to others.

Suppose, in the very early days of language, our hominid ancestors possessed a lexical protolanguage, with a large learned lexicon of meaningful words, but no complex syntax (cf. Bickerton 1990; Jackendoff 2002). Most communication would then proceed in terms of ‘holophrases’, i.e. one-unit communicative acts, where the signals used had no internal structure (Tomasello 2008). If our ancestors had already developed an early form of metarepresentational theory of mind capacity, enabling a degree of intention-reading, such one-unit communicative acts could be quite rich in content (i.e. speaker meaning). The following scenario is borrowed from Dennett (1991: 197). It is originally a story about how autostimulation (inventing new paths of internal communication) might enhance cognitive organisation but it can be adapted to illustrate how polysemy might have developed as well. Dennett invites us to imagine a hominid one day idly drawing two parallel lines on the floor of his cave. “[W]hen he looked at what he had done, these two lines reminded him, visually, of the parallel banks of the river that he would have to cross later in the day, and this reminded him to take along his vine rope, for getting across” (ibid). Now, adapting this scenario, imagine the hominid (Peter) had a friend sitting next to him (Mary), watching him draw the two parallel lines. Suppose there was already a word, *pak*, for riverbank in Peter and Mary’s language. After having become aware that the two lines reminded him of the riverbank, Peter looks at Mary and utters ‘Pak!’, intending to remind her that they have to cross the riverbank later that day. Mary recalls their impending expedition and correctly attributes to Peter the intention to draw her attention to this fact, and perhaps she will also draw the inference that Peter thinks it’s time to start preparing for the trip. Peter then repeats his utterance

of 'Pak!' in a more playful manner, together with a pointing at the drawing to communicate to Mary that it was this drawing that reminded him of the riverbank and so of their upcoming venture. Mary responds 'Pak!' back at Peter communicating to him that she recognises the visual similarity between the drawing and the banks of the river. Peter might then proceed by drawing a vertical line that crosses the parallel ones while uttering the word for vine, with the intention to inform Mary that they have to bring with them a vine rope to cross the banks. (Notice how this communicative exchange is inferential and metarepresentational through and through).

This communicative exchange may have an additional consequence. Suppose that after having crossed the riverbank, having walked for a while, Peter and Mary decide (using holophrastic communication) to make camp for the evening. Peter gathers some wood to make a fire, which he throws rather carelessly into a pile on the ground. Mary, who doesn't think this is a good way to arrange the wood, sits down by the pile and starts arranging the logs into a stack. She takes a pair of logs, places them side-by-side with a distance between them, then takes another pair of logs and places them side-by-side on top of the previous two, in the alternate direction. She looks at Peter so as to capture his attention, imitates the distance between the logs with her hands and utters 'Pak!'. Having Peter's previous use of the word fresh in mind (his showing her the visual analogy between the river banks and the parallel lines drawn on the floor of the cave, thereby communicating to her information about their crossing the river banks), assuming that Peter also remembers this use of the word, Mary intends to communicate to Peter that the logs should be arranged in parallel (alternating) lines when making a fire. Peter, who remembers their earlier communicative exchange (that is, he remembers his previous use of the word and is simultaneously aware of Mary's familiarity with this use), infers that this is indeed what Mary intends to communicate to him. This may give rise to an extended use of the word *pak*, as a term describing not only the parallel banks of a river, but one that also has a more general and abstract sense, 'parallel lines'. The relation between *pak* (riverbank) and *pak* (parallel lines) would, of course, be transparent to Peter and Mary but it might be opaque to new learners who encounter the meanings in distinct contexts.

This little story is intended to show how the development of polysemy could arise naturally in an environment where the communicators are equipped with a

metarepresentational theory of mind ability, which allows for the code to underspecify their intended meanings. In contrast, in an environment where the communicators lack this ability, as in many forms of animal communication, there appears to be no reason, indeed no possibility, for polysemy to evolve. So, we may hypothesise that the direction of development of linguistic signals in a non-metarepresentational environment might be the opposite of their development in a metarepresentational context. In the absence of metarepresentation there might be a drift towards univocality (e.g. the development of acoustically distinct alarm calls in vervet monkeys for each distinct situation), while in the presence of metarepresentation there seems to be a drift toward polysemy, resulting from the more general phenomenon of underdetermination of speaker meanings by the language.

3.6 Implications of the pragmatic approach

In this chapter, I have argued that polysemy arises as a result of encoded lexical meanings being massively underdetermining of speaker-intended concepts.⁷⁸ Given this, I argued that the comprehension of polysemy calls for an analysis in terms of the workings of the pragmatic inferential system, rather than in terms of linguistic operations over lexically encoded content. At the deepest level, then, polysemy is, using Sperber and Wilson's (1998: 197) words, "the outcome of a pragmatic process whereby intended senses are inferred on the basis of encoded concepts and contextual information". An upshot of this pragmatic approach is that it provides a basis for explaining how polysemy arises and develops, and what motivates its proliferation in natural language. I suggested, drawing on arguments from Sperber (2000) and Carston (2002a), that our metarepresentational theory of mind ability might provide the cognitive basis for polysemy, and that it might have developed as a result of the co-evolution of the capacities for language and metarepresentation. More specifically, I

⁷⁸ This view, is, in principle, compatible with two different positions on the relation between encoded concepts and speaker-intended, communicated concepts: (i) encoded concepts usually underdetermine the speaker-intended concepts, but in some cases, the encoded concept is also the one that is communicated by the speaker (e.g. *happy* in 'How many *happy* people are there in the world today?'); (ii) encoded word meanings are highly abstract, schematic entities (rather than being actual concepts they act as 'pointers' to a conceptual domain, or bundles of information), from which the concept communicated must be pragmatically derived on each occasion of use. (i) is the position taken by Sperber and Wilson (1998), while the more radical view in (ii) is suggested by Carston (2002b: Chapter 5; 2010), as well as several others (Recanati 2004; Bosch 2007; Pietroski 2008; Pritchard 2009).

proposed that the component that allows us to infer speaker meanings from encoded linguistic meanings provides the motivation for polysemy in language. On this pragmatic account, polysemy poses no paradox, but is what we would expect from having languages that are not capable of fully encoding our thoughts.

An implication of this pragmatic approach is that all instances of polysemy should, in principle, be traceable back to the operation of a pragmatic process. This, of course, is not in opposition to the existence of conventional polysemy (in a language community or in the lexicon of individual speakers), for which identifying the pragmatic process on the basis of which the polysemy was originally derived would be a matter of diachronic analysis. Code-based models of communication, which treat language comprehension as a mechanical process of encoding and decoding of linguistic signals, and polysemy to be generated according to linguistic rules that operate over semantic representations (as on Pustejovsky's account), can at best give a synchronic account of polysemy (although, as I have argued, this is far from clear); they have nothing to say about the development of polysemy. The pragmatic inferential model of communication, however, which gives a more 'organic' account of the process of language comprehension, provides a framework in which the prospect for explaining polysemy development is much more promising.

A further implication of the claim that polysemy is a fundamentally communicative phenomenon is that, in principle, we should expect not just words but any simple ostensive stimulus to be susceptible to polysemy (e.g. manual and facial gestures). One example might be the use that car drivers make of a single flash of their two headlights to another driver, which could (in England), mean at least the following: (a) go ahead in front of me; (b) thanks for giving way to me; (c) watch out – there is traffic police up ahead. The general meaning of this signal might be something like 'friendliness indication', with the more specific meanings derivable on the basis of context (positioning of the cars on the road, whether or not one of the drivers has already flashed his/her lights, etc.). Another example might be smiles, which, depending on the context, can communicate a range of (related) feelings: amusement, affection,

sympathy, etc. (see Wharton 2009 for further examples of the use of gestures as ostensive/communicative stimuli).⁷⁹

In this chapter, I have taken an essentialist position on linguistic underdeterminacy, espousing Carston's (2002b) claim that the complexity and fine-grainedness of the thoughts we can entertain and communicate make them impossible to fully encode, against a view that treats underdeterminacy as a matter of 'convenient abbreviation'. However, I think that the essentialist position on linguistic underdeterminacy is fully compatible with the assumption that polysemy often *is* a matter of effort-saving convenience for the speaker, albeit not in a sense that entails the existence of eternal sentences. The very high level of polysemy in languages suggests that it is apparently easier for people to take old words and extend them to new meanings than to invent new words. This may be because, as Sperber and Wilson (1998: 198) have pointed out, the stabilisation of a new word in a language is a slow and relatively rare process that has to be coordinated over a large group of individuals over time. The kind of small-scale coordination involved in communication is a less elaborate affair, and with our pragmatic abilities enabling us to form hypotheses about the speaker's meaning on the basis of her linguistic utterances and highly accessible contextual information, there would, in most cases, be no need for a new word to describe something that may just as well be described by using an already existing word (with an extended meaning). For instance, many cases of metaphorically generated polysemy have the feel of being shorthand for concepts that could in principle be paraphrased (albeit not in a fully encoding way). Consider (17):

- (17) a. My dog always holds a ball in its *mouth*.
b. If you blow across the *mouth* of a bottle you can often get a note.
c. The village is located near the *mouth* of the river.
d. They built a fire at the *mouth* of the cave.
e. A car accident happened at the *mouth* of the tunnel.

The 'aperture' sense of *mouth* in (17)b.-e. appears to have been derived from the sense in (17)a. by way of metaphorical extension. Although clearly tokenings of this

⁷⁹ Thanks to Robyn Carston for pointing this out to me.

general sense, each of (17)b.-e. expresses a different concept that could be paraphrased; for instance, ‘the opening for filling or emptying the bottle’, ‘the place where the river enters the sea’, etc. However, the metaphorical sense of *mouth* may be used to communicate these concepts in a considerably more succinct and efficient way. So far, I have not considered the details of the specific pragmatic processes involved in the derivation of new senses (hence giving rise to polysemy), which include the process of metaphorical extension of lexical meanings. It is to this topic I turn in the next chapter.

Chapter 4

RELEVANCE-THEORETIC LEXICAL PRAGMATICS: THE *HOW* OF POLYSEMY

4.1 Introduction

In this chapter, I address the nature of the lexical pragmatic processes that are involved in the construction of polysemy. How do word meanings get extended to have different senses?

In the previous chapter, we discussed the issue of linguistic underdeterminacy; the systematic gap between sentence meanings and intended speaker meanings. Bridging this gap during utterance comprehension can be seen as a matter of solving a variety of pragmatic subtasks. On the relevance-theoretic approach, where intended speaker meanings are taken to be derived inferentially from encoded linguistic meanings, these tasks can broadly be divided into the following: (i) constructing an appropriate hypothesis about the intended explicit content, or the *explicature* of the utterance ('what is said'); (ii) constructing an appropriate hypothesis about the intended contextual assumptions, or the *implicated premises*;⁸⁰ and (iii) constructing an appropriate hypothesis about the intended contextual assumptions, or *implicated conclusions* (Wilson and Sperber 2004: 615). As mentioned in the previous chapter, arriving at a hypothesis about the explicit content of an utterance involves, in addition to a decoding of the linguistic expressions, resolution of ambiguities and indexical reference assignment, the application of other pragmatic enrichment processes such as resolution of ellipses (e.g. developing an utterance of 'On the top shelf' into a full proposition), interpretation of unarticulated constituents (e.g. specifying the location from where the person is leaving in an utterance of 'Jane's leaving'), and lexical adjustment (e.g. specifying the appropriate interpretation of *open* in an utterance of

⁸⁰ The set of premises used in the interpretation of an utterance constitutes the *context*, which, in relevance theory, is seen as a psychological construct, as "a subset of the hearer's assumptions about the world", and includes assumptions derived from the observation of the physical environment, the hearer's encyclopaedic knowledge, memories and beliefs, as well as the preceding linguistic context (Sperber and Wilson 1986/1995: 15).

'John *opened* the car').⁸¹ How the hearer performs the task of adjusting lexical meanings in the construction of explicatures is the topic of this chapter.

The last three decades have seen a growing interest in the interaction between semantic and pragmatic aspects of word meanings. A central question is how, and to what extent, linguistic underdeterminacy applies at the level of individual words. The field of study that has come to be known as *lexical pragmatics* explores the pragmatic processes that enable us to derive contextual meanings from encoded meanings of individual lexical items, and has attracted the attention of scholars from various disciplines, including philosophers (Grice [1967] 1989; Lewis 1979; Searle 1983; Travis 1985; Recanati 1995, 2004), computational linguists (Pustejovsky 1995a; Copestake and Briscoe 1996; Lascarides and Copestake 1998; Asher forthcoming), pragmatists (Blutner 1998, 2002; Carston 1997, 2002b; Wilson and Carston 2007), and psychologists (Barsalou 1987; Gibbs 1994; Glucksberg 2001; Murphy 1997a, 2002). As we will see in this chapter, the data for lexical pragmatics coincides in large part with the polysemy data.

My aim in this chapter is to show how polysemy arises as a result of the on-line construction of occasion-specific senses by pragmatic processes that apply at the level of individual words, within the overall process of forming hypotheses about the explicit content of utterances. More specifically, I will outline how the relevance-theoretic approach to lexical pragmatics, according to which there is a single pragmatic process that adjusts the meanings of individual words in different directions (Carston 1997, 2002b; Wilson and Carston 2006, 2007; Vega-Moreno 2007; Sperber and Wilson 2008), can account for the construction of occasion-specific senses during utterance comprehension. The chapter is structured as follows. In section 4.2, I outline the relevance-theoretic account of lexical pragmatics, focusing on two varieties of lexical adjustment, narrowing and broadening of conceptual content, which I claim are involved in giving rise to polysemy. In section 4.3, I suggest a reanalysis of Pustejovsky's

⁸¹ Relevance theory does not, however, claim that the different subtasks involved in utterance comprehension are sequentially ordered (i.e. that the hearer first decodes the linguistic meaning of the utterance, then computes the explicit content, then decides on which contextual assumptions he should bring to bear, and then, finally, derives the implicated conclusions). Rather, it is assumed that utterance comprehension proceeds through a process of 'mutual parallel adjustment' of hypotheses about explicatures, contextual assumptions and contextual implications (Wilson and Sperber 2004: 617). I return to the issue of 'mutual adjustment' in section 4.2.1.

(1995a) cases of polysemy in terms of the relevance-theoretic lexical pragmatic account (section 4.3.1). I argue that this approach provides a simpler, more unified account, and avoids some of the problems associated with the generative lexicon theory, as discussed in Chapters 2 and 3. Then, I discuss some further implications of the relevance-theoretic approach to polysemy, concerning the notion of sense relations and the role of pragmatic narrowing and broadening in lexical semantic change (section 4.3.2). Finally, I discuss the issue of prepositional polysemy, which is a major topic in the cognitive linguistics literature but has received little attention in relevance theory, and sketch a possible direction for a relevance-theoretic analysis of the multiple (related) meanings of one particular preposition (section 4.3.3).

4.2 The pragmatic adjustment of lexical meaning

In this section, I consider two varieties (or outcomes) of the contextual adjustment of encoded word meanings, lexical narrowing and broadening, and outline the relevance-theoretic approach to lexical adjustment.

Lexical narrowing ('enrichment', 'strengthening', 'specification', 'precisification') is the case where the concept expressed by the use of a word has a more restricted denotation than that of the linguistically-encoded concept; that is, it denotes a proper subset of the linguistically-specified denotation (Atlas 1989; Bach 1994; Recanati 1995, 2004; Carston 1997, 2002b, 2010; Levinson 2000; Wilson 2003; Wilson and Carston 2007). Some examples are given below:

- (1) Susan needs a *knife* so she can start eating.
- (2) The blacksmith made a set of new *shoes* for the horse.
- (3) a. I've started *collecting* stamps.
b. The books are *collecting* dust on the shelves.
- (4) *cut* the grass/one's hair/a cake, etc.
- (5) It's *cold* outside.
- (6) *red* car/apple/watermelon/hair/eyes, etc.
- (7) Sally drives/walks/types *fast*.

The noun *knife* in (1) is used to denote a specific kind of knife ('table knife'), and thus picks out a subset of the more general, linguistically-encoded meaning of the noun. (2) is a similar case in which *shoes* is used to denote a specific kind of shoe ('horseshoe'), a concept that is clearly narrower than the one linguistically-encoded. If the encoded meaning of the verb *collect* in (3) is taken to be 'gather', the meanings that are expressed in (3)a., 'systematically seek and acquire', and (3)b., 'accumulate', could both be seen as denoting subsets of the linguistically-encoded denotation. The verb *cut* in (4) conveys different actions depending on the object denoted by the complement (and the wider context of utterance), each of which would have a more restricted denotation than the encoded meaning. The adjective *cold* in (5) encodes a scalar concept. The degree of coldness expressed on a given occasion is contextually determined, and could be narrowed down to different degrees, depending on the situation of utterance. Each increase in the degree of coldness would contribute to a further restriction of the linguistically-specified denotation (e.g. the communicated meaning of *cold* in the expression *cold weather* uttered on a winter's day in Norway would be narrower than the one communicated if uttered in the context of a holiday in Greece). In (6) the adjective *red* denotes different hues of red in describing a car, an apple, a watermelon, someone's hair, eyes, etc., In each case, *red* communicates a concept that is more specific than that encoded by this adjective. Finally, in (7), the adverb *fast* expresses different concepts when it describes the action of driving, walking, typing, etc. In each case, the concept communicated has a narrower denotation than that linguistically-encoded by *fast*.

Lexical *broadening* (or 'loosening') is the case where the concept expressed by use of a lexical item is more general than that which is linguistically-encoded; it denotes a proper superset of the linguistically-specified denotation (Sperber and Wilson 1985/1986, 2008; Gibbs 1994; Recanati 1995, 2004; Carston 1997, 2002b; Glucksberg 2001; Wilson 2003; Wilson and Carston 2006, 2007; Vega-Moreno 2007). Relevance theorists talk about *approximation*, *hyperbole*, *metaphor* and *category extension* within the general type of lexical broadening. Here are some illustrations:

- (8) The U.S. states Wyoming and Colorado are *rectangular*.
- (9) It's *boiling* outside.

- (10) Bill is a *wizard* in finance.
- (11) The novel was *slaughtered* by critics.
- (12) a. The weightlifter complained that his weights were *unbalanced*.
 b. UN officials argued that the report was *unbalanced*.
 c. I knew how *unbalanced* Paula had been since my uncle died.
- (13) Joan has a cold and needs a *Kleenex*.

The adjective *rectangular* in (8) is used to express ‘approximately rectangular’, representing a small departure from the linguistically-encoded concept (on the assumption that the concept which is encoded by *rectangular* is a strict geometrical concept). In (9) the adjective *boiling* is used hyperbolically to express ‘extremely hot (with regard to normal weather conditions)’, representing a further departure from the linguistically-encoded concept. The metaphorical use of the noun *wizard* in (10) represents a radical departure from the linguistically-encoded concept, contributing to the interpretation that Bill is very skilled and imaginative in financial matters. Similarly, the metaphorical use of *slaughtered* in (11), conveying that the novel was heavily criticised, also involves a radical broadening of the concept linguistically-encoded by the verb. In (12), the concept encoded by the adjective *unbalanced* is broadened when used to describe a report, as in (12)b. (expressing ‘biased’), or a person, as in (12)c. (expressing ‘mentally ill’). These broadened concepts are generally thought of as metaphorical extensions of the concept expressed in (12)a. (which is identical to a very close approximation of the concept encoded by *unbalanced*), but the broadenings have taken different directions. The example in (13) can be seen as an instance of ‘category extension’, in which the brand name *Kleenex* is used to denote a broader category of disposable tissues, including items from less salient brands.

Lexical narrowing and broadening are typically given different treatments in the semantics and pragmatics literature. Moreover, it is not uncommon to analyse the different varieties within each type in different ways. Levinson (2000: 114), who works within a neo-Gricean framework, analyses a range of instances of frequently occurring lexical narrowing (e.g. the narrowing of *and* to mean ‘and as a consequence’ or *if* to ‘if and only if’ or gender stereotypes for ‘secretary’ or ‘fire fighter’) as cases of default inference. He takes these to be governed by an informativeness principle (*I-principle*,

for short), which includes the following enrichment rule: “Amplify the informational content of the speaker’s utterance, by finding the most *specific* interpretation, up to what you judge to be the speaker’s m-intended point”. On the basis of this rule, semantically general expressions automatically undergo narrowing in the derivation of a (generalised) implicature and are only restored to the encoded concept if the resulting implicature clashes with a salient contextual assumption. A different kind of approach is that of Pustejovsky (1995a), who, as we have seen in the previous chapters, analyses the contextual specification of evaluative adjectives such as *good*, *fast*, and so on, as being derived through a lexicon-internal process of ‘selective binding’, enabling them to modify only an aspect of the linguistically specified meaning of the head noun (e.g. *good* in *good knife* comes out as ‘good for cutting’). Within more philosophically oriented approaches, such as the indexicalist semantic account espoused by Szabó (2001) and others, the logical form of an evaluative adjective such as *good* is seen as containing a variable which needs to be contextually saturated.

Lexical broadening, on the other hand, is standardly given different analyses depending on subtype (approximation, hyperbole, metaphor, category extension). For instance, following Lewis (1979), some philosophers of language analyse approximations of the kind in (8) as involving the contextual fixation of a standard of precision parameter, contained in the logical form of such expressions. Lasersohn (1999) proposes a different analysis, on which the denotation of linguistic expressions such as *rectangular* above are, relative to a given context, associated with what he calls a ‘pragmatic halo’. The pragmatic halo of a given expression in a context *C* is the set of objects of the same logical type as the denotation itself (which is included in the set), each of which “is understood to differ from the denotation only in some respect that is pragmatically ignorable in that context” (ibid. 526), allowing, in this way, a deviation from what is strictly speaking true. On the other hand, hyperboles and metaphor are still quite widely analysed along Gricean lines (Grice [1967] 1989), as blatant violations of Grice’s first maxim of Quality (‘Do not say what you believe to be false’) which trigger the hearer’s search for a related implicature. A radically different approach is taken by cognitive linguists, who analyse metaphorical interpretations in terms of mappings from one conceptual (source) domain to another conceptual (target) domain (Lakoff and Johnson 1980, 1999; Kövecses 1986; Lakoff 1987, 1993; Lakoff and Turner 1989;

Fauconnier and Turner 2002). The assumption is that knowledge from one conceptual domain is used to reason about another conceptual domain – this is called a ‘conceptual metaphor’. Much-cited examples are ARGUMENT IS WAR, TIME IS MONEY, LOVE IS A JOURNEY, etc. Such conceptual metaphors are taken to underlie linguistic metaphors (so, for instance, metaphorical utterances such as *We’re at a crossroads in our relationship*, *Our relationship has hit a dead-end*, *We are going separate ways*, etc. are instantiations of the underlying conceptual metaphor LOVE IS A JOURNEY). Although it has been extremely influential as a theory of metaphor in cognitive linguistics as well as in several other disciplines, the conceptual metaphor theory has received considerable criticism (Murphy 1996, 1997b; Stern 2000; McGlone 2001; Haser 2005; Vega-Moreno 2007). In particular, it has been pointed out that its implications for the representation of lexical concepts and on-line processing are unclear.⁸²

4.2.1 A single process of *ad hoc* concept formation

Relevance theory stands out from the approaches mentioned above by its proposal of a unified account of lexical narrowing and broadening (including approximation, hyperbole and metaphor), treating them as outcomes of a single pragmatic inferential process that adjusts the meanings of individual words, in accordance with the hearer’s expectations of relevance (Carston 1997, 2002b, 2010; Wilson and Sperber 2002; Wilson 2003; Wilson and Carston 2007; Sperber and Wilson 2008). The occasion-specific senses that are constructed on the basis of this pragmatic process are called *ad hoc* concepts, a term adopted (and adapted somewhat) from the psychologist Larry Barsalou (1983, 1987).⁸³ Thus, in lexical narrowing, the *ad hoc* concept will be more specific than the one

⁸² For instance, does it follow from the conceptual metaphor theory that many concepts are not understood directly but only with reference to a concept from a different domain? And how do speakers and hearers determine which of the properties and inferences associated with the source concept that are mapped onto the target, and how are they selected during on-line comprehension? Recently there have been some suggestions of how the cognitive linguistic and relevance-theoretic approaches to metaphor can be combined, despite some fundamental differences between them (Gibbs and Tendahl 2006; Tendahl and Gibbs 2008; Wilson 2009). I return to consideration of the cognitive linguistic approach, as well as the suggestion for a combination of the two theoretical frameworks, in Chapter 6 in connection with the discussion of metonymy.

⁸³ In a series of experiments, Barsalou (1987) found that people spontaneously construct novel, context-dependent (*ad hoc*) categories in working memory (e.g. TOURIST ACTIVITIES TO PERFORM IN BEIJING, THINGS TO PACK IN A TOILET BAG, WHAT TO TAKE FROM ONE’S HOME DURING A FIRE, etc.). Once they were structured, the *ad hoc* categories functioned as coherent categories and exhibited typicality effects in a similar way as stable categories do.

linguistically-encoded, denoting a proper subset of the linguistically-specified denotation, while in lexical broadening (including approximation, hyperbole, metaphor and category extension), the *ad hoc* concept is broader, denoting a proper superset of the linguistically-specified denotation. And there are some *ad hoc* concepts that involve both narrowing and broadening of the encoded concept so that their denotation merely overlaps with that of the encoded concept. Before I turn to the details of this approach a brief consideration of the relevance-theoretic stance on encoded lexical meaning is in order.

The general claim about lexical meanings in relevance theory is that they are not uniform entities. More specifically, Sperber and Wilson (1998) distinguish between three types of linguistic encodings: (i) words that do not encode concepts at all (but rather ‘procedures’, which constrain, or indicate, how the process of pragmatic inference should proceed, e.g. indexicals and discourse connectives, cf. Blakemore 1987; 1992; 2002); (ii) words that appear to encode so-called pro-concepts, i.e. words that have some (underspecified) conceptual content but whose semantic contribution must be contextually specified (e.g. words like *my*, *have*, *near*, *long*); and (iii) words that may be said to encode full-fledged concepts (e.g. *dog*, *horse*, *water*). Their central assumption, however, is that most words, whether or not they encode full concepts, behave *as though* they encoded pro-concepts, since the concepts they are used to convey on different occasions are pragmatically determined on the basis of contextual relevance.

For those words that encode full-fledged concepts, the claim is that they encode atomic concepts (in the sense of Fodor 1981, 1998). Relevance theory takes the Fodorian critique against decompositional (specifically, definitional) accounts of lexical meaning to be decisive (cf. Chapter 2), and the most plausible alternative to be a simple mapping from lexical form to mental concept (e.g. the lexical form *rabbit* encodes the unstructured concept RABBIT). A mentally represented concept, a constituent of the ‘language of thought’ (Fodor 1975, 2008), is seen as an address (or entry) in memory that may give access to three types of information: (i) the *logical* properties of the concept (e.g. one-way meaning postulates, such as RABBIT → KIND OF ANIMAL);⁸⁴ (ii) a set

⁸⁴ Recall the discussion of meaning postulates in connection with Fodor’s conceptual atomism in Chapter 2. While Fodor (1998) denied the existence of meaning postulates (based, in large part, on Quine’s (1999 [1953]) pessimism about the possibility of drawing a principled analytic/synthetic distinction), Horsey

of assumptions, or *encyclopaedic* information, about the denotation of the concept, that is, conceptually represented assumptions and beliefs, including stereotypes and culture-specific information about the denotation, and also, in many cases, imagistic and/or sensory-perceptual representations (e.g. RABBITS HAVE FUR, RABBITS ARE EDIBLE, RABBITS HAVE FOUR LEGS, RABBITS LOOK LIKE THIS: [MENTAL IMAGE], etc.);⁸⁵ (iii) the *lexical* (i.e. phonological and syntactic) information connected with the linguistic form that encodes the concept (Sperber and Wilson 1986/1995: 85-93).

In relevance theory, lexical adjustment resulting in the construction of narrowed or broadened *ad hoc* concepts involves taking the encoded concept and its associated logical and encyclopaedic information, together with a set of contextual assumptions, as input to the inferential process of constructing hypotheses about the speaker-intended meaning. Recall that relevance theory takes on-line utterance comprehension to involve the resolution of a number of pragmatic subtasks, including the construction of hypotheses about the explicit content, implicated premises, and contextual implications of an utterance. Typically, lexical adjustment in the form of the construction of *ad hoc* concepts is the result of a process of 'mutual parallel adjustment' of tentative hypotheses about explicit content, implicated premises and contextual implications, which are incrementally modified against the background of the hearer's context-specific expectations of relevance.

As a brief illustration of how this mutual adjustment process is assumed to work in an instance of lexical narrowing, consider the following example given by Carston and Powell (2006):

(2006) argued in favour of a psychological notion of meaning postulates, which accommodates people's intuitions about analyticities. This is the view espoused by relevance theory, where the logical entry for a concept is seen as consisting of a set of deductive rules (or meaning postulates), which form part of our mental logic (our 'deductive device', cf. Sperber and Wilson 1986/1995; Horsey 2006), and which apply reliably to any representation that meets their input conditions (e.g. the concept RED triggers the inference RED → COLOUR, the concept MOTHER triggers the inference MOTHER → PARENT, etc.). Such meaning postulates are thought to be content-constitutive of a concept, in the sense that someone could not be said to fully master a concept without also having completed the logical entry for that concept. It should be noted, however, that among relevance theorists there is not a perfect consensus as to the status and role of meaning postulates within the theory, nor as to whether the distinction between logical and encyclopaedic knowledge is warranted.

⁸⁵ With its distinction between logical and encyclopaedic information associated with concepts (Sperber and Wilson 1986/1995; Horsey 2006), relevance theory thus distinguishes itself sharply from the 'purist' atomist position now held by Fodor (1998), who, as already mentioned, denies any role for an analytic/synthetic distinction (in fact, denies that there *is* any such distinction).

(14) A (to B): Be careful. The path is *uneven*.

The second part of A's utterance is likely to achieve relevance for B by offering some kind of explanation as to why, or in what way, he should be careful. The concept encoded by *uneven* is very general (it is probably true of every path), but as B is looking for a particular kind of implication (e.g. he might trip over, he should take small steps, he should keep his eye on the path) he will narrow down the encoded concept UNEVEN so as to arrive at an interpretation that warrants the implication(s) that satisfy his context-specific expectation of relevance. The result is the construction of the *ad hoc* concept UNEVEN*, whose denotation is significantly narrower than that of the encoded concept. This process can be seen to rely on 'backward' inference from expected contextual implications to the explicitly expressed proposition containing the *ad hoc* concept.⁸⁶

Thus, in relevance theory, lexical adjustment in terms of the construction of narrowed or broadened *ad hoc* concepts is seen as emerging as a by-product of the hearer's search for an interpretation that satisfies his expectation of relevance. As mentioned in the previous chapter, utterance understanding is taken to employ a dedicated comprehension heuristic, which provides the hearer with the following procedure for constructing a hypothesis about the speaker's intended meaning (Wilson and Sperber 2004: 615):

- (i) Follow a path of least effort in deriving cognitive effects: Test interpretive hypotheses (disambiguation, reference resolutions, implicatures, etc.) in order of accessibility.
- (ii) Stop when your expectations of relevance are satisfied.⁸⁷

⁸⁶ See Wilson and Sperber (2002) and Wilson and Carston (2006) for more detailed accounts of the mutual adjustment process and the role of 'backward inference' in lexical interpretation.

⁸⁷ The comprehension heuristic is jointly suggested by the Communicative Principle of Relevance, 'Every act of ostensive communication communicates a presumption of its own optimal relevance', and the definition of optimal relevance: An utterance is optimally relevant if (a) it is at least relevant enough to be worth processing, and (b) it is the most relevant one compatible with the speaker's abilities and preferences. Thus, an optimally relevant utterance should achieve enough cognitive effects to be worth processing, and avoid causing the hearer any unjustifiable effort in achieving those effects. A consequence of this ban on wasted processing effort is that the first satisfactory interpretation is the only satisfactory interpretation, that is, if an utterance has a highly accessible interpretation which is relevant in the expected way, this is the only one that the hearer is justified in choosing (Wilson 2008).

The comprehension heuristic states that when the hearer arrives at an interpretation of the utterance that satisfies his expectations of relevance, this is where he will stop. Unless there is evidence to the contrary, this interpretation will be the most plausible hypothesis about the speaker's intended meaning. However, given the non-demonstrative nature of pragmatic inferences, this hypothesis could turn out to be false, but it is the best the hearer can do.

As an illustration, consider the exchange in (15), uttered in the context of the extensive vaccination of the population against swine flu:

(15) Bill: Has John had the vaccine yet?

Susan: He's already had *flu*.

A schematic outline of how Bill might apply the comprehension heuristic in the interpretation of Susan's utterance is given in the table below (where Bill's interpretive hypotheses are found to the left and his basis for deriving them to the right):⁸⁸

(16)

(a) SUSAN HAS SAID THAT [HE _x HAS ALREADY HAD FLU].	Decoding of Susan's utterance. Embedding of the decoded logical form into a description of Susan's ostensive behaviour.
(b) Susan's utterance will be optimally relevant to Bill.	Expectation raised by the recognition of Susan's utterance as an act of ostensive communication.
(c) Susan's utterance will achieve relevance by answering Bill's question.	Expectation raised by (b), and the fact that this answer would be the most relevant to Bill at this point.
(d) Someone who has already had swine flu doesn't need to get the vaccine.	First contextual assumption activated that satisfies the expectation in (c). Accepted as an implicit premise of Susan's

⁸⁸ See Wilson and Sperber (2004) and Sperber and Wilson (2008) for more detail and further examples of how the comprehension heuristic is thought to work.

	utterance.
(e) John has already had FLU* ('swine flu') ⁸⁹	First interpretation of the explicit content (explicature) of Susan's utterance, involving a narrowing of the encoded meaning of FLU to FLU* ('swine flu'). The explicature in (e) would combine with (d) to satisfy (c).
(f) John doesn't have to get the vaccine.	Implicature of Susan's utterance, inferred from (d) and (e). Accepted as the intended implicated conclusion of Susan's utterance.

Bill decodes Susan's utterance, and takes this decoded content (an incomplete logical form) as input to the inference process to follow, and assumes that her utterance is optimally relevant to him. Bill is interested in an answer to his question, that is, whether John has got the vaccine against swine flu, and thus expects Susan's utterance to achieve relevance by providing him with this information. The first contextual assumption to come to Bill's mind that satisfies this expectation is that 'Someone who has already had swine flu doesn't need to get the vaccine'. Bill takes this as an implicit premise of Susan's utterance. The first interpretation of the explicit content of Susan's utterance involves (in addition to the assignment of John as the referent of 'He') a narrowing of the encoded content of FLU to the *ad hoc* concept FLU*, denoting 'swine flu', i.e. a subset of the linguistically-specified denotation. The encoded concept would not achieve relevance in the expected way, as someone who has had flu but not swine flu would not be exempt from taking the vaccine. Together, the explicature 'John has already had FLU*' and the contextual assumption that 'Someone who has already had swine flu doesn't need to get the vaccine' would satisfy the expectation that Susan's utterance should provide an answer to Bill's question. They would also form the basis for the derivation of the implicated conclusion of her utterance, 'John doesn't need to get the vaccine', which Bill accepts as Susan's implicated conclusion (the implicature of her utterance).

⁸⁹ *Ad hoc* concepts are, by convention, marked with an asterisk.

As this shows, relevance theory treats the process of lexical adjustment as arising within the overall process of arriving at a confirmed hypothesis about the speaker's intended meaning, a process which is driven by the hearer's search for an interpretation that satisfies his occasion-specific expectations of relevance. Often, as already mentioned, the hearer may bring to the comprehension process a specific expectation of how the utterance will be relevant to him (i.e. the cognitive effects its processing is expected to yield). Here Bill expected Susan's utterance to provide an answer to his question. This expectation contributed, via a 'backwards' inference process, to his identification of the explicature and implicated premise.

The example in (15) involved lexical narrowing. In what follows I will illustrate how lexical broadening, including metaphor, is treated along exactly the same lines on the relevance-theoretic approach. Consider the following exchange:

(17) Stephen: I wonder if John's firm is going to make it through the crisis.

Kate: I wouldn't worry too much. John is a *wizard* in finance.

Below is a schematic illustration of how Stephen might interpret the last part of Kate's utterance, 'John is a *wizard* in finance' (again, the interpretive hypotheses are given to the left of the tableau, the basis for deriving them to the right).

(18)

(a) KATE HAS SAID THAT [JOHN _x IS A WIZARD IN FINANCE].	Decoding of Kate's utterance. Embedding of the decoded logical form into a description of Kate's ostensive behaviour.
(b) Kate's utterance is optimally relevant to Stephen.	Expectation raised by the recognition of Kate's utterance as an act of ostensive communication.
(c) Kate's utterance will achieve relevance by providing an explanation for Kate's previous assertion and as responding to Stephen's worry.	Expectation raised by (b), and the fact that such an explanation (and response to Stephen's worry) would be most relevant to Stephen at this point.
(d) John's firm might not make it through	Background assumption already activated

the current financial crisis.	by Stephen's previous utterance.
(e) A wizard can make seemingly impossible things happen.	Assumption activated by the linguistically-encoded concept WIZARD and the assumption in (d). Accepted as an implicit premise of Kate's utterance.
(f) John can make seemingly impossible things happen in finance, and will be able to take his firm through the crisis.	Implicit conclusion derivable from (e), together with an appropriate interpretation of Kate's utterance which satisfies the expectation of relevance raised by (c), by contradicting and (possibly) eliminating the assumption in (d). Tentatively accepted as an implicit conclusion of Kate's utterance.
(g) John is a WIZARD* in finance.	Interpretation of the explicit content of Kate's utterance, where the meaning of the encoded concept WIZARD is broadened to include people who can make seemingly impossible things happen, which together with (e) would imply (f). Accepted as the explicit meaning of Kate's utterance.
(h) John is a WIZARD* in finance and will be able to take his firm through the crisis.	First overall interpretation of Kate's utterance, which includes a specification of its explicit content and implicature, that satisfies Stephen's expectations of relevance raised by (b) and (c). Accepted as Kate's intended meaning.

As this shows, the metaphorical meaning of *wizard*, WIZARD*, having as its denotation people who make seemingly impossible things happen, is arrived at via a mutual adjustment of contextual assumptions, explicit content, and contextual implications guided by expectations of relevance. The broadened concept in (17) is inferred in

exactly the same way as the narrowed concept in (15).⁹⁰ Thus, lexical narrowing and broadening are not seen as two functionally different kinds, but as involving the same processes of meaning construction, sometimes leading to a narrowing of the encoded concept, sometimes to a broadening. Furthermore, narrowing and broadening often combine in the construction of a single *ad hoc* concept. Consider Carston's (2002b: 328) example in (19):

(19) John's a real *bachelor*.

The denotation of the *ad hoc* concept BACHELOR* might include married men who behave like stereotypical bachelors (and so be broader than the linguistically-encoded concept), but exclude other men who are unmarried but do not satisfy certain stereotypical assumptions about bachelors, e.g., the pope (and so be narrower than the linguistically-encoded concept).⁹¹

4.2.2 Implications of the relevance-theoretic account

The unified relevance-theoretic account of lexical narrowing and broadening has several implications. First, it makes it possible to see different types of broadening as forming a continuum, ranging from literal uses through approximation and category extension to 'figurative uses', such as hyperbole and metaphor. Strictly literal interpretations and metaphorical uses occupy each of the extreme ends of the continuum, with the other varieties of broadening falling into the space between them (with hyperbole being closer to the metaphor end, involving a greater departure from the linguistically-encoded denotation, and approximation closer to the literal end). For instance, the verb phrase in (20) below can be interpreted literally, as a hyperbole, or as a metaphor (Wilson 2003):

(20) That book *puts me to sleep*.

⁹⁰ The outlines of the on-line interpretive process leading to a narrowing of the linguistically-encoded concept in (15) and a broadening in (17) given here are, of course, very schematic, and there is obviously no claim that the hearer necessarily goes through exactly these inferential steps with exactly those premises and conclusions or that the process takes place in this particular order.

⁹¹ In fact, Sperber and Wilson (2008: 95) claim that a combination of narrowing and broadening is involved in most metaphors, as in the example below:

(i) Joan is an *angel*.

Here they see the communicated meaning of *angel*, ANGEL* as involving both a narrowing of the linguistically-specified denotation to include only prototypically good angels, excluding harmful angels, angels of wrath or fallen angels, and a broadening to describe all very good, caring people.

And second, following from this, the relevance-theoretic account of approximation, hyperbole and metaphor as varieties of a single phenomenon of lexical broadening involves an abandonment of the traditional distinction between literal and figurative meaning, as it is conceived in classical rhetoric and in Gricean pragmatics.⁹² The claim is that the interpretations of approximation, hyperbole and metaphor require the exact same interpretive mechanism that is involved in the comprehension of literal utterances.⁹³ Third, it follows from the unified account of lexical pragmatic processes proposed by relevance theory that lexical narrowing and broadening, including ‘figurative uses’ such as hyperbole and metaphor, contribute to the truth-conditional content of utterances, that is, to their explicit content (explicature, what is said or asserted). This view (which Recanati 1993 calls ‘truth-conditional pragmatics’) represents a radical departure from the standard Gricean view, which posits a minimal semantics where the processes contributing to ‘what is said’ are confined to saturation of indexicals and disambiguation of linguistic ambiguity. Recently, however, there appears to be a growing agreement that processes of lexical adjustment contribute to the explicit content of utterances (Recanati 1995, 2004; Carston 1997, 2002b; Kintsch 2000, 2001; Stern 2000; Glucksberg 2001; Wilson and Carston 2006, 2007; Sperber and Wilson 2008). Wilson and Carston (2007) provide several arguments for a ‘truth-conditional account’ of lexical pragmatic processes. A first argument is the existence of neologisms that are derived via *ad hoc* concept construction.⁹⁴ If such clearly pragmatically derived

⁹² However, this does not deny the fact that metaphors are often used as expressions of style, in order to create *poetic effects*, in relevance-theoretic terms. A poetic effect is “the peculiar effect of an utterance which achieves most of its relevance through a wide array of weak implicatures” (Sperber and Wilson 1986/1995: 222). Metaphors are seen as particularly well suited to create such poetic effects.

⁹³ The view that there are no interpretive mechanisms specific to metaphor, that it requires the same mechanisms as those involved in the interpretation of literal utterances, is shared by several other scholars, working within different theoretical frameworks (see, for instance, Lakoff and Johnson 1980; Gibbs 1994; Kintsch 2000, 2001; Glucksberg 2001; Atlas 2005; Evans 2009).

⁹⁴ An example of such a neologism is the expression *thorbjorned* in the following utterance (attested: <http://www.nytimes.com/2009/10/11/opinion/11brenner.html?ref=opinion>):

(i) Wow, those Scandinavians completely *thorbjorned* my hard-earned political capital.

This utterance occurred in a *New York Times* commentary about the 2010 Nobel Peace Prize laureate Barack Obama, to whom the utterance was attributed. The journalist was describing what she imagined to be Obama’s reaction to “receiving such a surprising and politically tone-deaf honor”. The coined verb *thorbjorned* (communicating the *ad hoc* concept THORBJORNED*, roughly paraphrasable as ‘destroyed’, or ‘ruined’) is derived from the first name of the leader of the Nobel Peace Prize committee, Thorbjørn Jagland. As the verb *thorbjorned* has no meaning other than the pragmatically constructed neologism above, there would be no proposition expressed by this utterance unless the *ad hoc* concept THORBJORNED* is taken to contribute to its explicit content.

concepts are not taken to contribute to the proposition expressed by the utterances in which they occur, it is hard to see that there would be any proposition expressed at all by such utterances. A second argument is the evidence that it is pragmatically adjusted meaning, and not linguistically-encoded meaning, that falls within the scope of sentence operators such as negation, conditionals, disjunctions, etc. (e.g. 'If John has already had FLU*, he doesn't need the vaccine', cf. (15) above). A third argument comes from the fact that lexical narrowing and broadening may lead to semantic change, so that for some people, a pragmatically derived lexical meaning may have become lexicalised (e.g. the metaphorical meaning of *wizard* in (17) above), requiring disambiguation, whose contribution to the explicit content of an utterance is widely agreed on. For other people, even within the same speech community, the same word may have only a single encoded sense and the others derived through *ad hoc* concept construction. Wilson and Carston argue that it seems implausible that one proposition would be expressed in the case where the hearer has two lexicalised concepts for a lexical form, and a different one in the case where he has only a single concept stored for the same lexical form, which points to *ad hoc* concepts also contributing to the proposition expressed. A final argument is that if lexical narrowing and broadening can combine in the interpretation of a lexical concept, as shown by (19) above, it seems implausible that only one of them (i.e. narrowing) should contribute to the proposition expressed (cf. Carston 1997).⁹⁵

Finally, the relevance-theoretic account takes the processes involved in lexical adjustment to be genuinely *inferential*, in as much as the hearer's goal is to arrive at a warranted conclusion about the speaker's intended meaning, with the encoded meaning playing an evidential role. For instance, as shown by (18) above, the move from the premise 'Mary has said that John is a WIZARD in finance' to the conclusion 'Mary meant that John is a WIZARD* in finance' rests on a non-demonstrative inference process that takes a set of premises as input and yields as output a set of conclusions which are wholly derivable from the premises. This is perhaps the most contentious aspect of the

⁹⁵ The truth-conditional account of lexical pragmatic processes is supported by a number of experimental studies of metaphor comprehension. Specifically, these studies suggest that people understand metaphorical meanings as quickly and automatically as they understand literal meanings (Glucksberg and Keysar 1990; Glucksberg, McGlone, and Manfredi 1997; Gibbs 1992, 1994; Glucksberg 1998, 2001). This goes against the view that metaphorical interpretations depends on the prior computation of a 'literal meaning' or 'what is said', which is then rejected, triggering a second stage of retrieval of the metaphorical meaning.

relevance-theoretic approach to lexical pragmatic processes, in particular its implications for metaphor comprehension, which, on most (non-Gricean) accounts, is analysed as an associative (i.e. non-inferential) process, involving some kind of domain mapping (i.e. systematic correspondences between different cognitive domains), feature or domain activation, etc. (see, for instance, Lakoff and Johnson 1980; Lakoff and Turner 1989; Recanati 1995, 2004; Kintsch 2000, 2001; Evans 2009).⁹⁶ A central argument in favour of an inferential account of lexical pragmatic processes, and specifically of metaphor, is that it would be considerably more constrained than a non-inferential account, which, it is claimed, could not avoid overgeneration (Wilson and Carston 2007).⁹⁷ I cannot get into this debate in more detail in this chapter, and I will broadly follow the relevance-theoretic assumption that lexical adjustment (including metaphor comprehension) is an inferential process. However, I return to consideration of this issue in Chapter 6, in connection with my proposal for a relevance-theoretic account of metonymy.

4.3 A lexical-pragmatic account of polysemy

In the previous chapter, I argued that polysemy arises as a result of encoded lexical concepts being massively underdetermining of speaker-intended concepts, and that, at

⁹⁶ An objection that is sometimes raised against inferential accounts of metaphor comprehension is that they are unable to account for so-called ‘emergent properties’, that is, properties that are conveyed by the metaphorical expression but which are not associated with any of the linguistically-encoded concepts in the utterance. A much-cited example is ‘My surgeon is a *butcher*’, in which the property conveyed by the concept BUTCHER* (that the surgeon is extremely incompetent, dangerous, etc.) is not standardly associated with either SURGEON or BUTCHER in isolation (see discussion in Wilson and Carston 2007). For proposals for a solution to this apparent problem within relevance theory, see Carston (2002b), Vega-Moreno (2007) and Wilson and Carston (2006).

⁹⁷ This issue is part of a broader debate in pragmatics concerning whether ‘what is said’ by an utterance is derived by pragmatic processes that are fundamentally associative (Recanati 1995, 2002a, 2004, 2007) or inferential (Sperber and Wilson 1986/1995; Carston 2002b, 2002a, 2007). The philosopher François Recanati distinguishes between primary and secondary pragmatic processes, where the former are associative and local, and the latter are properly inferential and global in the Gricean sense. Metaphor, on Recanati’s account would be a primary pragmatic process, hence local and associative. On the other hand, relevance-theorists take utterance interpretation to be a wholly inferential process, with a unitary, on-line pragmatic processing system which derives explicit content (‘what is said’) and conversational implicatures in parallel, driven by the hearer’s search for an interpretation that meets his occasion-specific expectations of relevance. As we have seen, this often involves a process of ‘mutual adjustment’, where a hypothesis about an implicature can both precede and shape a hypothesis about an explicature. Relevance theorists have claimed that Recanati’s distinction between primary and secondary pragmatic processes commits him to a sequential view of the derivation of implicatures, and hence makes him unable to account for the mutual adjustment phenomenon.

the deepest level, polysemy is the outcome of a pragmatic process whereby intended senses are inferred on the basis of encoded concepts and contextual information (Sperber and Wilson 1998). More specifically, then, given the account of lexical pragmatic processes outlined in this chapter, the claim would be that generally, polysemy is the outcome of a pragmatic process of *ad hoc* concept construction, resulting in communicated concepts that have narrower or broader denotations than those linguistically-encoded. The construction of *ad hoc* concepts (i.e. polysemy) takes place within the overall inferential process of forming hypotheses about speaker-intended meanings, which is driven by the hearer's occasion-specific expectations of relevance. As argued in Chapter 3, this is the very same overall inferential process that is responsible for the disambiguation of conventional polysemy (i.e. where two or several senses have become established in the mental lexicon of speakers). However, in the case of on-line lexical adjustment, this process is genuinely constructive; it takes encoded concepts, logical and encyclopaedic information, together with a set of contextual assumptions as input, and outputs an occasion-specific concept, which is either narrower or broader than the linguistically-encoded concept or overlaps with it (i.e. is both narrower in some respects and broader in others).

4.3.1 A re-analysis of the Pustejovsky (1995a) cases

The framework for analysing lexical adjustment in terms of *ad hoc* concept construction outlined here allows us to fine-tune the pragmatic account of Pustejovsky's (1995a) cases of polysemy (cf. Chapters 2 and 3). The relevant examples are repeated in (21)-(23) below:

- (21) a. John *baked* a potato.
b. John *baked* a cake.
- (22) a. Every chef needs a *good* knife.
b. 'American Pastoral' is a *good* book.
- (23) a. Mary *began* reading a book.
b. Mary *began* a book.

As we saw in Chapter 2, Pustejovsky analyses the verb *bake* as linguistically-encoding the change-of-state meaning which occurs in (21)a., and which is adjusted to the

creation sense in (21)b. by the generative process of ‘co-composition’. The application of this process is dependent on the semantics of the complement noun (which, in (21), turns on a semantically represented distinction between natural kinds and artefacts). However, on the pragmatic account pursued here, the creative sense of *bake* in (21)b. could be analysed as being arrived at through lexical broadening of the linguistically-encoded change-of-state sense, in which its denotation is extended to include the various pre-cooking activities required to bring something into existence (e.g. cake, pizza, bread, etc.).⁹⁸ In this case, an important step in the hearer’s construction of the *ad hoc* concept BAKE* will be his accessing of the encyclopaedic entry for the concept CAKE, which will contain assumptions about the way in which cakes come into being (i.e. as a result of a ‘creative’ process of baking). A crucial difference between the relevance-theoretic pragmatic analysis and that given by Pustejovsky is thus the status of the information used to derive the creative interpretation of *bake*: whether it is seen as being encoded as part of the semantics of the noun or merely contingent, stored as part of the hearer’s encyclopaedic knowledge. As argued in the previous chapter, a clear advantage of the latter analysis is that, instead of postulating default interpretations that can be overridden only in the case of explicit contextual evidence pointing to a different interpretation, it allows for the necessary flexibility in lexical interpretation (which is constrained, however, by the hearer’s expectations of relevance).

A possible objection to the analysis of the creative sense of *bake* as an instance of on-line *ad hoc* concept construction may be that this sense appears, in fact, to have become conventional. If this is correct, the linguistic form *bake* would have two linguistically-encoded senses, with distinct denotations. This may well be so (at least for some people), in which case the pragmatic analysis given above can be seen as providing an account of the semantic relation between the two senses: the creation sense may have become conventional as a result of frequent and routinised broadening of the linguistically-encoded change-of-state sense. I return to the issue of semantic change and stabilisation of senses in the next section.

⁹⁸ According to *The Oxford Dictionary of English Etymology* (1966/1985. Oxford: Clarendon Press), the basic sense of *bake* is ‘cook by dry heat’; thus, Pustejovsky’s assumption that the change-of-state sense is primary seems reasonable.

Turning now to Pustejovsky's examples of 'selective binding' in (22), these can be analysed as instances of lexical narrowing giving rise to *ad hoc* concepts (as hinted at in the previous chapter). As I said there, adjectives such as *good* can be seen as encoding highly general concepts, which have to be narrowed into more specific *ad hoc* concepts on each occasion of use (e.g. 'GOOD* knife', 'GOOD** book', 'GOOD*** mum', 'GOOD**** student', etc.). Depending on the context, *good* may express different *ad hoc* concepts in describing one and the same thing (e.g. a '*good* job' could be one that is well paid, offers interesting tasks, has an inclusive social environment, offers special benefits to the employees, gives a certain social status, and so on, each of which might involve a distinct *ad hoc* concept). As in the case of *bake* above, the pragmatic inferential process will take as input information stored in the encyclopaedic entry of other lexical concepts in the utterance in the derivation of the *ad hoc* concept. In the case of *good*, encyclopaedic information stored about the concept encoded by the head noun is of particular importance. However, this information is not restricted to a particular, linguistically-specified purpose (e.g. 'knives are for cutting'), but might include any information relevant to the interpretation of the speaker-intended meaning of the adjective, as shown by the following examples:

- (24) a. To become a member of Billy's exclusive gang you had to have a *good* knife.
 b. This is a *good* knife for people with wrist arthritis.

Here it appears that information such as 'knives can be used for stabbing people', 'knives can be designed in different ways', etc. play an important role in the derivation of the *ad hoc* concepts communicated by *good*. In this way, by taking lexical interpretation to be a matter of adjusting the interpretation of individual words in accordance with one's context-specific expectations of relevance, the relevance-theoretic approach accounts for, indeed predicts, that *good* may be used to communicate a number of different occasion-specific senses, far beyond those predicted by Pustejovsky's account.

Finally, let us consider the 'type shifting' constructions of the kind in (23). As already mentioned in Chapter 3, on the present account the speaker-intended event associated with the VP (e.g. 'begin *reading* a book') would be derived entirely by means of a pragmatic process, albeit not through *ad hoc* concept construction. I mentioned two

possibilities for a pragmatic analysis of ‘type shifting’ predicates. Either they could be analysed as instances of ‘free’ enrichment, where a constituent not expressed in the linguistic form of the utterance is supplied. Or, as De Almeida and Dwivedi (2008) suggest, the pragmatic derivation of speaker-intended meanings for type shifting predicates could be analysed as being structurally driven, that is, as a form of saturation. On this approach (which I tend to favour, due to the mandatory character of this process), the syntactic structure of sentences such as ‘Mary began a book’ are seen as containing an extra VP with an empty verbal head, as shown by (25) (de Almeida and Dwivedi 2008: 316):

(25) Mary began [_{VP} [_{V⁰} e] [_{NP} the book]]

Thus, the verbal gap that remains in the LF of such constructions has to be filled in (saturated) using information from the discourse context.⁹⁹ This process, which is linguistically mandated and consists in supplying a missing constituent to the proposition expressed, would be very different from the (pragmatically induced) construction of *ad hoc* concepts from underspecified concepts output by the linguistic system.¹⁰⁰ In fact, such ‘type shifting’ constructions appear to be significantly different from the other instances of adjustment of lexical meaning discussed by Pustejovsky (1995a), as well as those discussed in this chapter (to the extent that I am not convinced that they should be treated as instances of polysemy at all).

An important motivation for a re-analysis of the Pustejovsky cases is the fact that the generative lexicon, in spite of its incorporation of considerable amounts of world knowledge into the lexicon (in the form of, *inter alia*, ‘qualia structures’), still leaves a lot of work for the pragmatic system to do in finding the interpretation that was intended

⁹⁹ At no stage in the interpretation process, therefore, is it assumed that a ‘default’ interpretation (e.g., according to which Mary began reading a book) is computed and then cancelled by context.

¹⁰⁰ Recall my criticism of Pustejovsky’s type coercion mechanism in Chapter 2, where it was pointed out that the generative mechanism wrongly predicts that a VP such as *begin a car* should be interpreted as ‘begin to drive a car’ due to telic information stored as part of the lexical representation for *car* (e.g. cars are for driving). On the relevance-theoretic account presented here, there would, of course, be no such prediction, but this account would also require an event to be supplied in cases such as *begin a car* (e.g. an utterance of ‘Bill *began a car*’ could be interpreted as ‘Bill began repairing a car’ in the context of a garage). However, a speaker using the VP *begin a car* to describe a situation in which someone began driving a car would not be optimally relevant, as the choice to use this expression rather than the conventional *start a/the car*, would, due to the extra effort of processing it would induce, send the hearer off searching for additional effects, which would not be part of the speaker’s intended meaning.

by the speaker (specifically, in overriding ‘default’ interpretations in contexts where another, ‘non-default’ interpretation was clearly intended). I claim that the wholly pragmatic, relevance-theoretic account is able to do this work, *as well as* that part of the interpretive work that a generative lexicon does adequately. The question, then, as I pointed out in Chapter 2, becomes whether we in fact need a generative lexicon or whether anything is to be gained by deriving some interpretations in one way (via generative linguistic mechanisms) and others in a distinct way (via pragmatic mechanisms). Certainly, considerations of theoretical economy would favour the unitary pragmatic approach.

4.3.2 Polysemy, sense relations and lexical semantic change

A consequence of adopting the relevance-theoretic atomist stance on word meanings is that there are two (complementary) ways of describing the representation of polysemy, which were outlined in Chapter 2: (i) In the cases where the polysemous senses have become conventional (i.e. stored in the mental lexicons of speakers), they would be represented as distinct lexical entries with the same linguistic form (e.g. BAKE₁, BAKE₂), on a par with homonymy. (ii) In all other cases, there would be a single encoded meaning of the word, and the differences in meaning between its uses in different contexts would be pragmatic adjustments to the encoded meaning (e.g. GOOD*, GOOD**, GOOD***). As we saw in Chapter 2, this model is compatible with the empirical evidence on polysemy representation.¹⁰¹

A widely held objection to atomistic theories of word meaning is that they cannot capture intuitively semantic relations such as synonymy, analyticity, entailment and polysemy, which are thought to require complex lexical semantic representations. Focusing on polysemy here,¹⁰² the question is: if intuitions about semantic relations between the senses of conventionally polysemous words do not stem from the overlap

¹⁰¹ However, as we discussed, it is far from the only model compatible with the experimental evidence.

¹⁰² In relevance theory, at least analyticity and entailment can be accounted for in terms of meaning postulates attached to concepts. See Borg (forthcoming: Chapter 5) for a different suggestion of how such semantic relations can be handled within an atomist framework without meaning postulates.

or sharing of features in their lexical semantic representations, where do such intuitions come from?

I suspect there might not be a single answer to this question. First, in the large number of cases where the polysemy takes the form of pragmatically derived *ad hoc* concepts, the meaning relations between the distinct *ad hoc* concepts would be transparent to speakers/hearers. That is, we perceive a relation in meaning between the different uses of *good* (e.g. ‘good knife’, ‘good mum’) because they are derived (inferentially) from the same encoded concept. Second, in the case of conventional polysemy (which, on this account, would be stored in much the same way as homonymous senses), some intuitions may stem from our encyclopaedic information about the denotations of the different senses. For instance, if it is true that the two senses of the linguistic form *bake* are represented as independent entries, *BAKE*₁ and *BAKE*₂, the reason we perceive them to be closely related in meaning could be to do with the kind of information we associate with the two activities; both involve putting some initially inedible potential foodstuff into an oven and applying heat to it at a high enough temperature for a long enough time that it becomes edible. Other intuitions about sense relations could be reflective. For instance, the fact that we perceive a meaning relation between the different senses of a conventionally polysemous lexical item such as *mouth*, may be because we know that the *mouth* in (say) ‘*mouth* of a river’ has been metaphorically derived from the *mouth* in ‘*mouth* of a dog’; we are able to reflectively reconstruct the process of metaphorical extension from the primary sense to the extended sense. Finally, I think that the mere fact that the different senses are encoded by a single lexical form often leads us to anticipate that there is a relation between them. A kind of ‘null hypothesis’ appears to be that the different senses of a lexical form are related, which is revised only when we are presented with evidence that the multiple encoding is accidental (for instance, it seems possible that many non-linguists would think that there is a relation between the two meanings of *bank*).¹⁰³

¹⁰³ This point perhaps becomes clearer if we consider a cross-linguistic example: For instance, in Norwegian, there is a noun, *måne*, which is conventionally polysemous between the sense ‘moon’ and the sense ‘bald spot’. The latter is a broadened use, derived on the basis of a perceived resemblance between the shape of a bald spot and that of the moon. While the relation between these two senses is a matter of course for speakers of Norwegian, this may not be so for speakers of English (in which these two concepts are encoded by different lexical forms) unless they are made aware of it.

A related issue is the role of polysemy and pragmatic adjustment of lexical meanings in semantic change. Much fruitful research on semantic change has been conducted within the context of grammaticalisation theory (e.g. Fleischman 1982; Sweetser 1990; Heine, Claudi, and Hünnemeyer 1991; Hopper and Traugott 1993/2003; Bybee, Perkins, and Pagliuca 1994; Traugott and Dasher 2002). The term ‘grammaticalisation’ refers to the process whereby content words acquire grammatical functions, involving a shift in meaning from the specific to the schematic (e.g. the shift from purposive *be going (to...)* to auxiliary *be going to*), or existing grammatical expressions acquire further grammatical functions. A central focus in the literature on grammaticalisation is the general directionality observed for some kinds of semantic changes, and on developing accounts in which possible and impossible changes or directions of change can be predicted.¹⁰⁴

A hypothesis about semantic change in grammaticalisation theory is that its main driving force is pragmatic; that it is motivated by speaker-hearer interactions and communicative strategies (Hopper and Traugott 1993/2003; Traugott and Dasher 2002). Traugott and Dasher (2002) argue that pragmatic inferences that arise in specific contexts may come to be reanalysed as part of the conventional meaning associated with a given construction.¹⁰⁵ Given the pragmatic account of adjustments to lexically-encoded concepts outlined here, we should expect pragmatic narrowing and broadening of lexical meanings to play a role in semantic change. In historical linguistics, the categories of narrowing and broadening have long been part of traditional classifications of types of lexical semantic change. Some scholars consider them to be the principal kinds of semantic change, and other kinds (including hyperbole and metaphor) to be subtypes of them. Consider the examples in (26)-(34), given by Campbell (1998/2004: Chapter 9):

¹⁰⁴ For instance, according to the hypothesis of unidirectionality of grammaticalisation, grammatical forms are taken to have developed from lexical forms, but the reverse change does not occur (cf. Hopper and Traugott 1993/2003).

¹⁰⁵ Traugott and Dasher’s (2002) emphasis is the typical direction of certain kinds of semantic change (so-called ‘regular’ changes, which are observed cross-linguistically and repeatedly within a single language), and the pragmatic mechanisms (which they describe as ‘invited inferencing’ and ‘subjectification’) that they take to be responsible for their development.

Lexical narrowing:

- (26) *Meat*. The original meaning of this word was 'food' in general, which got narrowed into its current meaning, 'food of flesh'.
- (27) *Hound*. Its current meaning 'dog of a breed used for hunting (especially one able to track its prey by scent)' comes from the Old English word *hund*, which meant 'dog' in general.
- (28) *Starve*. The current meaning of this verb in English, 'to suffer or perish from hunger' has developed as a narrowing of the meaning of the Old English word *steorfan*, 'to die'.

Lexical broadening:

- (29) *Dog*. This word first appeared with the more specific meaning '(specific) powerful breed of dog', which generalised to include all breeds or races of dogs.
- (30) *Cupboard*. In Middle English, the meaning of this word was 'table upon which cups and other vessels were placed, a piece of furniture to display plates, a sideboard', which got broadened into its current meaning in British English, 'closet or cabinet with shelves for keeping cups and dishes'. In American English, its meaning has been further broadened into denoting any 'small storage cabinet', and in parts of Canada, it is used to denote 'wardrobe' or 'clothes closet'.
- (31) *Terribly, horribly, awfully, etc.* These words can all be said to have the meaning 'very' in modern English, being used to intensify the meaning of the adjective they modify. As a result of being used as hyperboles, they have developed into having no synchronic semantic connection with their origins, *terror, horror, awe, etc.*
- (32) *Grasp*. The meaning 'understand' is seen as a metaphorical extension of the original meaning, 'seize'.¹⁰⁶
- (33) *Dispose of, liquidate, terminate, take care of, eliminate, etc.* Their 'to kill' meanings are metaphorically derived.

¹⁰⁶ Sweetser (1990) showed that there is a cross-linguistic tendency for verbs of seeing and grasping to be metaphorically extended into verbs of understanding.

- (34) *Chill*. The informal meaning ‘to relax, calm down’, is metaphorically derived from the original meaning ‘to cool’.

Moreover, it is recognised by many scholars that semantic change must go through a stage of polysemy, in which related meanings of a word that emerged at historically different periods, coexist over time in a language, for individual speakers as well as for language communities (a phenomenon known as ‘layering’ in grammaticalisation theory, cf. Hopper 1991). The assumption is that a shift in the meaning of a word form *a* from meaning A to meaning B, goes through an intermediary stage in which both A and B coexist, as schematised in (35) below (adapted from Campbell 1998/2004: 266-267):

- (35) Stage 1: *a* means ‘A’
Stage 2: *a* means ‘A’ and ‘B’ (‘A’ > ‘A’, ‘B’)
Stage 3: *a* means ‘B’ (‘A’, ‘B’ > ‘B’)

This assumption about the progression of the process of semantic change is quite compatible with the relevance-theoretic account of lexical adjustment (and hence of polysemy), pursued in this thesis, where what starts out as an *ad hoc* concept in Stage 1, may become stabilised or conventional over time for individual speakers or within a language community, as in stage 2, as a result of frequent adjustment of the lexical meaning of a word in a specific direction (e.g. the broadening of the change-of-state sense of *bake* into the creative sense, as described in section 4.3.1 above). In such a case the construction of the *ad hoc* concept becomes progressively more routinised (a ‘pragmatic routine’ or inferential shortcut develops, cf. Vega-Moreno 2007). A possible development from there is that the *ad hoc* concept takes over from the originally encoded concept, as in Stage 3 (e.g. as in modern English, where the ‘food of flesh’ sense of *meat* has replaced the general ‘food’ sense, or, to take a more recent case, where the derived sense of the word *gay*, ‘homosexual’ has almost entirely taken over from the original sense, ‘lighthearted and carefree’).¹⁰⁷

¹⁰⁷ A very interesting question is what factors contribute to pragmatically adjusted senses ‘catching on’ in a speech community, which ultimately lead to their stabilisation. The philosopher Ruth Millikan proposes an evolutionary account of the process of stabilisation of senses. In Millikan’s (1984, 2005) framework, linguistic items, or ‘language devices’, as she calls them, have proper functions, or purposes (in a similar way as tools have functions or purposes). A proper function is what accounts for the continuing

Furthermore, at the synchronic level, individual speakers may differ with regard to which senses of a word they have stored in their mental lexicons. For instance, for some speakers of English, the broadened (metaphorical) meaning of *wizard*, WIZARD*, may have become conventional (thus stored in their mental lexicons), and recognising this concept as the one intended by a speaker on an occasion of use would be a matter of disambiguation rather than *ad hoc* concept construction (cf. Chapter 3). For other people, even within the same speech community, *wizard* may have only one encoded sense (WIZARD), and the metaphorical sense (WIZARD*) would be derived through *ad hoc* concept construction. Thus, inferred senses may be ephemeral or stable, shared by few or many or by a whole language community. The construction of a particular *ad hoc* concept may be a first-time affair for one communicator and a routine inferential pattern for another, and the reason this is possible is our pragmatic ability. As Wilson and Carston write:

One of the most important functions of pragmatic inference is to compensate for grammatical and lexical differences among members of a speech community, so that addressees with different encoded senses can end up with the same interpretations, albeit via different routes. (Wilson and Carston 2007: 241)

4.3.3 Prepositional polysemy

The cases of polysemy that we have looked at so far have been open-class lexical items (nouns, main verbs, adjectives and adverbs). However, within the cognitive linguistics tradition, the main focus of research on polysemy has been closed-class lexical items,

reproduction and proliferation of a language device. She distinguishes between two types of proper function, ‘direct stabilising proper functions’ and ‘derived proper functions’. A direct stabilising function is one that, “when performed, tends both to encourage speakers to keep using the device and hearers to keep responding to it with the same (with a stable) response” (Millikan 2005: 94). The stabilising function of a linguistic device is to contribute to its ‘conventional meaning’. It is what accounts for the ‘survival’ of the expression type in a language community; it is historically responsible for its reproduction. For instance, a direct stabilising function of the word *elephant* would be to produce representations of elephants in hearers’ minds. A derived proper function of a linguistic device token is derived from a speaker’s intention or purpose in using the linguistic device, and may involve the use of metaphor, irony, or other figures of speech. Using the word *elephant* to describe someone either big boned or not particularly graceful would be such a case. If, on Millikan’s account, there is repeatedly a conflict between the stabilising function and speaker purpose (or, speaker-intended meaning, on the account pursued in this thesis) for a linguistic device, this will result in either a change in function or the addition of a new function. If the same speaker-hearer pattern is reproduced long enough, the new use will become conventional; the result is a new stabilising function for a word. According to Millikan, a new meaning has been established in a public language “when the new meaning would continue to proliferate even if the old meaning were to die out” (2005: 193).

specifically prepositions (among many others, Lakoff 1987; Brugman 1988; Brugman and Lakoff 1988; Kreitzer 1997; Tyler and Evans 2003; Evans 2009). There has been little, if any, work on prepositions within the relevance theory framework. However, we might expect the underdeterminacy issue to be even more acutely manifested in the case of such closed-class items (which are few in number and new items cannot usually be added), than in the case of the open-class items we have discussed so far. Since speakers have a very limited set of linguistic resources at hand for expressing a very large (virtually unlimited) set of spatial relations (in addition to the set of abstract relations that they are used to express), pragmatic inference must presumably play a crucial role in the comprehension of prepositional expressions.

Two questions in particular arise from this: First, given the analysis of the polysemy of open-class items as the outcome of an inferential process of *ad hoc* concept construction, can the (presumably massive) underdeterminacy relation that exists between the class of prepositions and the spatial and abstract relations that can be expressed by use of them be resolved in terms of the same type of pragmatic process? Second, given the three types of linguistic encodings specified in section 4.2.1 above (procedural meanings, pro-concepts, full-fledged concepts), what do prepositions encode? While it seems clear that their encodings are in some way conceptual (they are constituents of conceptual representations and, in most cases at least, they affect the truth-conditions of utterances in which they occur), they seem to be much more schematic and abstract than the conceptual representations encoded by lexical items such as *flu*, *wizard*, *rectangular*, *bake*, etc., which, on this account, are claimed to encode full-fledged concepts. In this respect, prepositions have more the flavour of so-called pro-concepts, i.e. schematic concepts that need to be pragmatically fleshed out into full concepts.

As the issue of prepositional meaning could easily serve as the topic of a whole thesis, I will obviously not be able to do more than scratch the surface of it here. In the rest of this section, I will (briefly) consider the ‘Principled Polysemy’ approach to prepositional polysemy proposed by Tyler and Evans (2001, 2003). This approach, although different from the relevance-theoretic position in important ways, shares the assumption that the constructive role played by pragmatic inference in giving rise to polysemy must be built into an account of the phenomenon. I will end the section with a

tentative suggestion for a direction that a relevance-theoretic analysis of prepositional meaning might take.

A reason for the interest in prepositions taken by cognitive linguistics is that they are considered particularly clear illustrations of the so-called ‘embodiment thesis’ (Johnson 1987, 1997; Lakoff and Johnson 1999). This is the claim that meaning is grounded in our bodily experiences and interactions with the world, which allows us to conceptualise abstract areas of experience in terms of the familiar and concrete.¹⁰⁸ As a consequence, linguistic meaning is also seen as being largely grounded in spatio-physical experience, of which prepositions, expressing spatial relations that give rise to a range of non-spatial abstract senses, are prime examples. Recall from Chapter 2 the paradigmatic example of prepositional polysemy, *over*, repeated in (36) below:

- (36) a. The bird flew *over* the house. (‘above and across’)
b. The painting is *over* the couch. (‘above’)
c. The truck ran *over* the rabbit. (‘across’)
d. Sarah lives *over* the hill. (‘on the other side of’)
e. Mary nailed a board *over* the hole in the ceiling. (‘covering’)
f. I will read the papers *over* the weekend. (‘temporal’)
g. John has a strange power *over* Mary. (‘control’)

As we saw in Chapter 2, Brugman (1988; Brugman and Lakoff 1988) and Lakoff (1987) analysed *over* as a radial category composed of a range of distinct but related senses organised around a prototypical, or central sense (which, in their view, was the ‘above and across’ sense) in a lexical network structure. Their approach provided a full specification of the range of possible senses of *over*, all of which were taken to be stored in the long-term semantic memory of speakers. A common criticism of this full-specification approach is its failure to distinguish between those aspects of meaning that

¹⁰⁸ Thus, a central assumption about metaphors in cognitive linguistics is that they are shaped and constrained by our bodily experiences (Lakoff and Johnson 1980). For instance, spatial concepts of verticality are taken to be the ‘source domain’ of conceptual metaphors such as HAVING CONTROL OR FORCE IS UP, BEING SUBJECT TO CONTROL OR FORCE IS DOWN, whose physical basis is our experience that “physical size typically correlates with physical strength, and the victor in a fight is typically on top” (ibid. 15).

are part of the word meaning proper and those that result from its interaction with the context (which Sandra 1998 called the 'polysemy fallacy').

The 'Principled Polysemy' approach proposed by Tyler and Evans (2003) espouses the view that polysemous lexical items are represented in terms of sense networks centred around a prototypical sense, as well as the cognitive linguistic claim about embodied meaning. Their approach differs from the Lakoff-Brugman approach by distinguishing between established senses and uses of a polysemous word, thus allowing for many of the senses to be pragmatically inferred rather than fully specified in the lexicon.¹⁰⁹

Tyler and Evans (2003) seek to provide a methodology for distinguishing between distinct senses of a preposition, and for establishing the prototypical sense. In order to determine whether a particular use of a preposition counts as a distinct sense, they propose the following two criteria (Tyler and Evans 2003: 42-43; Evans and Green 2006): A sense is distinct if (i) it involves non-spatial meaning and/or a spatial configuration between the trajector (TR) and the landmark (LM) which is distinct from that found in the word's *protoscene* (i.e. the primary sense of the word, represented in terms of an idealised spatio-functional configuration);¹¹⁰ and (ii) there are instances of the sense that are context-independent, that is, which cannot be inferred from another sense and the context in which it occurs. To illustrate, consider the uses of *over* in (37) and (38):

- (37) a. The hummingbird is hovering *over* the flower.
b. The helicopter is hovering *over* the city.
- (38) a. John nailed a board *over* the hole in the ceiling.
b. John nailed a board *over* the hole in the wall.

In (37), *over* is used to express a spatial relation in which the TR (encoded by *the hummingbird/the helicopter*) is located higher than the LM (encoded by *the flower/the*

¹⁰⁹ Evans (2009) provides a detailed analysis of the 'state' concepts expressed by the English prepositions *in*, *on*, and *at*, within the framework of his newly developed *LCCM Theory (Theory of Lexical Concepts and Cognitive Models)*. I cannot get into the details of this analysis here, but I return to consideration of the LCCM theory as concerns its account of metonymy in Chapter 6.

¹¹⁰ In Langacker's (1987) terminology, the trajector (TR) is the primary focal element which follows a trajectory, and landmark (LM) the secondary focal ('backgrounded') element.

city). The same basic TR-LM configuration is expressed in both instances, and neither of them conveys any additional non-spatial meaning. Thus, they fail the criterion in (i) and should not be treated as distinct senses. However, according to Tyler and Evans, the sense of *over* expressed in (38) appears to be distinct. First, there is a primary non-spatial element of ‘covering/obscuring’ to this sense, in which the LM (*the hole in the wall/ceiling*) is obscured from view by the TR (*the board*). Furthermore, the spatial relations expressed between the TR and LM are not the same as that expressed by *over* in (37); in (38)a., the TR and LM are horizontal with respect to the vantage point and the TR is physically below the LM, while in (38)b., the TR and LM are vertical with respect to the vantage point, and the TR is physically next to the LM (ibid. 43). Thus, criterion (i) is satisfied. Second, Tyler and Evans argue that the covering sense of *over* in (38) cannot be pragmatically derived from the primary sense (or *protoscene*), if the primary sense is taken to involve a spatial configuration between TR and LM in which the TR is *higher* than the LM. Thus, the criterion in (ii) is also satisfied.¹¹¹ The claim is that unless we have the covering/obscuring sense already stored in semantic memory, the uses of *over* in (38) cannot be explained.

The methodology for distinguishing between distinct senses of a preposition above relied on a primary sense (a *protoscene*) being established for the word. In order to establish the primary sense in a polysemy network, Tyler and Evans (2003: 45-50) propose the following set of criteria: (i) earliest attested meaning; (ii) predominance in the semantic network (the primary sense will be the one that is most frequently involved in or related to the other distinct senses); (iii) use in composite forms (failure to participate in composite forms, e.g. *overcoat*, is taken to be suggestive of a non-primary sense); (iv) relations to other prepositions (the sense that participates in a contrast set, e.g. prepositions of verticality, is a likely candidate as a primary sense), and (v) ease of predicting sense extensions (the primary sense should be the best predictor of other senses in the network). Based on these criteria, Tyler and Evans (2003: 65) identify the

¹¹¹ However, Tyler and Evans (2003: 44) claim that in an utterance such as ‘The tablecloth is *over* the table’, the covering/obscuring sense can in fact be inferred from “the fact that the tablecloth is *over* and hence higher than the table, in conjunction with our knowledge that tablecloths are larger than tables and that we typically view tables from above the top of the table”. This would presumably explain how this sense came to be derived in the first place (which subsequently established as a distinct sense in the semantic network through routinisation).

'above' sense in (37) as the primary sense, or protoscene, of *over*. More specifically, the protoscene is described as a "spatial relation in which the TR is higher than but within potential contact of the LM". In addition, they identify another 14 distinct senses (thus 15 senses in total) that form part of the semantic network for *over* (see Tyler and Evans 2003: 80 for details).

Occasion-specific senses of *over* are constructed on the basis of the proto-scene, or any of the other distinct senses.¹¹² An instance of such a contextually-specified meaning is the 'above and across' sense, illustrated by (39) below (Tyler and Evans 2003: 69):

(39) The cat jumped *over* the wall.

In this example, the 'above and across' trajectory expressed by *over* should not be seen as being stored in semantic memory (as claimed by Lakoff 1987), but as being inferentially derived on the basis of the protoscene and our understanding of the action of jumping (e.g. that it involves motion and a trajectory), encyclopaedic assumptions about cats (e.g. their physical abilities), walls (e.g. that they provide vertical, impenetrable obstacles), as well as our knowledge about force dynamics (e.g. gravity).¹¹³

In spite of some fundamental differences, Tyler and Evans's Principled Polysemy approach is in many ways consonant with the relevance-theoretic approach to lexical adjustment, and with the analysis of polysemy proposed in this thesis. As mentioned above, the two accounts share the assumption that speaker meanings are, quite generally, radically underdetermined by linguistically-encoded meanings, and that pragmatic inference therefore plays a crucial role in the understanding of utterances. They also share the view that the hearer uses his encyclopaedic knowledge associated with the encoded concepts in the process of pragmatically inferring speaker-intended

¹¹² The precise details of this process of constructing occasion-specific senses are not entirely clear to me. Is the assumption that the hearer must first disambiguate the meaning of the preposition (i.e. choosing the speaker-intended established meaning), and then construct a context-specific sense on the basis of the disambiguated meaning?

¹¹³ Tyler and Evans (2003: 57) describe the inferential strategy used to arrive at this sense as one of 'best fit', in which the speaker chooses that preposition which "best fits the relevant (i.e. salient) conceptual spatial relation between the TR and LM at one point in the cat's trajectory, which will, in turn, prompt the [hearer's construction of] the appropriate entailments or inferences". In relevance-theory, this strategy is captured by the presumption of optimal relevance.

senses of lexical items. In this way, both accounts emphasise the central pragmatic aspect of polysemy. However, while on the relevance-theoretic account this pragmatic aspect is *fundamental* to the existence of polysemy, on Tyler and Evans's account it is ultimately grounded in the embodiment of meaning, that is, in our spatio-physical experience, which is taken to provide the fundamental conceptual structure from which other concepts are constructed.

It is with regard to their positions on linguistic meaning in particular that the two approaches part company. First, relevance-theory does not subscribe to the claim that linguistic meaning is embodied. Rather, it takes a Chomskyan/Fodorian modular view of the linguistic system, and treats linguistic meanings as incomplete conceptual representations output by the linguistic system, and taken as input to the pragmatic inferential system. Second, as we have seen, the relevance-theoretic account does not postulate semantic networks to account for polysemous senses of a word. Although I take it that it is, in principle, possible that we store many distinct senses for a preposition such as *over* in our mental lexicons, I wonder if the range of spatial and abstract relations that are expressible by this preposition actually requires the postulation of 15 distinct senses for this word. Recall that Tyler and Evans's (2003) argument for treating the covering/obscuring sense of *over* in examples such as 'John nailed a board *over* the hole in the ceiling' as distinct was that it could not be derived pragmatically from the primary 'above' (i.e. 'the TR is higher but within potential contact of the LM') sense of *over*, because the spatial relation expressed is, strictly speaking, that the board is *below* the hole in the ceiling. However, it seems that the contextual derivation of the covering sense is made impossible primarily by the decision to take the 'above' sense to be primary (and thus the assumption that the 'covering' meaning must be derived from this sense).

A similar point is made by Fieke Van der Gucht and De Cuypere (2007), in their discussion of the Principled Polysemy approach. In contrast to Tyler and Evans they maintain a monosemy approach to the meaning of *over*, taking it to encode a more schematic meaning, a relation between a spatially 'superior' TR and an 'inferior' LM, which can be instantiated in more than one possible way in extralinguistic reality (ibid. 746). The preposition thus exhibits a 'structured polyvalence' at the experiential level, in

the sense that the range of possible relations it can be used to refer to is constrained by the linguistic meaning, which is invariant in all its uses. They write,

To assess the relation between the invariable meaning proper and the polyvalence of *over*, it is imperative not to interpret the specific configuration which holds between the TR and LM in a single – call it ‘prototypical’ – way, excluding other possibilities which then have to be accounted for by invoking polysemy or certain kinds of ‘extension’. (Fieke Van der Gucht and De Cuypere 2007: 746).

Furthermore, Fieke Van der Gucht and De Cuypere question the view that prepositions encode concepts in the same way as verbs, nouns, and adjectives do. Their claim is that in isolation, a preposition such as *over* does not *mean*, e.g., ‘on the other side of ‘covering’, ‘control’ or ‘past’, rather, these meanings arise when *over* is interpreted in combination with other words in a sentence. Therefore, prepositions should rather be treated as encoding ‘instrumental meanings’, that is, relational meanings that can only be instantiated in grammatical combinations with other words.

I suggest that this is also the direction that a relevance-theoretic approach to prepositional polysemy would probably take. Traditionally, there has been a distinction in relevance theory between conceptual and procedural meaning (Blakemore 1987; Wilson and Sperber 1993). While most lexical items encode concepts (e.g. *flu*, *wizard*, *rectangular*, *bake*) which figure in conceptual representations and affect truth-conditions of utterances (i.e. explicit content), some expressions do not encode concepts (e.g. *but*, *so*, *however*, *also*, *anyway*, and, arguably, pronouns and various mood indicators and sentential particles); they are typically non truth-conditional and their meanings are procedural; they encode constraints on the inferential phase of comprehension. An intermediate case, as mentioned in section 4.2.1, is words that encode so-called pro-concepts, i.e. a schematic representation that requires contextual specification in order to make a fully conceptual contribution to the proposition expressed. Prepositions, then, could be seen as encoding such pro-concepts, say, in the form of a schematic representation of a spatial relation, which, in context, would trigger a process of construction of a relational meaning (an *ad hoc* concept), based on the preposition’s combination with other lexical concepts and their associated encyclopaedic entries, as well as accessible contextual assumptions (which, in spatial

uses, would often include assumptions derived from the immediate physical context). Sometimes the concept communicated will be narrower than the encoded conceptual meaning (which would arguably be the case in most spatial uses), sometimes it will be broader (as in more abstract uses). Consider the ‘covering’ uses of *over* again in (40):

- (40) a. The tablecloth is *over* the table.
b. John nailed a board *over* the hole in the ceiling.
c. John nailed a board *over* the hole in the wall.

Let us assume, for the sake of illustration, that *over* encodes a relation between a spatially ‘superior’ TR and an ‘inferior’ LM, as suggested by Fieke Van der Gucht and De Cuypere (2007). It is, I think, possible to argue that such a relation does in fact underlie all the uses of *over* in (40), but that, given our real-world knowledge about the objects that *over* is used to express relations between (i.e. between tablecloths and tables in (40)a., boards and holes in the ceiling in (40)b., and boards and holes in the wall in (40)c.), this relation is construed in different ways in the different uses in (40), so that, for instance, from the point of view of the speaker/hearer the board would strictly speaking be *below* the hole in the ceiling in (40)b., and *next to* it in (40)c. The broadened covering sense, then, could be seen as being inferred on the basis of the encoded meaning of *over* (from the fact that the TR is somehow ‘superior’ to the LM), in combination with encyclopaedic knowledge about the objects it relates (e.g. that tablecloths are typically larger than tables and are used to cover them, that a board can be used to cover up a hole in the ceiling/wall, etc.).

The abstract ‘control’ sense of *over* (cf. Lakoff 1987), repeated in (41), could be seen as a further broadening of the encoded spatial sense, involving an abstract relation between a ‘superior’ TR (John) and an ‘inferior’ LM (Mary):

- (41) John has a strange power *over* Mary.

In deriving this interpretation, the hearer will make use, *inter alia*, of encyclopaedic information associated with the concept POWER, specifically, his knowledge that in the

domain of human relations, it involves the ability to direct or influence the behaviour of others.¹¹⁴

4.4 Conclusion

In this chapter, I have outlined the relevance-theoretic account of lexical pragmatics, according to which there is a single process that adjusts the meanings of individual words in different directions (giving rise to *ad hoc* concepts with either narrower or broader denotations than the linguistically-encoded denotations), and claimed that it provides an account of how new senses for a word (giving rise to polysemy) are constructed during on-line utterance comprehension. The overall claim was that polysemy is not such a paradox in the context of a fully functioning pragmatic system.

I suggested a reanalysis of the Pustejovsky (1995a) cases of polysemy in terms of the relevance-theoretic lexical pragmatic account, and argued that, in contrast to the generative lexicon, the pragmatic approach predicts that lexical items are used to express a variety of occasion-specific concepts, which include but go far beyond the default senses predicted by Pustejovsky's account.

I then went on to consider the nature of sense relations in connection with polysemy, and suggested that our intuitions about relations between the different senses of a polysemous word may come from a variety of sources. I also considered some historical lexical data, which, combined with insights about semantic change from grammaticalisation theory, suggests that lexical narrowing and broadening may, via an intermediate stage of conventional polysemy, lead to semantic change. Finally, I discussed the issue of prepositional polysemy, a major topic in the cognitive linguistics research, which has received little attention in relevance theory, and sketched a possible direction for a relevance-theoretic analysis of prepositional meaning. A proper relevance-theoretic account of prepositional polysemy would, of course, have to be based on extensive empirical investigation into the range of possible uses of specific prepositions (cf. Tyler and Evans 2003), and should, ideally, also include some cross-linguistic investigation. I hope to return to this topic in future research.

¹¹⁴ In Tyler and Evans's framework (2003: 101), this 'control' sense comes out as a distinct sense, not derivable from the context.

I have claimed in this chapter that polysemy is the outcome of a pragmatic process of *ad hoc* concept construction, resulting in communicated concepts that have narrower or broader denotations than those linguistically-encoded. However, there are two important kinds of polysemy that I have not yet looked at, and for which it remains to be demonstrated that they can be given appropriate treatments within the present pragmatic framework. One is metonymy, which is the topic of Chapter 6, and the other is a set of cases which are often described as ‘systematic’ or ‘regular’ polysemy. This includes the kind of sense alternation that seems to be generated by a count or a mass use of a noun (e.g. the animal/meat senses of *rabbit*), and thus provides a good candidate for a linguistic-semantic analysis. I suggested in Chapter 3 that even in these apparently ‘systematic’ cases of polysemy, pragmatic inference is indispensable, but I did not suggest how a pragmatic analysis of systematic polysemy might proceed. This is what I aim to do in the next chapter.

Chapter 5

SYSTEMATIC POLYSEMY AND THE COUNT-MASS DISTINCTION

5.1 Introduction

In the previous chapters I have argued in favour of a wholly pragmatic account of polysemy. The view developed there is that polysemy arises as a result of the operation of a pragmatic inferential process of *ad hoc* concept construction, which yields hypotheses about speaker-intended senses on the basis of encoded senses, contextual information and relevance-based interpretive constraints. As we saw in the previous chapter, the outcome of this process may be a concept that has a narrower denotation than the one linguistically-encoded ('John OPENED* the bottle'), or a broader one ('John is a WIZARD* in finance'). Compared with other more semantically oriented approaches, this pragmatic approach drastically downplays the contribution of the linguistic system in the production and interpretation of polysemy, whose role is seen as that of providing a minimal input or clue which the inferential system uses as evidence to yield hypotheses about occasion-specific, speaker-intended senses.

A set of cases that may be seen as presenting a possible challenge to the pragmatic approach is so-called 'systematic polysemy', where the related senses of a word (usually a noun) are predictable on the basis of a general pattern of sense alternation observed for words denoting objects of the same category. Nouns are the most salient (and most discussed) examples of this type of polysemy, and I will concentrate on these in what follows.¹¹⁵ The generalisations involved in systematic polysemy have the following form:

- (1) If an expression has a use of type A, it also has a use of type A'.

This type of polysemy goes by various names, including 'regular polysemy' (Apresjan 1974; Kilgarriff 1992, 1995), 'semantic transfer rules' (Leech 1990 [1974]), 'lexical

¹¹⁵ As to other parts of speech claimed to exhibit a type of 'systematic' polysemy, see for instance the recent paper by Kennedy and McNally (2010), suggesting that colour adjectives alternate between a gradable and a non-gradable sense, and Jackendoff's (2002: Chapter 11) analysis of the causative alternation (e.g. The door *opened*/John *opened* the door) as a form of polysemy. See also the discussion of Jackendoff's position in Chapter 2.

implication rules' (Ostler and Atkins 1992), 'logical polysemy' (Pustejovsky 1991, 1995a), 'transfers of meaning' (Nunberg 1996, 2004), 'sense extensions' (Copestake and Briscoe 1992, 1996), and 'conversion' (Gillon 1999). Here are some examples:

- (2)
 - a. A *rabbit* jumped over the fence.
 - b. We're having *rabbit* for dinner.
 - c. The model wore *rabbit* on the catwalk.
 - d. After a tractor had run over the body, there was *rabbit* splattered all over the yard.
- (3)
 - a. We have a *pine* in our garden.
 - b. This table is made of *pine*.
- (4)
 - a. Susan decorated the cake with a *cherry*.
 - b. When the kids left, there was *cherry* all over the kitchen floor.
 - c. Jill and Joan have a *cherry* in their garden.
 - d. This table is made of *cherry*.
- (5)
 - a. The *newspaper* is lying on the coffee table.
 - b. The *newspaper* is up on the web.
 - c. The *newspaper* announced staff redundancies.
- (6)
 - a. *Oslo* is hosting the 2011 Nordic World Ski Championships.
 - b. *Cambridge* voted conservative.
 - c. Mary is reading *Dostoyevsky* on the tube.

The noun *rabbit* is used to express the animal in (2)a., its meat in (2)b., its fur in (2)c., and unspecified 'rabbit stuff' in (2)d. The noun *pine* is used to refer to a tree in (3)a. and to the wood from the tree in (3)b. The noun *cherry* is used to denote the fruit in (4)a., stuff of the fruit in (4)b., the tree carrying the fruit in (4)c. and the wood from the tree carrying the fruit in (4)d. The noun *newspaper* is used to denote a copy of a newspaper in (5)a., its informational content in (5)b., and the organisation that publishes it in (5)c. In (6), the name for a place is used to refer to the organisers of an event taking place at that place in (6)a., the name for a place is used to refer to its inhabitants in (6)b., while the name of an artist is used to denote his work in (6)c. On the basis of these examples,

we may identify the following lexical alternation patterns, described by means of the formula in (1):

- (7) If an expression has an 'animal' use, it also has a 'meat'/'fur'/'animal stuff' use.
- (8) If an expression has a 'tree' use, it also has a 'wood' use.
- (9) If an expression has a 'fruit' use, it also has a 'fruit stuff'/'tree' use.
- (10) If an expression has a 'publication' use, it also has a 'content'/'organisation' use.
- (11) If an expression has a 'place' use, it also has an 'organisers'/'inhabitants' use; If an expression has a 'writer' use, it also has an 'oeuvre' use.

And there are many more such lexical alternation patterns (see e.g. Lakoff and Johnson 1980: 38-39; Nunberg 1996: 117 for further examples). Much work in computational semantics has emphasised the need for an adequate formalisation of these observed regularities. The literature contains various suggestions of how they can be incorporated into the lexicon as a set of lexical inference rules, in this way avoiding a listing of all options for all words (Pustejovsky 1991, 1995a; Ostler and Atkins 1992; Copestake and Briscoe 1992, 1996; Boguraev and Levin 1993; Kilgarriff 1992, 1995; Kilgarriff and Gazdar 1995; Blutner 1998).

On the face of it, the arguments for treating systematic polysemy as being governed by linguistic rules seem well-founded. The processes in question appear to be regular, productive and available in many languages. Their application appears to be restricted by the existence of synonymous terms in the lexicon, so-called pre-emption or blocking (cf. Blutner 1998, 2002), for instance, the existence of the form *beef* in English 'blocks' the use of the form *cow* to refer to the meat of the animal. Some of them seem to affect syntactic behaviour, for instance, lexical alternations that rest on the count-mass distinction in English and in other languages with count-mass syntax, gender marking on the fruit for tree alternation in Romance languages (e.g. the distinction in Spanish between *aceituna* ('olive') and *aceituno* ('olive tree'), and *guinda* ('cherry') and *guindo* ('cherry tree')). The challenge for the pragmatic approach, then, is how to account for the apparent systematicity of these processes of sense extension. Are they best handled as instantiations of lexical rules, as derived mainly through pragmatic inference, or perhaps as involving a combination of the two?

Systematic polysemy is often seen as having an affinity with metonymy, the process whereby an expression that conventionally denotes one object or property is used to denote another object or property that stands in a certain relation to it.¹¹⁶ A famous example is Nunberg's (1978: 22) 'The *ham sandwich* is sitting at table 20', where *ham sandwich* refers to a customer in the context of a restaurant. Some of the sense alternations described as instances of systematic polysemy arguably involve metonymic relations between the different senses, (cf. (4)c., (5)c. and all the examples in (6)). For this reason, formal semantic accounts have proposed to treat not only these but also more creative cases of metonymy as having a non-pragmatic component (Copestake and Briscoe 1996). There is some evidence that creative cases may have grammatical effects, e.g., 'The *french fries* is getting impatient' (Nunberg 1996), where the agreement between the subject and the VP is determined by the referent of the NP *french fries* rather than by its syntax.¹¹⁷ From a communicative point of view, however, unlike the examples in (2)-(6), these creative cases of metonymy seem largely confined to referential uses (in a strong sense) and bear some relation to certain nicknames (e.g. 'Four Eyes is always reading math books'), a connection which I will pursue in the next chapter.

In this chapter I have two aims. The first is to consider a well-developed formal account of systematic polysemy, proposed by Copestake and Briscoe (1996), and assess whether systematic polysemy can be adequately accounted for as instantiations of lexical rules. I conclude that, although such an account captures the regularities involved, it does not provide the interpretive flexibility required to handle the full range of data. Thus, my second aim is to show how the data can be reanalysed from the perspective of relevance theory, using the machinery laid out in the previous chapter. I will defend the view that, although the sense alternations that give rise to systematic polysemy clearly have a linguistic component, the contribution of the linguistic system to its generation

¹¹⁶ This assumption goes back to Apresjan (1974: 16), who saw regularity as a distinctive feature of metonymy, and 'irregular polysemy' as arising from metaphor. I will come back to this issue in Chapter 6.

¹¹⁷ The cognitive grammar tradition (Lakoff and Johnson 1980; Lakoff 1987 and many others) usually makes no distinction between systematic polysemy that involves metonymic sense relations and metonymy more generally; the label 'metonymy' subsumes them both. In his pragmatic account, Nunberg (1996) takes the process of 'predicate transfer' to be responsible for the generation of systematic polysemy as well as creative metonymy. These accounts will be considered again in Chapter 6.

and interpretation is less central than is often thought. My conclusion will be that even in cases of systematic polysemy, the major interpretive work is done by pragmatics.

The focus of this chapter will be instances of systematic polysemy that (in English) rest on the distinction between count and mass uses of nouns (and thus are perhaps the clearest candidates for a linguistic analysis), as exemplified by (2), (3) and (4) above. Towards the end of the chapter, I will look at a set of cases that appears to involve alternation in meaning between different aspects of the object denoted (e.g. between a physical object sense and an information type sense, *newspaper*, *book*, *DVD*, etc.), and suggest how they can be treated within the relevance-theoretic account. Metonymy, however, both of the systematic and the creative kind, will be the subject of Chapter 6. Since it is an unresolved issue in relevance theory (but see Papafragou 1996 for an early relevance-theoretic account), I will devote the whole of that chapter to a proposal for a relevance-theoretic account of systematic and creative metonymy.

5.2 Computational semantic accounts: The case of ‘grinding’

Copestake and Briscoe (1992; 1996: 36) propose a formal semantic framework in which *sense extensions*, that is, “predictable creation of different but related senses”, are represented as lexical rules. One such rule is ‘universal grinding’ (Pelletier 1975). The effect of this rule is to create from a count noun (denoting a physical object) a mass noun with properties appropriate for an unindividuated substance.¹¹⁸ This is a general, abstract rule that yields the mass senses of *rabbit* and *cherry* in (12) and (13) (cf. (2)d. and (4)b. above), as well as the uses of *sun* and *shopping centre* in (14) and (15) to denote unindividuated substances:

(12) After a tractor had run over the body, there was *rabbit* splattered all over the yard.

(13) When the kids left, there was *cherry* all over the kitchen floor.

¹¹⁸ Pelletier’s (1975: 456-457) original thought experiment involves a machine, the ‘universal grinder’ (not unlike a meat grinder), which chops up “any object no matter how large, no matter how small, no matter how hard”. Then, the idea is that we could take any object corresponding to any (apparent) count noun, let’s say *man*, put it into the grinder and ask what is on the floor at the other end of the grinder, and get the following answer: ‘There is man all over the floor’. Given that “there can be made a *prima facie* case that nothing is immune from the grinder treatment”, Pelletier concludes that count nouns having physical objects as their extensions can always be given a mass sense.

- (14) We got quite dark from all the *sun*.
(15) That's a lot of *shopping centre* for a small town (Nunberg and Zaenen 1992).

Copestake and Briscoe further posit a set of conventionalised sub-cases of the grinding rule, including a specialised 'meat-grinding' function that forms food-denoting mass nouns from animal-denoting count nouns, and one that forms fur-denoting mass nouns from animal-denoting count nouns. Kilgarriff (1992, 1995) proposes a similar specialised rule that yields the wood senses of tree-denoting count nouns. Together, these rules account for the examples in (16)-(18) below:

- (16) We're having *rabbit* for dinner (*chicken, turkey, moose, frog, etc.*).
(17) The model wore *rabbit* on the catwalk (*mink, beaver, calf, lizard, crocodile, etc.*).
(18) This table is made of *pine* (*cherry, oak, chestnut, birch, etc.*).

So the claim is that there is a set of lexical rules (the universal grinder and its various conventionalised sub-cases) stored in the lexicons of individual speakers that are responsible for generating systematic polysemy of the kind illustrated above. In other words, the shift in meaning that is observed in these cases (e.g. from an animal sense to a food sense), is assumed to have a wholly linguistic basis. Copestake and Briscoe do point out, however, that pragmatics may contribute to further contextual specification of the denotation, for instance, by providing the information that rabbit meat usually excludes the bones while the meat of whitebait does not.¹¹⁹

What are the main arguments for taking this approach, and what are its advantages? Perhaps the most obvious advantage of the rule-based account is the explanation it offers for the apparent productivity of the sense alternations in question, i.e. the fact that they can be extended immediately to words that speakers encounter for the first time. For instance, if someone tells you that he saw an okapi in the Wild Animal

¹¹⁹ In the discussion to follow, it should be borne in mind that the primary goal of computational semantic accounts such as that of Copestake and Briscoe and others may not be to model how these sense alternations are manifested in actual communication, but to provide a formal account in which the process of operationalisation may be an end in itself. As such, the computational semantic accounts and the cognitive pragmatic account of utterance comprehension pursued in this thesis rest on substantially different methodological approaches. However, since the computational approach has proved to be so influential in accounts of systematic polysemy of the kind discussed in this chapter, it seems in my view appropriate to also assess its explanatory potential in accounting for how systematic polysemy is manifested in actual communication.

Park in San Diego last summer, and this was the first time you heard of the animal as well as the word for it, *okapi*, you would have no difficulty understanding an utterance of ‘*okapi* tastes a bit like horse’ (where *okapi* refers to the meat of the animal), although you would never have come across this use before. The computational semanticist view is that this suggests the existence of a productive rule, rather than several stored senses for a word (which they appear to see as the only alternative view).

However, the apparent productivity of such sense extensions need not have an entirely linguistic basis, although it is indeed a possible explanation for it. An alternative explanation would be to take the example above to be a case of the hearer exploiting his knowledge about certain regularities in the world (e.g. that animals are in general edible, that an animal’s flesh is considered ‘meat’, etc.) in his interpretation of *okapi* as referring to the meat. Once he has learned that *okapi* is an animal, the inference from animal to meat would come quite naturally. This would also serve to explain the conventional cases. However, since this real-world distinction coincides with the grammatical distinction between count and mass nouns in English, there would also be a clear linguistic clue (i.e. that the noun occurs without a determiner) to this interpretation. I return to this point in section 5.3.2.

Another seemingly compelling argument in favour of the rule-based account is the parallel one might draw between the conventionalised cases of grinding and derivational morphological processes. In particular, what Copestake and Briscoe see as a striking similarity is that both processes appear to be ‘blocked’ by the existence of an underived synonymous lexical form (Aronoff 1976; Briscoe, Copestake, and Lascarides 1995), (Clark and Clark 1979 describe this as ‘pre-emption by synonymy’). On this view, the existence of lexical forms such as *veal*, *pork* and *mutton* blocks the application of the meat-grinding rule, and thus explains why the following sentences strike us as odd:

- (19) ?Joan likes to eat *calf* (veal).
- (20) ?We’re having *pig* for dinner (pork).
- (21) ?Matt is preparing *sheep* for our anniversary (mutton).

Such lexicalised forms can be analysed as lexical exceptions to the meat-grinding rule (Ostler and Atkins 1992; Copestake and Briscoe 1996), by analogy with exceptions to

morphological processes. For instance, Aronoff (1976) notes that the form **gloriosity* does not appear in English (contrary to e.g. *curiosity*, *atrociousness*, etc.), which he attributes to the fact that it would be synonymous with the existing form *glory* and therefore blocked from being generated. However, unlike most exceptions to morphological patterns, which are usually considered ill-formed and do not occur in normal language use, the derived forms of animal terms such as *calf*, *pig*, *sheep* and *cow* seem to co-exist happily with the lexicalised forms *veal*, *pork*, *sheep*, and *beef*. Consider the following examples:

- (22) There were five thousand extremely loud people on the floor eager to tear into roast *cow* with both hands and wash it down with bourbon and whiskey (passage from Tom Wolfe's *The Right Stuff* (1979), cited by Copestake and Briscoe 1996: 38, my italics).
- (23) Hindus are forbidden to eat *cow* (?beef) (Nunberg and Zaenen 1992).

Copestake and Briscoe explain such cases in terms of non-synonymy of the ground and the lexicalised forms; in (22) and (23) *cow* is possible because it is not synonymous with *beef*. Their proposal is that when such uses occur, “they either convey a negative attitude to the consumption of the meat on the part of the speaker or an entailment of extended denotation, where more of the cow ... than is normally considered ‘meat’ is being treated as food” (1996: 38-39). They suggest that the utterances in (22) and (23) can be analysed as apparent violations of the Gricean maxim of Manner, that is, as choices of a less common or interpretable form from among different ways of expressing the same meaning, which, as they put it, “carries the (discourse) implication that the terms are not synonymous”. This gives rise to the interpretation that the speaker has expressed a negative attitude towards the consumption of the meat in (22), and the interpretation that the interdiction concerns the animal as a whole (due to its status in the Hindu religion) and not just its meat in (23).¹²⁰ This use is clearly more informative than the

¹²⁰ Nunberg and Zaenen (1992) propose a similar account, where such uses are explained in terms of apparent violation of the Gricean maxim of Quantity (roughly: *Say as much as and no more than the communicative circumstances require*) In the cases where the speaker has chosen to use the word *cow* instead of *beef*, the hearer is entitled to infer that she has some reason for using this vaguer term, as in (23). Unlike the computational semanticists, however, Nunberg and Zaenen do not think that the blocking phenomenon can be explained in entirely linguistic terms. Blutner (1998), working within a

possible utterance containing ‘beef’, which would not (at least not as easily) give access to this information.

However, there are many uses of derived forms which appear to be difficult to explain in terms of apparent maxim violation. For instance, in (24) both *pork* and *pig* would be acceptable, and in (25) they seem to be used interchangeably:

(24) Jews and Muslims don’t eat *pork/pig*.

(25) As a general rule, I don’t eat *pork*. This can be awkward – I often go for the veggie option, to avoid having to explain why I don’t eat *pig*.¹²¹

This does not mean that English speakers may not perceive the uses of *pig* in the examples above as being somehow more ‘marked’ than the conventional expression *pork* (although considerably less so than the uses of *cow* in (22) and (23)), to the extent that they may give rise to some additional “effects”. However, these uses seem so natural that, in my view, it is counterintuitive to assume that they represent (apparent) violations of a conversational norm (contrary to e.g. figures of speech, where this is more evident). Moreover, the uses of *sheep* and *pig* in (26) and (27) below are entirely conventional:

(26) Kate [*with a plate of food in front of her*]: This roast *sheep* is the best I’ve ever had.

(27) I love roast suckling *pig*.

Copetake and Briscoe could argue that although these uses are conventional, they carry with them an entailment that it is the whole animal that has been roasted, and not just any unindividuated portion of its meat. This may be right, but in that case it seems that the entailment cannot be explained as resulting from the use of a derived form instead

neo-Gricean framework, treats examples like (23) above as instances of Horn’s (1984) ‘division of pragmatic labour’ (unmarked forms tend to be used for unmarked situations and marked forms for marked situations), explained in terms of a Q-principle (corresponding to the first part of Grice’s Quantity maxim (*make your contributions as informative as required*)) and an I-principle (subsuming the second part of Grice’s Quantity maxim (*do not make your contribution more informative than required*), the maxim of Relation and (possibly) all of the Manner maxims).

¹²¹ Attested: <http://www.inventio.co.uk/pigs.htm>

of the lexicalised one, as the following derived uses may, given the appropriate context, carry the same kind of entailment (i.e. involving a whole roast chicken/turkey):

(28) We're having roast *chicken/turkey* for dinner.

This is not to say that the lexicon may not be sensitive to the frequency of use of such (quasi-) synonymous forms as *cow/beef*, *pig/pork*, etc., to the extent that the use of e.g. *cow* instead of the more frequent *beef* to denote the meat of the animal may be interpreted as carrying some additional meaning. However, the effect does not have to come from a semantic restriction on the use of *cow* due to the existence of the lexicalised form. An alternative approach is to analyse the so-called blocking phenomenon in terms of conventions of use (what Morgan 1978 terms 'conventions about language') where the use of a non-lexicalised form in cases where there exists a lexicalised one (e.g. *beef*, *veal*, *mutton*) induces a greater processing load on the hearer's pragmatic system which is then offset by the derivation of extra cognitive effects, rather than treating it as a semantic restriction on the application of a lexical rule. I return to this issue in section 5.3.2.

Another matter pointing to a pragmatic explanation of meat grinding cases is that whether a speaker chooses to use a ground or an unground form to refer to the meat of an animal appears to be governed by real world knowledge. Consider the examples in (29)-(31):

(29) Wolves eat *lambs* (?lamb).

(30) Snakes eat *mice* (?mouse).

(31) In Africa, lions eat *wildebeest*, *warthogs*, *zebras*, *buffalo* and different types of *antelope*.¹²²

Here the nouns are all used in their plural (i.e. count) senses, which is the conventional way of talking about what animals eat. The uses of *lamb* and *chicken* in (32), however, are straightforward cases of meat-grinding:

(32) Indians eat *lamb* and *chicken* (?lambs and chickens).

¹²² <http://www.catalogs.com/info/gadgets/what-do-lions-eat.html>

An intuitive explanation for the difference between these examples is the different ways in which humans and other animals get hold of their food. Unlike other carnivorous animals, humans usually do not hunt down the animals they eat themselves in their daily lives and, even when they do, they don't devour the uncooked flesh. It seems that the use of a ground sense somehow entails that the meat has been prepared in some human-like fashion (which is why e.g. 'wolves eat *mutton*' would sound odd). This does not seem to be the case for the 'unground' forms in (29), (30) and (31), where although it is the meat/flesh sense that is being communicated, the count form of the nouns allows the hearer to access his representations of the whole animals (which may give access to assumptions about the animals being hunted down as prey). The question, then, is this: Given identical linguistic environments, is it reasonable to assume that the meat senses in (32) result from the application of a rule of meat grinding, while the meat senses in (29)-(31) are pragmatically derived?¹²³ This asymmetry is, of course, possible, but one might ask if lexical inferences that are clearly influenced by world knowledge in this way are not better treated along pragmatic lines.¹²⁴

An argument that is often given in favour of treating systematic polysemy in terms of lexical rule application is the availability of default interpretations in uninformative (or 'null') contexts (Pustejovsky 1995a; Copestake and Briscoe 1996). For

¹²³ Ostler and Atkins's (1992: 84) discuss grinding applied to words denoting food items. This is only possible, they claim, when "the units of that food are not evident", as in e.g. *some egg, some crab, some salmon, some potato*, typically referring to some food substance on a plate, etc. They further note that words denoting pulses appear to resist grinding, so rather than saying 'Have some *pea/bean/lentil*' a speaker would say 'Have some *peas/beans/lentils*'. This they treat as a semantically-based constraint on the rule that turns food item denotations into food mass denotations (although, as far as I can tell, they do not specify what aspect of the semantics of pulse denoting word is responsible for the existence of the constraint). However, exceptions to this 'constraint' regularly occur. In appropriate contexts, the 'ground' senses of words denoting pulses are unproblematic, as in utterances such as 'I just love refried *bean* smeared on a taco' (<http://eatingindallas.wordpress.com/2009/07/26/battling-breakfast-tacos>), 'I'd like the grilled salmon fillet with mashed *pea*', 'Can I have a side-dish of *dhal*, please?', etc. How would these be explained in a rule-based account? As exceptions to a constraint on the application of a rule? In my view, the tendency to prefer the plural form of the pulse denoting noun over the mass form may rather have its explanation in the size of the food items in question - we seldom eat a single pea/bean/lentil/etc. - or anyway in some other real-world fact about pulses. In other words, this tendency appears to have a pragmatic explanation; one which obviously doesn't exclude the possibility of such words occurring with a mass interpretation, as in the examples above.

¹²⁴ The examples in (29)-(31) are, of course, not exceptions to the rule of meat-grinding, as the direction of the lexical inference goes from animal-denoting count nouns to food denoting mass nouns, and not the other way around (i.e. from a food denotation to the use of a mass term), and so the rule of meat-grinding could still be valid. However, the examples *do* show that the basis for animal terms taking on food denotations cannot be entirely linguistic.

instance, the most accessible interpretation of (33) below seems to be the one according to which *rabbit* denotes ‘rabbit meat’ (and thus that Sam enjoyed but later regretted *eating* the rabbit):

(33) Sam enjoyed but later regretted *the rabbit* (Copestake and Briscoe 1996: 42).

Recall Pustejovsky’s analysis of such cases as type coercions, where a verb that subcategorises for an NP or a progressive VP syntactically, semantically requires a complement with an eventive interpretation. In the case where this requirement is not satisfied by the surface syntactic structure, a type-coercion operator changes the denotation of the NP from an entity into an event consistent with eventive information (represented in the form of a ‘telic role’) stored in the lexical representation for the noun (the so-called qualia structure). Copestake and Briscoe’s account of the semantic processing of (33) is similar, and involves the selection of an appropriate aspect of the meaning of the complement (in this case its telic role: rabbits are for eating (by humans)). This, they claim, explains the ‘default’ character of the eating interpretation and militates against a pragmatic analysis. As they write, “the *meat-grinding* sense of *rabbit* provides a telic role which allows the eating interpretation to be constructed. However, if the lexicon does not propose such a sense, it is unclear what it is about the context which allows pragmatic specialization of the interpretation” (1996: 42).

I have two points to make about this claim. First, I cannot see how this explains the default character of the eating interpretation, when, on Copestake and Briscoe’s account, there are in fact three possible interpretations of *rabbit* made available by the lexicon: (i) the general ground sense (‘rabbit stuff’); (ii) the meat-grinding sense; and (iii) the fur-grinding sense. It is not clear to me why, on this lexicon-based account, it is the meat-grinding sense that is selected (which then gives rise to the eating interpretation), and not any of the other possible senses.¹²⁵ Second, I think the argument

¹²⁵ This shows a consequence of the account that Copestake and Briscoe themselves point out (1996: 56); that the lexical rules for sense extension lead to overgeneration. For instance, on their account, the sentence ‘*Rabbit* is expensive these days’ is, given the universal grinder, the rules of animal-meat grinding and animal-fur grinding, three-ways ambiguous between the general ‘rabbit stuff’ sense, the meat sense and the fur sense. While the overgeneration is in itself problematic (it does not seem psychologically plausible that all the senses are generated), there is also a question of how the ambiguity is resolved, that is, how hearers decide when one rule has prevalence over the others. A possibility would be to say that all three senses are initially generated by the linguistic system and then leave it to pragmatics to select the

that the availability of default interpretations in uninformative contexts shows the need for a semantic analysis considerably underestimates the fact that hearers rarely (if ever) come to the interpretation process ‘empty handed’, as it were. There are hardly any entirely context-free interpretations. Recall the relevance-theoretic view that one of the tasks the hearer has to solve in utterance comprehension is constructing a hypothesis about the contextual assumptions to be brought to bear in the process. This set of assumptions – a subset of his assumptions about the world – could include assumptions derived from the observation of the physical environment, encyclopaedic knowledge, memories and beliefs as well as the preceding linguistic context (Sperber and Wilson 1986/1995; Wilson and Sperber 2004). When the assumptions that the hearer may derive from the linguistic and extra-linguistic context are scarce, as in (33), he will have to rely more on information stored in his long-term memory in the interpretation of the utterance. The linguistic meaning of the utterance in (33) allows him to access information stored in long-term memory about the concept RABBIT, among which the assumption that ‘rabbits are delicious to eat’ would be highly accessible. This assumption, in combination with assumptions that the other concepts in the sentence may give access to (for instance, that ‘eating delicious food is an enjoyable activity’ which may have been made accessible by the concept ENJOY), leads to the eating interpretation being much more accessible than any of the other possible interpretations. This is what gives it a ‘default character’. However, its availability can be given a straightforward pragmatic explanation.

5.3 A relevance-theoretic account

In my proposal to follow for a (mainly) pragmatic account of systematic polysemy, I would like to start by considering a claim that Fodor and Lepore (2002: 117) make in their *Compositionality Papers* apropos the animal-meat alternation. Here they say, “Opportunities for polysemy arise from *how things are in the world* (or, anyhow, from how we take them to be)”. Fodor and Lepore are sceptical about the possibility of words being *lexically* polysemous. On their view, the reason that, e.g., *lamb* may be polysemous

correct one in the particular context. This would mean, however, that in the derivation of the ‘default’ eating interpretation of (33), pragmatics would have to intrude to yield the meat-grinding sense of *rabbit*, which in turn would allow for the eating interpretation to be constructed.

between the animal and the meat is that “lamb-the-meat comes from lamb-the-animal”, hence rests on a *real relation* among things in the world. Fodor and Lepore do not develop their view in any detail, but I think their claim is a deep one, which touches on what I see as the very essence of the polysemy phenomenon. I hope what I mean by this will become clear in the following sections.¹²⁶

The type of systematic polysemy giving rise to the animal-meat, animal-fur, tree-wood (and so forth) alternations in meaning is, as we saw in the previous section, closely linked to the well-studied grammatical distinction between count and mass expressions in English (and other languages with count-mass syntax). There is a vast literature on this topic, the details of which I cannot get into here. However, in the next section I will give a brief overview of a set of issues that have been raised concerning the count-mass distinction, before I present the relevance-theoretic account in section 5.3.2.

5.3.1 The count-mass distinction and systematic polysemy

According to common intuition, ‘count’ expressions describe individual entities (e.g. The *horses* were bought in the spring), or kinds of individual entities (e.g. *Horses* are animals), while ‘mass’ expressions describe portions of quantities (e.g. There’s some *milk* in the fridge), or kinds of quantities (e.g. *Milk* is healthy). Syntacticians have been interested in the morpho-syntactic characteristics underpinning the count-mass distinction. The usual view is that only count nouns admit of a contrast between the singular and the plural (e.g. *horse/horses*); mass nouns are almost always singular (e.g. *milk, rice, advice*, etc.). While count nouns can be modified by cardinal numbers (e.g. *two horses*) and by the quantifiers *each, every, many, several, few* and the stressed *some*; mass nouns occur with the quantifiers *little, much*, the unstressed *some*, as well as with measure phrases (e.g. *a bottle of milk*). Syntactically, common nouns are usually distinguished into count nouns and mass nouns by means of the feature [\pm count] (see for instance Gillon 1992, 1999).

Many philosophers (from Quine 1960 and onwards) have been interested in the semantic aspect of the count-mass distinction, *viz.* the ways in which count and mass

¹²⁶ In fact, the main implication of Fodor and Lepore’s claim, that what gets called polysemy is either a pragmatic phenomenon based entirely on how the world works, or it is standard homonymy, that is, a truly linguistic phenomenon where the same word form encodes two or several distinct meanings, is very much in line with the relevance-theoretic approach to polysemy maintained in this thesis.

expressions differ with respect to what they denote. Two criteria, those of cumulativity and divisivity of reference, have been proposed to distinguish the semantics of mass expressions from that of count expressions. The former criterion dates back to Quine (1960: 91), who noted that “*mass* terms like ‘water’ ... have the semantical property of referring cumulatively: any sum of parts which are water is water”. Although describing a characteristic feature of mass expressions, it is not distinguishing of them as it also applies to plurals (i.e. count nouns).¹²⁷ The second criterion of ‘divisive reference’, proposed by Cheng (1973: 286-287), states that for any stuff/entity referred to by a mass noun, any part of that stuff/entity is also denoted by that same mass noun. Although this criterion clearly has some intuitive appeal, Gillon (1992, 1999) points out that it seems to be refuted by Quine’s (1960: 99) earlier observation that “there are parts of water, sugar and furniture too small to count as water, sugar, furniture”: The denotations of many apparent mass nouns have parts that do not fall under the denotation of the same noun. Furthermore, there is a large class of words that appear to be count nouns, yet satisfy the criterion of divisivity of reference (e.g. *stone, rock, cord, rope*, etc.). So having divisive reference can neither be a necessary nor sufficient characterisation of a mass expression.

The distinction between mass and count syntax raises two psychological issues, concerning (i) how the distinction maps onto conceptual representation, and (ii) how children acquire the link between count-mass syntax and non-linguistic ontological categories. Traditionally, the grammatical count-mass distinction has been linked to the conceptual distinction between objects and substances (Macnamara 1982). Physical objects are typically count (e.g. *horse, man, rock*, etc), while substances are typically mass (e.g. *water, rice, air*, etc.). However, several authors have noted that the grammatical count-mass distinction only partly corresponds to the ontological distinction between objects and substances (see e.g. Ware 1975; Link 1983; Gordon 1985; Gillon 1992; Chierchia 1998; Pelletier and Schubert [1989] 2001). In English, for instance, count nouns also refer to abstract entities (e.g. *promise, emotion*) and *events* (e.g. *competition, journey*). Mass nouns also refer to abstract entities (e.g. *love, beauty*)

¹²⁷ For instance, if the animals in my garden are rats and the animals in my neighbour’s garden are also rats, then the animals in our two gardens are rats (cf. Link 1983).

and groups of unindividuated objects (e.g. *furniture, jewellery*). A widely held view, which better captures the range of possible denotations of count and mass expressions, is that count-mass syntax corresponds to a distinction between the kinds of entities we perceive as individuals, and those we perceive as unindividuated entities (e.g. Bloom 1994b) (where the notion of individual corresponds approximately to ‘discrete bounded entity’).¹²⁸

Already Quine (1960) addressed the problem of acquisition in connection with count-mass syntax and its link to non-linguistic ontological categories. He saw the ontological distinction between objects and substances as a cultural construction, and count-mass syntax as the means by which children learn to discriminate objects from substances. A logical consequence of this view is that there would be substantial differences in conceptual representation between children and adult speakers, and between speakers of English (and other languages with count-mass syntax) and speakers of languages that do not have count-mass syntax, as in classifier languages such as Chinese and Japanese (i.e. languages that contain morphemes whose function is to indicate the formal or semantic class to which items belong).¹²⁹ Against such linguistic determinism, psychologists have argued that the acquisition of count-mass syntax presupposes rather than gives rise to such basic ontological categories (Bloom 1994a, 1999; Papafragou 2005; Barner and Snedeker 2005, 2006), which may indeed have an innate basis.

There has been much less focus on the contribution of pragmatics to the interpretation of an expression as having a count or mass denotation. Pelletier and Schubert ([1989] 2001: 274), following Ware (1975), reject the possibility of developing a pragmatic account of the count-mass distinction, on the grounds that “most times there is nothing that the speaker has in mind which would allow us to classify the use of

¹²⁸ Some scholars do not think that this distinction adequately captures the semantics of mass terms either, as there seem to be cases where these denote discrete countable objects (e.g. *furniture, jewellery, money*). An alternative view is to treat the conceptual entity picked out by a mass term as being ‘unspecified’ as regards quantification and measurement (Gillon 1992, 1999; Barner and Snedeker 2005, 2006; Li, Barner, and Huang 2008).

¹²⁹ However, there is considerable controversy as to whether such languages in fact lack count-mass syntax (cf. Krifka (1995) and Chierchia’s (1998) view that all Chinese nouns are mass) or whether this distinction also appears in classifier languages under a different guise (c.f. Li, Barner and Huang’s (2008) distinction between ‘count-classifiers’ and ‘mass-classifiers’).

a noun phrase as either +count or +mass”. While I doubt that this is in fact correct, Ware’s argument is worth considering.

Ware’s point was the following. We cannot draw the count-mass distinction on the basis of speaker intentions, because (i) a speaker may have a determinate intention to refer to either an individual or substance of some kind, but the hearer need not recognise this intention for communication to be successful (e.g. the hearer may respond to a request for *the candy* without knowing whether the speaker is individuating or amassing), and (ii) the speaker herself may have an indeterminate intention, and so could use a definite singular NP with the intention to refer to something without a commitment to a count or mass denotation (e.g. if one Beethoven sonata was played, and the speaker says she liked *the Beethoven*, she does not have to be either using a count or a mass noun).

Although both Ware’s scenarios indeed seem possible, they require that the contexts are such that making the distinction would not be considered relevant (i.e. it would be too costly in processing terms or not contribute to the cognitive effects of the utterance). For instance, (i) could occur in a context where a bowl of candy is manifest to both speaker and hearer, and where the speaker asks the hearer to pass her *the candy*. Since the NP *the candy* is ambiguous with respect to whether it is count (in the singular) or mass, the hearer cannot be certain if the speaker considers what is in the bowl a portion of candy (mass) or a group of individual pieces of candy, but this is irrelevant to his ability to respond appropriately to her request. Similarly, in the example involved in scenario (ii), where the speaker says she liked *the Beethoven*, the communicative context does not require of either the speaker or the hearer that the denotation of *the Beethoven* is specified as either mass (‘the music by Beethoven that was just played’) or as count (‘the sonata by Beethoven that was just played’), since there would be no doubt on the part of the hearer what the speaker is referring to, regardless of whether he construes the referent in the one or the other way (or indeed, leaves the issue indeterminate).¹³⁰

¹³⁰ The situation where a distinction is not grasped but where the communication is still considered good enough for practical purposes is in fact quite common. Another sort of example is the use of an NP such as *The children* where ‘children’ can be thought of in distinction from adults or as ‘children of so and so’ (even when they are adult) and there can be instances where we would not need to know which of these ways the speaker was construing some particular individual children when she says ‘*The children* will be

In addition, Ware (1975) describes a number of cases in which a count or a mass occurrence of a noun does not make a noteworthy difference to its interpretation (this concerns in particular abstract terms, e.g. ‘needing *much more theory/many more theories*’, ‘finding *more truth/truths*’, etc.), as well as cases where there seems to be no obvious real world fact to explain the tendency to use a noun with a count or mass determiner (e.g. we say ‘much *whiskey/gin and orange*’, but ‘many *martinis/orange blossoms*’). Ware’s overall conclusion (endorsed by Pelletier and Schubert), appears to be that since there are cases in which the count-mass distinction is not reflected in the speaker-intended meaning of an utterance, pragmatics is not useful in making the distinction.

Of course, if the aim is to classify a noun as being (lexically specified as) either [+count] or [+mass], which is, as I understand it, Pelletier and Schubert’s aim, then I agree that pragmatics does not take us very far in the above examples. However, this aim appears to rest on the assumption that it is meaningful to speak of a count-mass distinction in the language existing independently of the real-world distinction between countable (individuated) entities and substances (unindividuated entities). I am not convinced that it is. In my view, rather than showing the inadequacy of a pragmatic explanation, Ware’s examples indicate that pragmatics may in fact have quite an important role to play in distinguishing between count and mass *uses* of words (or, as in the above cases, in not making a distinction at all). Recall the analysis of the definite singular NP *the rabbit* in (33) (‘She enjoyed but later regretted *the rabbit*’), which involved the activation of highly accessible encyclopaedic knowledge about rabbits (e.g. that rabbits are delicious to eat), which, together with assumptions activated by other concepts encoded by the sentence, gave rise to the interpretation that the rabbit was eaten, involving a mass use of *rabbit*. Note that the linguistic meaning of (33) is not at odds with a count interpretation of *rabbit* (e.g., that she enjoyed *having* the (live) rabbit but later regretted it, for instance, in a context where the referent has decided to buy a rabbit for her children - at first she enjoys having it (it’s cute, cuddly, etc.), but later she regrets it due to the amount of work involved in caring for it). In the absence of any

pleased’, that is, it won’t matter with regard to the hearer’s ability to recognise which individuals she is referring to.

syntactic clues, a crucial aspect of distinguishing between count and mass uses of nouns involves applying encyclopaedic knowledge about the entities they denote (and contextual information more generally) in their interpretation.

This brings us back to our starting point. As I briefly touched on above, many syntactic approaches take the count-mass distinction to arise from common nouns being marked as either [+count] or [-count] in the lexicon. Underlying this assumption is the idea that some nouns *are* count nouns, while others *are* mass nouns; a noun either falls into the one or the other category and the basis for this view is rarely questioned.¹³¹ However, several authors have noted that, in fact, most nouns (if not all) can be both count nouns *and* mass nouns (Pelletier 1975; Ware 1975; Krifka 1995; Gillon 1999; Pelletier and Schubert [1989] 2001). For instance, the mass noun *water* can be used as a count noun, e.g. in the context of a restaurant ('Two *waters*, please'), abstract terms like *love* and *hope* can be used both as mass nouns and as count nouns ('She needs some *love*'/'She's the *love* of my life', 'There is not much *hope* left'/'My primary *hope* is that I will finish this thesis'), count nouns like *sun*, *shopping centre* and *chicken* can be used as mass nouns ('We had a lot of *sun* at the beach', 'That's a lot of *shopping centre* for a small town', '*Chicken* is my favourite kind of food'). As Ware (1975: 384) puts it, "there are words with a double life – sometimes occurring as count nouns and sometimes occurring as mass nouns". This raises several issues: Given the ability of nouns to occur as both count and mass, does it make sense to operate with a syntactic distinction between count and mass *nouns*? Shouldn't the distinction rather be treated as pertaining to occurrences of entire noun phrases, where these are interpreted as being semantically count or mass, depending on whether they denote individuals or unindividuated entities (what Pelletier and Schubert [1989] 2001 term the 'semantic occurrence approach') (and in many cases, unspecified with regard to this distinction)? If so, where do our intuitions about single nouns being either count or mass (e.g. *horse* is count, *water* is mass) come from?

¹³¹ In light of this, it is easy to see the motivation behind rule-based accounts of the systematic polysemy that arises from the alternation between count and mass senses of a noun. If a noun is syntactically marked as either [+count] or [-count], then a lexical rule may operate to change the value of the feature, thus rules like 'universal grinding' and 'portioning' (a rule that converts food or drink denoting mass nouns into a portion of that substance, e.g. *three beers*) (Copestake and Briscoe 1996), or 'conversion' of count nouns into mass nouns and vice versa (Gillon 1999).

In the next section, I will outline a pragmatic, relevance-theoretic account of the class of systematic polysemy that rests on this ability of nouns to take on both count and mass interpretations. The account will be based on something like a ‘semantic occurrence approach’ to the count-mass distinction, as briefly described above. This allows for entire noun phrases to be encoded as having either count or mass denotations (or as being unspecified with respect to the distinction, as in (33) above), thus capturing what is clearly a morpho-syntactic dimension to the distinction, while not having to assume that nouns in themselves are syntactically either count or mass.¹³² The aim will be to show that pragmatics, often aided by syntactic clues, has a very constructive role to play in the interpretation of systematic polysemy involving alternations between count and mass interpretations of a single word. I will also suggest a possible answer to the question of where our intuitions about single nouns being either count or mass come from.

¹³² Pelletier and Schubert ([1989] 2001) is a survey paper in which they identify and discuss a range of problems associated with the count-mass distinction, and the ways in which these are manifested in the various syntactic and semantic approaches to it. As they see it, there are only two viable alternatives for an account of the count-mass distinction: (i) a ‘syntactic expression approach’, where COUNT and MASS are syntactic features which attach to nouns, and where ‘lexical extension’ rules may operate to change the count/mass feature and alter the semantic representation of the noun accordingly, and (ii) a ‘semantic occurrence approach’, where occurrences of noun phrases are interpreted as being (semantically) count or mass (depending on whether the entity it refers to is true of stuff or things, which is determined by the syntactic constructions) (ibid. 314). The first, syntactic account has the advantage of accounting for our intuitions about many nouns being ‘count nouns’ and many nouns being ‘mass nouns’. However, Pelletier and Schubert note that against this syntactic account it may be held that it is unclear what syntactic work the features actually do, when, “for every sentence which has a mass term in a given location, there is another sentence which has a count term in that position” (and vice versa); no constructions will be ruled out (ibid. 322). The only role of the count/mass features seems to be to direct the semantic translation process, which points to a semantic distinction rather than a syntactic one. The semantic occurrence account, on the other hand, in which Pelletier and Schubert take it that all nouns have a ‘comprehensive extension’, that is, being “true of kinds (of stuff or things), true of conventional kinds of servings, true of quantities (of stuff) and true of objects coinciding with quantities of stuff” (ibid. 314), and that their precise denotations are determined by the syntactic constructions they enter into (e.g. the determiner *a* allows semantic operators to convert this ‘basic’ meaning of a noun into one that is true of individuals and kinds of individuals), our intuitions about the ‘count’ or ‘mass’ nature of many nouns is not accounted for. Although I think this account is certainly on the right track, I take it a step further in section 5.3.2, where, in the pragmatic account that I propose, much of the work of deciding on the precise denotation of occurrences of nouns is left to pragmatic inference. This allows us to assume that instead of having ‘comprehensive extension’, many nouns encode concepts that are mentally represented as denoting (kinds of) individuals or (kinds of) unindividuated entities (the count-mass distinction is thus primarily a cognitive-conceptual one), and that pragmatics, in combination with syntactic clues, operates to yield occasion-specific senses of nouns.

5.3.2 A (mainly) pragmatic account of systematic polysemy

First some preliminaries. In relevance theory, which, as we have seen, is fundamentally a cognitive account of utterance interpretation, the distinction between linguistic semantics and pragmatics is seen as corresponding to different processes involved in utterance comprehension: linguistic decoding and pragmatic inference (Sperber and Wilson 1986/1995; Wilson and Sperber 1993; Carston 2002b). Assuming a modular language input system in the sense of Fodor (1983), the process of decoding linguistic utterances is performed by a language perception module (or parser), which takes the linguistic stimulus as input, executes a series of grammatical computations (or mappings) and delivers a semantic representation as output (called ‘logical form’ in relevance theory),¹³³ which then feeds the pragmatic inferential system.¹³⁴ A central characteristic of Fodor’s input systems (in addition to, *inter alia*, their automaticity and informational encapsulation) is that they deliver shallow outputs. The structured set of concepts that constitute the logical form of an utterance is rarely (if ever) fully propositional. Carston describes the logical form of an utterance as a ‘template’ or ‘schema’ for a range of possible propositions, which contain slots that have to be filled – a process that requires pragmatic inference (see Carston 2002b, Chapter 1, section 1.5.2 for more detail). A consequence of the assumption that linguistic decoding and pragmatic inference combine in verbal comprehension is that there is no need to assume that all words necessarily encode full concepts, or that every concept communicated by an utterance *could* in principle have been encoded (Wilson 1998; Carston 2002b).

Recall from the previous chapter that a mentally represented concept, a constituent of the ‘language of thought’ (Fodor 1975, 2008), is seen as an address (or entry) in memory that may give access to three types of information: (i) the *logical* properties of the concept (e.g. a one-way meaning postulate such as HORSE → KIND OF ANIMAL); (ii) a set of assumptions, or *encyclopaedic* information, about the denotation of

¹³³ This is, of course, a kind of abstraction since in actual on-line processing the decoded words/morphemes are delivered rapidly to the pragmatic processing system (which does not ‘wait’ to get the logical form as a whole).

¹³⁴ As mentioned in Chapter 3, in relevance theory, the pragmatic inferential system is also seen as being modular, perhaps constituting a sub-module of the more general theory of mind module (Sperber and Wilson 2002b; Wilson 2005), though in a somewhat weaker sense of ‘modular’ than Fodor’s.

the concept, that is, conceptually represented assumptions and beliefs, including stereotypes and culture-specific information about the denotation, and also, in many cases, imagistic and/or sensory-perceptual representations (e.g. HORSES HAVE MANES, HORSES ARE EDIBLE, HORSES HAVE FOUR LEGS, HORSES ARE USED FOR RIDING, HORSES LOOK LIKE THIS: [MENTAL IMAGE], etc.), and (iii) the *lexical* (i.e. phonological and syntactic) information connected with the linguistic form that encodes the concept (Sperber and Wilson 1986/1995: 85-93). As we saw in Chapter 4, the construction of *ad hoc* concepts (i.e. occasion-specific senses, narrower or broader than the linguistically-encoded senses) involves taking the encoded concept and its associated logical and encyclopaedic information as input, together with a set of contextual assumptions, using the relevance theoretic comprehension heuristic to derive warranted conclusions about the speaker's meaning.

Let us assume that our ability to discriminate between kinds of individuals and unindividuated entities is independent of count-mass syntax. This seems plausible, as there is little reason to think that speakers of e.g. Chinese or Japanese represent the world differently with respect to this distinction. There is also substantial evidence that early individuation capacities predate the acquisition of count-mass syntax in development (e.g. Papafragou 2005; Barner and Snedeker 2005).¹³⁵ If we assume that the ability to distinguish between individuals and unindividuated entities is one of the means by which we carve up and make sense of the world around us, it seems reasonable that many of our concepts will be mentally represented as either individual-denoting or unindividuated entity-denoting, in accordance with how we perceive (or categorise) the things they refer to in the world (independently of their natural-language counterparts). Given the short description of a mentally-represented concept I gave above, we can see the concept HORSE as being represented in the following way:

¹³⁵ In psychology, a dominant theory of the nature of human object knowledge is the *core knowledge hypothesis*, proposed by Spelke and her colleagues (Spelke 2000; Spelke et al. 1994). This hypothesis states that our knowledge of the physical world stems from innate principles for reasoning about inanimate physical objects. This knowledge shares many of the properties of a Fodorian module by being experience-independent, domain-specific and informationally encapsulated. A possible hypothesis, then, is that our ability to recognise countable objects versus undifferentiated masses (like water, snow, etc) is part of this 'core knowledge'. I return to this issue shortly in the discussion of 'meaning postulates' as a possible source of a conceptual count-mass distinction.

(34) HORSE

Lexical entry: +N (... etc.)

Logical entry: HORSE → ANIMAL OF A CERTAIN KIND

Encyclopaedic entry: IS USED FOR RIDING, IS OFTEN DOMESTICATED, IS POPULAR AMONG YOUNG GIRLS, LOOKS LIKE THIS [MENTAL IMAGE], ... etc.

The concept HORSE in (34) is associated with three types of information stored in its lexical, logical and encyclopaedic entries. The lexical entry contains information about the natural language counterpart of the concept (e.g. that the lexical form *horse* belongs to the category N). The logical entry associated with the concept consists of a meaning postulate (in the form of an inference rule), which applies reliably to any representation that meets its input conditions (i.e. the concept HORSE).¹³⁶ The encyclopaedic entry contains a set of assumptions that represent the individual's real world knowledge about horses (e.g. that horses are used for riding, are often domesticated, are popular among young girls, and so on). Similarly, the concept WATER may be represented (schematically) along the lines of (35), as associated with the tree types of information:

(35) WATER

Lexical entry: +N (... etc.)

Logical entry: WATER → LIQUID OF A CERTAIN KIND

Encyclopaedic entry: IS USED FOR DRINKING, IS NECESSARY FOR LIFE ON THE PLANET, FREEZES AT ZERO DEGREES CELSIUS, LOOKS LIKE THIS [MENTAL IMAGE] ... etc.

On the basis of the representations in (34) and (35), there are two different ways in which we could see the concepts as being specified as denoting individuals or unindividuated entities.¹³⁷ Either it could be seen as (i) following from the meaning postulate(s) associated with the concept, or (ii) being derivable from real-world knowledge of the denotation of the concept, stored in the encyclopaedic entry associated with it. I will consider each possibility in turn.

¹³⁶ Of course, a concept can be associated with more than a single meaning postulate.

¹³⁷ Having argued that concepts, and not their natural language counterparts, may be mentally represented as denoting individuals or unindividuated entities, the lexical entries associated with the concepts would not contain any information that would allow us to make such a distinction.

As already mentioned, relevance theory has adopted the Fodorian, informational atomist stance on lexical meaning. On the view expressed in (i), the logical information (or meaning postulates) associated with a concept is hypothesised to be what allows us to determine that it denotes individuals or unindividuated entities. Explaining how this might work requires some elaboration. Unlike Fodor, who, in recent work, denies that any of the inferences a concept enters into is constitutive of its content (e.g. Fodor 1998), relevance theorists take concepts to be associated with content-constitutive meaning postulates (see, in particular Horsey 2006 for a detailed defence of the need for supplementing Fodor's informational semantics with content-constitutive rules of inference).¹³⁸ That a concept is associated with a meaning postulate means that a tokening of the concept, e.g. HORSE, will cause the activation of the associated meaning postulate(s), 'HORSE \rightarrow ANIMAL OF A CERTAIN KIND', thereby tokening the concept ANIMAL OF A CERTAIN KIND. Assuming that the concept ANIMAL is associated with the meaning postulate 'ANIMAL \rightarrow LIVING KIND', the concept LIVING KIND would equally be tokened as a result of a tokening of the concept HORSE). This is one of a large array of concepts that Horsey (2006: section 4.3.2) describes as 'intuitive', that is, concepts for which we have an intuitive grasp (e.g. HORSE, ANIMAL, WATER, LIQUID, etc.).¹³⁹ He distinguishes between two types of intuitive concepts: those that are 'perceptual' and those that are 'inferential'. Assuming a modular view of mental architecture (cf. Fodor

¹³⁸ Horsey argues that the distinction between logical and encyclopaedic information associated with a concept should be seen as a psychological one, which does not coincide with the analytic/synthetic distinction in philosophy: "[T]he question of whether a given inference is supported by a meaning postulate or is represented in encyclopaedic knowledge is an empirical question which is to be determined by psychological investigation." (ibid. 75).

¹³⁹ The notion of 'intuitive concepts' comes from Sperber (1997), who distinguishes between 'intuitive' and 'reflective' beliefs. Intuitive beliefs are beliefs derived from perception (including introspection), or are derived directly or indirectly from these on the basis of intuitive inference mechanisms, e.g. in the form of meaning postulates attached to concepts (e.g. HORSES ARE ANIMALS/WATER IS A LIQUID). They can also be acquired via communication (someone might tell you that e.g. OKAPIS ARE ANIMALS), which, in such cases can be seen as playing the role of experience by proxy (the individual might herself have formed this belief via perception or spontaneous inference, had she been placed in a position to experience its perceptual basis). Reflective beliefs are those that are entertained by virtue of being embedded in a *validating context* (which may involve reference to authority, to divine revelation, explicit argument or proof, etc.), and may be derived, *inter alia*, from communication (resulting in, e.g., a belief of the form: MARY [WHO I FIND RELIABLE] SAID THAT *P*) or conscious thought (e.g. THERE IS EVIDENCE IN SUPPORT OF *P*). Parallel to this is a distinction between 'intuitive concepts' (i.e. those which may appear in intuitive beliefs and inferences, e.g. concepts such as HORSE, ANIMAL, WATER, LIQUID) and 'reflective concepts' (i.e. those of which we have no intuitive grasp, but understand because we have (reflective) beliefs about them, e.g. the concept H₂O, and which "are introduced by explicit theories which specify their meaning and the inferences that can be drawn on their basis").

1983; Sperber 1994a), perceptual concepts are the outputs of sensory stimuli modules (broadly identified with so-called ‘basic level categories’, cf. Rosch et al. 1976), while inferential concepts are taken to be innate concepts which are introduced by inference rules (meaning postulates), and which are not part of the vocabulary of the perceptual mechanisms. As an example, Horsey gives the concept *LIVING KIND*, which can be introduced via the spontaneous inference ‘*ANIMAL* → *LIVING KIND*’, and which Sperber (1994a), on the basis of strong commonalities in living-kind classification across cultures, suggests treating as a domain-specific cognitive module. Much research in cognitive science indicates that there are several different cognitive domains that licence spontaneous inferences about entities relevant to them, and that the domains have many of the properties of a Fodorian module. It is argued that such domains plausibly include a naïve physics (Spelke 2000), a naïve biology (Atran 2002) and naïve psychology (Leslie 1987), and that they are part of the cognitive apparatus which makes knowledge acquisition possible (Sperber 1994a) (cf. the *core knowledge hypothesis*, mentioned in footnote 135). Along these lines, we may speculate that the concepts *INDIVIDUAL* and *UNINDIVIDUATED ENTITY* may be such inferential concepts, perhaps constituting submodules of the more general domain of naïve physics. For instance, the concept *WATER*, which, according to the meaning postulate given in (35) causes tokening of the concept *LIQUID OF A CERTAIN KIND*, may activate the concept *UNINDIVIDUATED ENTITY* via the spontaneous inference ‘*LIQUID OF A CERTAIN KIND* → *UNINDIVIDUATED ENTITY*’. In the case of the concept *HORSE* in (34), which is argued to activate the inferential concept (or conceptual module) *LIVING KIND* as a result of the tokening of the concept *ANIMAL*, there may be a further spontaneous inference activated by the concept *ANIMAL*, ‘*ANIMAL* → *INDIVIDUAL*’. In this way, the distinction between concepts that denote individuals and those that denote unindividuated entities can be seen as arising from meaning postulates attached to the concepts.

A second option is to take the distinction between individuals and unindividuated entities to arise from encyclopaedic information stored about the denotations of concepts. As I have said above, the encyclopaedic entry attached to a concept contains information about the objects, events and/or properties which instantiate it. Typically, this entry contains information about prototypical instances of the concept (cf. the discussion of ‘prototypes’ in Chapter 2, section 2.2.2). For instance,

the encyclopaedic information about the denotation of the concept HORSE may contain information about what constitutes a typical horse, e.g. the kind of horse commonly used for riding (or the sport/show horse), and information about less typical instances, such as wild horses and small pony-sized horses (e.g. the Icelandic horse). For many concepts, we may also have an imagistic representation of the typical instances of the entities/properties they denote as part of the encyclopaedic entries attached to them (cf. the suggestion that physical object words may be associated with 3D model representations in Jackendoff 2002, as discussed in Chapter 2, section 2.2.4). It may be argued, then, that the fact that we perceive the concepts HORSE and WATER in (34) and (35) as denoting individuals and unindividuated entities respectively may stem from imagistic representations of typical instances of horses and water stored in the encyclopaedic entries attached to the concepts. We 'see' horses out in the world as bounded countable individuals, and water as an unbounded entity. On this view (which I tend to favour over (i), due to its intuitive plausibility), the count-mass distinction has a perceptual basis.

By arguing that the conceptual distinction between individuals and unindividuated entities can be seen as arising from either (i) the logical information attached to a concept, or (ii) an imagistic representation stored in the encyclopaedic entry of a concept, I hope to have shown that we can account for the intuition that some nouns *are* count nouns, while others *are* mass nouns, without having to assume that this is a linguistic property of nouns (thus avoiding the problems associated with a syntactic account) but rather to do with the kinds of things that they denote in the world. In my view, this offers a more psychologically plausible account: it explains that children and Chinese speakers may perceive horses as bounded countable entities and water as an unindividuated entity despite their 'lack of' count-mass syntax, and allows for many concepts, for which the classification in terms of individuals and unindividuated entities seems less relevant (as is the case for, e.g., abstract terms for emotions such as *love*, *hope*, *grief*, *anger*, and so on, plurals such as *blues*, *wares*, *clothes*, etc.), to be unspecified with regard to the distinction.

However, this approach does not deny (the obvious fact) that, in languages that have it, count-mass syntax plays a central role in determining the denotations of nouns. But rather than taking count-mass syntax as being (somehow) prior to the conceptual

distinction between individuals and unindividuated entities, count-mass syntax can be seen as an instance of a conceptual (ontological) distinction being reflected in a syntactic distinction at the level of NPs. Let us assume, then, that count-mass syntax (as it is manifested at the level of NPs) maps onto semantic representation as schematised in (36):¹⁴⁰

- (36) a. Count syntax → (KINDS OF) INDIVIDUALS
b. Mass syntax → (KINDS OF) UNINDIVIDUATED ENTITIES

These mappings allow a hearer to infer, on the basis of syntactic clues, that the entity denoted by the noun is either an individual or unindividuated (a substance).¹⁴¹ But instead of taking the nouns themselves to be either count or mass (which, on the basis of that syntactic feature, select specific determiners, can appear in the plural form, etc.) the view is now that it is concepts that are perceived as denoting (kinds of) individuals or unindividuated entities, and, independently of this, the language can encode NPs as denoting (kinds of) individuals or unindividuated entities.¹⁴² The normal case is, of course, an overlap between these two situations (with count-mass syntax allowing for the modification of nominal denotations by specifying cardinality, quantity, portion, amount, etc.), and so the concepts encoded by nouns such as *horse*, *man*, *pencil*, being mentally represented as denoting individuals, are most frequently used with count syntax, while the concepts encoded by nouns such as *water*, *rice*, and *snow*, being mentally represented as denoting unindividuated entities, are most frequently used with mass syntax. However, whenever a noun that encodes a concept which is perceived as

¹⁴⁰ These mappings are thought to model the linguistic knowledge of adult language users. I remain neutral as regards the very contentious issue of how such mappings are acquired (for discussion, see Bloom 1994a, 1994b, 1999; Papafragou 2005; Barner and Snedeker 2005, 2006).

¹⁴¹ For instance, if someone hears an utterance of 'There's much *glurf* these days' without having ever come across the noun *glurf* before, he or she can infer that the denotation of the NP *much glurf* is an unindividuated entity.

¹⁴² Krifka (1995) also considers the possibility of not distinguishing syntactically between count and mass nouns, but concludes that it must be done in order to account for what appears to him to be a syntactic difference, that mass-denoting nouns can be straightforwardly used as names of kinds (e.g. 'Wine contains alcohol'), while count-denoting nouns need a definite article or pluralization to perform this task (e.g. 'The bear/bears can be aggressive'). However, another way to see this is that count-denoting nouns can in fact be used without a determiner to name a kind, only that this would force a mass interpretation upon them (for instance, an utterance of 'Bear is delicious' would be possible given the appropriate context). So it seems that it is not whether the noun is syntactically defined as count or mass but rather how it occurs in the sentence that is the determining factor here. Occurring without a determiner and plural marker yields a mass interpretation.

denoting a (kind of) individual is used with mass syntax, or a noun that encodes a concept which is perceived as denoting a (kind of) unindividuated entity is used with count syntax, this gives rise to polysemy. Consider (37) and (38):

- (37) a. *The horses* are grazing in the field.
b. Yesterday, I had *horse* for the first time.
- (38) a. The plants need *some water*.
b. [To waiter]: I'll have the hamburger and *a water*, please.

In (37)a., the noun *horse* is used in a count NP to denote individuals, and in (38)a., the noun *water* is used in a mass NP to denote an unindividuated entity. There is thus a correspondence between the conceptual representations HORSE and WATER (being perceived as denoting individuals and unindividuated entities respectively) and the syntax of the NPs in which their natural language counterparts occur. In (37)b., where *horse* is used in an NP with mass syntax, the language parser outputs a concept, let's say HORSE{unindividuated entity}, which comes with an instruction to the pragmatic system that its denotation is constrained to unindividuated entities. This concept provides a highly underspecified input to pragmatic processing. The relevance-theoretic comprehension heuristic then operates to yield the speaker-intended concept (which is one among the various sub-tasks it performs in developing the logical form into a fully propositional string of concepts) resulting in the construction of the *ad hoc* concept HORSE*, paraphrasable as 'horse meat', which is narrower than the linguistically specified concept. The hearer constructs this concept on the basis of highly activated encyclopaedic information associated with the concept HORSE (e.g. HORSES ARE EDIBLE) and contextual information derived from the utterance situation (e.g. that the speaker went to a fancy restaurant last night). Similarly, in (38)b., where *water* is used with count syntax, the language parser outputs a semantic representation that is specified as an individual, let's say WATER{individual}. The pragmatic processor then takes over to construct an *ad hoc* concept on the basis of the encoded concept, highly activated encyclopaedic assumptions associated with it (e.g. WATER IS DRUNK FROM CONTAINERS), and other contextual assumptions derivable from the utterance situation, say, from the fact that the speaker and hearer are at a restaurant where people are served glasses of

water. The communicated concept, WATER*, paraphrasable as ‘glass of water’, is narrower than the linguistically specified concept, denoting a subset of the set of individual containers of water.

We are now (finally) in a position to account for the examples of systematic polysemy that we set out in the beginning of this chapter, under (2), (3) and (4), repeated below as (39), (40) and (41):

- (39) a. A *rabbit* jumped over the fence.
b. We’re having *rabbit* for dinner.
c. The model wore *rabbit* on the catwalk.
d. After the tractor had run over the body, there was *rabbit* splattered all over the yard.
- (40) a. We have a *pine* in our garden.
b. This table is made of *pine*.
- (41) a. Susan decorated the cake with a *cherry*.
b. When the kids left, there was *cherry* all over the kitchen.
c. Jill and Joan have a *cherry* in their garden.¹⁴³
d. This table is made of *cherry*.

In (39), the unindividuated senses of *rabbit* are indicated by the mass syntax of the NPs, but their narrower, contextually appropriate interpretations, i.e. ‘rabbit fur’ (39)b., ‘rabbit meat’ (39)c., and ‘rabbit stuff’ (39)d., are all derived by pragmatic inference, on the basis of encyclopaedic assumptions associated with the concept RABBIT (e.g. RABBITS ARE EDIBLE, RABBITS HAVE FUR, RABBITS ARE OF FLESH AND BLOOD, etc.). The wood sense of *pine* in (40)b. and of *cherry* in (41)d. arise as a result of the nouns occurring in NPs with mass syntax, encoding unindividuated pine and cherry concepts, combined with pragmatic narrowing. Importantly, this relevance-theoretic pragmatic account claims that the reason we can use *rabbit* to denote the senses in (39)b., (39)c., and (39)d., and *pine* and *cherry* to denote kinds of wood in (40) and (41) is first and foremost a result of what we know about the world (or what we perceive the world to be like), and not due

¹⁴³ The metonymic extension from the meaning of *cherry* in (41)a. above to the tree sense of *cherry* is, of course, not accounted for here but will be addressed in Chapter 6. The alternation between the tree sense and the wood sense, however, follows the pattern described above.

to any such information being stored in the lexical entries for the words, although mass syntax provides an important clue to the interpretations. Examples like (33) above, however, where the syntax of the NP is unspecified as to whether it encodes a count or a mass concept, show that the choice of a count or mass interpretation for a concept can have an entirely pragmatic basis.¹⁴⁴

However, although they are pragmatically derived, the fact that we perceive such polysemies as being ‘systematic’ or ‘regular’ is, I think, due to several (independent) factors: One concerns the syntactic distinction between count-denoting and mass-denoting NPs. The fact that our linguistic knowledge tells us that most (if not all) nouns can be used with both count and mass syntax, and that this syntactic distinction is paired reliably with a semantic distinction between individuals and unindividuated entities, gives rise to a sense of regularity. A second factor is the way in which this type of polysemy reflects certain regularities in the world. Fodor and Lepore make this point:

Suppose it’s right that ‘lamb’ is polysemous between the animal and the meat. Surely that’s because lamb-the-meat comes from lamb-the animal. Surely there just *couldn’t be* a word that’s polysemous between *lamb-the-animal* and (say) *beef-the-meat*? Or between *lamb-the-animal* and *succotash-the-mixed-vegetable*? That there couldn’t may itself sound like a deep fact of lexical semantics. But no; it’s just the truism that, the less one can see what the relation between X and Y might be, the more one is likely to think of an expression that is X/Y ambiguous as homonymous rather than polysemous. (Fodor and Lepore 2002:117)

Polysemous relations of the animal-meat, the animal-fur or the tree-wood kind indeed reflect highly regular and predictable states of affairs in the world. Because we know that there is an inherent relation between an animal and its meat/fur, and between a tree and its wood, we can easily infer, upon encountering a new kind of animal or tree, that the relation also applies to this instance (and in certain cases that it does not, for instance, if the animal in question has no fur). Since words for animals and trees are usually polysemous with respect to these relations, we can infer that the word for the new instance can also be used in this way. The view, then, is that the polysemy that the language exhibits in these cases is dependent on, and arises from our knowledge of

¹⁴⁴ Presumably, this would be crucial in classifier languages such as Chinese and Japanese, which do not distinguish syntactically between count-denoting and mass-denoting NPs.

predictable real world relations, rather than being generated by the linguistic system on the basis of stored lexical information.

A third factor that contributes to the perceived ‘regularity’ of this type of polysemy is, of course, that some of its instances have become conventional (e.g. *lamb*, *chicken*, *oak*, etc.), as a result of frequency of use and possibly other factors contributing to the stabilisation of senses (cf. Chapter 4). In such cases, the conventionalised sense may have acquired a conceptual address of its own, and would be associated with its own logical and encyclopaedic entries.

An advantage of the pragmatic approach is that it allows for considerable flexibility of interpretation. Consider (42):

- (42) a. A *chicken* pecked the ground in the backyard.
b. John brought home a *chicken* for supper.
c. Corn-fed *chicken* is difficult to find these days.
d. Corn-fed and inexpensive *chicken* is difficult to find these days.

In (42)a., *chicken* is used in the stereotypical way to denote the live animal. In (42)b. however, it is used in an NP with count syntax to refer to an individual chicken prepared to be eaten. This ‘meat’ sense of *chicken* represents a pragmatic narrowing of the linguistically-encoded meaning, by picking out a subset from among the set of chickens (i.e. the set of whole chickens prepared to be eaten). In fact, depending on the level of specificity required by the context, it could be narrowed down further to communicate, e.g., ‘a dead unplucked chicken’, ‘a frozen plucked chicken’ (like the ones you find in any Sainsbury freezer cabinet), ‘a cooked chicken ready to be eaten’ (as from a delicatessen), and so on. It is difficult to see how this ability of *chicken* (used in its individual-denoting sense) to take on both animal and meat senses could be analysed in an account based on lexical rules for sense extension. It is possible that the proponents of such an account would say that deriving the appropriate sense of *chicken* in (42)b. above is a pragmatic matter, and that they never intended their rules to account for such flexibility in meaning. If so, this would mean that only the rule-generated meat sense (e.g. the one that results from ‘meat-grinding’ on Copestake and Briscoe’s (1996) account) could be considered as ‘semantic’, while the meat sense of *chicken* in (42)b.

would come about as a result of pragmatic inference. This asymmetry does not arise in the pragmatic account, which I think points towards a more empirically adequate account. Concerning the mass uses of *chicken* in (42)c. and (42)d., they show that a property primarily associated with the live animal ('corn-fed') can be used with the meat sense of *chicken*, and that such a property can be conjoined with a property primarily associated with the meat sense ('inexpensive'). The fact that we do not find these utterances the least incoherent may be due to logical and encyclopaedic information activated during our processing of them (e.g. a tokening of the concept CHICKEN causes tokening of the concept ANIMAL, as well as the activation of encyclopaedic assumptions about chicken-the-animal, e.g. CHICKENS EAT CORN; CORN-FED CHICKENS HAVE MORE FLAVOUR THAN OTHER CHICKENS, etc.).

Finally, consider the uses of *mink* in (43) below:

- (43) a. A *mink* lives in the attic of our mountain cabin.
 b. I opened the cage and suddenly there was *mink* all over the place
 c. *Mink* is expensive these days.
 d. I wore my jeans and she wore a *mink*.¹⁴⁵

Let us assume that *mink* conventionally encodes two senses, the animal sense, MINK₁, and the fur sense, MINK₂ (where the MINK₂ sense has become conventional as a result of the frequent narrowing of the mass use of the animal sense ('mink stuff') into the 'fur' sense). Thus, the MINK₁ concept would be perceived as denoting a (kind of) individual, while the MINK₂ concept would be perceived as denoting a (kind of) unindividuated entity. In (43)a., *mink* is used in the conventional way (in an NP with count syntax) to denote the live animal (MINK₁), with a possibly narrower communicated meaning, paraphrasable as 'wild mink' (MINK₁*). In (43)b., *mink* is used in the same sense, only with mass syntax to refer to an unindividuated entity of *live* minks, communicating (possibly) the *ad hoc* concept MINK₁** ('farmed mink'), which is pragmatically constructed on the basis of the semantic representation MINK₁{unindividuated entity} output by the linguistic system. In (43)c., and (43)d. the noun *mink* is used to denote the fur sense, MINK₂, but in different ways: According to our analysis of the fur sense as

¹⁴⁵ Lyrics of 'She Ain't Pretty' from the album *Snow in June* (1990) by the band Northern Pikes.

being established, in (43)c., *mink* is used in the conventional way (in an NP with mass syntax) to denote mink fur ($MINK_2$). However, the count syntax of the NP *a mink* in (43)d. allows it to map onto an individual-denoting concept ($MINK_2\{\text{individual}\}$), while pragmatic narrowing yields the communicated meaning, e.g. ‘mink stole’ (or the *ad hoc* concept $MINK_2^*$). The $MINK_2^*$ concept can be seen as denoting a subset of the set of all individual mink furs. Although (43)b. (43)d. are not conventional uses of the noun *mink*, they do suggest that all that count-mass syntax encodes is an instruction that the concept output by the linguistic parser is specified as being either an individual or unindividuated.

Returning to the ‘blocking’ cases discussed in section 5.2, these can be analysed within the current, relevance-theoretic account as involving an increase in the effort of processing on the part of the hearer, which is offset by an increase in the cognitive effects he may derive from the utterance. According to the Communicative Principle of Relevance (cf. Chapter 3) every utterance (or other ostensive stimulus) “creates a presumption of its own optimal relevance” (Sperber and Wilson 1986/1995: §3). To be ‘optimally relevant’, the utterance should be (a) at least relevant enough to be worth the hearer’s processing effort, and (b) the most relevant one compatible with the speaker’s abilities and preferences. So, a speaker aiming at optimal relevance should try to do at least the following: achieve enough cognitive effects to make the utterance worth processing (cf. (a)), and avoid causing the hearer any unjustifiable effort in achieving those effects (cf. (b)). This ban on processing effort has two consequences. First, it implies that the first interpretation that the hearer finds satisfactory (for instance, if there is a highly accessible interpretation which is relevant in the expected way), this is the only one he is justified in choosing. A second consequence is that the extra processing effort demanded by any indirectness in an utterance should be offset by extra (or different) effects, which would not have been achieved by use of a more direct utterance. Consider again the examples in (22) and (23), repeated as (44) and (45) below:

- (44) There were five thousand extremely loud people on the floor eager to tear into roast *cow* with both hands and wash it down with bourbon and whiskey (Copestake and Briscoe 1996: 38).

(45) Hindus are forbidden to eat *cow* (beef?) (Nunberg and Zaenen 1992).

In both cases, the use of *cow* with mass syntax instead of the conventional expression *beef* demands an extra (albeit relatively minor) effort of processing of the hearer, which is offset by extra cognitive effects (compared to the alternative utterances containing *beef*): In (44), it allows the hearer to draw implications about the speaker's somewhat derogatory attitude towards the people consuming the meat, and in (45) about the status of the cow in Hindu religion. As to the utterances in (19)-(21), repeated below in (46)-(48), it can be argued that the reason the utterances seem odd is that, without a more specific context, the extra processing effort that they require of the hearer is not offset by any extra (or different) cognitive effects that he may derive from them:

(46) ?Joan likes to eat *calf* (veal).

(47) ?We're having *pig* for dinner (pork).

(48) ?Matt is preparing *sheep* for our anniversary (mutton).

Finally, in (49) and (50) ((24) and (25) above), the reason the uses of *pig* seem less marked than the one in (47) may be that in both cases the 'pig' sense has already been primed by the (linguistic) contexts (by the subject *Jews and Muslims* in (49) and the general subject matter in (50), 'the speaker does not eat pork'). In this way, the uses of *pig* do not necessarily require much extra effort to process, and hence the effects that might be drawn from them are quite minimal:

(49) Jews and Muslims don't eat *pork/pig*.

(50) 'As a general rule, I don't eat *pork*. This can be awkward – I often go for the veggie option, to avoid having to explain why I don't eat *pig*.'

(51) Kate [*with a plate of food in front of her*]: This roast *sheep* is the best I've ever had.

(52) I love roast suckling *pig*.

As to the utterances in (51) and (52) ((26) and (27) above), the uses of *sheep* and *pig* to refer to the meat of the animals are entirely conventional; there would thus be no extra processing effort required, and no extra effects obtained.

5.3.3 The *book/window* cases

In this section, I will discuss a set of cases that seem to involve a regular alternation between distinct senses, but where the intuition is that the senses (somehow) belong to a single conceptual unit, and where there is no syntactic/morphological difference corresponding with different senses. I call these the *book/window* cases for convenience. As we are going to see, these cases raise some difficult issues. My aim here is not to provide a full account of the *book/window* cases, but rather to suggest a possible approach to them. Consider (53)-(56):

- (53) a. The *book* is sitting on the coffee table.
b. John found the *book* interesting.
c. That *book* with the gorgeous cover is really good.
- (54) a. John broke the *window*.
b. Mary crawled through the *window*.
d. The *window* is the most important part of a bedroom.
- (55) a. The *school* needs a refurbishment.
b. The *school* announced that its GCSE results were the best it had ever received.
c. John made many friends at *school*.
- (56) a. Ken spilled coffee on the *newspaper*.
b. The *newspaper* told a story about an armed robbery that took place in London.
c. The *newspaper* announced that redundancies would be made.

In (53)a., the noun *book* refers to a physical object, in (53)b. it refers to an information type, and in (53)c. it denotes both. This alternation, between a physical object sense and an information type sense, extends to many other words denoting objects of similar type, e.g. *DVD*, *pamphlet*, *newspaper*, *encyclopaedia*, *piano*, *radio*, *film*, and so on. In (54), the noun *window* refers to a physical object in (54)a., an aperture in (54)b., and the two combined in (54)c. We find this alternation for the noun *door*, as well as for a number of objects similar to windows and doors, including *porthole*, *jalousie*, *dormer*, *entrance*, *entry*, *exit*, *portal*, *hatch* and so on. In (55), the alternation between a physical object (building) sense of *school*, as in (55)a., and an institution sense, as in (55)b.,

which may also combine, as in (55)c., is exhibited by a number of other nouns, including *church*, *parliament*, *castle*, *bank*, etc. Finally, the uses of *newspaper* in (56) are interesting, as they involve an alternation between a physical object sense, as in (56)a., and an information type sense, as in (56)b., as exhibited by *book* above, as well as an alternation between a physical object sense and an institution (/organisation) sense, as exhibited by *school* in (55). Note that the three senses cannot combine (i.e. the combination of the senses in (56)a. and (56)c. has no extension). Other nouns that exhibit this kind of sense alternation are *journal*, *magazine*, *periodical*, *radio*, etc. The patterns of sense alternation exhibited by (53)-(56) can be schematised as in (57):

- (57) a. physical object/informational content
b. physical object/aperture
c. physical object/institution

Let me start by considering what all these instances have in common. First, they can all be characterised as ‘regular’, in the sense that the patterns of sense alternation range over a large set of words, and can be extended to new words denoting objects of similar type. Second, they all have a physical object sense which alternates with a sense denoting a property (more or less) intrinsic to the physical object denotation of the concept. Third, the sense alternations are largely constrained by real-world knowledge of the denotation of the nouns, as shown by the examples below involving the nouns *book*, *radio*, and *computer*, exhibiting the sense alternation in (57)a.:

- (58) a. John found the *book/?radio/?computer* interesting.
b. The *book/radio/?computer* told a story about an armed robbery that took place in London.
c. The *book/computer/?radio* contained compromising information.

An utterance of ‘John found the *radio/computer* interesting’, as in (58)a., cannot be used to communicate that John found the informational content transmitted by the radio/stored on the computer interesting, but would rather be interpreted as pertaining to their technical details (in some way or other). The example in (58)b. shows that a book or a radio can tell a story about something, while a computer usually can’t,

presumably because we think of computers primarily as containers and not as transmitters of information. Hence (58)c. is fine with *computer* and *book*, but not with *radio*, which is an information transmitter, not a container.

The *book/window* cases raise several questions. First, there is a question concerning whether they involve a single lexical entry or distinct lexical entries. For instance, Cruse (1995: 36) points out that the senses may give rise to independent truth-conditions, as shown by (59) and (60):

(59) A: Do you like the *book*?

B: (i) No, I find the sentimentality nauseating.

(ii) Yes, it's magnificently produced – a pity the poems are such rubbish.

(60) A. Do you like the *church* in this town?

B: (i) No, it is in desperate need of refurbishment.

(ii) Yes, it's got a liberal position on same-sex marriage.

It is possible to conceive of a situation in which both of B's answers in (59) and (60) were simultaneously true. Recalling Quine's logical test for ambiguity mentioned in Chapter 1, this should indicate that the different senses of *book* and *church* are discrete, and hence involve distinct lexical entries. However, the following instances of 'co-predication' are fully acceptable, that is, both senses of *book* and *church* are simultaneously modified by different predicates without giving rise to 'zeugma':

(61) Roth's new *book* has a clever plot and an eye-catching cover.

(62) The *church*, albeit in desperate need of refurbishment, announced budget cuts next year.

Cruse takes these diverging results to show that the senses of *book* (and hence of *church*) are discrete but the word is not ambiguous, that is, it is represented as a single lexical entry. Cruse's view is that there is no sharp dichotomy between polysemy and monosemy but rather a continuum of degrees of distinctness, ranging from senses that are not distinct at all (e.g. the lexical form *teacher*, which could have either a female or a male referent), to senses that are fully distinct (e.g. homonyms such as *bank*), and presumably *book* and *church* fall into some middle position on this continuum.

Copestake and Briscoe (1996: 49) agree with Cruse that acceptable instances of co-predication suggest the availability of a single lexical entry but deny that a strong zeugmatic effect is a valid diagnostic of distinct lexical entries for different senses of a word form. Thus, in their view, the uses of *door* and *window* below are instances of a 'single lexical structure':

- (63) a. John went through the *door* and closed it behind him.
b. John opened the *door* and walked through it.
c. ? John walked through the *door* and slammed it behind him.
d. ?? John took the *door* off its hinges and walked through it.
- (64) a. The *window* was broken so many times that it had to be boarded up (Nunberg 1979: 150).
b. Mary opened the *window* and crawled through it.
c. ? Mary crawled through the *window* and slammed it shut.
d. ?? Mary pulled the *window* off its hinges and crawled through it.

Here, both the physical object and the aperture senses of *door* and *window* are modified by different predicates, giving rise to different degrees of acceptability. For instance, while the utterances in (63)a. and (63)b. both appear acceptable (with perhaps a difference in the degree of acceptability), (63)a. being slightly more acceptable than (63)b.), the utterance in (63)c. appears less acceptable and the one in (63)d. gives rise to a zeugmatic effect. Similar differences in degrees of acceptability can be observed for the examples involving *window* in (64). What causes such differences in the acceptability of co-predication? Intuitively, the difference appears to arise from the different types of events that the utterances describe. For instance, the sequence of events described in (63)a., that of walking through a doorway and closing the door behind oneself, is highly stereotypical, one that a person performs perhaps several times a day. This sequence of events may be stored as a 'chunk' in long-term memory, becoming highly accessible as a result of the activation of the concept DOOR. Thus, no anomaly is perceived in the processing of the utterance. Similarly, although *window* strictly speaking refers simultaneously to the glass and the aperture in (64)a., the fact that the utterance describes a highly conceivable sequence of events makes it seem perfectly natural. At the

other end of the scale, in (63)d., where the utterance is intended to describe the quite unusual event where John (for some unknown reason) goes through the trouble of pulling the physical door off its hinges before walking through the aperture it leaves, it seems that once the scenario where the physical door is being taken off its hinges is evoked and the referent assigned to *door* is the physical door (which the hearer is perhaps picturing as being placed facing the wall next to the aperture), it can no longer refer simultaneously to the aperture. The anomaly of (63)d. (and similar remarks apply to (64)d.), results from the ‘salience’ of the sequence of events evoked; in other words, from the application of world knowledge during the processing of the utterance. The degrees of acceptability of the utterances in (63) and (64) can also be seen as correlating with the amount of processing effort they require on the part of the hearer: The more stereotypical the sequence of events is, the less processing effort is demanded, whereas the more ‘marked’ (or salient) cases come with an increase in the cost of processing.

I therefore agree with Copestake and Briscoe in that a zeugmatic effect does not necessarily indicate the presence of distinct lexical entries for the different senses of a word form, since whether an utterance (or sentence) is perceived as zeugmatic or not seems to depend to a large extent on the hearer accessing general world knowledge during his processing, and, furthermore, is a matter of degree (cf. my comment on this issue in Chapter 1). However, although I am generally sceptical about the ability of such ambiguity tests to determine whether a word form is associated with single or multiple lexical entries (as these tests cannot always distinguish between effects that are due to lexical encoding and those that arise from pragmatic processing), I also agree that the acceptability of co-predication, as shown by (61), (62), (63) and (64) above, is at least suggestive of the availability of a single lexical entry.

So, assuming there is only a single lexical entry at play in the *book/window* cases (this also seems to be the dominant view in the literature, cf. Nunberg 1979; Cruse 1986; Pustejovsky 1995a; Copestake and Briscoe 1996), a second issue that these cases raise is what we could reasonably take their lexical representations to look like in order to give access to the different senses. The literature contains various suggestions. Pustejovsky (1995a: 91-95, 149-156) treats the *book/window* cases as so-called dot objects. For instance, the lexical representation of *book* is taken to include both the ‘physical object’ sense and the ‘informational content’ sense, as well as their combination, forming what

he terms a 'lexical conceptual paradigm'. Other word forms exhibiting the same type of polysemy are seen as subtypes of this lexical conceptual paradigm (e.g. *DVD*, *encyclopaedia*, *novel*, etc.). Assuming a system of lexical inheritance, a word like *novel*, describing a subtype of book, inherits these properties (as well as a set of other properties associated with *book*), and thus exhibits the same type of sense alternation. As discussed in other chapters, there are a number of general problems with Pustejovsky's generative account of polysemy, which I will not get back into here. More specifically, with regard to this idea that words such as *book*, *window*, *newspaper*, and so forth encode lexical conceptual paradigms, it has been argued that Pustejovsky seems to be doing little more than listing different senses under a single lexical entry, which it is one of the explicit goals of his theory to avoid (Fodor and Lepore 2002). However, although apt, this critique is somewhat unfair, since, unlike the Katz style sense enumeration lexicons, the system of lexical inheritance, together with the notion of dot objects, does allow Pustejovsky to capture the apparent systematicity and productivity of the sense alternations under discussion. His account predicts that words that are subsumed under a single lexical conceptual paradigm should exhibit the same kinds of sense alternations, and it predicts, by virtue of this information being encoded as part of single lexical representations for the words, the acceptability of the co-predications in (61), (62), (63)a., and (64)a.,. At the same time, however, it wrongly predicts that the zeugmatic sentences in (63)b. and (64)b. should be acceptable. If both the aperture and the physical object senses of *door* and *window* are encoded in the lexical representations of the words as part of a lexical conceptual paradigm, then there is no reason why the utterances should be unacceptable. Furthermore, in the case of *newspaper* (and words exhibiting a similar kind of sense alternation), for which Pustejovsky takes it that there must be a complex lexical representation enabling the derivation of the physical object sense, the informational content sense and the organisation sense, there is nothing in the theory to prevent the linguistic system from also generating the combination of the three senses, for which there is no corresponding denotation (i.e. there exists no object which is a physical object, an information type as well as an organisation, which may be called a *newspaper*).

Cruse (1995, 2004) proposes a similar account of the encoded meanings of the *book/window* cases, where these are seen as giving access to a specific kind of bundle of

semantic properties, called ‘facets’, described as discrete components of a single sense. For instance, *book* encodes a [TEXT] facet and a [TOME] facet; *church* encodes a [PREMISES] facet and an [INSTITUTION] facet. These are independent in the sense that they can give rise to different truth-conditions by allowing only one of the facets to be interpreted as relevant, as in (59) and (60), and they allow modifying adjectives to attach to one or another of them, giving rise to ambiguous phrases (e.g. *a new book* could mean either ‘a new text’ or ‘a new tome’). However, words encoding facets are not in themselves ambiguous, cf. the acceptability of co-predication, as in (61) and (62).¹⁴⁶ While Cruse’s account provides a nice description of how some of the *book/window* cases may be associated with different senses (or ‘facets’), it is not particularly explanatory – for instance, it does not give any indication of what it is about such nouns that gives rise to their encoding of such ‘facets’, or about how these are processed/activated in communication.

Pustejovsky’s and Cruse’s accounts represent what seems to be a dominant view in the linguistic semantic literature: that the sense alternations in question are to be given a linguistic explanation in terms of aspects of encoded word meanings. How might these cases be approached within relevance theory, on the assumption that encoded concepts are atomic, unstructured mental representations? Vicente and Martínez Manrique (2010) argue that the relevance-theoretic commitment to atomism and to inferentially derived *ad hoc* concepts encounters problems faced with ‘encoded polysemy’ of the kind I have been discussing in this section. They take *window* as an example, and claim that in order to account for its different senses in ‘He crawled through the *window*’, ‘The *window* is broken’, ‘The *window* is rotten’, ‘The *window* is the most important part of the bedroom’, there are only two possibilities available to the relevance theorist, neither of which is satisfactory: Either she could say that the word *window* encodes at least four different atomic concepts (e.g., WINDOW₁, WINDOW₂, WINDOW₃, WINDOW₄), an option which they claim requires a departure from the

¹⁴⁶ It should be noted that Cruse (1986: 65-66) sees the *window/door* cases, as exemplified in (54), as being ambiguous, due to the zeugma that arises in examples like (63)b. and (64)b. He thinks that the apparent compatibility of the different senses as shown by (63)a. and (64)b., should not be accepted as evidence for a single encoding, rather, “each case must be examined carefully to determine whether there are special factors preventing the appearance of zeugma”. I think this point could equally well go the other way: there may be special factors that contribute to the appearance of zeugma, cf. the discussion of (63)b. above.

commitment to the view that words encode atomic concepts. Or she has to assume that a different *ad hoc* concept is communicated in each case (e.g., WINDOW*, WINDOW**, WINDOW***, WINDOW****), an option which fails, they claim, since the polysemy is ‘encoded’ and therefore “does not depend on pragmatic processing”. According to Vicente and Martínez Manrique, the only way to account for the polysemy exhibited by *window*, is via a lexical decomposition approach of the kind proposed by e.g. Pustejovsky (1995a). If we take the lexical representation of *window* to contain the concept PANE, this would provide direct access to its meaning in an utterance of ‘The *window* is broken’.

There are several comments to make about this critique from Vicente and Martínez Manrique. First, it is not clear what they take to be the evidence for the underlying premise that the apparent polysemy observed for nouns such as *window* does not depend on pragmatic processing; as far as I can tell, they simply assume this. Until they can provide some evidence to support this claim, their rejection of the possibility of giving a pragmatic account does not appear to have much force. However, there is a sense in which Vicente and Martínez Manrique are partly right when they claim that an analysis of the different senses of *window* in terms of *ad hoc* concept construction fails. As we have seen, the outcome of *ad hoc* concept construction is either a concept with a narrower denotation than the linguistically-encoded concept (a proper subset of it) or a broader denotation (a proper superset of it). If it is the case (which appears to be what is assumed by Vicente and Martínez Manrique) that in an utterance of ‘The *window* broke’, the concept expressed by the noun is WINDOW PANE, then it is true that this concept could not have been arrived at via narrowing or broadening - its denotation would be a *sub-part* of the encoded denotation, but not a proper sub- or superset of it - hence this would not count as an instance of *ad hoc* concept construction. In the next chapter, I suggest that sometimes (in metonymy, in particular) the outcome of *ad hoc* concept construction may not be a narrower or a broader concept but an entirely different concept. This may well turn out to be what happens in some of the *book/window* cases as well (which would account for the truth-conditional differences observed in some cases). I will briefly return to this option in the next chapter, where, despite the fact that the *book/window* cases are not cases of inferential

narrowing/broadening in the strict sense, I will consider how an *ad hoc* approach for them might go.

Concerning the first option, that is, to take the word *window* to encode at least four different atomic concepts, I agree with Vicente and Martínez Manrique insofar as postulating different atomic concepts as being encoded by *window* is not a viable option for the relevance theorist; not because it would necessarily be at odds with conceptual atomism,¹⁴⁷ but because it seems highly implausible that speakers should have a distinct concept for each of the uses of *window* above. As already mentioned, there is a strong intuition that the *book/window* cases involve a single conceptual unit,¹⁴⁸ and there is therefore some intuitive appeal to the view that the different senses stem from complex representations associated with single lexical items. But does this have to be the case? Let us consider a different example, (which resembles the *window* case above):

- (65) a. The *car* is parked out back. (whole object)
b. The *car* is green. (frame)
c. The *car* has broken down. (engine)
d. The *car* is unlocked. (door)
e. The *car* smells of cigarette smoke. (interior)
f. The groceries are in my *car*. (trunk)
g. The *car* is leaking gasoline/oil/break fluid. (gasoline tank/oil tank/breaks)

Each use of *car* in (65) can be said to denote a distinct object, all of which are part of a car (except (65) where the whole object is denoted), in much the same way as different

¹⁴⁷ In her response to Vicente and Martínez Manrique, Carston (2010) points out that in relevance theory, the view is not that the mapping between words and atomic concepts is strictly bi-directional (that is, one-to-one; for every lexical item there is a single atomic concept and vice versa); rather, the view is that the repertoire of atomic concepts in our language of thought well exceeds the set of lexicalised atomic concepts. The mapping is more appropriately seen as unidirectional, that is, from a lexical item to a single atomic concept, but not vice versa. In the case of conventionalised ('encoded', 'semantic') polysemy, the consequence of this view is, as we have seen, that the different senses will be represented in much the same way as fully homonymous forms, that is, as distinct lexical entries (perhaps with a considerable degree of overlap in the encyclopaedic information they give access to). However, if we see the process of sense conventionalisation (or establishment) as being gradual, it is likely that there will be a stage at which a single lexical item can be said to 'encode' more than one atomic concept, which then, over time as a result of frequency of use, gradually part company and develop into separate mental files, but I cannot see why this should be an impossible situation for the conceptual atomist, so long as the concepts themselves are fully atomic.

¹⁴⁸ In fact, it seems that it is mainly linguists that are aware of the different senses of the *book/window* cases; they normally pass unnoticed to ordinary speakers (unlike, for instance, many metaphors).

aspects of a window are denoted in ‘He crawled through the *window*’ (aperture), ‘The *window* is broken’ (pane), ‘The *window* is rotten’ (frame), ‘The *window* is the most important part of the bedroom’ (whole object). Do we want to say that *car*, like *window*, is an instance of ‘encoded polysemy’, that the word is *polysemous* between all of these different senses? I think not. The uses of *car* in (65) seem to involve a kind of ‘economy of expression’: Instead of saying, e.g., ‘The *car*’s frame is green’, ‘The *car* door is unlocked’ or ‘The *car*’s interior smells of cigarette smoke’, and so forth, which would all be possible, but more costly ways of expressing the same meanings, the speaker can rely on the hearer’s ability to quickly and effortlessly figure out which part of the car she is referring to, which he does on the basis of *what he knows about what actual cars are like*, rather than on the basis of what he knows that the word *car* means.¹⁴⁹ It also seems hard to think of any of the parts communicated by use of *car* in (65) without at the same time activating a representation of the whole car. Cruse (2000: 47; 2004: 115) sees the source of such nuances in meaning as different ‘ways of seeing’, or ‘perspectives’, being built into the lexical representation of a word (not unlike Pustejovsky’s qualia roles). He writes:

[P]erspectives [can be explained] by analogy with looking at an everyday object from in front, from the sides, from behind, from on top, etc. All these different views are perceptually distinct, but the mind unifies them into a single conceptual unity. Something similar happens with meaning.

I think this observation is essentially right. However, I do not think that this happens at the level of lexical representation but rather at a level at which linguistic knowledge, perceptual knowledge and general world knowledge are integrated. The sense alternations observed for *book*, *window*, *school* and *newspaper* as shown by (53), (54), (55) and (56) could be seen as arising from our encyclopaedic knowledge about actual books, windows, schools and newspapers, and the perspectives from which we see them, and not from our linguistic knowledge connected with the natural language items that encode representations of them. We know that tokens or instances of books are physical objects and have an abstract informational content of some kind, that most windows

¹⁴⁹ Similarly, in the case of *window*, the speaker could have chosen to say ‘He crawled through the aperture of the window’, ‘The window glass is broken’, ‘The window frame is rotten’, but this would be an unnecessarily laborious way of expressing these meanings, as the hearer will immediately access these interpretations in the absence of any evidence that a different interpretation might be intended.

can be opened, are usually made of glass and have a frame, that prototypical schools have premises (buildings) and are a kind of institution, and that prototypical newspapers are physical objects with an abstract informational content published by an organisation/company. This knowledge is part of our encyclopaedic information about the denotations of the concepts BOOK, WINDOW, SCHOOL, and NEWSPAPER, along with other information about the objects (e.g. that a book's contents can be of high/poor quality, that books take a long time to write, are published by publishing houses, and so on and so forth). However, what gives this information a special status (which I think leads to the widely held view that it must somehow be lexically encoded) is the fact that it represents knowledge about properties that are, to some extent, intrinsic to (or defining of) them, certain aspects of which are often 'highlighted' (e.g. 'The *book* has an interesting plot' (content)).

In addition to this conceptual knowledge of their denotations, it is, as I mentioned in section 5.3.2 above, likely that we have stored spatial/imagistic information about the objects in our long-term memory. Recall Jackendoff's (1992b: 44; 2002: 346-350) talk of *spatial structure* (SpS) in connection with physical object denoting lexical items. The SpS component of an object word's meaning can be thought of as an imagistic representation of the prototypical instance of a category, in the form of a 3D-model, in a shape that allows for its recognition at different distances and from different perspectives. The SpS component is also thought to represent the moment-by-moment appearance of objects as well as the integration over time of the shape, motion, and layout of objects in space, and possibly the forces among them. I think there is something very appealing about Jackendoff's idea that object denoting words are represented partly in terms of an SpS structure. From a relevance-theoretic point of view, we can see this spatial information about an object being stored as part of the encyclopaedic information associated with the concept that encodes it.¹⁵⁰ Some of the sense alternations we have been discussing in this section may have as their primary basis spatial information associated with the concepts encoded. For instance, the concept WINDOW, due to the activation of an SpS stored in its encyclopaedic entry,

¹⁵⁰ As we saw in Chapter 2, however, Jackendoff's view is that SpS is an actual component of lexical meaning, and not merely associated with it.

allows the hearer to ‘zoom in’ on different aspects of the SpS during the processing of utterances such as ‘The *window* is broken’, ‘The *window* is rotten’, etc. (a process that would, of course, be relevance-driven). Similarly, utterances such as ‘The *book* is heavy’, ‘The *book* has a beautiful design’, ‘The *school* needs a refurbishment’, ‘Ken spilled coffee on the *newspaper*’, plus all the utterances involving *car* in (65) will activate spatial (imagistic) encyclopaedic information associated with the encoded concepts, allowing the hearer to ‘perspectivise’ his representation of the concepts.¹⁵¹

On this highly suggestive approach, the meaning variation observed for *window* (and many of the other *book/window* cases) is not treated as an instance of *ad hoc* concept construction. Instead, *window* is treated as having a constant referential meaning across these different contexts, with some of its encyclopaedic information receiving extra activation in each case, resulting in the perception of the object being viewed from a specific ‘perspective’. This perspective can be derived from the linguistic meaning of the utterance, or from assumptions made available by the broader utterance context. The speaker does not have to encode the perspective from which she ‘sees’ the object she is referring to (for instance, by producing utterances such as ‘The informational content of the book is interesting’, ‘The school personnel/institution announced budget cuts’, ‘The window glass is broken’); she can rely on the hearer quickly accessing his representation of the object in which the relevant part would receive more activation, due to the real world knowledge he has stored about the object. So, in this sense, perspective taking is part of the overall pragmatic process of arriving at the intended speaker meaning (although it is, in many cases at least, not necessary in order to arrive at the proposition expressed).¹⁵²

¹⁵¹ The notion of ‘perspective’ discussed here has much in common with Langacker’s notion of ‘active zones’ (1984, 1991). In Langacker’s terms, the part of a word’s *profile* (i.e. the entity or relation designated by the word) that is relevant or active in a given context is called its *active zone*. However, a crucial difference between the two is that while Langacker posits no clear-cut boundary between linguistic meaning and encyclopaedic knowledge, the relevance-theoretic view advocated here takes the perspectivising to happen at the level of encyclopaedic knowledge, and not at the level of encoded meaning.

¹⁵² The reader will have noticed that this suggestive account implies a move away from the claim that polysemy is an instance of *ad hoc* concept construction as concerns the *book/window* cases. This is partly right. A reason for this is that I am not certain that we have to do with genuine polysemy in many of these cases (i.e. clearly distinct senses of a word). However, although it seems clear to me that *the book/window* cases cannot be analysed as instances of lexical narrowing in the sense of Wilson and Carston (2007), they do have an affinity with narrowing, as the different perspectives highlight more specific parts of the

In Chapter 3 (section 3.3.1), I discussed Carston's (2002b) arguments for the ineffability of thought and for taking an essentialist position on the question of linguistic underdeterminacy ('for any given proposition/thought there is no sentence which fully encodes it'). Here a 'thought' was understood in its individualist, psychological sense (a more fine-grained notion of thought than 'proposition expressed') and such thoughts did not seem amenable to linguistic encoding. It was argued that the thoughts we have about ourselves, of people and of things around us, of our temporal and spatial locations at any given moment, will be different from the thoughts other people may have about the same entities, that is, our relation(s) to the referents will determine our mental representations of them. It is hard to see how such private representations could be encoded in natural language. Recanati (1993: 97) describes the kind of thought expressed by an utterance containing a referential expression as a *de re* thought. A *de re* thought has two dimensions: It has an objective truth-conditional content, that is, the singular proposition that the utterance expresses, plus a subjective content, 'a mode of presentation of the reference', which affects the cognitive significance of the object but not its truth conditions. So, for instance, relative to an utterance of 'The *car* is parked out back', two people may have *de re* thoughts that differ with respect to the mode of representation of the referent 'car', depending on their relation to it. I would like to suggest that the perspectives that the speaker invites the hearer to take on the referents of *book*, *window*, *car*, etc. in many of the examples we have been discussing in this section, have much in common with Recanati's *de re* modes of presentation. In many cases, they do not seem to be truth-conditionally significant (e.g., the referent of *book* is a book even though we are talking of its contents)¹⁵³, but they are cognitively significant. Perspectives, however, although clearly private in one sense, are partly shared between interlocutors, and, to some extent, they can be encoded (they are thus not entirely ineffable) (cf. the possibility of saying, e.g., 'The window *frame* is rotten', 'The car *door* is unlocked'). Consider an utterance of 'The *window* is rotten' for which the hearer probably focuses in on the wooden frame of the window during the interpretation. We

objects. The *book/window* cases could also be seen as having an affinity with metonymy, by involving types of part-whole relations. I return to this possibility briefly in the next chapter.

¹⁵³ However, in the case of *newspaper* (and the set of similar cases) the distinction between the physical object sense and the organisation sense is clearly truth-conditionally significant.

can see the hearer's activation of this perspective as resulting from the speaker exploiting what she can take to be a stable mental representation of the hearer's.¹⁵⁴ That is, the speaker can take for granted that the perspective she takes on the object she is referring to, is grasped by the hearer by virtue of the (encyclopaedic) information he has about the object, thus there is no need for a more specific encoding. It seems likely that we would find this sort of property highlighting for any physical object denoting lexical item of a certain frequency in the language. This is a way in which speakers can minimise the effort of production – we do not have to encode something that we can take for granted that people know about the objects we are referring to – which in turn reduces the processing load on hearers.

Finally, consider the uses of the indexicals in (66)-(70) (similar examples are discussed by Nunberg 1978, 1993; Sag 1981):

- (66) [*Pointing at a copy of The Sun*]: Murdoch has owned *that* since 1969.
- (67) [*Pointing at a copy of a periodical*]: *That* is up on the web now.
- (68) [*Holding up a copy of a book in a bookshop*]: *This* is the best I have ever read.
- (69) [*Pointing at the frame of a door, where the actual door has been removed for replacement*]: The police smashed *that* open.
- (70) [*Pointing at the local church*]: *They* are liberal on same-sex marriage.

In these utterances, the speaker is using an indexical in a way that exploits the hearer's encyclopaedic information about the relevant object. In (66), the speaker is pointing at a copy of *The Sun*, and expects the hearer to infer that she is intending to refer to the organisation that publishes it, in (67) she is pointing at a copy of the periodical while referring to its abstract informational content; similar situations arise in (68)-(70). In these cases, it seems clear that once the physical objects are (ostensively) evoked in the contexts (e.g. a copy of *The Sun*, a periodical, a book, etc.), they give immediate access to stored information about them, otherwise it would not have been possible for the

¹⁵⁴ Recanati (1993: Chapter 7) distinguishes between two types of *de re* modes of presentation: egocentric concepts and encyclopaedia entries. Egocentric concepts exist only by virtue of a 'fundamental epistemic relation' (i.e. one that presupposes acquaintance with the object) that holds between the individual and the object at time *t*, and cease to exist when the relation ceases to hold (e.g. 'The man who keeps staring at me'). They are thus of an unstable and temporary character. In contrast, encyclopaedia entries are stable mental representations, which do not presuppose acquaintance with the object (e.g. 'My sister's piano teacher').

speaker to use the indexicals in this deferred way and expect the hearer to understand what she is referring to. Furthermore, they show that the object itself activates this kind of knowledge quite independently of any natural language encoding. This is, of course, not an argument against the possibility of this knowledge being linguistically-encoded. Decompositionalists could say that it is doubly represented, once as general (language-independent) encyclopaedic knowledge about the object and once as part of the lexical representation of the natural language item for the object.¹⁵⁵ However, it is possible to argue against this possibility for reasons of economy; there is in principle *no need* for the lexicon to encode this information, and, furthermore, there is, as far as I can tell, no *evidence* for it being encoded either. Once mental representations of the objects are activated (be it via ostension, perception or natural language encoding), they give access to a range of information about the objects, which speakers and hearers take advantage of in communicating with each other.

The deferred uses of the indexicals in (66)-(70) raise difficult questions concerning the distinctness of what we have called the different ‘perspectives’ we might take on objects such as newspapers, periodicals, books, doors and churches. In (66), (67) and (68), tokens (of a newspaper, periodical, book) are used to stand for their types, on the basis of which the hearer infers the intended referents. In (66), the type that the token is an instance of (‘the newspaper called The Sun’) is clearly distinct from the (type of) intended referent (‘Murdoch’s News Corporation’). Since the word *newspaper* can be used to refer to them both (albeit not simultaneously) I think that it is reasonable to assume that we have to do with genuinely distinct (metonymically derived) senses of the word in this case, and I will return to such cases in Chapter 6. In (67), however, it is not equally clear that the perspectives involve entirely distinct objects. In this case, the speaker is pointing to a (physical object) token of a periodical, which can be seen as ‘standing for’ the type, and succeeds in referring to its (electronic) informational content. This type has physical object instantiations as well as electronic instantiations both of which contain the same informational content, and can be said to constitute different realisations of the same abstract object. In (68), a physical object token of a

¹⁵⁵ Within encyclopaedic semantics (Langacker 2008), this would not be an issue since words are taken to encode encyclopaedic information; thus no clear-cut boundary is assumed between linguistic knowledge and general world knowledge.

book is used to stand for the informational content of the type, to which the speaker attributes the property of 'being the best she has ever read'. In this case it seems clear that we have to do with a single conceptual unit. In (69), the aperture of a door is used to stand for the removed panel, and the reason the speaker can do this, it seems, is that she can assume that the hearer considers the aperture, the panel and the frame of the door parts of one and the same object (i.e. the door).¹⁵⁶ Finally, in (70), the local church building is used to stand for the views of the church as an institution. This is a typical case of a metonymic relation (hence the senses should be referentially distinct) yet there is a feeling that the building and the institution are somehow different parts of a single (perhaps more abstract) object.

We have argued that 'perspectivising' of communicated meanings of words such as *book*, *newspaper*, *door*, *church* and so on arise from properties intrinsic to, or closely associated with, their denotations (information about which is stored as part of the encyclopaedic entries for the encoded concepts), rather than from information stored as part of their lexical representations. It may seem that underlying this argument is the assumption that these words have physical objects as their denotations (for instance, *book* encodes the concept BOOK, which has as its denotation a physical object, a book, and the encoded concept activates a range of encyclopaedic information about this denotation). There would, I think, be good reason to take this view on the encoded meanings of these words. The physical objects are the most accessible meanings; they are usually the first to come to mind when hearing the words in isolation (they are something like Copestake and Briscoe's (1996) 'default senses'), and they are the first to be acquired (Bloom 2000). It is also hard to think of, e.g., a book or a newspaper without at the same time thinking about their physical instantiations.

However, it also seems possible to think of the meanings of these lexical items in more abstract terms, where a word such as *book*, for instance, encodes an abstract type with different instantiations, e.g., bound books, audio books, e-books, etc., which, in view of the recent technological developments, may prove to be a more precise description of its denotation. A further, more extreme version of this abstract view of

¹⁵⁶ Although, in the case of *door* (and similarly in the case of *window* and their relatives) it seems likely that there will be disagreement among people as regards what constitutes a door: Is it the whole object, including the doorway, the panel as well as the aperture, is it only the panel, and so on.

the meaning of these words is that they do not encode concepts with denotations at all, but, rather, something much more schematic and underdetermined (a pointer), which gives hearers access to the conceptual material mentioned above, from which they construct the contextually appropriate concept (which does have a denotation). This is a very interesting option, which is currently being explored in relevance theory (Carston 2002b, 2010; Kjølth forthcoming) and by scholars working within different frameworks (Recanati 2004; Bosch 2007; Pietroski 2008; Pritchard 2009).

However, in the final instance, the determination of the precise denotation of a word is (if at all possible) an empirical issue. And, given that it is sometimes hard to determine exactly what parts constitute objects of a given category (cf. the discussion of the examples in (66)-(70)), it is not surprising that word meanings may reflect this uncertainty. For instance, although we can be said to have the concept DOOR, we may still be unable to determine whether a door consists of an aperture, a frame and a panel, or just the frame and panel, or the panel alone, and so on; we can very easily distinguish school buildings from the people running the school, but we may be uncertain as to whether a school building without anyone running it would still count as a school (cf. 'This used to be a *school*'), yet there can be no doubt that the majority of us have the concept SCHOOL. So perhaps this lack of precision of encoded word meanings is something that we (that is, linguists and philosophers) will have to live with. Underspecified lexical entries are, however, compensated for by an efficient pragmatic system, which allows speakers to tailor their utterances in accordance with the knowledge that they can expect hearers to have (thereby minimising effort – there is no need for the speaker to encode something she can expect the hearer to have immediate access to), and which allows hearers to infer speaker meanings on the basis of encoded content and contextual assumptions, a process which is guided and constrained by considerations of relevance.

5.4 Conclusion

In this chapter I have suggested a mainly pragmatic account of the systematic sense alternations involved in examples such as 'John shot a *rabbit*/had *rabbit* for lunch', usually seen as prime candidates for an analysis in terms of lexical rule application (where the effect of the lexical rules is to change the value of a linguistically marked

+count or +mass feature on a noun). I argued that such rule-based accounts do not provide interpretive flexibility required to handle the variety of interpretations that this kind of sense alternation may give rise to. Instead I proposed that our intuitions about some nouns being ‘count’ and other nouns being ‘mass’ stem from our mental representations of the concepts that are encoded by them, which may be perceived (by virtue of meaning postulates attached to the concepts or imagistic representations in their encyclopaedic entries) as denoting individuals or unindividuated entities (or indeed, be unspecified with regard to this distinction). I further suggested that the count-mass distinction as it is manifested in language could be seen as an instance of this conceptual (ontological) distinction being reflected in a syntactic distinction at the level of NPs (rather than at the level of individual nouns). The cases of ‘systematic’ polysemy would arise whenever a noun that encodes a concept perceived as denoting an individual is used with mass syntax, or when a noun that encodes a concept perceived as denoting an unindividuated entity is used with mass syntax. The pragmatic process of *ad hoc* concept construction would then operate to yield context-specific senses of these nouns (e.g. to specify whether a mass use of *rabbit* communicates ‘rabbit meat’, ‘rabbit fur’, ‘rabbit stole’, ‘rabbit remains’ etc.).

In the last part of this chapter, I have discussed the systematic polysemy of nouns such as *book* and *window*, and suggested that their different senses could be analysed in terms of activation of encyclopaedic properties associated with the concepts (so-called ‘perspectivising’). I have also pointed out that many of the examples that were discussed under this section appear to have an affinity with metonymic uses of words, a topic to which I turn in the next chapter.

Chapter 6

METONYMY

6.1 Introduction

This chapter is about metonymy as a source of polysemy. Metonymy (derived from Greek *metonumía*, meaning ‘change of name’) is the case where an expression that conventionally denotes one object or property is used to refer to (or ‘stand for’) something that falls outside its conventional denotation but with a clear (associative) relation holding between the conventional and the metonymic denotations. Here are some illustrations:

- (1) The *ham sandwich* left without paying.¹⁵⁷
- (2) Susan is reading another *Woolf*.
- (3) Jane is just a *pretty face*.
- (4) Peter is bringing the *loudmouth* with him to the cinema this evening.
- (5) John has married a *free ticket to the opera*.¹⁵⁸

In (1), uttered in the context of a restaurant, the speaker is referring to a customer who has ordered a ham sandwich, while the proper name *Woolf* in (2) is used to refer to a book written by the author Virginia Woolf. In (3) the predicate *pretty face* is used to describe Jane as being good looking (someone who ‘has’ a pretty face), and, helped by ‘just’, to convey the implicature that she is a shallow and/or unintelligent person. In (4), the noun phrase *the loudmouth* is used to refer to a person who tends to talk too much and/or too loudly (and as such it allows the speaker to convey a negative attitude towards this person). In (5), the noun phrase *a free ticket to the opera* refers to someone who regularly receives free tickets to the opera, and, as a whole, the utterance may implicate that John’s marriage is in a difficult state or that his primary interest in marrying the woman was her ability get him free tickets to the opera.

Such metonymies have been shown to pose a formidable challenge to semantic and pragmatic theory, and, although they have received increasing research attention

¹⁵⁷ Example due to Nunberg (1978).

¹⁵⁸ Example due to Sperber (2007 [1975]).

over the last couple of decades, the mechanisms underlying them are not well understood. In this chapter, my main concern will be to suggest an answer to the following question: How are metonymic interpretations derived within the inferential process of forming and confirming hypotheses about intended speaker meanings? Two hypotheses concerning metonymy interpretation will be investigated: According to the first hypothesis, metonymic interpretations are derived by means of the process of *ad hoc* concept construction, as discussed in the previous chapters. The second hypothesis takes referential metonymy to involve a creative use of ‘naming’, closely related to nicknaming, as was proposed by Anna Papafragou (1996) in an early relevance-theoretic account of metonymy. The overall aim of this chapter is to investigate the contribution of metonymy to polysemy.

While a topic of long-standing interest in rhetoric, poetics and literary criticism, the study of metonymy has a relatively recent history in philosophy of language, linguistics and cognitive science, dating back to the pioneering approaches by Nunberg (1978, 1979) and Lakoff and Johnson (1980). Compared with metaphor, which has been intensively studied over the last 30 years, metonymy has received only moderate attention, and mainly so in the cognitive linguistics tradition (Lakoff and Johnson 1980; Lakoff 1987; Lakoff and Turner 1989; Panther and Radden 1999; Barcelona 2000b; Fauconnier and Turner 2002; Taylor 1989/2003; Dirven and Pörings 2003), which analyses linguistic metonymies as reflections of underlying cognitive structures, in the form of ‘metonymic conceptual mappings’ (e.g. PRODUCER FOR PRODUCT (2), FACE FOR PERSON (3)). I come back to these approaches in section 6.2.2.

In recent years, however, there has been a growing interest in the phenomenon from the point of view of philosophy of language and pragmatics. The recognition that metonymy cannot be analysed along traditional Gricean lines ([1967] 1989), as involving implicatures derived on the basis of the hearer having ‘said’ something patently absurd/false (e.g. in (1) above, that the object made of bread and ham is getting impatient), has led to several proposals for how metonymy can be accommodated at the truth-conditional level, hence within a compositional semantics (Sag 1981; Stallard

1993; Pustejovsky 1995a; Stern 2000, 2006; Stanley 2005).¹⁵⁹ The challenge for more cognitively or pragmatically oriented approaches to metonymy (Nunberg 1978, 1979, 1996, 2004; Fauconnier 1994 [1985]; Papafragou 1996; Recanati 1995, 2004; Ward 2004; Evans 2009), which are less concerned with the problems metonymy raises for semantic compositionality, is to explain the nature of the pragmatic process that leads to metonymic interpretations. Central questions are: What are the circumstances under which a speaker may take an expression *a* which refers to *A* and use it to successfully refer to *B*? What are the constraints on the possible relations between *A* and *B*? What is the cognitive and communicative motivation for using a metonymic expression, instead of a literal expression with a similar meaning?

Most accounts distinguish between metonymy and metaphor. Traditionally, metonymy has been conceived of in terms of real-world contiguity relations (e.g. between ham sandwiches and their orderers), while metaphor has been taken to involve a resemblance relation between two entities (e.g. ‘Jane is (like a) princess’) (see e.g. Ullmann 1962). In the cognitive linguistics tradition, it is common to distinguish metonymy and metaphor in terms of mapping relations occurring within a single domain (metonymy) and across domain boundaries (metaphor). Some accounts, however, do not make the distinction; for instance, Nunberg (1996, 2004), treats both metonymy and metaphor as outcomes of a single process of ‘meaning transfer’. I return to this issue under section 6.2.

The connection between metonymy and polysemy has long been recognised (e.g. Apresjan 1974; Lakoff 1987; Taylor 1989/2003; Pustejovsky 1995a; Nunberg 1978, 1979, 1996, 2004; Bartsch 2002; Seto 2003). Apresjan (1974: 16) distinguished between ‘regular’ and ‘irregular’ polysemy, and claimed that being metonymically motivated was

¹⁵⁹ Unlike metaphor, which, as mentioned in Chapter 4, is still quite commonly analysed as involving a blatant violation (a flouting) of the maxim of Quality which triggers the search for a related implicature (although the view that it contributes to the proposition expressed by an utterance continues to gain support), philosophers of language and pragmatists now widely agree that metonymy is a local process that does not presuppose the prior identification of a proposition serving as input to the interpretive process. A possible exception may be Bach, who, according to Recanati (1993: 265) claims that “the absurdity of the ‘literal’ interpretation is what triggers the transfer from ham sandwich orderer”. See Bach (1994) for discussion. However, if metonymy contributes to the proposition expressed by an utterance (i.e. to its truth-conditional content) it violates semantic compositionality, that is, the principle that the meaning of a complex expression is determined by the meanings of its constituent parts, which is seen by many philosophers as the main problem in connection with metonymy.

a distinctive feature of the former, while the latter was typically metaphorically motivated. He described the polysemy of a word *a* with the senses A_i and A_j as being ‘regular’ if, in a given language, there exists at least one word *b* with the polysemous senses B_i and B_j , being semantically distinguished in exactly the same way as A_i and A_j (cf. Chapter 5). The polysemy of a word *a* was called ‘irregular’ if the semantic distinction between A_i and A_j was not exhibited by any other word in the language. The connection between metonymy and ‘regular’ polysemy on the one hand, and metaphor and ‘irregular’ polysemy on the other, arguably has some intuitive appeal. Recall from the previous chapter that many systematic sense alternations (such as ‘fruit for tree’, ‘publication for publisher’, etc.) appeared to involve metonymic relations between the encoded and the communicated concepts.¹⁶⁰ These obviously differ from poetic metaphors, which are usually one-off and thus clearly ‘irregular’ in Apresjan’s sense, but they also differ from ‘everyday’ metaphors (e.g. ‘Bill is a *bulldozer*’, ‘The critics *slaughtered* her new book’), which, although often conventional, cannot be said to involve any kind of ‘regularity’ in the above sense. However, there are numerous cases of so-called creative metonymy that are not obviously ‘regular’ either, as shown by the examples in (3), (4) and (5) above. Moreover, there are instances of metaphor, which, if not ‘regular’, appear to be at least semi-productive, such as for instance the metaphorical use of an animal denoting noun to describe some particular human characteristic (e.g. ‘Peter is a *lion/chicken/pig*’), or metaphorical uses of (human) body parts to describe analogous parts of an inanimate objects (e.g. *foot* of a mountain/tree/table, *mouth* of a bottle/river/cave).¹⁶¹ Furthermore, if the cognitive linguistic accounts (e.g. Lakoff and Johnson 1980 and many others) are right about the existence of metaphorical conceptual schemes (e.g. LOVE IS A JOURNEY, PSYCHOLOGICAL FORCES ARE PHYSICAL FORCES, etc., cf. Chapter 5), then there should in fact be a considerable amount of regularity in metaphorical use of language.

¹⁶⁰ These gave rise to examples such as the following: ‘Joan and Jill have a *cherry* in their garden’, ‘The *newspaper* sacked its editor-in-chief’, etc.

¹⁶¹ In fact, the apparent ‘productivity’ of the extension from animals to human characteristics has led some computational semanticists (e.g. Copestake and Briscoe 1996: 37) to suggest that, in addition to the pragmatic factors that clearly play a role in this process, there is a “core component to such processes which should be expressed in terms of a sense extension rule”. Although it should be clear from the previous chapter that I am sceptical about the feasibility of such an approach, I agree that there is a sense in which these metaphorical uses are different from more poetic metaphors as concerns their predictability in English (and probably also in a range of other languages).

In previous chapters, I have argued that polysemy arises as a result of a pragmatic process of *ad hoc* concept construction, which operates at the level of individual words and whose outcomes are concepts with either a narrower or a broader denotation than those linguistically-encoded. As we have seen, relevance theory takes utterance comprehension, including the construction of *ad hoc* concepts, to be a wholly inferential process, with a unitary, on-line pragmatic processing system which derives explicit content ('what is said') and conversational implicatures in parallel, driven by the hearer's search for an interpretation that meets his expectations of relevance. This often involves a process of 'mutual adjustment', where a hypothesis about an implicature may both precede and shape a hypothesis about an explicature. The process of mutual adjustment is, as we saw, a crucial component of the relevance-theoretic account of lexical interpretation (Carston 2002b, 2010; Wilson and Sperber 2002; Wilson and Carston 2006, 2007; Sperber and Wilson 2008). As already mentioned, my aim in this chapter is to suggest an account of metonymy that brings it within the scope of this inferential theory of utterance comprehension. Most accounts of the comprehension of metonymy take it to proceed through some kind of associative (i.e. non-inferential) relation between the encoded concept and the speaker-intended concept, in the form of 'pragmatic functions' (Nunberg 1978, 1979; Fauconnier 1994 [1985]; Ward 2004), 'domain mappings' (Lakoff and Turner 1989; Croft 1993; Gibbs 1994; Radden and Kövecses 1999), 'meaning transfer' (Nunberg 1996, 2004), or 'schema activation' (Recanati 1995, 2004). To my knowledge, apart from Papafragou's (1996) early relevance-theoretic account, there have been no attempts to investigate the extent to which the comprehension of metonymy may be explained inferentially, that is, as dependent on a non-demonstrative inferential process of constructing and confirming hypotheses about intended speaker meanings. As will be discussed in section 6.3 of this chapter, metonymy represents a challenge to such inferential accounts, since it appears to rely on arbitrary (i.e. associative) rather than inferential relations between linguistic meanings and communicated meanings, but the possibility of an inferential account is well worth investigating since it would provide both a more constrained account of how metonymy works and a more unified account of pragmatic processes overall.

With this as my overall goal, the relevance-theoretic analysis of metonymy will be divided into three parts. I will start by discussing Papafragou's (1996) formulation of

a preliminary account of metonymy in terms of metarepresentational (interpretive) use of language, in section 6.3.2. Then, I will suggest two possible ways in which a relevance-theoretic account of metonymy may proceed: In section 6.3.3, I sketch an account of metonymic interpretations as being derived by means of the process of *ad hoc* concept construction. In section, 6.3.4, I suggest a different account that builds on Papafragou's preliminary work, in which metonymy is treated as a form of creative naming, closely related to 'nicknaming', where a salient property of an individual is used to create a new name. The two accounts are compared, and their respective advantages and shortcomings discussed. However, before I turn to the relevance-theoretic account, I will look briefly at two other influential lines of approach to metonymy. The first analyses metonymy as an instance of (pragmatic) function application and dates back to Nunberg (1978, 1979), the second includes the wide range of approaches that have followed in the wake of Lakoff and Johnson's (1980) original formulation of the conceptual metaphor theory.

6.2 Previous approaches to metonymy

6.2.1 Metonymy as pragmatic function application

In a widely influential approach, Nunberg (1978, 1979) analysed metonymy as a subvariety of 'deferred reference', that is, as an instance of a more general process according to which a speaker may point at, name or describe some thing *A* with the intention to refer to some thing *B*. Consider (6)-(9):

- (6) *He [pointing at a large footprint in the sand] must be a giant.*
- (7) Murdoch has owned *that [pointing at a copy of The Sun]* since 1969.
- (8) *The newspaper* sacked its editor-in-chief.
- (9) *The ham sandwich* is getting impatient.

The idea is that each of the italicised expressions in (6)-(9) is interpreted via a *referring function*, which is "that function that the hearer (correctly) selects from among an infinitely large number of functions that take the demonstratum as arguments" (Nunberg 1979: 156), for instance, 'A was made by B' in (6), 'A is published by B' in (7) and (8), 'A was ordered by B' in (9). The correct referring function is chosen by a hearer

on the basis of its ‘cue-validity’, described as “the probability with which a given referent *b* can be identified as being in the value of a certain function *f* at a demonstratum *a*”, in more general terms, “the relative usefulness of a given description for purposes of identification” (ibid 160). Other things being equal, the hearer can take a rational speaker to have intended him to select the function that has the highest cue-validity for its referent, that is, the one that allows him most easily to identify the referent in terms of its relation to the demonstratum (which may well be its encoded meaning, as in e.g. ‘*The ham sandwich is being eaten*’), in which case the correct function is the identity function. Which function has the highest cue-validity on a given occasion depends, of course, on a range of pragmatic factors. For instance, in (7) above, the high cue-validity of the function ‘A is published by B’ stems from encyclopaedic knowledge about newspapers; we know that there are companies that publish them and that these are owned by powerful individuals among whom Rupert Murdoch, the subject of the utterance, is one of the best known. In (9), uttered in the context of a restaurant, the function that maps food items into customers has a high cue-validity against the set of contextual assumptions shared by the waiters; the link between customers and their food orders provides an efficient means of referent identification.

The idea that metonymy is generated by the application of pragmatic functions also shows up in the early work of Fauconnier (1994 [1985], in whose framework, pragmatic functions are called ‘connectors’), and in Sag (1981), who outlines a formal semantic account geared at accommodating the context-dependency of metonymy in a strict compositional semantics. In Sag’s view, a formal semantic account is the only viable alternative to a Gricean treatment of metonymy as post-semantic inferences from absurdities, which, as he rightly observes, fails. We don’t seem to first compute the absurd ‘literal’ interpretation and then infer the metonymic meaning; rather, metonymy appears to be a local process, which does not presuppose the prior identification of a proposition serving as its input. If this is right, Sag argues, there would be no input for the Gricean inferential system to operate on the interpretation of metonymy.¹⁶² Sag’s

¹⁶² However, a more modern kind of Gricean account (e.g. Bach 1994; Levinson 2000) may grant that metonymy contributes to ‘what is said’ but still claim that the absurd strict-and-literal interpretation is entertained *en route* to the metonymic interpretation.

proposal is to introduce a sense transfer function into the logical language (L),¹⁶³ which allows for a common noun (e.g. *ham sandwich*) to translate into L as a predicate (e.g. as denoting ‘ham sandwich customer’). This preserves semantic compositionality; once the denotation of the transferred sense is fixed, the compositional semantic process proceeds as usual. In (7) above, for instance, it would be the expression *ham sandwich* that has a shifted denotation, while the determiner *the* plays its normal semantic role. Sag emphasises that by itself, this procedure is almost entirely unconstrained and must therefore be paired with a pragmatic theory that places (perhaps drastic) constraints on which sense transfers are possible.

In practice, then, Sag’s proposal leaves the interpretive work entirely up to pragmatics. For this reason, its ability to explain what is going on in metonymy beyond the application of sense transfer functions is limited (e.g. what triggers their application?, what are the contextual factors that license them?, where do they come from?, and so on). It seems then that this account could be viewed as an early version of a truth-conditional pragmatic approach but with all the crucial detail of how it could work missing.¹⁶⁴ However, the idea that metonymy involves a kind of sense transfer rather than deferred reference was significant and has been taken over into recent treatments by Nunberg (1996, 2004). Consider:

(10) I am *parked out back*.

(11) The *ham sandwich* is getting impatient.

Nunberg now treats examples such as (10) and (11) as instances of a more general process of *meaning transfer*, where “the name of a property that applies to something in one domain can sometimes be used as the name of a property in another domain, provided the two properties correspond in a certain way” (1996: 111).¹⁶⁵ On this

¹⁶³ Natural language sentences are translated into this logical language L , which captures their semantics.

¹⁶⁴ However, Sag himself did not consider the possibility of developing a ‘truth-conditional pragmatic’ account (cf. Recanati 1993). This is the approach favoured by contemporary pragmatic frameworks (Recanati 1993, 1995, 2004; Romero and Soria 2005, 2010), including relevance theory (Papafragou 1996; Wilson and Carston 2007; Carston 2010), as well as the approach taken here.

¹⁶⁵ This process is taken to apply to a whole range of ‘figurative’ uses of language, including metaphor, synaesthesia, synecdoche and metonymy. A consequence of this account is that it blurs the distinction between metaphor and metonymy, which are both derived according to the same process. As Nunberg (1996: 113) puts it, “metaphors and metonymy aren’t different types of transfer, they’re different conditions that can support a functional relation between properties”.

analysis, the subject of the sentence in (10) (which, according to Nunberg, we tend to think of as the target of the metonymy, that is, as denoting the car and not the speaker), is used in its conventional sense to refer to the speaker, while it is the predicate *parked out back* that has a transferred meaning, contributing a property of persons, “the property they possess in virtue of the locations of their cars” (ibid.). In support of this analysis, he provides the following examples:

(12) a. I am parked out back and have been waiting for 15 minutes.

b. ?I am parked out back and may not start.

(13) We are parked out back.

The examples in (12) are taken to show that a predicate conjoined to the sentence uttered in (10) may well describe the speaker, as in (12)a., but not so easily the car, as illustrated by (12)b., suggesting that the appropriate referent of *I* in these cases is the speaker, and not the car (and thus that the predicate *parked out back* has a transferred meaning).¹⁶⁶ Similarly, a speaker would not utter (13) in a case where she has several cars parked out back.

As concerns (11), Nunberg argues, contrary to his previous analysis, that the transfer applies to the noun *ham sandwich* and not to the entire NP (the referentially used definite description), and so “this is a case of meaning transfer, rather than reference transfer” (1996: 115). In support of this analysis, he provides examples such as the following,

(14) That (*those) french fries is (*are) getting impatient.

(15) The three ham sandwiches in the corner has (*have) asked for his bill.

In (14), where the expression *ham sandwich* has been replaced by the plural *french fries* (denoting ‘french fries customer’) the demonstrative pronoun agrees with the number

¹⁶⁶ A similar example (originally discussed by Jackendoff 1992a) in support of the predicate transfer analysis involves the use of a reflexive pronoun (Nunberg 2004: 357):

(i) Ringo squeezed himself into a tight space.

Instead of taking the subject to refer to a person and the reflexive to refer to a car (and thus facing the problem of explaining that the reflexive is not coreferential with the subject in this case), Nunberg treats them both as referring to a person, and the predicate *squeeze x into a tight space* as having a transferred meaning, “denoting the relation that people enter into in virtue of the maneuvers they perform with their cars” (ibid. 358).

of the transferred meaning and not with the number that is linguistically-encoded, suggesting that it is the expression *french fries* (and not the whole NP) that has a transferred meaning. Similarly, in (15), where ‘three ham sandwiches’ refers to a single male customer, the verb and the possessive agree with the transferred meaning and not with the linguistically-encoded plural of the NP.

Nunberg’s meaning transfer account is quite a radical move from his previous account of metonymy in terms of deferred reference. First, the transfer is now taken to operate on the linguistic meaning of the expression thereby changing its denotation rather than being a case of a direct function on denotations. Nunberg (1996: 115) points out that an implication of the early analysis of cases like (11) (as a subtype of the more general phenomenon of deferred reference) was that the actual ham sandwich figured in the interpretation process. This point is perhaps best illustrated by cases of deferred indexical reference, as in (7), repeated as (16) below:

(16) *That* [*pointing at a copy of The Sun*] is owned by Rupert Murdoch.

In this case, the object of ostension (i.e. the copy of *The Sun*) provides the input to the referring function which yields the interpretation of *that* as referring to the company that publishes *The Sun*. Similarly, on the analysis of (11) as an instance of deferred reference, it is the actual ham sandwich that provides the input to the referring function which returns the customer as its output. On the current analysis of (11), however, the transfer operates from the linguistic meaning of *ham sandwich* to the contextual meaning ‘ham sandwich customer’, which then allows for the NP *The ham sandwich* to pick out the individual who has the property of being the ham sandwich customer. Second, meaning transfer is seen as a ‘linguistic’ process (presumably specified within the language faculty) with pragmatic constraints on its application, rather than an entirely pragmatic process, as was previously assumed.

In my view, Nunberg is probably right in distinguishing between deferred reference (cf. (6) and (7) above), involving the use of a demonstrative or indexical, and ‘meaning transfers’ that target the linguistic meaning of expressions to yield either a transferred meaning for a predicate, as in (10), or for a common noun, as in (11). Although the deferred cases and the meaning transfers may ultimately turn out to be

manifestations of a single cognitive mechanism (operating in the first type of case on a representation of the object referred to by the demonstrative or indexical, and in the latter case on the representation linguistically-encoded by the common noun or predicate), my focus in this chapter will be on the cases that Nunberg treats as ‘meaning transfers’, these being the ones from which polysemy may arise. On the other hand, I do not think there is sufficient evidence to warrant Nunberg’s claim that meaning transfer is a *linguistic* process. Recall that examples such as (14) and (15), where the metonymies appear to render the number clashes felicitous, are taken to show that the transfer applies to the meaning of the common nouns rather than to the reference of the whole NPs, and that we have to do with a linguistic (albeit pragmatically constrained) process. However, as Ward (2004: 268) shows, this cannot be the whole story, as on Nunberg’s account, (17)a. should be felicitous if the meaning transfer applies to the plural noun *sunglasses*:

- (17) a. *I wouldn’t mind going out with that *sunglasses* at Table 7.
b. I wouldn’t mind going out with the *sunglasses* at Table 7.

In this case, contrary to Nunberg’s predictions, it appears to be the number clash between the grammatically plural noun and the singular demonstrative determiner that renders the metonymy infelicitous, and that the absence of such a clash renders (17)b. felicitous. Thus, as Ward points out, it may be that the felicity of examples such as (14) and (15) is more likely to depend on “specific principles of English morphosyntax”, (as, in fact, Nunberg himself concedes (1996: 115)), than on the process of meaning transfer.¹⁶⁷

However, regardless of whether metonymy is treated as a kind of deferred reference or a meaning transfer (or as a more general process of pragmatic mapping, as Ward (2004) claims), the challenge for this type of account is to explain how the process that associates objects and properties in metonymic interpretations is constrained. I will end this section with a brief consideration of Nunberg’s (1996, 2004) recent suggestion

¹⁶⁷ It should be noted that Ward’s (2004) main concern is to provide evidence against Nunberg’s claim that cases such as (11) involve a transfer of the meaning of the common noun rather than a deferred reference for the whole NP. I remain neutral on this issue here, but I will return to the question of whether metonymy involves actual ‘meaning transfer’ or just reference substitution in section 6.3.

of ‘noteworthiness’ as a pragmatic condition on meaning transfer. This criterion states that a “transfer is only possible when the property contributed by the new predicate is ‘noteworthy’, which is to say one that is useful for classifying or identifying its bearer relative to the conversational interests” (2004: 349; see Clark and Clark 1979 for a similar view). Although the criterion is intuitively correct (for instance, the property of being a ham sandwich orderer in (11) above is clearly noteworthy in the sense that it is useful for the identification of the referent; indeed, all instances of successful referential metonymies will satisfy the noteworthiness criterion in this way), it is not particularly explanatory. What makes a property ‘noteworthy’ in the first place, that is, what is required for a property to be useful for the identification of a referent? And why do we find it ‘useful’ to refer to someone or something via a ‘noteworthy’ property instead of via a more literal encoding? These are, in my view, the questions that a pragmatic account of metonymy has to answer. As to Nunberg’s noteworthiness criterion, it is, as it stands, too vague and needs to be developed further in order to provide a proper explanation of what constrains metonymic uses of expressions. I return to this issue in section 6.3. My claim will be that being ‘noteworthy’ (i.e. useful for the classification or identification of a referent) is only one of the ways in which a metonymic expression may achieve relevance.

6.2.2 Cognitive linguistic approaches to metonymy

Cognitive linguistic accounts of metonymy build on the assumption, implicit in the accounts discussed in the previous section, that cultural/experiential factors license associations between different entities, and that such associations form the basis for metonymy in natural language. However, where Nunberg, Fauconnier, Sag and Ward left open the question of what the cognitive basis for such associations might be (whether in the form of ‘referring functions’, ‘connectors’ ‘sense transfer functions’, ‘meaning transfers’ or ‘pragmatic mappings’), cognitive linguists make a precise claim about their source: metonymic associations are one of the principles by which our knowledge is structured (Lakoff and Johnson 1980; Lakoff 1987 and many others).

A central claim of mainstream cognitive linguistics is that tropes like metaphor and metonymy, rather than being ‘figures of speech’ (that is, purely linguistic/rhetorical devices) understood by contrast to ‘literal’ expressions, are ‘figures of thought’ (Gibbs

1994, 1999) which structure how we conceptualise the world around us (Lakoff and Johnson 1980). Lakoff and Johnson claim that our conceptual system is structured in terms of conceptual metaphors such as ARGUMENT IS WAR, LOVE IS A JOURNEY, HAPPY IS UP, and so on. In addition, they claim, the conceptual system contains metonymic concepts, as illustrated by (18)-(22):

(18) PART FOR WHOLE

We need some new *faces* in this organisation.

(19) PRODUCER FOR PRODUCT

Susan is reading another *Woolf*.

(20) OBJECT USED FOR USER

The *violin* is in a bad mood today.

(21) INSTITUTION FOR PEOPLE RESPONSIBLE

In 1773, the *Parliament* passed the Tea Act.

(22) PLACE FOR PEOPLE RESPONSIBLE

Oslo is hosting the 2011 Nordic World Ski Championships.

Thus, the use of metonymic expressions in language is seen as reflecting underlying conceptual metonymies. However, metaphor and metonymy are seen as different kinds of processes: While metaphor involves conceiving one thing in terms of another, and is analysed as a mapping from a source domain to a target domain (Lakoff and Johnson 1980, 1999; Lakoff 1987, 1993; Lakoff and Turner 1989), metonymy is the process whereby one entity is used to stand for another, and is seen as relating a vehicle and a target concept within a single domain or domain matrix (Lakoff and Turner 1989; Goossens 1990; Croft 1993; Taylor 1989/2003; Gibbs 1994; Kövecses and Radden 1998; Radden and Kövecses 1999). As concerns their communicative functions, referential metonymy is generally seen as providing an efficient (effort-saving) route to the intended referent (e.g. Langacker 1993), and metaphor as a means to “provide understanding” (cf. Chapter 5, Lakoff and Johnson 1980; Lakoff 1987).

An approach that has been foundational to much of the research on metonymy within the cognitive linguistics paradigm is Lakoff's (1987) theory of knowledge organisation in terms of *Idealised Cognitive Models* (ICMs). An ICM is viewed as a

relatively stable mental structure that represents an individual's knowledge and beliefs about a particular domain of experience, and guides categorisation and reasoning relative to that domain.¹⁶⁸ For instance, as discussed in Chapter 2, Lakoff (1987: 70) suggests that the concept BACHELOR is defined against an ICM "in which there is a human society with (typically monogamous) marriage and a typical marriageable age". This allows us to judge whether someone is a typical or less typical representative of the category BACHELOR (e.g. that a single heterosexual man in his 20's is a typical bachelor while the Pope isn't). Lakoff sees metonymy as an important source of such typicality effects. As an example, he discusses the stereotype of a housewife mother, which defines an (old-fashioned) cultural expectation about what a mother is supposed to be like. Lakoff sees social stereotypes of this sort as cases of metonymy; in this case the 'metonymic ICM' HOUSEWIFE-MOTHER is used to 'stand for' or represent the category MOTHER as a whole.¹⁶⁹ Importantly, this is all seen as taking place at the level of thought; the metonymic model exists independently of any natural language encoding.

Croft (1993) takes a somewhat different approach and suggests that metonymy functions by *highlighting* one domain within a domain matrix (cf. Langacker 1987). Following Langacker, Croft takes the concept encoded by a word to be structured in terms of multiple domains, called the domain matrix of that concept. Consider (23):

- (23) a. *Woolf* committed suicide at the age of 59.
 b. *Woolf* is difficult to read.

According to Croft, the concept WOOLF is defined against a domain matrix containing, *inter alia*, the domain of persons and the domain of creative activity. The 'primary domain' of a concept is that which defines its most central aspects; in the case of WOOLF this would arguably be the domain of persons. This accounts for the 'literal' use of *Woolf* in (23)a. In the metonymic use of *Woolf* in (23)b., Croft sees the shift of reference from the person to her work to be paralleled by a shift of domain within the domain matrix

¹⁶⁸ An ICM is idealised in the sense that it is an abstraction over a range of experiences, and does not include knowledge about all possible real-world situations pertaining to its domain.

¹⁶⁹ Thus, Lakoff is working with a much broader and quite different notion of metonymy than the other accounts discussed to far. According to the definition of metonymy provided at the beginning of this chapter, the housewife stereotype would not count as an instance of metonymy, as it does not seem to involve an associative relation between disjoint sets (rather, it denotes a subset of the denotation of the concept MOTHER, i.e. it could be seen as an instance of narrowing).

for the concept, that is, from the domain of persons to the domain of creative activity. This, Croft argues, can be seen as a kind of ‘domain highlighting’, “since the metonymy makes primary a domain that is secondary in the literal meaning” (1993: 179).¹⁷⁰

Radden and Kövecses (1999), following Langacker’s (1993, 1999) claim about metonymy as a ‘reference point phenomenon’ (i.e. as a product of our “ability to invoke the conception of one entity for purposes of establishing mental contact with another” (1993: 5)), see metonymy as a mechanism for providing mental access to a desired target via another conceptual entity. They combine this view with Lakoff’s notion of ICMs in the following definition of metonymy:

Metonymy is a cognitive process in which one conceptual entity, the vehicle, provides mental access to another conceptual entity, the target, within the same idealized cognitive model. (Radden and Kövecses 1999: 21)

On this view, a metonymy provides a route of access to a particular target within a single domain. For instance, within the ‘Production ICM’, there is a metonymic relationship PRODUCER FOR PRODUCT (of which the metonymy ARTIST FOR OEUVRE is a subtype), which allows for the vehicle *Woolf* to provide access to the literary works of Virginia Woolf in (23)b.

The work of Radden and Kövecses is descriptively interesting. They provide a ‘typology’ of metonymy producing relationships (1999: 29-44) where these are described as occurring within specific ICMs (e.g. within the ‘Location ICM’ we find metonymies such as PLACE FOR INHABITANTS, PLACE FOR INSTITUTION, etc.). Furthermore, they describe a number of cognitive salience principles (e.g. HUMAN OVER NON-HUMAN, SUBJECTIVE OVER OBJECTIVE, etc.) that they take to govern the choice of metonymic vehicle (i.e. the particular access route), and which lead to ‘natural’ or ‘default’ cases of metonymy. In general, they suggest, the more such cognitive principles apply, the greater the cognitive motivation of a metonymy. Although producing a typology of various metonymy-producing relationships is clearly useful and an interesting goal in

¹⁷⁰ It is unclear to me if Croft means by this that proper names such as *Woolf* have all this information as part of their literal meaning. As a word, *Woolf* is just a name that can be used for anything or anyone. Croft seems to assume that in this particular case, the name gives immediate access to the individual concept in our minds of the author to whom this proper name is linked, which seems right, but he needs to explain how we make this link from among the many other possible individuals that this name may be linked to (e.g. the man Leonard Woolf, my pet spaniel Woolf, etc.).

itself, it provides perhaps more by way of description than of explanation of what is going on in metonymy comprehension. As to the cognitive salience principles, they may well be able to explain why we choose one metonymic vehicle over another, but they cannot be taken to provide a motivation for speakers' *use* of metonymy, as in most cases of metonymic use, there is an obvious and readily available paraphrase (e.g. 'The novels of Woolf', 'The ham sandwich orderer', etc.). I return to this issue in section 6.3, where I argue that an important factor motivating metonymic use is the reduction of processing effort it may involve.

There is no doubt that cognitive linguistic approaches to metonymy provide illuminating and insightful analyses of a range of natural language metonymies, and shed light on the role of association in human conceptualisation. Indeed, some of the cognitive linguistic perspectives on metonymy may prove complementary to a relevance-theoretic analysis. Recently, there have been some suggestions of how the cognitive linguistic and relevance-theoretic approaches to metaphor may be combined, despite some fundamental differences between the two theoretical frameworks (Gibbs and Tendahl 2006; Tendahl and Gibbs 2008; Wilson 2009). In section 6.3, I will suggest that combining them may also prove fruitful in an account of metonymy, in particular with regard to the more 'regular' cases.

However, as I see it, the cognitive linguistic approaches as described above equally have some shortcomings. The first concerns the fact that the distinction between metonymy and metaphor in most accounts rests on the notion of a cognitive domain or an 'ICM' (where, as we have seen, metonymy occurs within a single domain or ICM and metaphor between distinct domains), neither of which is well defined. How one decides to delineate a domain or ICM would therefore be crucial for the classification of a use as a metonymy or a metaphor.¹⁷¹ Second, as Papafragou (1996: 176) notes, the analysis of metonymy in terms of pre-existing cognitive structures disregards its connection to the outside world, and the fact that most 'creative' cases of metonymy involve a novel conceptualisation of an external entity, rather than an activation of a pre-existing

¹⁷¹ Also within the cognitive linguistics paradigm, the definition of metonymy as intra-domain mappings is subject to criticism. See for instance Haser (2005: 36ff.) and Peirsman and Geeraerts (2006) for discussions of the problem and suggestions for a solution to it, and Croft (2006) for a defence of the domain mapping account.

metonymic concept. Third, cognitive linguistic accounts generally pay little attention to the pragmatic aspects of metonymy, that is, why we use it and how metonymic interpretations are derived in verbal communication. As the reader will know, these are fundamental issues in the relevance-theoretic account, to which I will turn shortly. Before I do this, however, I would like to look in some detail at a recent approach to metonymy from the point of view of cognitive linguistics, proposed by Evans (2009), who addresses the issue of how metonymy (and tropes more generally) may be understood in language.

Evans's approach, termed the Lexical Concepts and Cognitive Models theory (the LCCM theory), stands out among the cognitive linguistic approaches by virtue of (i) its distinction between linguistically-encoded content in the form of 'lexical concepts', and encyclopaedic knowledge in the form of 'cognitive model profiles' to which lexical concepts afford access, and (ii) its proposal for an account of the semantic composition process, or 'linguistically mediated meaning construction', that is, of the selection of contextually appropriate lexical concepts and the integration of these in an overall interpretation of an utterance. In LCCM theory, linguistic forms encode lexical concepts, the contents of which are highly schematic (abstracted from across usage events). These provide access to a range of cognitive models, or 'cognitive model profiles', described as 'coherent bodies of multimodal knowledge of any kind' (cf. Lakoff's ICMs, Langacker's and Croft's domain matrices). The cognitive model profile of a lexical concept is seen as its 'semantic potential', and consists of a set of 'primary' and 'secondary cognitive models'. The primary cognitive models for a lexical concept are those with which it is directly associated; that is, the association areas that constitute its 'access site', while the secondary cognitive models are indirectly associated via connections with the primary ones. Although these are not directly associated with the lexical concept, they form part of its semantic potential by being linked to its access site (ibid. 207). As an illustration, consider the linguistic form *France* (ibid. 76-81). In LCCM theory *France* encodes the lexical concept [FRANCE]. The primary cognitive model profile for this lexical concept consists of (at the very least) the cognitive models GEOGRAPHICAL LANDMASS, NATION STATE, and HOLIDAY DESTINATION, relating to knowledge about the specific geographical region known as France (as reflected in, e.g., an utterance of '*France* is hexagonal'), the nation state France (e.g. '*France* is an

important member of the EU'), and about France as a holiday destination (e.g. 'We're doing *France* this year'). These cognitive models provide the *access site* for the lexical concept [FRANCE]. In addition, each of them provides access to further cognitive models, constituting the secondary cognitive model profile for the lexical concept, as a result of chaining within the conceptual system. These secondary models may include NATIONAL SPORTS, POLITICAL SYSTEM, CUISINE (all of which are chained to the primary cognitive model NATION STATE), and which may themselves be chained to further secondary cognitive models (e.g. POLITICAL SYSTEM giving access to ELECTORATE, HEAD OF STATE, etc.), as well as many others (see Evans 2009: 207-209 for more detail).

The focus of the LCCM theory account of figurative language (metonymy and metaphor) differs from that of most cognitive linguistic approaches; rather than being concerned solely with the cognitive representations underlying metaphors and metonymies in natural language, it aims to provide a theory of the understanding of figurative language. Like relevance theorists, Evans assumes that literal and figurative meaning arises from a single process of meaning construction. What distinguishes a literal and a figurative use of an expression/utterance "relates to that part of the semantic potential which is activated during the process of interpretation during the construction of a conception [i.e. a meaning]" (Evans 2009: 285). Consider (24):

- (24) a. *France* is a country of 26 regions.
b. *France* said it would send two planes with humanitarian aid to Haiti.

The utterance in (24)a. contains a literal use of the expression *France*, which, given the cognitive domain profile for the lexical concept [FRANCE] as outlined above, could be analysed in LCCM theory as activating a cognitive model NATION STATE within the access site of the lexical concept. In (24)b., where the expression *France* is used metonymically to stand for the French government, the lexical concept [FRANCE] would activate a cognitive model in its secondary cognitive model profile HEAD OF STATE (the access route would thus proceed from NATION STATE via POLITICAL SYSTEM to HEAD OF STATE). The distinction between literal and figurative meaning thus amounts to whether a lexical concept activates a cognitive model in the primary cognitive model profile (which leads to a 'literal' interpretation) or in the secondary cognitive model profile

(which leads to a ‘figurative’ interpretation). As a rule, the greater the length of the access route, the more figurative the use of an expression is likely to feel.¹⁷² The utterance comprehension process proceeds in the following (simplified) way: First, the process of ‘lexical concept selection’ ensures that the most appropriate lexical concepts associated with each linguistic item in the utterance are selected, on the basis of (linguistic and extra-linguistic) context (i.e. ambiguity resolution). Then, the process of ‘fusion’ (which involves two compositional processes working together) enables (a) ‘lexical concept integration’, that is, the construction of larger linguistic entities, driven by lexical knowledge, and (b) ‘interpretation’, that is, a matching between the cognitive model profiles of the lexical concepts which have undergone lexical concept integration, which is constrained by several principles. Among the range of cognitive principles that the LCCM theory proposes as governing the interpretation phase of language understanding (Evans 2009: Ch. 13) is the ‘Principle of Conceptual Coherence’, which requires that a non-match between the cognitive model profiles of lexical concepts undergoing interpretation is avoided, and the ‘Principle of Ordered Search’, which ensures that primary cognitive models are the first to undergo matching, and that, if a search in secondary cognitive models is necessary, it proceeds in a coherent way (i.e. in order of ‘distance’ from the point of lexical access). In (24)a., the literal interpretation of *France* results from a match between a cognitive model in the primary cognitive model profile activated by the lexical concept [FRANCE] and the meaning that has been assigned to the predicate ‘a country of 26 regions’ through integration and interpretation. However, the metonymic use in (24) results from a clash between all the cognitive models in the primary cognitive model profile for [FRANCE] and the meaning assigned to its predicate ‘said it would send two planes with humanitarian aid to Haiti’. Due to the Principles of Conceptual Coherence and Ordered Search, this triggers a search for a match in the secondary cognitive model profile for the lexical concept, where a match is

¹⁷² It is not entirely clear to me how this distinction between primary and secondary cognitive models is made: What makes, for instance, the cognitive models HOLIDAY DESTINATION and GEOGRAPHICAL LANDMASS more ‘primary’ than e.g. CUISINE? Which models are primary and which are secondary are likely to vary between individuals and between groups of language users, with the possible consequence that people will differ in their judgements of whether an utterance is an instance of ‘figurative’ use or not (given that the activation of a cognitive model or models within the primary cognitive models profile of a lexical concept is claimed to give rise to ‘literal’ interpretations and the activation of a cognitive model within the secondary cognitive model profile to ‘figurative’ interpretations).

found with the HEAD OF STATE cognitive model. *France* thus has a figurative meaning in (24)b. The choice of lexical concept selected for clash resolution is, of course, contextually determined, as stated by the Principle of Context-Induced Clash Resolution (ibid. 293):

(25) Principle of Context-induced Clash Resolution

In cases where clash resolution is required, the lexical concept whose secondary cognitive model profile is searched to resolve the clash is determined by context. This is achieved by establishing a figurative target and a figurative vehicle, on the basis of context. The lexical concept that is established as the figurative vehicle is subject to clash resolution.

In (24)b., the utterance context, i.e. the immediate aftermath of the Haiti earthquake, as well as the linguistic context, in particular, the use of the verb *say*, select the lexical concept encoded by *France* as the figurative vehicle, while the figurative target, ‘the political leaders of France’ is found in the secondary cognitive model profile for this concept.

Evans (2009: Ch. 14) analyses metaphor in the same way – as an instance of context-induced clash resolution activating cognitive model(s) in the secondary cognitive model profile of a lexical concept. On the LCCM theory account, then, metaphor and metonymy are simply different outcomes of a single process of meaning construction. The distinction between metaphor and metonymy relates to whether or not the figurative target and the figurative vehicle exhibit ‘alignment’: In metonymy, as exemplified by (24)b., the figurative target and the figurative vehicle are accessed via the cognitive model profile associated with a single lexical concept ([FRANCE]); there is thus an ‘alignment’ of the target and the vehicle in a single cognitive model profile, and the clash resolution site corresponds to the access route for the figurative target. In metaphor, however, at least those of the form ‘X is a Y’ (which are the only ones discussed by Evans), the situation is different: Here, Evans takes X to correspond to the figurative target and Y to the figurative vehicle, which is also the clash resolution site. He discusses the example ‘My boss is a pussycat’ (‘docile’). On the LCCM theory approach, the figurative target is the lexical concept [BOSS], while the figurative vehicle is

the lexical concept [PUSSYCAT], and the clash resolution takes place in the secondary cognitive model profile of this concept. However, referential metaphors such as ‘The wilting violet has finally gone home’ appear to present a problem for this way of distinguishing metonymy and metaphor, as in Evans’s terminology the figurative vehicle (‘the wilting violet’) and the target (the person of whom this property is predicated), would exhibit alignment in such cases; thus, they should be analysed as metonymies.¹⁷³

The LCCM theory represents, in my view, a new and important contribution to the cognitive linguistics enterprise by addressing the interaction between linguistic knowledge, encyclopaedic knowledge and utterance context in the processing of utterances. In particular, within this paradigm, it provides a new perspective on the nature of linguistic knowledge, by positing a principled distinction between the semantic information encoded by the language and non-linguistic encyclopaedic information, a perspective which has a clear affinity with the relevance-theoretic distinction between linguistically-encoded content and non-linguistic encyclopaedic information stored in long-term memory. As already mentioned, the LCCM theory also stands out by providing an account of figurative language *understanding* (and not just of the underlying conceptual structures which are held to provide the basis for figurative language), in which encyclopaedic knowledge (or cognitive models in LCCM theory terms) associated with encoded concepts play a central role. As such, I think LCCM theory provides a both valuable and insightful approach, which is, in many ways, consonant with the pragmatic account of lexical adjustment that I have been pursuing in the previous chapters. However, focusing on its account of figurative language understanding, the LCCM theory raises some problematic questions that might be taken into account in the further development of the theory. The first question concerns the

¹⁷³ Evans follows the common practice of taking the target or topic of metaphor/metonymy to be the literal element to which the figurative predicate applies, so in the metaphor ‘My boss is a pussycat’, being a PUSSYCAT* applies to the boss, and in a metonymy such as ‘The ham sandwich wants his sandwich now’, the property of being [in such and such a relation with a ham sandwich] applies to the customer. A (somewhat confusing) consequence of this view is that the notion of target needs to be distinguished from the notion of *communicated* meaning: although in the case of metonymy the target and the meaning communicated by the metonymically used expression are the same, in metaphor the target is another constituent of the sentence, thus distinct from the communicated meaning of the metaphorically used expression (e.g. PUSSYCAT*).

assumption that figurative interpretations arise as a result of a lack of conceptual coherence (a clash) between two or more primary cognitive models. Barring the issue of what it means in more precise terms for two or more cognitive models to be semantically (in)coherent, it is unclear to me how figurative expressions that do not appear to involve a clash in cognitive models would be explained on the LCCM theory, as illustrated by the metaphors below:

- (26) a. Mary is no *angel*.
b. Not all lawyers are *sharks*.

Presumably, the utterances in (26) would be interpreted as ‘literal’ on the LCCM theory account, activating cognitive models in the primary cognitive model profiles of their constituents, as they cannot be said to involve semantic incoherence of any kind. But such an analysis would miss the intended meanings of the utterances; that Mary isn’t a person of exemplary conduct, and that not all lawyers are vicious, aggressive and tenacious.¹⁷⁴ Similarly, utterances containing so-called regular metonymies are likely to be analysed as ‘literal’ on the LCCM theory account:

- (27) a. *Woolf* is difficult to read.
b. The local *church* is liberal on same-sex marriage.

In (27)a., it seems reasonable to assume that a cognitive model along the lines of AUTHOR would figure in the primary cognitive model profile for the lexical concept [WOOLF], that is, as being directly associated with it,¹⁷⁵ and that a cognitive model, say, INSTITUTION, is directly associated with the lexical concept [CHURCH] in (27)b. In view of my previous comments regarding such cases (cf. Chapter 5), treating them as literal may well be the correct analysis, but it is one that does not capture the intuition that there is (somehow) something less literal about instances like (27), and clearly literal utterances such as ‘Woolf is a famous Modernist writer’, ‘The local *church* is two blocks away’.

¹⁷⁴ In fact, this problem is similar to the well-known critique of the Gricean account of figurative language in terms of violations of the maxim of Quality (truthfulness); metaphorical utterances such as those in (26) are trivially true and involve no violation of any maxim of truthfulness (Wilson and Sperber 1981).

¹⁷⁵ This is, of course, hard to determine as long as we do not seem to have any precise criteria for distinguishing between primary and secondary cognitive models (cf. footnote 172).

The second question concerns the cognitive motivation for metonymy and metaphor in LCCM theory. As we have seen, previous cognitive linguistic accounts have claimed that metonymy and metaphor in language reflect underlying conceptual patterns and that their function is mainly to provide understanding (Lakoff and Johnson 1980; Lakoff 1987). Of metonymy, it has been claimed that it serves the function of providing an efficient (effort-saving) route to the intended referent (Langacker 1993; Kövecses and Radden 1998; Radden and Kövecses 1999). In LCCM theory, however, both metonymy and metaphor seem to imply an increase in the effort of processing compared to 'literal' utterances. They arise by virtue of a clash between primary cognitive models and result in the activation of secondary cognitive models associated with the lexical concepts encoded by the figuratively used expressions; thus, their interpretation involves a longer 'access route' than 'literal' interpretations do. It is, of course, correct that many poetic metaphors often demand an extra effort of processing on the part of the hearer, but it is not entirely clear to me what would be gained by inducing a longer access route, as claimed in LCCM theory, compared with using a literal expression which activates the intended domain directly (i.e. whose primary cognitive model profile includes the cognitive model intended for activation). This point is perhaps clearest in the cases of referential metonymies of the 'ham sandwich' kind, whose use is usually seen as motivated by the quick and easy identification of a referent in a given utterance situation that they enable. If, in fact, these involve an extra effort of processing on the part of the hearer, it is at least not obvious what motivates such referential uses in LCCM theory.

The third question concerns the fact that figurative language use is entirely a matter of association in LCCM theory. As we have seen, relevance theorists (Carston 1997, 2002b; Vega-Moreno 2004, 2007; Wilson and Carston 2006, 2007; Sperber and Wilson 2008) have provided convincing evidence that metaphorical interpretations are largely derived via inference, and often depend on a process of mutual adjustment where interpretive hypotheses about explicit and implicit content are computed in parallel. In the next section, I argue that metonymic interpretations also arise within, and are dependent on, this process of inferring intended speaker meanings.

6.3 A relevance-theoretic approach

As we have seen, relevance theorists claim that utterance comprehension, including the process of *ad hoc* concept construction, is inherently inferential, in the sense that it depends on a spontaneous inference process that takes as its premises the encoded linguistic meaning (embedded in ‘The speaker has said that...’) together with highly accessible contextual assumptions, and yields conclusions in the form of hypotheses about speaker-intended meanings, a process which is guided and constrained by (occasion-specific) expectations of relevance. Most accounts of metonymy (indeed, all of the accounts discussed so far), however, are *associative*, that is, they take metonymic senses to arise as a result of the activation of associative links between the object or property linguistically denoted and the speaker-intended object or property. As evidenced by the approaches discussed in section 6.2 above, metonymy lends itself easily to an analysis in terms of association. That this is so is hardly surprising; any two concepts could in principle be associated with one another provided that the context licenses the association. The challenge for associative accounts is to explain not just *how* a given association occurs (which, as we have seen, can easily be done by reference to pragmatic function application, domain mapping/highlighting/activation), but *why* it is that a particular association is selected in a given context and not any of the other possible associations that are equally licensed by that same context. Thus, they have to be paired with pragmatic constraints that are able to drastically reduce the number of possible associations on a given occasion.

6.3.1 Lexical pragmatic processes: associative or inferential?

Relevance theorists have criticised associative accounts of lexical pragmatic processes for being too unconstrained (Wilson and Carston 2007: 252-253). The view is that an inferential account, if empirically adequate, would be considerably more constrained and therefore preferable on theoretical grounds.¹⁷⁶ However, there is much that indicates that association plays an important role in many instances of metonymy, and

¹⁷⁶ The argument goes as follows: While all inferential relations are also associations between (constituents of) premises and (constituents of) conclusions, not all associations are inferential (that is, evidentially grounded or warranted). If only those associations that are inferential are exploited in lexical adjustment, this would considerably constrain which adjustments are possible on a given occasion.

relevance theorists have expressed doubts about the possibility of analysing metonymy in entirely inferential terms (Wilson and Carston 2007; Carston 2010).

The question whether metonymic interpretations are associatively or inferentially derived is part of a broader debate in pragmatics concerning the nature of the pragmatic processes that contribute to the explicit content of an utterance (what is said, the proposition expressed): are they purely associative (e.g. Recanati 1993, 1995, 2004) or inferential (e.g. implemented by a dedicated pragmatic mechanism that derives the explicit and implicit content of utterances in parallel, by forwards and backwards inference, resulting in a logical argument, as maintained by relevance-theorists)? For instance, Recanati (1993, 2004) makes a distinction between *primary* and *secondary* pragmatic processes. Primary pragmatic processes are responsible for computing ‘what is said’ (in Recanati’s extended sense of the term) by an utterance, and include ‘bottom-up’, linguistically mandated processes such as saturation (indexical resolution), as well as ‘top down’, optional processes such as free enrichment, loosening and semantic transfer (in relevance theory terms: free enrichment and *ad hoc* concept construction, where ‘semantic transfer’ is Recanati’s term for metonymy). These pragmatic processes are ‘pre-propositional’ (i.e. do not involve the prior computation of a proposition serving as input), unconscious (in the sense that normal “interpreters need not be aware of the context-independent meanings of the expressions used” (2004: 23), nor of making the associative links), and inherently associative. Recanati writes (2002a: 113-114), “*primary* pragmatic processes (...) need not involve an inference from premises concerning what the speaker can possibly intend by his utterance. *Indeed, they need not involve any inference at all*” (emphasis in original). Secondary pragmatic processes, however, responsible for the derivation of implicatures, are properly inferential in the Gricean sense of the term; they are reflective and ‘post-propositional’, that is, “they cannot take place unless some proposition *p* is considered as having been expressed” (2004: 23).¹⁷⁷ Secondary pragmatic processes satisfy the ‘availability condition’,

¹⁷⁷ Millikan (1984, 2005), however, maintains an entirely non-inferential theory that treats language understanding as a form of perception, including both ‘primary’ and ‘secondary’ pragmatic processes.

according to which interpreters must have conscious access to what is said, what is implied, and to the inferential relation between them.¹⁷⁸

The existence of ‘mutual adjustment’ is a central relevance-theoretic objection against Recanati’s distinction between primary and secondary pragmatic processes (Carston 2002a, 2007), as described above. Carston writes:

[H]ow can these two quite distinct types of process (local, associative, unconscious, subpersonal, in the one case; global, inferential, consciously accessible, person-level in the other case) interact in such a way as to effect adjustments to each other’s content? Specifically, how can implicatures, whose derivation takes ‘what is said’ as its input, affect the content of ‘what is said’? (emphasis in original) (Carston 2007: 24-25).¹⁷⁹

The metonymy phenomenon is particularly interesting in the light of this controversy between Recanati and the relevance theorists. As we have seen, it is widely agreed that metonymy contributes to the explicit content of an utterance. Thus, Recanati counts it among his primary pragmatic processes. In relevance theory, the prediction is that metonymic meanings are derived as part of the overall process of inferring speaker meanings on the basis of linguistic input. However, while relevance theorists have successfully shown that most of the pragmatic processes that Recanati counts as ‘primary’ can be described in inferential terms, and involve a process of mutual adjustment of explicit and implicit content (including narrowing, broadening, ‘free’ enrichment, as well as ‘linguistically mandated’ processes such as disambiguation and indexical reference resolution) (Sperber and Wilson 1986/1995; Carston 2002b), it is not clear, as noted above, how far metonymy can be analysed in inferential terms, and to what extent it may depend on a process of mutual adjustment of explicit and implicit content.

¹⁷⁸ This does not mean, Recanati (2004: 50) argues, that the inferential process itself must be conscious; it is sufficient that “the subject herself has the reflective capacities for making the inference explicit”, i.e. that she is, on reflection, capable of rationally justifying her interpretation.

¹⁷⁹ In response to this objection, Recanati (2004: Chapter 3; 2007) maintains that mutual adjustment of explicit and implicit content is compatible with his approach, because “in the actual mental process of the interpreter, there is no step-by-step, linear reasoning leading from the premises to the conclusion. Still (...) the inferential link is grasped by the interpreter, who can, on reflection, make it explicit” (2004: 50). Thus, the priority of explicit content is seen as a *logical*, and not a *temporal* matter. Carston (2007: 25) argues that this leaves Recanati without an account of the on-line derivation of conversational implicatures, since his secondary pragmatic processes turn out not to account for how they are derived in actual communication but rather for how interpreters may rationally reconstruct the interpretive process that led to the implicature derivation.

In the remainder of this chapter, I will discuss the possibility of developing an account of metonymy that accommodates it within a theory of verbal comprehension which takes speaker meanings to be derived inferentially on the basis of encoded linguistic meaning, contextual assumptions and considerations of relevance. As mentioned in the introduction to this chapter, I will approach this issue from two different angles: First, I will investigate whether metonymy can be brought within the scope of the relevance-theoretic account of lexical pragmatic processes (cf. Chapter 4), which takes occasion-specific lexical meanings to be understood in terms of *ad hoc* concept construction. Second, I will suggest how an early relevance-theoretic analysis of metonymy, proposed by Anna Papafragou (1996), which takes metonymic interpretations to arise from a creative process of ‘naming’, can be carried over to a more up-to-date relevance-theoretic account. As will become clear, both options have their advantages and, unfortunately, some shortcomings, which will be discussed. I will start, however, with a discussion of Papafragou’s early analysis.

6.3.2 Papafragou (1996)

Papafragou (1996) sketched an early relevance-theoretic account, where she analysed metonymy as a variety of metarepresentational use of language (cf. Noh 2000; Wilson 2000) involving the introduction of a new name for a referent.¹⁸⁰ She saw metonymy primarily as a way of economising on the processing effort required for assigning reference, but also, secondarily, as a means of achieving cognitive effects that would not be communicated by any literal paraphrase (e.g. by allowing the speaker to express a variety of attitudes toward the metarepresented content). For instance, economy of

¹⁸⁰ It should be noted that Papafragou (1996) talks about metonymy in terms of *interpretive* use, as opposed to *descriptive* use of language (cf. Sperber and Wilson 1986/1995), rather than in terms of *metarepresentation*. According to Sperber and Wilson’s definition (1986/1995: 224ff.), an utterance is used interpretively when its propositional form represents some other representation with a propositional form “in virtue of a resemblance between the two propositional forms”, where ‘interpretive resemblance’ of propositional forms is defined as a sharing of logical and contextual implications. However, given the more recent development within relevance theory of the notion of metarepresentation as a category subsuming interpretive uses (cf. Noh 2000; Wilson 2000), and the fact that although Papafragou’s examples of metonymy clearly involve linguistic metarepresentation of some kind, it is not obvious that they are in fact cases of interpretive use (i.e. that they involve propositional forms sharing a set of implications). It therefore seems more adequate to speak of metarepresentational use of language in connection with Papafragou’s discussion of metonymy, rather than of interpretive use. Thus, the reader should be aware that the talk of ‘metarepresentation’ or ‘metarepresentational use’ in this section replaces the use of the notion ‘interpretive use’ in Papafragou’s paper.

effort provides the motivation for a use of the expression *The ham sandwich* to refer to the ham sandwich customer in the context of a busy café, while achieving cognitive effects would be the main purpose of the metonymic expression in an utterance of ‘Watch out, *the smelly breath* is approaching’, when from the reference point of view, use of the person’s name (known to both speaker and hearer) would be the most efficient linguistic device to use. In this section, I will discuss the merits of this account, as well as some of the problems it raises.

As we saw in Chapter 3, metarepresentation (representation of another representation) involves a higher-order representation, usually an utterance or thought, with a lower-order representation, either in the form of a *public representation* (e.g. utterances), a *mental representation* (e.g. thoughts), or an *abstract representation* (e.g. sentences, propositions), embedded in it (see Wilson 2000 for a range of examples). (I will leave aside the abstract representations, since they are not central to my purposes here.) As an illustration of how public and mental representations can be metarepresented, consider (28):

- (28) a. Jane enjoyed the dinner party.
b. Jane told me that she enjoyed the dinner party.
c. ‘I’ll enjoy this dinner party’, Jane thought.
d. Jane enjoyed the ‘dinner party’.

In (28)a., the expression *the dinner party* is used literally as part of a proposition the speaker holds to be true about a state of affairs in the world; this is what Sperber and Wilson refer to as *descriptive use* of language (1986/1995: 224ff.). The remaining examples contain metarepresentational uses of language: In (28)b., a public representation is being metarepresented; the expression *she enjoyed the dinner party* contains an explicit report of an utterance of Jane’s. This type of metarepresentational use, which involves quotation, achieves its relevance primarily by informing the hearer about the content of the original. (28)c. contains a metarepresentation of a mental representation, i.e. of Jane’s thought ‘I’ll enjoy this dinner party’, while in (28)d. the NP *the dinner party* is implicitly attributed to Jane (as reporting an utterance or thought of hers); the metarepresentational use enables the speaker to simultaneously express an

attitude of dissociation toward the descriptive content of the NP (perhaps the event was, in the speaker's opinion, so meagre that it could hardly be called a dinner party).¹⁸¹ This type of use is termed *echoic*; it “achieve[s] relevance mainly by conveying the speaker's attitude to an attributed utterance or thought” (Wilson 2000: 432).¹⁸²

A further type of metarepresentational use of language involves the *naming* of an individual or object, either with the purpose of coining a new name for something, as in (29) below, or of informing someone, who doesn't already know it, of an established name or label for something, as in (30) and (31):

(29) I name thee ‘Tabitha Hermione Lovell-Smith’.

(30) That is called an ‘okapi’. (Compare with the descriptive use ‘An okapi ran past’.)

(31) ‘Motor neurons’ form synapses with the muscles and command movements.

These examples all involve ‘mention’ of referring expressions (as opposed to ‘use’, cf. Lyons 1977a). They are instances of linguistic metarepresentation but unlike the speech/thought reports in (28)b. and (28)c., and the echoic use in (28)d., where a content is attributed to someone else, they are non-attributive, in the sense that the recognition of the metarepresentational use as such does not involve the identification of the source of the metarepresented content (indeed, it may have no identifiable source).¹⁸³ While the naming in (29) involves metarepresentation of form only (as proper names arguably have no semantic content), there is a sense in which the semantic content of the metarepresented expressions in (30) and (31) plays a role in the naming process (e.g. x is called an *okapi* because it belongs to the category of okapis), although they too, of course, involve metarepresentation of linguistic *forms*. This is, if I

¹⁸¹ A more fundamental kind of interpretive relation is that between utterances and thoughts: Sperber and Wilson (1986/1995: 230) claim that “every utterance is used to represent a thought of the speaker”. By this they mean that the propositional form of an utterance resembles the propositional form of the thought it is used to communicate.

¹⁸² The relevance-theoretic view is that all varieties of metarepresentation, whether they involve public, mental or abstract representations, can be analysed in terms of a notion of *representation by resemblance* (Wilson 2000: 425). The resemblance relation may be of any type: perceptual, linguistic, logical, mathematical, conceptual, sociolinguistic, stylistic, typographic, etc. Thus, the attributive cases in (28)b., (28)c. and (28)d. all involve conceptual resemblance, that is, resemblance of content.

¹⁸³ Unlike the attributive cases, which require of the hearer the ability to attribute an utterance or thought to someone else (thus depend on his theory of mind capacity, cf. Chapter 3), the understanding of the cases of mention in (29)-(31) are not obviously linked to this intuitive metapsychological ability (see discussion in Wilson 2000: 437).

understand her correctly, why Papafragou (1996: 180) takes cases like (30) and (31) to communicate something like ‘individual that can be appropriately called *okapi*’, and ‘things that can be appropriately called *motor neurons*’; the entities in question can be ‘appropriately called’ such and such because of what they are and because of what the names introduced for them *mean*. The form of metarepresentation involved in such cases of ‘mention’ allows the speaker to focus the hearer’s attention on the word itself (instead of on the representation of the actual thing it encodes).

Another form of naming may occur when an expression is used to pick out an entity that falls outside of its linguistically specified denotation. This happens in spontaneous dubbings (cf. Kaplan 1989 [1977]):

(32) Hi, ‘gorgeous’!

(33) What can I do for you, ‘love’?

As with (30) and (31), the semantic content of the naming expressions in (32) and (33) plays a role; for instance, the use of ‘gorgeous’ in (32) can be unpacked as ‘you who can appropriately be called *gorgeous*’ (and so the speaker is, in a sense, attributing to himself the view that the addressee possesses the property denoted by the name and thus she can be appropriately called by it). However, unlike the examples in (30) and (31), the dubbings in (32) and (33) are referential, in that they serve to pick out the individuals who can appropriately be called ‘gorgeous’ and ‘love’.¹⁸⁴

Papafragou’s (1996: 181) claim is that creative metonymy is a variety of metarepresentational use of language that shares its central characteristics with the cases in (29)-(33) above: It involves the use of an expression to introduce a new name, as in (29)-(31), and the intended referent of this expression does not fall under its conventional denotation, as in (32) and (33). Consider the following examples, discussed by Papafragou:

(34) Where is *the brain* now that we need him?

(35) *The violin* is in a bad mood.

¹⁸⁴ They could thus be seen as instances of simultaneous use and mention of a word, as discussed in the philosophical literature on quotation.

In both cases, Papafragou argues, an expression is used metarepresentationally to introduce a new name ('the person who could appropriately be called *the brain*' and 'the person who could appropriately be called *the violin*'), and the intended referent of the expression used does not fall under its conventional denotation.

In general, according to Papafragou, metonymy (at least of the referential kind), is motivated by the more general human cognitive tendency to identify individuals or objects on the basis of salient properties that they possess, which serves the purpose of reduction of cognitive effort in accessing referents. Given the view that concepts (for individuals and objects) give access to encyclopaedic information about their denotations, we can take it that salient properties (derived from perception, for instance) of a given individual or object (e.g. has-a-big-nose, violin-playing) translate into their encyclopaedic entries as highly accessible assumptions (e.g. 'X has a big nose', 'X is a violin-player in the quartet'). The idea, then, is that metonymy involves the metarepresentational use of an expression denoting a (contextually) particularly salient object/property "to give access to an individuating conceptual representation of an individual (or another object) through a highly accessible encyclopaedic assumption including the two" (1996: 181-182).

More specifically, the comprehension of metonymy requires of the hearer that he first recognises the metarepresentational use of the expression, and then uses this as a starting point for pragmatically deriving the speaker-intended referent. Let us consider how he might proceed in the case of (35) (Papafragou 1996: 182-183): First, the hearer accesses the 'literal' interpretation, 'The violin is in a bad mood', as a sort of default hypothesis about the speaker's meaning, which he then rejects on the basis of considerations of relevance: a rational speaker aiming at optimal relevance could not have intended him to derive this interpretation. Then, on the basis of the type of predicate involved, the hearer will form a new (approximate) hypothesis about the speaker-intended meaning which includes the expectation that the referent of the definite description is a person; he thus recognises that the definite description is being used metarepresentationally. The proposition he entertains will have something like the following form (where the metarepresentationally used constituent appears within quotation marks):

(36) The person that could appropriately be called ‘the violin’ is in a bad mood.

This proposition requires additional fleshing out by the hearer in his search for a hypothesis about the speaker-intended meaning that satisfies his occasion-specific expectations of relevance; in particular, he must identify the individual whom the speaker is claiming is in a bad mood. Imagine that (35) is uttered at a concert with Anne-Sophie Mutter and the Trondheim Soloists. The assumption that ‘Anne Sophie Mutter is a violin player’ would be highly accessible (and it includes the salient property as well as the referent), and would thus warrant the hearer’s conclusion that the proposition expressed is:

(37) Anne-Sophie Mutter is in a bad mood.

On this account, then, the comprehension of referential metonymy in fact involves three steps: (i) accessing and rejecting the ‘literal interpretation’; (ii) identifying the metarepresentationally used expression, and (iii) forming a hypothesis about the speaker-intended referent.

Let me start by describing what I think are the merits of this account. First, I think it captures a set of highly central characteristics of the metonymy phenomenon: It preserves the strong intuition that metonymy contributes to the proposition expressed by an utterance (i.e. its truth-conditional content). It has an account of how salient properties or objects are turned into highly accessible encyclopaedic assumptions and used to derive the intended referent of a metonymically used expression, instead of having to postulate an inventory of (often arbitrary) associative links between concepts, whose cognitive status remains unclear (are they part of our innate cognitive machinery, are they learned, are only some of them innate, and so on). The account further aims at capturing the crucial aspect of effort reduction involved in metonymically used referring expressions, by assuming that referential metonymies achieve relevance primarily by representing the most cost-efficient way of identifying a referent (but see my comment below). Treating metonymy as a subtype of the metarepresentational use of language also captures the fact that many referential metonymies have an affinity with nicknames (cf. *the brain* in (34)a. above, and the nickname ‘Brainy’). Second, the account situates metonymy within the more general human metarepresentational capacity, which helps

to explain how this sort of use of linguistic expressions arises. Such uses do not have to be taught or learned as they are a natural consequence of this more general metarepresentational capacity (Wilson 2000). In this way, the spontaneous character of metonymy in production and comprehension is accounted for, and the view that metonymy should represent some kind of isolated deviation from a linguistic norm is rejected.

It should be clear from this that I think that Papafragou's proposal represents an important step in the direction of a cognitively plausible account of metonymy, which could potentially form the basis for an up-to-date relevance-theoretic account (a possibility which I explore in section 6.3.4). Its greatest merit is perhaps that it takes seriously the fact that an appropriate treatment of metonymy must be consistent with a general account of utterance interpretation, a point which is often missed, by cognitive linguistic accounts in particular (with the notable exception of Evans 2009). However, it also raises some problematic issues, to which I now turn.

As I mentioned earlier, the comprehension of metonymy on Papafragou's account appears to involve three steps, the first of which is to access, and then reject, the 'literal' interpretation of the metonymically used expression, based on the irrelevant (and usually false) proposition it would yield. This leads to the recognition of the metarepresentational use of the expression, which is the second step in the interpretation process. This way of construing the process seems to be at odds with a central claim about utterance comprehension in recent manifestations of relevance theory (e.g. Wilson and Sperber 2002, 2004); namely, that it involves a *mutual adjustment* of explicit content, contextual assumptions and contextual implications in accordance with the hearer's occasion-specific expectations of relevance (in the next section, I suggest how the process of mutual adjustment may contribute to the derivation of metonymic meanings). The view is that hypotheses about the explicit content of a given utterance, the contextual assumptions that should be brought to bear on its interpretation, and the cognitive effects it is expected to yield are incrementally modified, while a consequence of Papafragou's account seems to be that such hypotheses must be computed in sequence: the hearer must first reject the 'literal' interpretation in order to 'kick off' the inferential process, whose outcome is the identification of the speaker-intended referent for the metonymically used expression.

In commenting on this issue, however, Papafragou (1996: 185) denies that this is the case: “[T]he search for a possible referent for the metonymic expression is seen as a local, i.e. sub-propositional, process. ... my approach does not entail that the literal, i.e. descriptive, referent is necessarily computationally prior to the non-descriptive one.” The problem is that *in practice*, it seems that this is not what her account implies: Recall the analysis of (35). In this case, the hearer recognised the metarepresentational use of the referring expression on the basis of the incompatibility of its descriptive content with that of the predicate of the sentence. From this, he inferred that the speaker-intended referent of the expression was an individual. How could he have achieved this without the prior computation of the ‘literal’ meaning of the whole utterance? Papafragou continues, “Usually, of course, the descriptive interpretation needs to be computed and ruled out, if only because it corresponds to the encoded meaning and the latter is needed as a starting point for any further processing (...)” (ibid.). But, on Papafragou’s account, this computation of the descriptive content in fact seems *crucial* to the recognition of the expression as being an instance of metarepresentational use (and hence an instance of metonymy): How, in the case of metonymy, can a speaker be expected to recognise the expression as being metarepresentationally used if not by contrast with its descriptive content, and the irrelevant proposition it would yield (thus by taking the ‘literal’ meaning of the whole utterance into account)?

There is another, theory-internal issue that arises in connection with Papafragou’s three-step procedure for the interpretation of referential metonymy. Carston (1997; 2002b: Chapter 5) was the first in relevance theory to suggest a symmetrical treatment of lexical narrowing and broadening (including metaphor), as different outcomes of a single pragmatic process of *ad hoc* concept construction, contributing to the proposition expressed by an utterance. As we saw in Chapter 4, this is now the common relevance-theoretic analysis of lexical adjustment (Wilson and Carston 2006, 2007; Sperber and Wilson 2008). Previously, narrowing and broadening had been given quite distinct treatments: while narrowing was treated as contributing extra conceptual material into the proposition expressed (explicature), broadening (or ‘loosening’, as it was called at that time) was treated as a case where the speaker chose to produce an utterance which was a less-than-literal interpretation of the thought she intended to communicate, on the assumption that this would make her thought more

accessible to the hearer than a literally used utterance would (Sperber and Wilson 1986/1995: 231ff.). In these cases, the hearer was thought to recover the speaker-intended meaning of the utterance only as implicatures. As Carston (1997) pointed out, a consequence of this approach was that utterances containing metaphorically used expressions did not communicate the propositions they expressed, unlike utterances involving lexical narrowing.

On the 'old' relevance theory view, then, lexical narrowing, broadening and metonymy were treated as distinct processes (for instance, metonymy was seen as having more in common with irony and other clearly metarepresentational uses of language than with metaphor). However, on the assumption that the current relevance theory account of lexical meaning in context (including cases of metaphor and hyperbole), in terms of *ad hoc* concept construction is correct, one might at least question whether it is right that metonymy should involve the operation of an entirely different pragmatic process, given the many similarities between metonymy and those language uses that result in *ad hoc* concept construction. For one thing, both cases of lexical narrowing and (as has convincingly been argued) lexical broadening, contribute to the proposition expressed by an utterance (Carston 1997, 2002b; Wilson and Carston 2006, 2007; Sperber and Wilson 2008). As noted by Papafragou (and many others), this also appears to be the case for metonymy. For instance, it seems clear that the proposition expressed (i.e. the truth-conditional content) of an utterance of 'The *ham sandwich* is sitting at table 20' has a person and not a ham sandwich as a constituent. And, like narrowing and broadening, metonymy appears to contribute a concept to the mental representation of the proposition expressed. Second, as Papafragou (1996: 176) points out, there is much to indicate that metonymy and metaphor are related processes, given for instance that they can both be used in referential expressions and as predicates, with very similar effects, as shown by her examples in (38) and (39):

- (38) a. The pretty face just went out.
 b. The pretty doll just went out.
- (39) a. Maria is a divine voice.
 b. Maria is a nightingale.

Furthermore, it may sometimes be difficult to decide whether the use of an expression is an instance of a metonymy or a metaphor, as in (40):

(40) The primary goal of Olympic athletes is to *get on the podium*.

The metonymic interpretation would involve seeing the expression *get on the podium* as denoting a highly salient property of the whole event of winning a competition and so used to ‘stand for’ the communicated meaning ‘win an Olympic competition’. Alternatively, the expression could be interpreted metaphorically, as expressing the *ad hoc* concept [GET ON THE PODIUM]* (paraphrasable as ‘win an Olympic competition’), whose denotation would include actual instances of getting on the podium but also other manifestations of winning a competition that might not necessarily include actually appearing on the podium, such as ‘being in the best physical shape’, ‘being the mentally strongest’, ‘beating the other competitors’, etc.¹⁸⁵

In view of this, I think we might question the view that metonymy and metaphor should be the outcomes of entirely different pragmatic processes (although even on the differential treatment approach they would both, of course, be guided and constrained by the general pragmatic strategy of searching for an interpretation consistent with the presumption of optimal relevance). In particular, Papafragou’s account of metonymy in terms of a metarepresentational use of language and the recent relevance-theoretic account of metaphor in terms of *ad hoc* concept construction seem to imply that the two processes differ as to their computational complexity. While metaphor comprehension involves the development of a logical form constituent (output by linguistic decoding) into a communicated concept appearing in the proposition expressed, metonymy comprehension on Papafragou’s account proceeds from the decoded logical form to an ‘intermediary’, approximate hypothesis about the speaker’s intended meaning containing the metarepresented material (as illustrated by the proposition in (36) above), which is then ‘pragmatically unpacked’ into a hypothesis about the proposition expressed. Thus, the comprehension of metonymy would seem to involve three layers of

¹⁸⁵ However, a proponent of a differential processing view of metaphor and metonymy could argue that in this particular example, two different processing routes (the metaphorical broadening and the metonymic naming) eventuate in very similar results.

representation, while there would only be two such layers involved in the comprehension of metaphor.

Finally, Papafragou's three-step procedure for the interpretation of metonymy seems hard to reconcile with the view that referential metonymy represents the most cost-efficient way of identifying a referent. In my view, rather than reasoning from the descriptive content of the concept (x) to a metarepresentationally used concept with a 'naming' function ('the person/object that can appropriately be called 'x)'), on the basis of which a referent (y) is identified (as a result of the activation of encyclopaedic or contextual information involving x and y), it seems more likely that the hearer (somehow) is able to reason directly from the encoded salient property (x) to an accessible encyclopaedic or contextual assumption including the property and the referent (y) (e.g. 'the y that has an x'), on the basis of which the referent (y) is identified. Computing the intermediary proposition containing the 'metarepresented material' seems like a redundant (effort-demanding) step in the comprehension process.

To me, there seem to be two possible ways in which a relevance-theoretic account of metonymy may proceed: Either we could take Papafragou's early analysis as a starting point, and investigate whether her original idea that metonymic interpretations arise as a result of a creative process of 'naming' can be carried over to a more up-to-date relevance-theoretic account, one that avoids the problems described above. Or, we could take a different route and investigate the possibility of developing an entirely different relevance-theoretic account of metonymy in terms of *ad hoc* concept construction. I will now consider each possibility in turn, starting with the latter.

6.3.3 Metonymy as *ad hoc* concept construction?

As mentioned above, since Papafragou's (1996) formulation of a preliminary account of metonymy as an instance of metarepresentational (interpretive) use of language, the development of the relevance-theoretic account of *ad hoc* concept construction has resulted in a kind of asymmetry between the theory's treatment of lexical pragmatic processes such as narrowing and broadening (including metaphor), on the one hand, and its treatment of metonymy, on the other. In view of the affinity between metonymy and *ad hoc* concept construction (they target individual words or phrases, contribute to

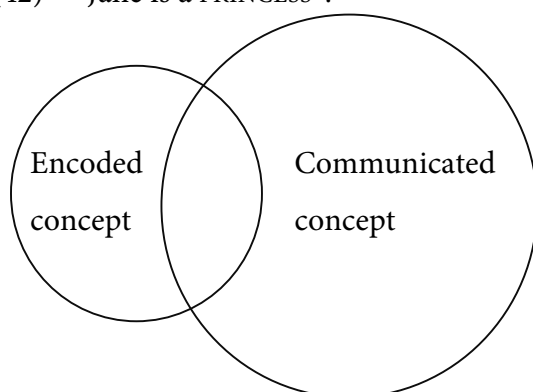
the proposition expressed (explicature), and give rise to a continuum of literal and figurative cases¹⁸⁶), and the fact that metonymy and metaphor may serve similar discourse functions (referential, predicative, etc.) and overlap in several ways (Taylor 1989/2003; Goossens 1990; Papafragou 1996; Barcelona 2000a; Radden 2000), a unified treatment of them seems desirable on both empirical and theoretical grounds. In this section, I will sketch such an account.

Previously, we have seen that relevance theorists describe the outcome of the process of *ad hoc* concept construction as either a concept with a narrower denotation than that which is linguistically-encoded (i.e. picking out a subset of it) or a concept that has a broader denotation (i.e. picking out a superset of the linguistically-specified denotation). While it seems clear that metonymy does not involve narrowing (e.g. the individuals who order ham sandwiches do not constitute a subset of the set of ham sandwiches), could it perhaps be analysed as a kind of broadening of the linguistically-encoded meaning? Consider the following metaphorical utterance:

(41) Jane is a *princess*.

The view is that in interpreting this utterance, the hearer constructs the broadened *ad hoc* concept PRINCESS*, communicating ‘spoiled’ and/or ‘arrogant woman’, which includes actual princesses who satisfy the description and excludes those who don’t (which is why most metaphors also involve an element of narrowing, cf. Carston 1997; 2002; Sperber and Wilson 2008). This can be illustrated as follows:

(42) Jane is a PRINCESS*.




¹⁸⁶ E.g. from the literal ‘*Woolf* is a famous modernist writer’, to the less literal ‘*Woolf* is on the top shelf’ to the figurative ‘The new *Woolf* is coming to the party tonight’ (used to describe an up-and-coming female novelist).

Now consider the ‘ham sandwich’ example again (repeated below as (43) for convenience):

(43) The *ham sandwich* is getting impatient.

An analysis of (43) in terms of *ad hoc* concept construction would imply that the hearer constructs the communicated concept HAM SANDWICH*, paraphraseable as ‘ham sandwich customer’, on the basis of the encoded concept HAM SANDWICH, contextual assumptions derived from the utterance situation (e.g. CUSTOMER X HAS ORDERED A HAM SANDWICH) and context-specific expectations of relevance (e.g. in the context of a busy restaurant, (43) might be expected to indicate to the waitress what she needs to attend to next). However, in this case, the *ad hoc* concept HAM SANDWICH* is not broader than the linguistically-encoded concept in the sense of having a denotation that includes or overlaps with the linguistically-specified one; rather, the concept communicated is entirely different from the one encoded, as can be illustrated in the following way:

(44) The HAM SANDWICH* is getting impatient.



So it seems that if metonymy is an instance of *ad hoc* concept construction, it contributes neither a narrowed nor a broadened concept (denoting a subset or a superset of the encoded lexical or phrasal concept nor a set that overlaps with it) but an entirely different concept to the proposition expressed by the utterance (Carston 2010).

However, in a brief comment on metonymy and its connection to *ad hoc* concept construction, Wilson and Carston (2007: 253-254) suggest that many instances of (what seem to be cases of) metonymy may in fact be analysable in terms of lexical broadening, based on Nunberg’s (1996, 2004) reanalysis of cases such as (45) below as involving an extended meaning for the predicate *parked out back* rather than the first person pronoun having a ‘deferred reference’ (cf. Section 6.2.1):

(45) I'm *parked out back*.

On a relevance-theoretic analysis, the predicate *parked out back* can be treated as communicating the *ad hoc* concept PARKED OUT BACK*, denoting the set of individuals who have their vehicles parked out back. However, this would not, I think, count as a genuine instance of broadening, as the denotation of the *ad hoc* concept is not a proper superset of the linguistically-encoded denotation (the set of things that are parked out back). Rather, the situation appears to be the one illustrated in (44) above (cf. the analysis of this example in Recanati 2004: 26).¹⁸⁷

Furthermore, Wilson and Carston's suggestion is that if Nunberg's analysis is correct, then many of the cases standardly treated as metonymies should be straightforwardly analysable in terms of lexical broadening based on *ad hoc* concept construction, including cases such as (46) and (47):

(46) Which wide body jets *serve dinner*?

(47) Nixon *bombed Hanoi*.

While it is common to treat the NPs *wide body jets* and *Nixon* as metonymies for 'airlines' and 'the U.S. Army', Wilson and Carston suggest instead that it is the predicates, *serve dinner* and *bombed Hanoi*, that have extended senses in (46) and (47). An argument in favour of this view is that in both cases, a conjoined predicate may take the encoded meaning of the NP as its subject, as shown by (48) and (49):

(48) Which wide body jets *serve dinner* and can land on water?

(49) Nixon *bombed Hanoi* and was later forced to resign from presidency over the Watergate scandal.

On the broadening analysis, only the predicates *serve dinner* and *bombed Hanoi* would have extended meanings, while the remaining constituents of the sentence would

¹⁸⁷ That is, unless we can conceive of a concept whose meaning is sufficiently abstract that it can incorporate entities that are in fact parked out back and individuals who are responsible for (strongly causally implicated in) bringing about the parked out back property of those entities. Such an approach would have an affinity with the superordinate concepts that Wilson and Carston (2006) suggest underlie the (basic) physical and (metaphorical) psychological senses of adjectives such as *hard*, *rigid*, *cold*, etc. The idea is that the superordinate concepts (e.g. HARD*, RIGID*, COLD*), which have both physical and psychological instances, have arisen through broadening of the basic physical sense of these adjectives.

receive more or less ‘literal’ interpretations. This also means that in (50) and (51) below, both predicates would have extended senses:

(50) Which wide body jets *serve dinner* and *offer a wide range of inflight entertainment*?

(51) Nixon *bombed Hanoi* and *invaded Cambodia*.

Although I see the motivation behind it, I am not entirely convinced that the analysis in terms of broadening of the linguistically-encoded meanings of the predicates, instead of treating the NPs as metonymies, is the best way to account for these examples. For one thing, the analysis does not capture what appears to be a productive pattern according to which proper names (e.g. ‘Woolf’, ‘Nixon’) can be used to refer to properties associated with their bearers. For another, there seems to be something counterintuitive about the analysis on which the predicates of the utterances in (50) and (51) have extended senses, when it is likely that few speakers would perceive them as such; they seem like straightforward, (more or less) literal uses of expressions. Third, if the predicates in the examples above did in fact communicate phrasal concepts with broader denotations than the ones assigned to them by the grammar, it is not entirely clear to me what these broadened concepts would be.¹⁸⁸ A problem with this is of course that if we take the subjects and not the predicates to have extended meanings in (48) and (49), it now seems that the predicates have different subjects, which does not seem likely in view of the ellipsis of the second subject constituent. However, it is not inconceivable

¹⁸⁸ Perhaps the content of the phrasal *ad hoc* concepts SERVE FOOD* and BOMBED HANOI* in (46) and (47) above could be paraphrased along the lines of ‘have staff that serve food’ and ‘ordered his armed forces to bomb Hanoi’. However, this would look more like (quite unusual) instances of pragmatically supplied unarticulated constituents rather than *ad hoc* concept construction (where a linguistically supplied constituent is pragmatically developed). Indeed, Romero and Soria (2005, 2010) suggest an account on which metonymies are treated as instances of unarticulated constituents that have to be filled in by the context, triggered by the semantic or contextual anomaly of the literal meaning (e.g., in the interpretation of the utterance ‘The *ham sandwich* is getting impatient’ the concept HAM SANDWICH is enriched into the concept HAM SANDWICH CUSTOMER). Carston (2010) points out that this would be quite an unorthodox case of a pragmatically supplied constituent since it would not involve the usual process of pragmatic modification of a linguistically-given value (as in the provision of the location for ‘It’s raining IN OSLO’, for instance), but rather the pragmatic provision of the head value (CUSTOMER) for the linguistically given modifier. Adjuncts like location constituents or quantifier domain constituents are optional in as much as the phrase could stand without them, while arguments like ‘the customer (who ordered a ham sandwich)’ or ‘The book (by Woolf)’ are obligatory elements of a phrase. Carston is doubtful about free enrichment supplying arguments for the uttered element to modify and says that, at the least, it needs detailed justification.

that in the interpretation of (48) and (49), we (somehow) construe the subject referents in a way that is sufficiently general for both predicates to apply to them. Perhaps, due to the fact that the encyclopaedic information that the metonymies exploit is highly accessible, and likely to be activated every time the concepts NIXON and WIDE BODY JETS are accessed (it could even be accessed in ‘chunks’, so that, for instance, assumptions about Nixon the person are activated together with assumptions about his time as president of the United States, and head of the U.S. Army), two predicates can make reference to different aspects of this encyclopaedic information without there being any incoherence (or ‘zeugmatic effect’) perceived by the hearer.¹⁸⁹ However, although I am sceptical about it, I remain open to the possibility that in the particular cases of (48) and (49), we may in fact have to do with extended meanings for the first predicates.

Nevertheless, as Wilson and Carston point out, even on the analysis of (45), (46) and (47) as broadenings of the predicate meanings, there would still remain cases “which seem to involve genuine reference substitution and which are not straightforwardly reducible to lexical narrowing or broadening” (2007: 254). Examples are the following:

(52) The *saxophone* walked out.

(53) *Downing Street* refused to give an interview.

In these cases, conjoining a predicate that makes reference to the linguistic meaning of the italicised expressions creates a strong ‘zeugmatic effect’:

(54) ? The *saxophone* walked out and looked newly polished.

(55) ? *Downing Street* refused to give an interview and was full of protesters.

So it seems that the shift in meaning targets the subjects and not the predicates. In what follows, I will assume that metonymies such as those in (54) and (55) represent a genuine category of cases that arise pragmatically through *ad hoc* concept construction. As argued above, I will take it that metonymy differs from narrowing and broadening in that the concepts (or phrasal concepts) communicated do not denote proper subsets or

¹⁸⁹ This possibility was also evoked towards the end of Chapter 5. The examples in (48) and (49) are in fact not entirely unlike the *book/window* cases that I discussed there. I return briefly to these cases towards the end of this section.

supersets of the linguistically-encoded denotations, but have denotations that fall entirely outside the ones encoded by the metonymically used expressions.

So, how would the construction of genuinely metonymic *ad hoc* concepts proceed during utterance comprehension? As I mentioned in section 6.3.2, it seems that in the interpretation of metonymy, the hearer is (somehow) able to ‘reason’ directly from the linguistically-encoded conceptual content of the metonymically used expression to an accessible encyclopaedic or contextual assumption which includes the linguistically-determined denotation and the speaker-intended meaning, on the basis of which the speaker-intended meaning is identified. As an illustration, consider the ‘ham sandwich’ example again (repeated here as (56) for convenience):

(56) The *ham sandwich* is getting impatient.

Imagine that (56) is uttered at a café by one waiter to another (let’s say from Jane the waitress to Sam the waiter) during lunchtime, a very busy time of the day at this café. The waiters are running around trying to serve the customers their correct orders in time. Against this background, an anticipated conclusion of Jane’s utterance would be that whoever (among the customers) is getting impatient should be served his or her food as quickly as possible. The linguistically specified concept HAM SANDWICH would provide additional activation to an already highly accessible contextual assumption about ham sandwiches as possible orders at this café, and by a process of spreading activation, about customers having ordered ham sandwiches. Let’s say that at the time of utterance of (56) there is only one customer waiting for his order of a ham sandwich. In this case, the encoded meaning of HAM SANDWICH would activate the contextual assumption CUSTOMER A HAS ORDERED A HAM SANDWICH. The interpretation of *ham sandwich* as communicating the *ad hoc* concept HAM SANDWICH* (person who ordered ‘ham sandwich’) allows Sam to identify customer A as the referent of the expression, and warrants the implicated conclusion that customer A should be served his food as quickly as possible. Below is a schematic outline of how Sam’s interpretation of Jane’s utterance in (56) might proceed (interpretive hypotheses to the left, the basis for deriving them to the right):

(57)

(a) Jane has said to Sam [THE HAM SANDWICH IS GETTING IMPATIENT].	Decoding of Jane's utterance. Embedding of the decoded logical form into a description of Jane's ostensive behaviour.
(b) Jane's utterance will be optimally relevant to Sam.	Expectation raised by the recognition of Jane's utterance as an act of ostensive communication.
(c) Jane's utterance will achieve relevance by providing information relating to their current task of serving customers their lunch.	Expectation raised by the overall context and the fact that this kind of information would be the most relevant to Sam at this point.
(d) A customer who is getting impatient should be served his order as quickly as possible.	First contextual assumption activated that satisfies the expectation in (c). Accepted as an implicit premise of Jane's utterance.
(e) A ham sandwich is a possible order at this café.	A highly accessible encyclopaedic assumption that receives additional activation by the concept HAM SANDWICH.
(f) Customer A has ordered a ham sandwich.	Accessible contextual assumption, activated by the encoded concept HAM SANDWICH via the encyclopaedic assumption in (e).
(g) The HAM SANDWICH* is getting impatient.	First interpretation of the explicit content of Jane's utterance, involving the <i>ad hoc</i> concept HAM SANDWICH* ('person who ordered ham sandwich'), derived associatively from the contextual assumption in (f). This property is assigned to the referent, 'customer A'. Accepted as the explicature of Jane's utterance.
(h) Customer A should be served his food.	Implicature of Jane's utterance. Satisfies

	the expectation in (c). Accepted as the intended implicated conclusion. Logically implied by (d) and (g).
--	---

Sam decodes Jane’s utterance and takes this decoded content (an incomplete logical form) as input to the inferential process to follow, and assumes that her utterance is optimally relevant to him. Given the particular context, Sam expects Jane’s utterance to achieve relevance by providing him with information relating to their current task of serving customers their lunch. The first contextual assumption to come to Sam’s mind that satisfies this expectation is that ‘A customer who is getting impatient should be served his order as quickly as possible’, which he takes as an implicit premise of Jane’s utterance. The encoded concept HAM SANDWICH provides additional activation of an already highly accessible encyclopaedic assumption, that ham sandwiches are possible orders at this café, which, together activates the accessible contextual assumption that ‘Customer A has ordered a ham sandwich’. On the basis of this assumption, the *ad hoc* concept HAM SANDWICH* (‘person who ordered ham sandwich’) is associatively derived and assigned to the referent ‘customer A’, as part of the first hypothesis about the explicit content of Jane’s utterance ‘The HAM SANDWICH* is getting impatient’ (where the concept IMPATIENT will probably also be adjusted (narrowed) to an *ad hoc* concept IMPATIENT*, specifying the kind/degree of impatience of the customer), which Sam accepts as the explicature of Jane’s utterance. Together, the contextual assumption in (d) and the explicature warrant the derivation of the implicature in (h), which satisfies the expectation in (c), and is accepted as the intended implicated conclusion of Jane’s utterance.¹⁹⁰

In (57), the interpretive step from (a) to (f) would be associative; the move from [the speaker has said] HAM SANDWICH to [the speaker meant] HAM SANDWICH* does not rest on a logical or evidential relation between the encoded HAM SANDWICH and the communicated HAM SANDWICH*. However, the overall inferential process leading to the derivation of the implicature in (h) severely constrains the range of possible associative

¹⁹⁰ In a case where this type of metonymy has become more of a coded or routine operation (which is indeed conceivable, e.g. among a group of employees at a busy café), there might be an immediate (associative) move from ‘The ham sandwich is getting impatient’ to ‘Customer A is getting impatient’.

relations that the encoded concept HAM SANDWICH may enter into. Moreover, the ‘working out’ of the metonymic meaning takes place within a process of mutual adjustment of explicit and implicit content. In this particular case, an occasion-specific expectation of relevance, satisfied by the implicature, preceded and shaped the interpretative hypothesis about the explicature, containing the metonymic *ad hoc* concept.

Consider another example, in which the mutual adjustment process also appears to play an important role in the derivation of the intended metonymic interpretation:

- (58) [The conductor]: Is everyone here?
[An orchestra member]: The *saxophone* just left.

In this case, assuming the orchestra member is aiming at optimal relevance, the conductor can take her utterance to achieve relevance by providing him with an answer to his question. A highly accessible assumption in the context would be that ‘if someone just left, everyone isn’t here’. The encoded meaning of *saxophone* (the concept SAXOPHONE) would activate assumptions about the person playing the saxophone in the orchestra, say, ‘Sally is the player of the saxophone’. The hypothesis that the explicit meaning (explicature) of the utterance is ‘The SAXOPHONE* just left’, involving the *ad hoc* concept SAXOPHONE* (‘player of the saxophone’, a property assigned to the referent Sally), warrants the derivation of the expected conclusion (i.e. provides an answer to the conductor’s question) ‘Everyone isn’t here’, that is, it allows for the interpretation to achieve relevance in the expected way. Thus, in this case, the hypothesis about the implicature (‘Everyone isn’t here’) precedes and shapes the hypothesis about the explicature, the derivation of which involves constructing the metonymic *ad hoc* concept.

An advantage of the relevance-theoretic approach is that it can also explain how certain accessible associations are ruled out in the interpretation process, and why some metonymies do not work (in a given context). Consider the example in (59):

- (59) A. So who’s Mary’s new boyfriend?
B. The *red shorts* in the corner.

Imagine that (59) is uttered at a party, where A has just been chatting with another party guest, C, and got into a discussion about football kits. During this discussion, A, an enthusiastic Liverpool supporter, tells C about the Liverpool children's kit she has bought her son for his birthday, upon which C, an ardent Chelsea supporter, says that she absolutely hates that red kit, and in particular, 'those ghastly red shorts!'. To A, therefore, the association between the concept RED SHORTS and the Liverpool football kit would have a high degree of activation at the time of B's utterance of 'The red shorts in the corner', and is, in fact, likely to be more highly activated than the association between the concept RED SHORTS and the guy standing in the corner wearing an 'ordinary' pair of red shorts. However, this association would yield a clearly wrong hypothesis about the speaker-intended meaning. The context-specific expectation of relevance that B's utterance comes with, that is, that it should provide A with an answer to her question (i.e. picking out to her Mary's new boyfriend), ensures that this initially accessed hypothesis is quickly rejected so that she tries out the next most accessible association, that is, the one between RED SHORTS and the guy standing in the corner (who is wearing red shorts), also made (perceptually) accessible in the context. This hypothesis is accepted as it contributes to B's utterance achieving relevance in the expected way (i.e. as having the explicature 'The RED SHORTS* ('guy wearing the red shorts') in the corner [is Mary's new boyfriend]'). Thus, in this case, a strong association (RED SHORTS - LIVERPOOL KIT) is ruled out as part of the overall inferential process of arriving at a hypothesis about the speaker-intended meaning; the hypothesis it would yield is disconfirmed by a context-specific expectation of relevance. The important point here is that the inferential framework that relevance theory provides ensures that there are strong constraints on the associative relations that will be accepted by the relevance-based comprehension procedure.

Consider also the following examples of metonymies that do not work in the given contexts:

(60) [Professor to his students]: ?There are many good *bodies* in this university.¹⁹¹

(61) [Student to her fellow students]: The *cream coat* is in a bad mood today.

¹⁹¹ Example due to (Croft 1993: 181).

In (60), the expression *bodies* is not a well-chosen metonymy if the speaker's intended referent is 'students'. The associative link between students and their bodies leads the hearers to assume that this property of students is somehow central to the overall meaning of the utterance. More specifically, it may give rise to the assumption that the professor is openly declaring an interest in the bodies of his students, which might contradict other strongly held assumptions that the students might have about him (or anyway about what he would feel free to say openly). On the current account, the inappropriateness of this metonymy could be explained as a result of the encoded concept BODIES giving access to encyclopaedic assumptions about the metonymic referent which, if taken as premises in the inferential process, would yield a wrong hypothesis about the speaker-intended meaning. Overall, this could be seen as a case in which the extra effort required to process the metonymy sets the hearer off searching for extra effects, and those he arrives at would not be intended by the speaker.¹⁹² Compare with a case in which the professor utters 'There are many good *minds* in this university', with the intention to refer to the students; this metonymy would probably also lead to miscommunication, as the most accessible association is likely to be that between good minds and the academic staff; and the professor would be taken to refer to his colleagues rather than to the students. As to the metonymy in (61), imagine that it is uttered by one student to another while they are waiting for their professor in her office, and her cream-coloured coat is hanging up next to the door. This metonymy seems to fail because it is neither effort-saving nor likely to have any special effects. There would probably be several other far more salient properties that the hearer would associate

¹⁹² The choice of properties in synecdoche seems, in many cases, to be largely governed by stereotypical assumptions about the referents, to the extent that they almost appear to have a normative function. Consider the examples below:

- (i) We need some new *heads* in the Philosophy department.
- (ii) We need some new *hands* in the home for the elderly.

In both cases, the italicised expressions are used as synecdoches to refer to people, but give rise to quite different assumptions about the referents. In (i), the use of *heads* evokes stereotypical assumptions about philosophers; that they are intelligent, talented, and so on, which may give rise to implicatures concerning the status of the department in need of such qualities. In (ii), the use of *hands* evokes stereotypical assumptions about caregivers in homes for the elderly; that they are caring, attentive, considerate, and so on, which may give rise to implicatures about the condition of the home in question (it may not be able to offer the required services to the patients, it may be in financial difficulties, and so on). Compare with 'We need some new *heads* in the home for the elderly', which would automatically activate assumptions about the management of the home, not about the caregivers. The choice of a particular property to be used as a synecdoche thus seems to be partly governed by assumptions about what the referents *ought to be like*, which is what gives some of them a normative flavour.

with the professor and so he might not be able to make an immediate connection between her and the coat, although this connection is made available by the immediate physical context. It is hard to see how such cases of miscommunication or non-workable metonymies could be properly explained within an entirely associative account.

Finally, I would like to consider briefly how the *ad hoc* concept construction approach might be applied to some of the systematic sense alternations discussed in the previous chapter (the so-called *book/window* cases, cf. section 5.3.3), which seemed to rest on a metonymic relation between a physical object sense and a more abstract (institution or organisation) sense. Consider (62) and (63):

(62) The *newspaper* announced staff redundancies.

(63) This *church* has got a liberal position on same sex marriage.

The ‘management of the publishing organisation’ sense communicated by *newspaper* in (62) could be analysed as being derived by means of *ad hoc* concept construction, on the basis of the encoded ‘physical object’ sense and encyclopaedic information activated by it, within the overall inferential process of forming a hypothesis about the speaker intended meaning. The *ad hoc* concept communicated, NEWSPAPER*, would have an entirely different denotation than the encoded concept. Similarly, the metonymic use of *church* in (63) could be treated as communicating the *ad hoc* concept CHURCH* (‘clerical leaders’), constructed on the basis of the encoded ‘building’ sense of the noun and encyclopaedic information that it gives access to, constrained by the hearer’s occasion-specific expectations of relevance.

However, with regard to the *window* cases, for which it was claimed that we ‘zoom in’ on more specific aspects of the denotation depending on the context, these cannot as easily be treated in terms of entirely different *ad hoc* concepts being communicated (as pointed out in the previous chapter); at least, such an analysis would not capture the part-whole relation that is crucially involved in such cases. Consider again a couple of the examples from Chapter 5, repeated in (64):

- (64) a. The *window* is broken.
b. The *window* is rotten.

The perspectives on *window* expressed in (64) cannot be analysed in terms of lexical narrowing of a general physical object sense of window, as the concepts communicated (let's say, 'window pane', 'window frame') would be subparts rather than proper subsets of the linguistically-specified denotation. In view of this, my suggestion was to account for the intuition about different senses being expressed in (64)a. and (64)b. in terms of certain encyclopaedic information receiving extra activation, and treating the referential meaning of *window* as being constant across these uses. However, I also mentioned the possibility of treating this type of word as encoding not a full-fledged concept, but a much more schematic representation (a pointer), which gives access to the aspects of meaning exemplified in (64). On this approach, *window* could be treated as expressing genuinely different *ad hoc* concepts in (64)a. and (64)b., each of which would be inferentially constructed on the basis of encyclopaedic information made accessible by the schematic encoded meaning. A consequence of this analysis would be that the process of *ad hoc* concept construction would become obligatory in the comprehension of these words in context. I cannot get further into this issue here, but I think an approach along these lines deserves serious consideration.

Some advantages and shortcomings of the account

The account of metonymy in terms of *ad hoc* concept construction has an obvious advantage in that it allows for a unified treatment of metonymy, narrowing and broadening, as different outcomes of a single pragmatic process operating at the level of individual words (narrowing contributing a concept to the proposition expressed (explicature) which is a proper subset of the linguistically-specified denotation, broadening contributing a proper superset of the linguistically-specified denotation and metonymy an entirely different concept), and, as such, it captures the commonalities between them. Given that metonymy is an important source of polysemy, the account is also consistent with the overall claim in this thesis that polysemy arises as a result of the operation of a pragmatic inferential process, which yields hypotheses about speaker-intended senses on the basis of encoded senses, contextual information and relevance-

based interpretive constraints. The *ad hoc* concept construction account further meshes well with Nunberg's (1996, 2004) claim that metonymy is an instance of 'meaning transfer' rather than 'deferred reference'; the metonymic meaning appears as an *ad hoc* concept in the proposition expressed by an utterance, and arises as a result of a radical 'adjustment' of the encoded concept (i.e. into an entirely different concept), on the basis of which the speaker-intended referent is picked out.

Furthermore, the 'regularity' often observed in connection with metonymic uses, for instance, that referring to customers via their food orders may be a standard customer identification procedure among the employees of a café (and similarly, referring to orchestra members via the instruments they play among members of an orchestra, or to people via salient perceptual properties they possess quite generally), is also quite compatible with an account in terms of *ad hoc* concept construction. As we have seen, cognitive linguists take this 'regularity' of many linguistic metonymies to stem from underlying cognitive structures in the form of 'conceptual metonymies', and claim that the same is true of metaphor (c.f. 'conceptual metaphors' such as LOVE IS A JOURNEY, ARGUMENT IS WAR, etc.). However, from a relevance-theoretic point of view, as indeed Wilson (2009) has suggested for metaphor, it is possible to see such regularities as cases where a repeated use of a linguistic metonymy that links different concepts together has set up a pattern of conceptual activation, or a 'pragmatic routine', with similar characteristics as the structures that cognitive linguists claim are instances of 'conceptual metonymies' (e.g. FOOD ORDERS FOR CUSTOMERS, MUSICAL INSTRUMENTS FOR ORCHESTRA MEMBERS, BUILDING FOR INSTITUTION, etc.). Thus, the regularity observed could be seen as originating in the construction of *ad hoc* concepts for communicative purposes (which, if repeated often enough, might result in the setting up of systematic correspondences), rather than as surface reflections of underlying conceptual metonymies (Wilson 2009).

However, a set of cases for which I do not think that the account in terms of *ad hoc* concept construction works very well is the metonymic use of proper names. If metonymy is a pragmatic process that targets the linguistic meaning of an expression and yields a context-specific concept on the basis of this meaning, then it is not clear what entity it operates on in the case of proper names, which arguably do not have any encoded semantic content. For *ad hoc* concept construction to work in these cases, a

referent has to be assigned to the proper name before the process can operate on a representation of this referent. This might work for fully established cases such as *Woolf*, *Dostoyevsky* and so on, but it is far from clear that it would work in cases where the referent of a name is unknown to the hearer and yet the metonymy goes through, as, for instance, in an utterance of ‘*Needham* is difficult to read’, where the hearer would interpret *Needham* as referring to ‘book(s) written by Needham’, regardless of whether or not he has ever heard about this author. In this case, the metonymy seems to be constructed on the basis of an established rule (e.g. AUTHOR FOR BOOKS) that might apply to any proper name given the appropriate contextual circumstances, rather than involving any kind of creative meaning construction. Note that this could not so easily be analysed as an instance of a pragmatic routine resulting from the repeated construction of *ad hoc* concepts linking authors and their books, as the new name itself would not be associated with any encyclopaedic information that might activate such systematic correspondences.

A further issue that arises in connection with the account of metonymy in terms of *ad hoc* concept construction is the cognitive and communicative motivation for it. On this account, narrowing can be seen as being motivated by considerations of economy (i.e. the reduction of processing effort); there is in principle no need for a more specific encoding so long as the speaker can rely on her hearer easily arriving at her intended meaning. Broadening, on the other hand, is more often associated with an increase in effort which is offset by extra effects; using a metaphor, for instance, may enable the speaker to convey a range of (weak) implicatures, some of which may not even be amenable to natural language encoding (cf. Chapter 3). Metonymy, however, is not so easily motivated within this account. Compared with a literal encoding (which is, in most cases possible), it seems in fact often to involve an *increase* in effort, with no obvious gain in effects. In those cases where the metonymy involves a reduction of processing effort, this seems to be due to an already existing rule or ‘code’ for referent identification which has become established among a group of language users (or has developed into a code that is shared by the language community as a whole, e.g. AUTHOR FOR BOOKS). It appears, then, that in most cases, novel uses of metonymies must have some other motivation than contributing to a reduction of processing effort. This fact, I think, is better captured by the account that treats metonymy as a creative use of

naming, as was originally proposed by Papafragou (1996). In the final section to follow, I will sketch a different account of metonymy that builds on the early analysis discussed above. When appropriately modified, this account seems better suited to capture both the crucial aspect of processing effort reduction involved in metonymic use, as well as an important aspect of metonymy that has so far been largely ignored; the fact that it often involves an expression of attitude towards the metonymic referent.

6.3.4 Metonymy as ‘naming’ revisited

The main contribution of Papafragou’s (1996) account was to view referential metonymy as a variety of metarepresentational use of language involving a form of ‘naming’, closely related to nicknaming, where a salient property of an individual is used to create a new name. In this section, I will suggest a way in which this idea can be carried over to a more up-to-date relevance-theoretic account.

Recall from section 6.3.2 that, on Papafragou’s account, the comprehension of metonymy requires that the hearer first recognise the metarepresentational use of the expression, and use this as a starting point for pragmatically deriving the speaker-intended referent. As I discussed, a problem with this account is that the recognition of the metarepresentational use (and hence of the metonymic use) seems to require the prior computation and rejection of the literal meaning of the utterance. This is what is supposed to trigger the construction of an ‘intermediary proposition’ containing the metarepresented material (i.e. ‘the *x* that can appropriately be called *y*’), on the basis of which the appropriate referent is identified. In my view, this problem can be avoided if we allow for mutual adjustment to play a role in the interpretation process. Consider the example in (58) again (repeated here as (65) for convenience):

(65) [The conductor]: Is everyone here?

[An orchestra member]: The *saxophone* just left.

As argued in the discussion of (58) above, the conductor can assume that the orchestra member’s utterance will achieve relevance by providing him with an answer to his question. A highly accessible assumption in the context would be that ‘if someone just left, then everyone isn’t here’. As in (58) above, the encoded meaning of *saxophone* (the concept *SAXOPHONE*) would activate assumptions about the person who plays the

saxophone in the orchestra, say, ‘Sally is the player of saxophone’. The recognition of the expression *The saxophone* as an instance of naming with the purpose of picking out Sally (as the individual who can appropriately be called ‘The saxophone’), which would give rise to the explicature ‘Sally has just left’, warrants the derivation of the expected conclusion ‘Everyone isn’t here’, that is, it allows for the interpretation to achieve relevance in the expected way. Thus, the recognition of the metarepresentational use could be seen as taking place within the process of mutually adjusting the explicit and implicit content of the utterance, that is, as taking place independently of any prior computation of the literal meaning of the whole utterance.¹⁹³

I will now suggest how the communicative functions of metonymic uses can be captured on the current account. As we have seen, Papafragou (1996) takes metonymy to be motivated by a more general cognitive tendency to identify individuals on the basis of salient properties they possess. Let us assume, then, that in metonymy, a speaker *names* a referent by use of the word or phrase for a property she takes the referent to possess (e.g. ‘The [person that can appropriately be called] ‘saxophone’ has left), and which she either assumes is also highly accessible to the hearer (and thus might allow for efficient identification of the referent), or wants to make the hearer aware of (thus inducing cognitive effects that a literal utterance would not). Consider the following examples:

- (66) The *red shorts* in the corner is Mary’s new boyfriend.
- (67) The *green tea* is coming over to talk to us.
- (68) Peter is bringing the *loudmouth* with him to the opera.

When uttered, for instance, in a context where Jane is wondering who Mary’s new boyfriend is, the metonymy in (66) can be seen as having a purely *referential function*. The speaker takes a property of the referent that is highly salient to her (i.e. his wearing red shorts), assumes that it is also highly salient to the hearer, and uses the linguistic expression for this property to *name* the referent, with the purpose of allowing the hearer to easily identify him (say among a group of people at a party). In (67), the

¹⁹³ I am not denying that in some cases the recognition of a metonymic use may in fact take place as a result of computing (and rejecting) the literal meaning of the utterance. However, I do not think that a general account of metonymy can be built on this assumption.

function of the metonymy *green tea* may be entirely referential, for instance, in the context of a market where vendors can be easily identified on the basis of the type of product they sell. It is also possible to conceive of a context in which this metonymy would involve an expression of attitude toward the referent, say, at a party where the speaker and hearer have overheard the person in question asking for a cup of green tea instead when offered a glass of champagne. In such a case, the main purpose of the metonymic naming would be to allow the hearer to draw conclusions from it, for instance, that the coming conversation might turn out to be boring, and so on. In (68), the primary function of the metonymy *loudmouth* is to draw the hearer's attention to a specific property of the referent (i.e. the property of talking too much in an offensive or tactless way). Using this expression to name the referent (by 'holding up', as it were, the property for the hearer to attend to it), allows the speaker to express her (negative) attitude toward him or her in an indirect way, and enables the hearer to infer that she wants him to know that she thinks that the referent possesses this property.

On this account, then, metonymy achieves relevance either by representing the most cost-efficient way to identify an individual, as in (66), compared with using a direct encoding (e.g., 'the man wearing the red shorts'). Or, in the cases where its primary function is to draw the hearer's attention to a particular property of the referent, as in (68), the extra effort that may be required to identify the referent of the description is offset by extra effects that would not be achieved by use of a more direct utterance. (However, there may also be an element of effort reduction involved in these cases, as the metonymy represents an efficient way of getting across this additional meaning.)

Broadly speaking, then, metonymies may serve two different communicative functions, which may give rise to two different situations. First, in cases where the naming serves the purpose of quick and easy identification of a referent, this typically gives rise to the establishment of productive 'mini-codes' for referent identification among sub-groups of language users. For instance, the members of an orchestra may have come to share the code INSTRUMENT FOR PLAYER, which they use as an efficient way of identifying each other. Further examples are given in (69)-(71):

- (69) Among café employees: FOOD ORDERS FOR CUSTOMERS (*The french fries and a coffee* has requested her bill').
- (70) Among hospital workers: DIAGNOSIS FOR PATIENT (*The arthritis* is in room 205').
- (71) Among university staff: ESSAY TOPIC FOR STUDENT (*The explicit/implicit distinction* will get a high mark'), DEGREE PROGRAMME FOR STUDENT (*The PhD* is getting on my nerves').

An example of a code for referent identification that has developed into one that is now shared by the (English) language community at large is that which allows us to use the name for an author to refer to the book(s) written by him or her.¹⁹⁴ The productive character of such codes is what allows us to understand *Needham* in the example above as referring to the books written by the author Needham (even though we might never have heard of this author before), provided that there is sufficient contextual information to warrant a metonymic interpretation.¹⁹⁵ It seems likely that a reason why some codes come to be shared by a language community as a whole is that the kind of conceptual relations they rest on are of more or less permanent character (such as, for instance, the relation between authors and their books, between copies of newspapers and the companies producing them, between churches and the clerical staff running them, etc., cf. Chapter 5), or the relation is an intrinsic one (for instance, that between humans and their body parts, between fruits and the trees producing them, etc.). It could thus be hypothesised that the more established a code is in a language community, the more we tend to think of the lexical items that activate it as being polysemous. Consider (72)-(75):

- (72) Mary is reading *Dostoyevsky* on the tube.
- (73) *Cambridge* voted conservative.
- (74) The *newspaper* sacked its editor-in-chief.
- (75) Jill and Joan have a *cherry* in their garden.

¹⁹⁴ As noted by many cognitive linguists, the code AUTHOR FOR OEUVRE/BOOKS may be an instance of an even more general coded rule PRODUCER/CREATOR FOR PRODUCT/CREATION.

¹⁹⁵ That such interpretations are largely code or rule-based finds support in a recent experiment on the processing of such metonymies, conducted by Frisson and Pickering (2007).

These examples have been described as instances of systematic polysemy with a metonymic basis. I think that an important reason why we perceive this type of polysemy as being systematic is the regularity of the real-world relations that they describe. Indeed, all of the metonymies above rest on properties which, if not definitional, are at least highly central properties of the denotations of the vehicle expressions, to the extent that it may be hard to think of their denotations without at the same time thinking of this property. For instance, in (72), the property of being a writer is probably *the* most salient property that speakers associate with the name *Dostoyevsky*; in fact, they may be more prone to access a representation of his work than a representation of the actual writer himself upon hearing it. In this way, there is a sense in which *Dostoyevsky* is not only something ‘one might appropriately call’ a copy of one of his books; it is, as it were, something that that book *is*, by virtue of being written by him. Similarly, in the other examples, a (more or less) intrinsic property of a city is that it has inhabitants; a highly central property of a newspaper is that it is being published by someone; an intrinsic property of a fruit is that it grows on a plant of some kind, usually a tree. These facts, I think, are the main reason why we feel that the metonymies in (72)-(75) remain close to the ‘literal’ end of language use; treating them as instantiations of codes for referent identification shared by the language community as a whole, provides a possible explanation for this intuition.

The account of metonymy as a form of naming can also deal with more creative cases of metonymy involving the use of proper names. Consider the following (attested) example:

(76) A. Hope *Leipzig* went well.

B. *Leipzig* hasn’t happened yet. (*Leipzig* = B’s talk at the Max Planck Institute in Leipzig).¹⁹⁶

In this case, A ‘names’ B’s forthcoming talk on the basis of a highly accessible property of it, namely, that it is going to take place in Leipzig. In this case, the function of the metonymy is clearly effort reduction (for both speaker and hearer); it works as a shorthand for a much more elaborate description.

¹⁹⁶ Thanks to Robyn Carston for providing this example.

Second, the cases where the main purpose of the naming appears to be to highlight, or make the hearer aware of, a property that the referent possesses (and in this way potentially expressing an attitude toward the referent), have, to my knowledge, apart from Papafragou's early discussion, so far received little attention in the research on metonymy. The main purpose of such metonymies appears to be to convey implicatures about the referent of the expression (cf. (68) above). Such cases are accounted for on the current account as metarepresentational uses in which the speaker is 'holding up' a property associated with the referent that she wants the hearer to attend to, and by the mere fact that she is doing so, she may succeed in conveying an attitude towards the referent in an implicit way. It is in these cases that the affinity between metonymy and nicknaming is particularly clear. My claim is that rather than giving rise to mini-codes for referent identification, metonymies thus motivated often seem to give rise to nicknames (e.g. 'I saw *Loudmouth* the other day', '*Four Eyes* is always reading math books', 'Where is *Brainy* now that we need him?'), which are sometimes precursors to proper names (e.g. the English and Scottish surname *Sellar*, originally used to name someone who worked in a cellar) and/or racial or social slurs (e.g. *albino*, *redneck*, *redskin*, *camel jockey* etc.), the latter cases invariably involving a strong expression of negative attitude.¹⁹⁷

An advantage of treating metonymy as a form of naming, compared with treating it as a case of *ad hoc* concept construction, is that metonymic uses appear to have a stronger motivation on this account (i.e. quick and easy referent identification contributing to reduction of processing effort/conveying extra effects that would not have been achieved by use of a literal encoding). The account also provides a straightforward explanation of the affinity between nicknames and metonymy, of why many slurs are metonymy-based, and (when appropriately modified) of the regular,

¹⁹⁷ Interestingly, there is another productive linguistic strategy in English (and in Norwegian, and probably in other languages as well) for naming individuals on the basis of properties they possess, as illustrated by the following examples:

- (i) I was just told by Ms. Politically Correct that I should mind my language on that topic.
- (ii) Miss Selfish decided to keep the goods to herself.
- (iii) Mr. Computer Geek over there can help you with any IT-related problem.

Here, the metonymic uses of the expressions are linguistically marked as dubbings by the presence of the titles, the capital letters (and special stress). The titles enable adjectives to be used directly as metonymic referring expressions. In addition to being referential, these cases often communicate an attitude toward the referent, for instance, (i) may imply that the referent is excessively politically correct.

code-like character of many metonymies. Furthermore, it can account for the same range of cases as the *ad hoc* concept construction account, as well as those involving proper names, which, as we saw in section 6.3.3, were not obviously analysable in terms of a process of meaning transfer targeting encoded word meanings.

However, recalling Nunberg's discussion in section 6.2.1, a consequence of pursuing the metonymy as naming account is a return to the view that the process is a matter of reference substitution rather than of meaning transfer. This would distinguish metonymy from other pragmatic processes that operate at the level of individual words and phrases, such as broadening and narrowing, which would involve the operation of a different pragmatic mechanism. In this way, some of the common characteristics between metonymy and *ad hoc* concept construction would not be captured.

A more serious problem, however, is that the process of creative naming is in fact unlikely to be a distinctive feature of metonymy, as it applies equally to referential metaphors (and possibly to other cases as well). Compare the metonymy in (77)a. and the metaphor in (77)b, which were discussed in section 6.3.2:

- (77) a. *The pretty face* just went out.
b. *The pretty doll* just went out.

Here, both the metonymy and the metaphor are used to name a referent with very similar effects. Moreover, just like metonymy, referential metaphors make good nicknames, as illustrated by the following description of the former French president, Jacques Chirac:

- (78) In the three decades of his political career, the man they call 'The Bulldozer' has forged his own path and proved a thorn in the side of more than one French government.¹⁹⁸

There is also a clear connection between referential metaphors (79) and metaphorically based proper names (80):

- (79) *The wilting violet* seems like she's ready to leave.

¹⁹⁸Attested: http://www.bbc.co.uk/bbcfour/documentaries/profile/profile_jacques_chirac.shtml

(80) Violet, Rose, Iris, Daisy, etc.

What conclusions may we draw from this? It seems that if we want to maintain an account of metonymy in terms of a creative process of naming, we must assume that we have to do with a broader category of cases, including referential metaphors (in the form of metarepresentational *ad hoc* concepts), and possibly other cases as well. On this view, metonymy would not be a natural kind, but rather a variety of uses of creative naming. As I see it, the most promising option would be to devise an account that is able to combine the constructive insights of the two approaches discussed here, that is, one that is able to capture the common features of metonymy, broadening (including metaphor) and narrowing, the regularity involved in many metonymic uses, as well as what seems to be an intrinsic connection between 'figurative' referential uses and metarepresentation, nicknaming and expressions of attitude. Much more needs to be said about this. I must leave the working out of such an account to future investigations within the relevance-theoretic framework.

Chapter 7

CONCLUSION: A PRAGMATIC ACCOUNT OF POLYSEMY

My aim in this thesis has been to show how the phenomenon of polysemy is intimately connected to, in fact overlaps with, cases of on-line contextual modulation of meanings due to pragmatic processes operating at the level of individual words. In this concluding chapter, I provide an overview of the central arguments presented in this thesis and discuss some of their implications for future work on polysemy/lexical pragmatics. I will also address some unresolved issues raised by my pragmatic account of polysemy.

This thesis started out by identifying several issues that polysemy raises in semantic theory and semantic applications. The main questions were:

- (i) How should the phenomenon of polysemy be defined, and (if desirable) distinguished from homonymy (accidental multiple encoding), on the one hand, and contextual modulation of meaning, on the other?
- (ii) How are polysemous lexical forms represented in the mental lexicon?
- (iii) What is the cognitive-communicative motivation for the proliferation of polysemy in natural languages?
- (iv) How do lexical meanings get extended into several different meanings?

I also pointed out that, in striking contrast to these difficult issues, polysemy is largely unproblematic from the point of view of communication; the ambiguity or indeterminacy that it creates is something that we handle effortlessly and unconsciously, for the vast majority of time. This has led some scholars to talk of a ‘polysemy paradox’ (Ravin and Leacock 2000; Taylor 2003). In what follows, I will briefly summarise my suggestions for how the questions I raised can be answered within the pragmatic account of polysemy that I have outlined in this thesis. Taken together, these suggested answers provide a possible solution to the paradox.

Chapter 2 discussed the problem of polysemy representation. A first question was whether an adequate account of polysemy representation requires decompositional word meanings. A widely held view among proponents of this view is that word meanings must consist of complex representations in order to capture meaning

relations between words (e.g. synonymy, analyticity, polysemy). I started the chapter by considering several influential decompositional theories of lexical semantics, including Katz's (1972) semantic theory, Lakoff's lexical network theory (1987), Pustejovsky's (1995a) generative lexicon account, and Pinker's (1989) and Jackendoff's (2002) theories of partial decomposition of lexical meanings. Their implications for polysemy representation were discussed. I then considered the conceptual atomist account maintained by Fodor (1998), and identified two ways of describing polysemy within this kind of framework, which posits unstructured (atomic) word meanings: either polysemous senses could be represented as separate lexical entries ($PAPER_1$, $PAPER_2$), with the same linguistic form (*paper*), or there could be a single encoded meaning of the word (*cut* - CUT), and the differences in meaning perceived between its uses in different contexts could be treated as instances of pragmatic adjustment of the encoded meaning (CUT*, CUT**, CUT***).

I concluded that the existence of polysemy does not provide a very strong argument in favour of lexical decomposition. First, it is not always clear that polysemous sense relations are in fact captured by all decompositional theories. Second, for many of those decompositional theories where polysemous sense relations are captured, the problem is often that much of what we would consider aspects of lexical meaning in *use* is incorporated into the lexical semantics. Third, independent of the issue of polysemy representation, decompositional theories of word meaning are associated with a range of largely unresolved problems (e.g. the incompleteness of many decompositions, the vagueness of many concepts, the fact that we may be ignorant or mistaken about the properties we take the instances of a concept to have, etc.). Finally, in view of my discussion of Fodor's atomist account, there is no obvious reason why a conceptual atomist approach combined with an adequate pragmatic theory should not be at least as able as decompositional theories to provide an account of polysemy. The view taken in this thesis was that word meanings are represented as unstructured atomic concepts, thus avoiding most of the problems associated with decompositional accounts, and that pragmatics can account for most of the variation in lexical meanings across contexts.

A second question was whether polysemous lexical forms are represented as multiple lexical entries (sense enumeration) or as a single entry (core meaning). I examined some relatively recent psycholinguistic studies that apply experimental

methods in order to glean evidence from on-line processing that may bear on the question of whether polysemous lexical items are represented as multiple or single entries in the lexicon. The experimental evidence, although far from conclusive, indicated that not all instances of polysemy are represented in the same way. Rather than storing only a single core meaning or a total list of all the possible distinct senses, the lexicon may store some reasonable number of senses, and other senses are constructed in context on the basis of the stored senses. This result was compatible with the atomist approach to lexical meanings taken in this thesis.

Chapter 3 examined the cognitive-communicative motivation for the proliferation of polysemy in natural languages. I argued that a solution to the problem of polysemy motivation has to take as its starting point the fact that polysemy is largely unproblematic from a communicative point of view. My hypothesis was that a pragmatic account that is able to explain how and why this is so, is also likely to shed light on the issue of what motivates the proliferation of polysemy in natural languages in the first place. I proposed a pragmatic account of polysemy within the framework of relevance theory (Sperber and Wilson 1986/1995; Carston 2002b), in which polysemy arises as a result of encoded lexical meanings being massively underdetermining of speaker-intended meanings and is grounded in our pragmatic inferential ability to construct occasion-specific senses on the basis of encoded meanings and contextual information. I argued against more semantically-oriented approaches that although there is, of course, an important linguistic aspect to polysemy, it is, at the deepest level, a consequence of how communication works. I further suggested, drawing on arguments from Sperber (2000) and Carston (2002b), that our metarepresentational theory of mind ability might provide the cognitive basis for polysemy, and that it might have developed as a result of the co-evolution of the capacities for language and metarepresentation. More specifically, I argued that polysemy is fundamentally grounded in the component that allows us to infer speaker meanings from encoded linguistic meanings; this is what provides the motivation for polysemy in language. According to this pragmatic account polysemy poses no paradox, but is what we would expect from the fact that our languages do not fully encode our thoughts.

In Chapter 4, I addressed the nature of the pragmatic processes that are involved in the construction of polysemy. I claimed that new senses for a word (giving rise to

polysemy) are constructed during on-line utterance comprehension by means of a single process of *ad hoc* concept construction, which adjusts the meanings of individual words in different directions (giving rise to *ad hoc* concepts with either narrower or broader denotations than the linguistically-encoded denotations, e.g. ‘John CUT* the cake’, ‘The steak is RAW*’, ‘It’s BOILING* outside’, ‘The book was SLAUGHTERED* by reviewers’), in accordance with the hearer’s occasion-specific expectations of relevance (Carston 2002b; Wilson and Carston 2007). At the end of the chapter, I briefly discussed the issue of prepositional polysemy, and sketched how an analysis of the multiple (related) meanings of the preposition *over* in terms of *ad hoc* concept construction might proceed. It is a subject for further research whether such an analysis would be workable for the range of uses of *over*, and whether it could be extended to the other prepositions of English (or to the prepositions of other languages).

In Chapter 5, I applied this relevance-theoretic approach to lexical pragmatics to cases of so-called systematic polysemy. The focus was the kind of sense alternation that appears to depend on whether a noun is used with count or mass syntax (e.g. John shot a *rabbit*/had *rabbit* for lunch). This type of polysemy is usually seen as a prime candidate for an analysis in terms of lexical rule application, where the effect of the lexical rules is to change the value of a linguistically marked [+count] or [+mass] feature on a noun. I argued that such rule-based accounts do not provide the interpretive flexibility required to handle the variations in meaning that this sense alternation may give rise to. Instead I proposed that our intuitions about some nouns being ‘count’ and other nouns being ‘mass’ stem from our mental representations of the concepts that are encoded by them. These may be specified, in terms of meaning postulates attached to the concepts or imagistic representations in their encyclopaedic entries, as denoting individuals or unindividuated entities, or, indeed, be unspecified with regard to this distinction. I further suggested that the count-mass distinction as it is manifested in language can be seen as an instance of this conceptual (ontological) distinction being reflected in a syntactic distinction at the level of NPs (rather than at the level of individual nouns). The cases of ‘systematic’ polysemy would arise whenever a noun that encodes a concept perceived as denoting an individual is used with mass syntax, or when a noun that encodes a concept perceived as denoting an unindividuated entity is used with mass syntax. The pragmatic process of *ad hoc* concept construction would

then operate to yield context-specific senses of these nouns (e.g. to specify whether a mass use of *rabbit* communicates ‘rabbit meat’, ‘rabbit fur’, ‘rabbit stole’, ‘rabbit remains’ etc.). This way of approaching the count-mass type of systematic polysemy is largely unexplored in the literature. A possible direction for future research on this topic would be to investigate the feasibility of this mainly pragmatic account when applied to a larger set of data.

In Chapter 5 I also discussed the systematic polysemy of nouns such as *book* and *window*, and suggested that their different senses (e.g. ‘The *book* is sitting on the coffee table/John found the *book* interesting’, ‘Mary broke the *window*/Mary crawled through the *window*’) could be analysed in terms of activation of encyclopaedic properties associated with the concepts (so-called ‘perspectivising’), the referential meanings of the nouns remaining constant across contexts. I also briefly evoked the possibility of an analysis of these cases in terms of *ad hoc* concept construction. Pursuing this option would require that these nouns be seen as encoding not full-fledged concepts, but more schematic representations, on the basis of which their different senses could be derived via pragmatic specification in terms of *ad hoc* concept construction. The possibility that all words encode nothing more than pointers to bundles of information (hence do not encode concepts) is currently being explored in relevance-theory (Carston 2002b, 2010; Kjøllo forthcoming) and elsewhere (Recanati 2004: Chapter 9; Bosch 2007; Pietroski 2008; Pritchard 2009). Although I have chosen to subscribe to the ‘traditional’ relevance-theoretic stance here, according to which most words encode atomic concepts, I remain open to the possibility that word meanings may in fact prove to be much less specific than this. For instance, with regard to the so-called *book/window* cases, treating them as encoding schematic representations might provide a compelling option, as it captures the fact that the precise denotation of such concepts is often surprisingly hard to pin down. A consequence of treating word meanings as schematic abstract entities in this way would be that the process of *ad hoc* concept construction would become both entirely general and obligatory in the comprehension of words in context; it would be triggered by underdetermined lexical-semantic representations in order to construct full-fledged concepts to the proposition expressed by an utterance. The question of whether words encode concepts or representations that are much more schematic is thus still very much an open issue in need of further investigation.

A cross-cutting topic of Chapters 2-5 has been whether the aspects of meaning that are involved in the construction of polysemy have a primarily linguistic or non-linguistic basis. The question has been to what extent polysemy is a result of the operation of linguistic processes, and to what extent is it governed by pragmatic inferential processes applying at the level of individual words. As the question indicates, it is possible to see this as a difference of degree. While some approaches maintain that a considerable amount of linguistic knowledge is involved in the generation of polysemy (e.g. computational semantic accounts), other approaches, such as the pragmatic account of polysemy that I have been arguing for here, downplay the linguistic aspect and claim that polysemy mainly arises as a result of the operation of pragmatic processes over encoded word meanings, taking contextual information and encyclopaedic assumptions about the denotation of the concept involved as input to the inferential process. However, more than being a matter of degree, the difference between these two types of approaches is, in reality, to do with two fundamentally different conceptions of what a language is. The computational semantic accounts see the language as providing a rich code that enables speakers to encode and decode their thoughts in much detail, and pragmatics as a useful add-on to this linguistic capacity, operating primarily when some interpretation other than the default interpretation generated by the linguistic system is indicated by the context. By contrast, the pragmatic account pursued here sees the role of the linguistic system as being that of providing a minimal input or clue, which the inferential system uses as evidence to yield hypotheses about occasion-specific, speaker-intended meanings. As I have argued throughout this thesis, the assumption that a large part of the interpretive work involved in the processing of polysemy should be attributed to the linguistic system itself requires independent justification, given that we have an independently motivated pragmatic interpretation system, automatically activated by verbal utterances, which is capable of rapidly generating new senses in contexts.

Finally, in Chapter 6 I discussed the phenomenon of metonymy, an important source of polysemy, where an expression that conventionally denotes one object or property is used to refer to or 'stand for' something that falls outside its conventional denotation, but with a clear associative relation holding between the conventional and the metonymic denotations (e.g. 'The *ham sandwich* is getting impatient', 'Woolf is on

the top shelf'). I suggested two possible ways in which a relevance-theoretic account may proceed. First, I sketched an account according to which metonymic interpretations arise within the process of *ad hoc* concept construction, and where the range of possible associative relations that the encoded concept may enter into are constrained by the hearer's occasion-specific expectations of relevance. Second, I sketched a different account in which metonymy is treated as a form of creative naming, closely related to 'nicknaming', where a salient property of an individual is used to create a new name. As the discussion of metonymy revealed, it remains much of an open issue within the relevance-theoretic framework. Nevertheless, the accounts sketched go some way in explaining how metonymic interpretations are derived (and constrained) within the overall inferential process of utterance understanding. In particular, the 'naming' element, including the expression of attitude that often follows from it, involved in metonymic uses needs to be further explored. It can be explored either as part of an account of metonymy in terms of *ad hoc* construction, or as part of an account that treats metonymy as a sub-variety of a broader phenomenon of creative naming, including at least the referential metonymies and metaphors discussed in this chapter, and possibly other cases as well.

As I mentioned in Chapter 6, recently, some researchers have begun to compare the relevance-theoretic and cognitive linguistic accounts of metaphor and to consider how they might be combined (Gibbs and Tendahl 2006; Tendahl and Gibbs 2008; Wilson 2009). Gibbs and Tendahl, for instance, suggest that, despite some fundamental differences, the two approaches should be seen as providing complementary rather than contradictory perspectives on metaphor. An interesting direction for future research would be to investigate whether a relevance-theoretic approach to the comprehension of metonymy might also be fruitfully combined with the cognitive linguistic approach to metonymy, in particular the cognitive linguistic claim about the existence of underlying metonymic concepts. A central question to be investigated is what role, if any, such underlying metonymic concepts play in the comprehension of metonymy.

Finally, I would like to return briefly to the problem of polysemy definition addressed at the beginning of this thesis. The problem concerned how to distinguish polysemy from homonymy on the one hand, and pragmatic adjustment of lexical meanings, on the other. Given my account of polysemy as a result of pragmatic

adjustment of lexical meanings, the question becomes whether we really need the label 'polysemy' at all. Perhaps it would be sufficient to posit a distinction between ambiguity (including homonymy and conventional polysemy), on the one hand, and pragmatically adjusted meanings, on the other, as this is the only distinction that has an impact on the way in which an utterance is processed, involving a process of disambiguation in the first case and *ad hoc* concept construction in the second case. Given the conceptual atomist stance on word meanings taken in this thesis, this would imply that speakers' intuitions about sense relations would not be directly reflected at the level of lexical semantic representation, contrary to the common assumption in traditional linguistic semantics and in cognitive linguistic frameworks (whose account of polysemy is built around the existence of sense relations at the level of lexical semantic representation). Instead, as I suggested in Chapter 4, these intuitions may be due to several factors: While some sense relations may be transparent, as with clearly contextually derived senses, other intuitions about sense relations may be caused either by an overlap in the encyclopaedic information associated with concepts, by our ability to reflectively reconstruct the process of sense extension from a primary to a derived sense (as with many conventional metaphors), or by the mere anticipation of a relation between senses encoded by a single lexical form (cf. homonyms). In my view, the idea that sense relations are not represented as part of the lexical representation of a polysemous lexical form deserves serious consideration in future work on the topic. While nobody would dispute the fact that we have intuitions about sense relations, there is, as far as I can tell, no clear-cut evidence that these are directly reflected in the lexical representation of polysemous words. Until such evidence can be provided, it is necessary to consider other possible explanations for their existence. Furthermore, operating with a distinction between lexical ambiguity and *ad hoc* concept construction meshes well with the psycholinguistic evidence suggesting that not all instances of polysemy are necessarily represented in the same way. The assumption that there need not be a single representational format for the data falling under the label 'polysemy' is, I think, supported by the instances of polysemy investigated in this thesis, which, as we have seen, are to some extent heterogeneous.

In conclusion, I hope to have shown in this thesis that the pragmatic/communicative aspect of polysemy is crucial to its existence and proliferation

in natural languages. While probably few would deny that pragmatics plays an important role in giving rise to different senses of a lexical form, this has, as we have seen, rarely been the focus of investigation in the research on polysemy, which tend to see it either as a linguistic reflection of how cognitive categories are structured or as a result of linguistic processes operating at the level of individual words (and pragmatics as useful add-on to these processes). While both these lines of research have provided important insights into the nature of polysemy, I think future work on this subject would benefit from taking the fundamental pragmatic aspect of the phenomenon more seriously. As I have argued in this thesis, the pragmatic perspective allows us to question common assumptions about the representation of polysemous lexical forms, and to re-examine the role of the linguistic system in giving rise to several senses for a lexical item, which may be less central than previously thought. Ultimately, of course, these are empirical issues. I am optimistic that future experimental studies will provide us with a deeper understanding of how polysemy is represented and constructed during on-line pragmatic processing.

References

- Allerton, D. J. 1979. *Essentials of Grammatical Theory: A Consensus View of Syntax and Morphology*. London: Routledge & Kegan Paul.
- Allwood, J. 2003. Meaning potentials and context: Some consequences for the analysis of variation in meaning. In *Cognitive Approaches to Lexical Semantics*, edited by H. Cuyckens, R. Dirven and J. R. Taylor. Berlin: Mouton de Gruyter. 29-66.
- Anderson, R. C., and A. Ortony. 1975. On putting apples into bottles - A problem of polysemy. *Cognitive Psychology* 7 (2):167-180.
- Apresjan, J. D. 1974. Regular polysemy. *Linguistics* 14 (2):5-32.
- Armstrong, S. L., L. Gleitman, and H. Gleitman. 1999 [1983]. What some concepts might not be. In *Concepts: Core Readings*, edited by S. Margolis and E. Laurence. Cambridge, Massachusetts: The MIT Press. 225-259.
- Aronoff, M. 1976. *Word Formation in Generative Grammar, Linguistic Inquiry Monographs 1*. Cambridge Massachusetts: The MIT Press.
- Asher, N. forthcoming. *Lexical Meaning in Context*. Cambridge: Cambridge University Press.
- Asher, N., and A. Lascarides. 1995. Lexical disambiguation in a discourse context. *Journal of Semantics* 12 (1):69-108.
- . 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.
- Atlas, J. D. 1989. *Philosophy without Ambiguity: A Logico-Linguistic Essay*. Oxford: Clarendon Press.
- . 2005. *Logic, Meaning and Conversation*. Oxford: Oxford University Press.
- Atran, S. 2002. Modular and cultural factors in biological understanding: an experimental approach to the cognitive basis of science. In *The Cognitive Basis of Science*, edited by P. Carruthers, S. Stich and M. Siegal. Cambridge: Cambridge University Press. 41-72.
- Bach, K. 1994. Conversational implicature. *Mind and Language* 9 (2):124-162.
- Bach, K., and R. Harnish. 1979. *Linguistic Communication and Speech Acts*. Cambridge, Massachusetts: The MIT Press.

- Barcelona, A. 2000a. On the plausibility of claiming a metonymic motivation for conceptual metaphor. In *Metaphor and Metonymy at the Crossroads: A Cognitive Perspective*, edited by A. Barcelona. Berlin: Mouton de Gruyter. 31-58.
- , ed. 2000b. *Metaphor and Metonymy at the Crossroads: A Cognitive Perspective*. Berlin: Mouton de Gruyter.
- Barner, D., and J. Snedeker. 2005. Quantity judgments and individuation: evidence that mass nouns count. *Cognition* 97 (1):41-66.
- . 2006. Children's early understanding of mass-count syntax: individuation, lexical content, and the number asymmetry hypothesis. *Language Learning and Development* 2 (3):163-194.
- Barnes, J., ed. 1984. *The Complete Works of Aristotle. Vol 1*. Princeton: Princeton University Press.
- Baron-Cohen, S. 1995. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, Massachusetts: The MIT Press.
- Baron-Cohen, S., A. M. Leslie, and U. Frith. 1985. Does the autistic child have a theory of mind? *Cognition* 21 (1):37-46.
- Barsalou, L. W. 1983. Ad hoc categories. *Memory & Cognition* 11:211-227.
- . 1987. The instability of graded structure: implications for the nature of concepts. In *Concepts and Conceptual Development*, edited by U. Neisser. New York: Cambridge University Press. 101-140.
- Bartsch, R. 2002. Generating polysemy: Metaphor and metonymy. In *Metaphor and Metonymy in Comparison and Contrast*, edited by R. Pörings and R. Dirven. Berlin: Mouton de Gruyter. 49-74.
- Beretta, A., R. Fiorentino, and D. Poeppel. 2005. The effects of homonymy and polysemy on lexical access: An MEG study. *Cognitive Brain Research* 24 (1):57-65.
- Bickerton, D. 1990. *Language and Species*. Chicago, Illinois: Chicago University Press.
- Blakemore, D. 1987. *Semantic Constraints on Relevance*. Oxford: Blackwell.
- . 1992. *Understanding Utterances*. Oxford: Blackwell.
- . 2002. *Relevance and Linguistic Meaning: The Semantics and Pragmatics of Discourse Markers*. Cambridge: Cambridge University Press.

- Bloom, P. 1994a. Possible names: The role of syntax-semantics mappings in the acquisition of nominals. *Lingua* 92 (1-4):297-329.
- . 1994b. Syntax-semantics mappings as an explanation for some transitions in language development. In *Other Children, Other Languages: Issues in the Theory of Language Acquisition*, edited by Y. Levy. Hillsdale, New Jersey: Erlbaum. 41-76.
- . 1999. The role of semantics in solving the bootstrapping problem. In *Language, Logic and Concepts: Essays in Memory of John Macnamara*, edited by R. Jackendoff, P. Bloom and K. Wynn. Cambridge, Massachusetts: MIT Press. 285-310.
- . 2000. *How Children Learn the Meaning of Words*. Cambridge, Massachusetts: The MIT Press.
- Blutner, R. 1998. Lexical pragmatics. *Journal of Semantics* 15 (2):115-162.
- . 2002. Lexical semantics and pragmatics. In *Semantics*, edited by F. Hamm and T. E. Zimmermann. Hamburg: Helmut Buske Verlag. 27-58.
- . 2004. Pragmatics and the lexicon. In *Handbook of pragmatics*, edited by L. R. Horn and G. L. Ward. Oxford: Blackwell. 488-514.
- Boguraev, B., and B. Levin. 1993. Models for lexical knowledge bases. In *Semantics and the Lexicon*, edited by J. Pustejovsky. Dordrecht: Kluwer Academic Publishers. 325-340.
- Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.
- . forthcoming. *Pursuing Meaning*. Oxford: Oxford University Press.
- Borowsky, R, and M. E. J. Masson. 1996. Semantic ambiguity effects in word identification. *Journal of Experimental Psychology. Learning, Memory, and Cognition* 22 (1):63-85.
- Bosch, P. 2007. Productivity, polysemy and predicate indexicality. In *Proceedings of the Sixth International Tblisi Symposium on Language, Logic and Computation*, edited by H. Zeevat and B. ten Cate. Heidelberg/Berlin: Springer.
- Bréal, M. 1924 [1897]. *Essai de sémantique: science des significations*. Paris: Gérard Monfort.

- Briscoe, T., A. Copestake, and A. Lascarides. 1995. Blocking. In *Computational Lexical Semantics*, edited by P. Saint-Dizier and E. Viegas. Cambridge: Cambridge University Press.
- Brown, R. 1958. How shall a thing be called? *Psychological Review* 65 (1):14-21.
- Brown, S. W. 2008. Polysemy in the mental lexicon. *Colorado Research in Linguistics* 21.
- Brugman, C. 1988. *The Story of over: Polysemy, Semantics, and the Structure of the Lexicon*. New York: Garland.
- Brugman, C., and G. Lakoff. 1988. Cognitive topology and lexical networks. In *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology and Artificial Intelligence*, edited by S. Small, G. Cottrell and M. Tannenhaus. San Mateo, California: Morgan Kaufman. 477-507.
- Bybee, J. L., R. Perkins, and W. Pagliuca. 1994. *The Evolution of Grammar: Tense, Aspect, and Modality in the Languages of the World*. Chicago, Illinois: University of Chicago Press.
- Byrd, R. J., N. Calzolari, M. S. Chodorow, J. L. Klavans, M. S. Neff, and O. A. Rizk. 1987. Tools and methods for computational lexicology. *Computational Linguistics* 13 (3-4):219-240.
- Call, J., and M. Tomasello. 2005. What do chimpanzees know about seeing revisited: An explanation of the third kind. In *Joint attention: Communication and Other minds: Issues in Philosophy and Psychology*, edited by N. Eilan, C. Hoerl, T. McCormack and J. Roessler. Oxford: Oxford University Press. 45-64.
- Campbell, L. 1998/2004. *Historical Linguistics. An Introduction*. Edinburgh: Edinburgh University Press.
- Cappelen, H., and E. Lepore. 2005. *Insensitive Semantics*. Oxford: Blackwell.
- Caramazza, A., and E. Grober. 1976. Polysemy and the structure of the subjective lexicon. In *Semantics: Theory and Application*, edited by C. Rameh. Washington: Georgetown University Round Table on Languages and Linguistics. 181-206.
- Carston, R. 1988. Implicature, explicature, and truth-theoretic semantics. In *Mental Representations: The Interface between Language and Reality*, edited by R. M. Kempson. Cambridge: Cambridge University Press. 155-181.
- . 1997. Enrichment and loosening: complementary processes in deriving the proposition expressed. *Linguistische Berichte* 8:103-127.

- . 2002a. Linguistic meaning, communicated meaning and cognitive pragmatics. *Mind & Language* 17 (1/2):127-148.
- . 2002b. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell Publishers.
- . 2007. How many pragmatic systems are there? In *Saying, Meaning and Referring. Essays on François Recanati's Philosophy of Language*, edited by M. J. Frápolli. Basingstoke: Palgrave Macmillan. 18-48.
- . 2010. Explicit communication and 'free' pragmatic enrichment. In *Explicit Communication: Robyn Carston's Pragmatics*, edited by B. Soria and C. Romero. Basingstoke: Palgrave Macmillan.
- Carston, R., and G. Powell. 2006. Relevance theory - new directions and developments. In *The Oxford Handbook of Philosophy of Language*, edited by E. Lepore and B. Smith. Oxford: Oxford University Press. 341-360.
- Cheney, D. L., and R. M. Seyfarth. 1990. *How Monkeys See the World: Inside the Mind of Another Species*. Chicago, Illinois: The University of Chicago Press.
- Cheng, C. Y. 1973. Response to Moravcsik. In *Approaches to Natural Language*, edited by J. Hintikka, J. Moravcsik and P. Suppes. 286-288.
- Chierchia, G. 1998. Reference to kinds across language. *Natural language semantics* 6 (4):339-405.
- Clark, E. V., and H. H. Clark. 1979. When nouns surface as verbs. *Language* 55 (4):767-811.
- Clark, H. H. 1992. *Arenas of Language Use*. Chicago, Illinois: The University of Chicago Press.
- . 1996. *Using Language*. Cambridge: Cambridge University Press.
- Coleman, L., and P. Kay. 1981. Prototype semantics: the English verb *lie*. *Language* 57 (1):26-44.
- Copetake, A., and T. Briscoe. 1992. Lexical operations in a unification-based framework. In *Lexical Semantics and Knowledge Representation, Proceedings of the First SIGLEX Workshop*, edited by J. Pustejovsky and S. Bergler. Berlin: Springer-Verlag. 101-119.

- . 1996. Semi-productive polysemy and sense extension. In *Lexical Semantics: The Problem of Polysemy*, edited by J. Pustejovsky and B. Boguraev. New York: Clarendon Press. 15-67.
- Croft, W. 1993. The role of domains in the interpretation of metaphors and metonymies. *Cognitive Linguistics* 4:335-370. Reprinted in *Metaphor and Metonymy in Comparison and Contrast*, edited by Dirven, R. and R. Pörings. Berlin: Mouton de Gruyter, 161-206.
- . 2006. On explaining metonymy: Comment on Peirsman and Geeraerts, "Metonymy as a prototypical category". *Cognitive Linguistics* 17 (3):317-326.
- Cruse, D. A. 1986. *Lexical Semantics*. Cambridge: Cambridge University Press.
- . 1995. Polysemy and related phenomena from a cognitive linguistic viewpoint. In *Computational Lexical Semantics*. Cambridge: Cambridge University Press. 33-49.
- . 2000. Micro-structure of word meanings. In *Polysemy: Theoretical and Computational Approaches*, edited by Y. Ravin and C. Leacock. Oxford: Oxford University Press. 30-51.
- . 2004. *Meaning in Language: An Introduction to Semantics and Pragmatics*. Oxford: Oxford University Press.
- Cuyckens, H., and B. Zawada, eds. 1997. *Polysemy in Cognitive Linguistics: Selected Papers from the Fifth International Cognitive Linguistics Conference*. Amsterdam: John Benjamins.
- de Almeida, R. G. 1999. *The Representation of Lexical Concepts: A Psycholinguistic Inquiry*. PhD Thesis. Rutgers, The State University of New Jersey.
- . 2004. The effect of context on the processing of type-shifting verbs. *Brain and language* 90 (1-3):249-261.
- de Almeida, R. G., and V. D. Dwivedi. 2008. Coercion without lexical decomposition: Type-shifting effects revisited. *Canadian Journal of Linguistics* 53 (2/3):301-326.
- Dennett, D. C. 1991. *Consciousness Explained*. London: Penguin Books.
- Dirven, R., and R. Pörings, eds. 2003. *Metaphor and Metonymy in Comparison and Contrast*. Berlin: Mouton de Gruyter.
- Dunbar, G. 2001. Towards a cognitive analysis of polysemy, ambiguity, and vagueness. *Cognitive Linguistics* 12 (1):1-14.

- Evans, V. 2005. The meaning of time: polysemy, the lexicon and conceptual structure. *Journal of Linguistics* 41 (1):33-75.
- . 2009. *How Words Mean: Lexical Concepts, Cognitive Models and Meaning Construction*. Oxford: Oxford University Press.
- Evans, V., and M. Green. 2006. *Cognitive Linguistics: An Introduction*. Edinburgh: Edinburgh University Press.
- Faarlund, J. T. 1978. *Verb og predikat. Ein studie i norsk verbalsemantikk*. Oslo: Universitetsforlaget.
- Falkum, I. L. 2007. A relevance-theoretic analysis of concept narrowing and broadening in English and Norwegian original texts and translations. *Languages in Contrast* 7 (2):119-142.
- Fauconnier, G. 1994 [1985]. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge: Cambridge University Press.
- Fauconnier, G., and M. Turner. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Fieke Van der Gucht, K. W., and L. De Cuyper. 2007. The iconicity of embodied meaning: Polysemy of spatial prepositions in the cognitive framework. *Language Sciences* 29:733-754.
- Fillmore, C. J. 1982. Towards a descriptive framework for spatial deixis. In *Speech, Place, and Action: Studies in Deixis and Related Topics*, edited by R. J. Jarvella and W. Klein. New York: Wiley.
- Fiorentino, R., and D. Poeppel. 2004. Decomposition of compound words: An MEG measure of early access to constituents. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*, edited by R. Alterman and D. Kirsch. Mahwah, NJ: Erlbaum. 1342.
- Fitch, W. T. 2010. *The Evolution of Language*. Cambridge: Cambridge University Press.
- Fleischman, S. 1982. *The Future in Thought and Language: Diachronic Evidence from Romance*. Cambridge: Cambridge University Press.
- Fodor, J. 2004a. Having concepts: A brief refutation of the twentieth century. *Mind & Language* 19 (1):29-47.
- . 2004b. Reply to commentators. *Mind & Language* 19 (1):99-112.

- Fodor, J. A. 1970. Three reasons for not deriving 'kill' from 'cause to die'. *Linguistic Inquiry* 1 (4):429-438.
- . 1975. *The Language of Thought*. Cambridge University Press: Harvard University Press.
- . 1981. The present status of the innateness controversy. In *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, Massachusetts: The MIT Press. 257-316.
- . 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Massachusetts: The MIT Press.
- . 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- . 2008. *LOT2: The Language of Thought Revisited*. Oxford: Oxford University Press.
- Fodor, J. A., M. F. Garrett, E. C. Walker, and C. H. Parkes. 1980. Against definitions. *Cognition* 8 (3):263-267.
- Fodor, J. A., and E. Lepore. 1998. The emptiness of the lexicon: Reflections on Pustejovsky's 'Generative Lexicon'. *Linguistic Inquiry* 29:269-288.
- . 2002. *The Compositionality Papers*. Oxford: Clarendon Press.
- Fodor, J. D. , J. A. Fodor, and M. F. Garrett. 1975. The psychological unreality of semantic representations. *Linguistic Inquiry* 6 (4):515-531.
- Frazier, L., and K. Rayner. 1990. Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language* 29 (2):181-200.
- Frisson, S., and M. J. Pickering. 2007. The processing of familiar and novel senses of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes* 22 (4):595-613.
- Geeraerts, D. 1993. Vagueness's puzzles, polysemy's vagaries. *Cognitive Linguistics* 4 (3):223-272.
- . 1994. Polysemy. In *The Encyclopedia of Language and Linguistics*. Oxford: Pergamon. 3227-3228.
- . 1997. *Diachronic Prototype Semantics: A Contribution to Historical Lexicology*. Oxford: Oxford University Press.

- Gentner, D., and S. Goldin-Meadow. 2003. *Language in Mind*. Cambridge, Massachusetts: The MIT Press.
- Gibbs, R. W. 1992. Categorization and metaphor understanding. *Psychological Review* 99 (3):572-577.
- . 1994. *The Poetics of Mind*. Cambridge: Cambridge University Press.
- . 1999. Speaking and thinking with metonymy. In *Metonymy in Language and Thought*, edited by K.-U. Panther and G. Radden. Amsterdam: John Benjamins. 61-76.
- Gibbs, R. W., and M. Tendahl. 2006. Cognitive effort and effects in metaphor comprehension: Relevance theory and psycholinguistics. *Mind & Language* 21 (3):379-403.
- Gillon, B. S. 1992. Towards a common semantics for English count and mass nouns. *Linguistics and philosophy* 15 (6):597-639.
- . 1999. The lexical semantics of English count and mass nouns. In *The Breadth and Depth of Semantic Lexicons*, edited by E. Viegas. Dordrecht: Kluwer Academic Publishers. 19-37.
- Glucksberg, S. 1998. Understanding metaphors. *Current Directions in Psychological Science* 7:39-43.
- . 2001. *Understanding Figurative Language: From Metaphors to Idioms*. New York: Oxford University Press.
- Glucksberg, S., and B. Keysar. 1990. Understanding metaphorical comparisons: Beyond similarity. *Psychological Review* 97:3-18.
- Glucksberg, S., M. S. McGlone, and D. A. Manfredi. 1997. Property attribution in metaphor comprehension. *Journal of Memory and Language* 36:50-67.
- Goddard, C. 2000. Polysemy: a problem of definition. In *Polysemy: Theoretical and Computational Approaches*, edited by Y. Ravin and C. Leacock. Oxford: Oxford University Press. 129-151.
- Goossens, L. 1990. The interaction of metaphor and metonymy in expressions for linguistic action. *Cognitive Linguistics* 1 (3):323-340. Reprinted in *Metaphor and Metonymy in Comparison and Contrast*, edited by Dirven, R. and R. Pörings. Berlin: Mouton de Gruyter, 349-377.

- Gordon, P. 1985. Evaluating the semantic categories hypothesis: The case of the count/mass distinction. *Cognition* 20 (3):209-242.
- Grice, H. P. [1957] 1989. Meaning. In *Studies in the Way of Words*. Cambridge, Massachusetts: Harvard University Press. 213-223.
- . [1967] 1989. Logic and conversation. In *Studies in the Way of Words*. Cambridge, Massachusetts: Harvard University Press.
- Hall, A. 2008. Free enrichment or hidden indexicals? *Mind & Language* 23:426-456.
- Halle, M., and A. Marantz. 1993. Distributed morphology and the pieces of inflection. In *The View from Building 20*, edited by K. Hale and S. J. Keyser. Cambridge, Massachusetts: The MIT Press. 111-176.
- Haser, V. 2005. *Metaphor, Metonymy, and Experientialist Philosophy: Challenging Cognitive Semantics*. Berlin: Walter de Gruyter.
- Hauser, M. D., N. Chomsky, and W. T. Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298 (5598):1569-1579.
- Heine, B., U. Claudi, and F. Hünemeyer. 1991. *Grammaticalization: A Conceptual Framework*. Chicago, Illinois: The University of Chicago Press.
- Hopper, P. J. 1991. On some principles of grammaticization. In *Approaches to Grammaticalization, Vol. 1*, edited by E. C. Traugott and B. Heine. Amsterdam: John Benjamins. 17-35.
- Hopper, P. J., and E. C. Traugott. 1993/2003. *Grammaticalization*. Cambridge: Cambridge University Press.
- Horn, L. R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicatures. In *Meaning, Form and Use in Context*, edited by D. Schiffrin. Washington DC: Georgetown University Press. 11-42.
- Horse, R. 2006. *The Content and Acquisition of Lexical Concepts*. PhD Thesis, University of London.
- Horton, M. S., and E. M. Markman. 1980. Developmental differences in the acquisition of basic and superordinate categories. *Child Development* 51:708-719.
- Hurford, J. 2007. *The Origins of Meaning*. Oxford: Oxford University Press.
- Huttenlocher, J., and L. V. Hedges. 1994. Combining graded categories: Membership and typicality. *Psychological Review* 101 (1):157-165.

- Jackendoff, R. 1983. *Semantics and Cognition*. Cambridge, Massachusetts: The MIT Press.
- . 1990. *Semantic Structures*. Cambridge, Massachusetts: The MIT Press.
- . 1992a. Mme. Tussaud meets the binding theory. *Natural Language and Linguistic Theory* 10:1-31.
- . 1992b. What is a concept, that a person may grasp it? In *Languages of the Mind: Essays on Mental Representation*. Cambridge, Massachusetts: MIT Press. 21-52.
- . 1997. *The Architecture of the Language Faculty*. Cambridge, Massachusetts: The MIT Press.
- . 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- . 2010. *Meaning and the Lexicon: The Parallel Architecture 1975-2010*. Oxford: Oxford University Press.
- Jastrzemski, J. E. 1981. Multiple meaning, number of related meanings, frequency of occurrence, and the lexicon. *Cognitive Psychology New York, N. Y.* 13 (2):278-305.
- Jastrzemski, J. E., and R. F. Stanners. 1975. Multiple word meanings and lexical search speed. *Journal of Verbal Learning and Verbal Behavior* 14 (5):534-537.
- Johnson, M. 1987. *The Body in the Mind: The Bodily Basis of Meaning, Reason and Imagination*. Chicago, Illinois: The University of Chicago Press.
- . 1997. Embodied meaning and cognitive science. In *Language Beyond Postmodernism: Saying and Thinking in Gendlin's Philosophy*, edited by D. M. Levin. Evanston, Illinois: Northwestern University Press. 148-175.
- Kamp, H., and B. Partee. 1995. Prototype theory and compositionality. *Cognition* 57:129-191.
- Kaplan, D. 1989 [1977]. Demonstratives. In *Themes from Kaplan*, edited by J. Almog, J. Perry and H. Wettstein. Oxford: Oxford University Press. 481-564.
- Katz, J. J. 1972. *Semantic Theory*. New York: Harper & Row.
- . 1976. A hypothesis about the uniqueness of natural language. *Annals of the New York Academy of Sciences* 280 (1):33-41.
- . 1981. *Language and Other Abstract Objects*. Oxford: Blackwell.

- Katz, J. J., and J. A. Fodor. 1963. The structure of a semantic theory. *Language* 39 (2):170-210.
- Katz, J. J., and P. M. Postal. 1964. *An Integrated Theory of Linguistic Descriptions*. Cambridge, Massachusetts: The MIT Press.
- Kempson, R. M. 1977. *Semantic Theory*. Cambridge: Cambridge University Press.
- Kennedy, C., and L. McNally. 2010. Color, context, and compositionality. *Synthese* 174 (1):79-98.
- Kilgarriff, A. 1992. *Polysemy*. PhD Thesis, University of Sussex.
- . 1995. Inheriting polysemy. In *Computational Lexical Semantics*, edited by P. Saint-Dizier and E. Viegas. Cambridge: Cambridge University Press. 319-335.
- Kilgarriff, A., and G. Gazdar. 1995. Polysemous relations. In *Grammar and Meaning: Essays in Honour of Sir John Lyons*, edited by F. R. Palmer. Cambridge: Cambridge University Press. 1-25.
- King, J. C., and J. Stanley. 2005. Semantics, pragmatics, and the role of semantic content. In *Semantics Versus Pragmatics*, edited by Z. G. Szabó. Oxford: Oxford University Press. 111-164.
- Kintsch, W. 1974. *The Representation of Meaning in Memory*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- . 2000. Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review* 7 (2):257-266.
- . 2001. Predication. *Cognitive Science* 25:173-202.
- Kjøll, G. forthcoming. *Words, Concepts and Ontology from the Perspective of Radical Pragmatics and Semantic Externalism*. PhD Thesis, University of Oslo.
- Klein, D. E., and G. L. Murphy. 2001. The representation of polysemous words. *Journal of Memory and Language* 45 (2):259-282.
- . 2002. Paper has been my ruin: conceptual relations of polysemous senses. *Journal of Memory and Language* 47 (4):548-570.
- Klepousniotou, E. 2002. The processing of lexical ambiguity: Homonymy and polysemy in the mental lexicon. *Brain and Language* 81 (1-3):205-223.
- . 2007. Reconciling linguistics and psycholinguistics: On the psychological reality of linguistic polysemy. In *The Cognitive Basis of Polysemy: New Sources of*

- Evidence for Theories of Word Meaning*, edited by M. Rakova, G. Petho and C. Rakosi. Frankfurt: Peter Lang Verlag. 17-46.
- Klepousniotou, E., and S. R. Baum. 2007. Disambiguating the ambiguity advantage effect in word recognition: An advantage for polysemous but not homonymous words. *Journal of Neurolinguistics* 20 (1):1-24.
- Klepousniotou, E., D. Titone, and C. Romero. 2008. Making sense of word senses: the comprehension of polysemy depends on sense overlap. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34 (6):1534-1543.
- Kreitzer, A. 1997. Multiple levels of schematization: a study in the conceptualization of space. *Cognitive Linguistics* 8 (4):291-325.
- Krifka, M. 1995. Common nouns: a contrastive analysis of Chinese and English. *The generic book*:398-411.
- Kövecses, Z. 1986. *Metaphors of Anger, Pride and Love: A Lexical Approach to the Structure of Concepts*. Amsterdam: John Benjamins.
- Kövecses, Z., and G. Radden. 1998. Metonymy: developing a cognitive linguistic view. *Cognitive Linguistics* 9 (1):37-77.
- Lakoff, G. 1970. *Irregularity in Syntax*. New York: Holt, Rinehart & Winston.
- . 1976. Toward generative semantics. In *Notes from the Linguistic Underground*, Vol. 7, edited by J. D. McCawley. New York: Academic Press. 43-61.
- . 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago, Illinois: The University of Chicago Press.
- . 1993. The contemporary theory of metaphor. In *Metaphor and Thought*, edited by O. A. Cambridge: Cambridge University Press. 202-251.
- . 1999 [1987]. Cognitive models and prototype theory. In *Concepts: Core Readings*, edited by E. Margolis and S. Laurence. Cambridge, Massachusetts: The MIT Press. 391-421.
- Lakoff, G., and M. Johnson. 1980. *Metaphors We Live by*. Chicago, Illinois: The University of Chicago Press.
- . 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.
- Lakoff, G., and M. Turner. 1989. *More Than Cool Reason: A Field Guide to Poetic Metaphor*. Chicago, Illinois: The University of Chicago Press.

- Langacker, R. W. 1984. Active zones. *Proceedings of the annual meeting of the Berkeley Linguistics Society* 10:172-188.
- . 1987. *Foundations of Cognitive Grammar, Vol. 1: Theoretical Prerequisites*. Stanford: Stanford University Press.
- . 1988. A usage-based model. In *Topics in Cognitive Linguistics*, edited by B. Rudzka-Ostyn. Amsterdam, Philadelphia: John Benjamins. 127-161.
- . 1991. *Concept, Image, and Symbol: The Cognitive Basis of Grammar*. Berlin/New York: Mouton de Gruyter.
- . 1993. Reference-point constructions. *Cognitive Linguistics* 4 (1):1-38.
- . 1999. *Grammar and Conceptualization*. Berlin: Mouton de Gruyter.
- . 2008. *Cognitive Grammar: A Basic Introduction*. Oxford: Oxford University Press.
- Lascarides, A., and A. Copestake. 1998. Pragmatics and word meaning. *Journal of Linguistics* 34:387-414.
- Lascarides, A., A. Copestake, and T. Briscoe. 1996. Ambiguity and coherence. *Journal of Semantics* 13 (1):41-66.
- Lasersohn, P. 1999. Pragmatic halos. *Language* 75 (3):522-551.
- Laurence, S., and E. Margolis. 1999. Concepts and cognitive science. In *Concepts: Core Readings*, edited by E. Margolis and S. Laurence. Cambridge, Massachusetts: The MIT Press. 3-81.
- Leech, G. 1990 [1974]. *Semantics: The Study of Meaning*. Harmondsworth: Penguin Books.
- Leibniz, G. W. 1996 [1765]. *New Essays on Human Understanding*. Edited by P. Remnant and J. F. Bennett, *Cambridge Texts in the History of Philosophy*. Cambridge: Cambridge University Press.
- Leslie, A. M. 1987. Pretense and representation: The origins of “theory of mind.”. *Psychological Review* 94 (4):412-426.
- Levin, B. 1993. *English verb classes and alternations: A preliminary investigation*. Chicago, Illinois: The University of Chicago Press.
- Levin, B. 1985. Lexical semantics in review: An introduction. In *Lexical Semantics in Review: Lexicon Project Working Papers, Vol. 1*, edited by B. Levin. Cambridge, Massachusetts: MIT Center for Cognitive Science.

- Levinson, S. C. 2000. *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge, Massachusetts: The MIT Press.
- Lewis, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8 (1):339-359.
- . 2002 [1969]. *Convention: A Philosophical Study*. Oxford: Blackwell.
- Li, P., D. Barner, and B. H. Huang. 2008. Classifiers as count syntax: Individuation and measurement in the acquisition of Mandarin Chinese. *Language Learning and Development* 4 (4):249-290.
- Link, G. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. In *Meaning, Use and the Interpretation of Language*, edited by R. Bäuerle, C. Schwarze and A. von Stechow. Berlin: Walter de Gruyter. 303-323.
- Locke, J. 1975 [1689]. *An Essay concerning Human Understanding* Edited by P. H. Nidditch. Oxford: Oxford University Press.
- Lyons, J. 1977a. *Semantics: Volume 1*. Cambridge: Cambridge University Press.
- . 1977b. *Semantics: Volume 2*. Cambridge: Cambridge University Press.
- Macnamara, J. 1982. *Names for Things: A Study of Human Learning*. Cambridge, Massachusetts: The MIT Press.
- Margolis, E. 1998. How to acquire a concept. *Mind and Language* 13 (347-369).
- Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman.
- McCawley, J. D. 1968. Lexical insertion in a transformational grammar without deep structure. In *Papers from the Fourth Meeting of the Chicago Linguistic Society*, edited by B. J. Darden, C. J. Bailey and A. Davison. Chicago, Illinois: Department of Linguistics, The University of Chicago. 71-80.
- McElree, B., S. Frisson, and M. J. Pickering. 2006. Deferred interpretations: Why starting Dickens is taxing but reading Dickens isn't. *Cognitive Science: A Multidisciplinary Journal* 30 (1):181-192.
- McElree, B., Matthew J., M. J. Traxler, M. J. Pickering, R. E. Seely, and R. Jackendoff. 2001. Reading time evidence for enriched composition. *Cognition* 78:17-25.

- McGlone, M. S. 2001. Concepts as metaphors. In *Understanding Figurative Language: From Metaphors to Idioms*, edited by S. Glucksberg. New York: Oxford University Press.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, Massachusetts: The MIT Press.
- . 1987. What Peter thinks when he hears Mary speaks. In *Précis of Relevance: Communication and Cognition. Behavioral and Brain Sciences* 10:725-726.
- . 2005. *Language: A Biological Model*. Oxford: Oxford University Press.
- Millis, M. L., and S. B. Button. 1989. The effect of polysemy on lexical decision time: Now you see it, now you don't. *Memory & Cognition* 17 (2):141-147.
- Morgan, J. 1978. Two types of convention in indirect speech acts. In *Syntax and semantics, Vol. 9: Pragmatics*, edited by P. Cole. New York: Academic Press. 261-280.
- Murphy, G. L. 1996. On metaphoric representation. *Cognition* 60:173-204.
- . 1997a. Polysemy and the creation of novel word meanings. In *Creative Thought: An Investigation of Conceptual Structures and Processes*, edited by T. B. Ward, S. M. Smith and J. Vaid. Washington DC: American Psychological Association. 235-265.
- . 1997b. Reasons to doubt the present evidence for metaphoric representation. *Cognition* 62:99-108.
- . 2002. *The Big Book of Concepts*. Cambridge, Massachusetts: The MIT Press.
- Nerlich, B. 2003. Polysemy: past and present. In *Polysemy: Flexible Patterns of Meaning in Mind and Language*, edited by B. Nerlich, Z. Todd, V. Herman and D. D. Clarke. Berlin: Mouton de Gruyter. 49-76.
- Nerlich, B., Z. Todd, V. Herman, and D. D. Clarke. 2003. *Polysemy: Flexible Patterns of Meaning in Mind and Language*. Berlin: Mouton de Gruyter.
- Noh, E.-J. 2000. *Metarepresentation: A relevance-theory approach*. Amsterdam: John Benjamins.
- Norrick, N. R. 1981. *Semiotic Principles in Semantic Theory*. Amsterdam: John Benjamins.
- Nunberg, G. 1978. *The Pragmatics of Reference*. Bloomington, Indiana: Indiana University Linguistics Club.

- . 1979. The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy* 3 (2):143-184.
- . 1993. Indexicality and deixis. *Linguistics and Philosophy* 16:1-43.
- . 1996. Transfers of meaning. In *Lexical Semantics: The Problem of Polysemy*, edited by J. Pustejovsky and B. Boguraev. Oxford: Clarendon Press. 109-132.
- . 2004. The pragmatics of deferred interpretation. In *Handbook of Pragmatics*, edited by L. R. Horn and G. L. Ward. Oxford: Blackwell. 344-364.
- Nunberg, G., and A. Zaenen. 1992. Systematic polysemy in lexicology and lexicography. In *Proceedings of Euralex II*, edited by H. Tommola, K. Varantola, T. Salmi-Tolonen and J. Schopp. Tampere: University of Tampere. 387-398.
- Osherson, D. N., and E. E. Smith. 1981. On the adequacy of prototype theory as a theory of concepts. *Cognition* 9 (35-58).
- Ostler, N., and B. T. S. Atkins. 1992. Predictable meaning shift: Some linguistic properties of lexical implication rules. In *Lexical Semantics and Knowledge Representation, Proceedings of the First SIGLEX Workshop*, edited by J. Pustejovsky and S. Bergler. Berlin: Springer. 87-100.
- Owen, G. E. L. 1960. Logic and metaphysics in some earlier works of Aristotle. In *Plato and Aristotle in the Mid-Fourth Century*, edited by I. During and G. E. L. Owen. Göteborg: Almqvist and Wiksell. 163-190.
- Panther, K., and G. Radden, eds. 1999. *Metonymy in Language and Thought*. Amsterdam: John Benjamins.
- Papafragou, A. 1996. On metonymy. *Lingua* 99:169-195.
- . 2005. Relations between language and thought: Individuation and the count/mass distinction. In *Handbook of Categorization in Cognitive Science*, edited by H. Cohen and C. Lefebvre. Amsterdam: Elsevier. 256-277.
- Peirce, C. S. 1955. Logic as semiotic: The theory of signs. In *Philosophical Writings of Peirce*, edited by J. Buchler. New York: Dover Press. 98-119.
- Peirsman, Y., and D. Geeraerts. 2006. Metonymy as a prototypical category. *Cognitive Linguistics* 17 (2):269-316.
- Pelletier, F. J. 1975. Non-singular reference: some preliminaries. *Philosophia* 5 (4):451-465.

- Pelletier, F. J., and L. K. Schubert. [1989] 2001. Mass expressions. In *Handbook of Philosophical Logic. Vol. 10*, edited by D. Gabbay and F. Guentner. Dordrecht: Reidel, 2nd edition. 249-335.
- Pickering, M. J., and S. Frisson. 2001. Processing ambiguous verbs: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27 (2):556-573.
- Pickering, M. J., B. McElree, and M. J. Traxler. 2005. The difficulty of coercion: A response to de Almeida. *Brain and Language* 93:1-9.
- Pietroski, P. M. 2008. Minimalist meaning, internalist interpretation. *Biolinguistics* 2 (4):317-341.
- Pinker, S. 1984. *Language Learnability and Language Development*. Cambridge, Massachusetts: Harvard University Press.
- . 1989. *Learnability and Cognition: the Acquisition of Argument Structure*. Cambridge, Massachusetts: The MIT Press.
- . 2007. *The Stuff of Thought: Language as a Window into Human Nature*. New York: Viking Penguin.
- Plato. 1987. *Theaetetus*. London: Penguin Books.
- Predelli, S. 2005. *Contexts. Meaning, Truth, and the Use of Language*. Oxford: Clarendon Press.
- Premack, D, and G Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1:515-526.
- Prinz, J. J. 2002. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, Massachusetts: The MIT Press.
- Pritchard, T. 2009. *The Pattern of Word Meaning: How Words Enable us to Express Content*. PhD Thesis, Kings College London.
- Pustejovsky, J. 1991. The generative lexicon. *Computational Linguistics* 17 (4):409-441.
- . 1995a. *The Generative Lexicon*. Cambridge, Massachusetts: The MIT Press.
- . 1995b. Linguistic constraints on type coercion. In *Computational Lexical Semantics*, edited by P. St. Dizier and E. Viegas. Cambridge: Cambridge University Press. 71-97.
- . 1998a. Generativity and explanation in semantics. Reply to Fodor and Lepore. *Linguistic Inquiry* 29:289-311.

- . 1998b. The semantics of lexical underspecification. *Folia Linguistica* 32 (3-4):323-347.
- Pylkkänen, L., R. Llinás, and G. L. Murphy. 2006. The representation of polysemy: MEG evidence. *Journal of Cognitive Neuroscience* 18 (1):97-109.
- Pylkkänen, L., A. Stringfellow, E. Flagg, and A. Marantz. 2000. A neural response sensitive to repetition and phonotactic probability: MEG investigations of lexical access. Paper read at Proceedings of Biomag 2000, 12th International Conference on Biomagnetism, at Espoo, Finland.
- Quine, W. V. O. 1960. *Word and Object*. Cambridge, Massachusetts: The MIT Press.
- . 1999 [1953]. Two dogmas of empiricism. In *Concepts: Core Readings*, edited by E. Margolis and S. Laurence. Cambridge, Massachusetts: The MIT Press. 153-169.
- Radden, G. 2000. How metonymic are metaphors? In *Metaphor and Metonymy at the Crossroads: A Cognitive Perspective* edited by A. Barcelona. Berlin: Mouton de Gruyter. 93-108.
- Radden, G., and Z. Kövecses. 1999. Towards a theory of metonymy. In *Metonymy in Language and Thought*, edited by K.-U. Panther and G. Radden. Amsterdam: John Benjamins Publishing. 17-59.
- Ravin, Y., and C. Leacock. 2000. Polysemy: An overview. In *Polysemy: Theoretical and Computational Approaches*, edited by Y. Ravin and C. Leacock. Oxford: Oxford University Press. 1-29.
- Recanati, F. 1987. Contextual dependence and definite descriptions. *Proceedings of the Aristotelian Society* 87:57-73.
- . 1989. The pragmatics of what is said. *Mind & Language* 4:295-329.
- . 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- . 1994. Contextualism and anti-contextualism in the philosophy of language. In *Foundations of Speech Act Theory*, edited by S. Tsohatzidis. London and New York: Routledge. 156-166.
- . 1995. The alleged priority of literal interpretation. *Cognitive Science* 19:207-232.
- . 2002a. Does linguistic communication rest on inference? *Mind & Language* 17 (1/2):105-126.
- . 2002b. Unarticulated constituents. *Linguistics and Philosophy* 25 (3):299-345.

- . 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- . 2007. Reply to Carston. In *Saying, Meaning and Referring. Essays on François Recanati's Philosophy of Language*, edited by M. J. Frápolli and R. Carston. Basingstoke: Palgrave Macmillan. 49-54.
- Rey, G. 1993. The unavailability of what we mean: A reply to Quine, Fodor and Lepore. In *Holism: A Consumer Update*, edited by J. A. Fodor and E. Lepore. Atlanta: Rodopi. 61-101.
- Rips, L. J., E. J. Shoben, and E. E. Smith. 1973. Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behaviour* 12:1-20.
- Rodd, J., G. Gaskell, and W. Marslen-Wilson. 2002. Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language* 46 (2):245-266.
- Romero, E., and B. Soria. 2005. "I" as a pure indexical and metonymy as language reduction. In *Modeling and Using Context: CONTEXT 2005, Lecture Notes in Artificial Intelligence*, edited by A. Dey. Berlin, Heidelberg: Springer. 436-450.
- . 2010. Phrasal pragmatics in Robyn Carston's programme. In *Explicit Communication: Robyn Carston's Pragmatics*, edited by B. Soria and E. Romero. Basingstoke: Palgrave Macmillan.
- Rosch, E. 1973. On the internal structure of perceptual and semantic categories. In *Cognitive Development and the Acquisition of Language*, edited by T. Moore. New York: Academic Press. 111-144.
- . 1975a. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General* 104 (3):192-233.
- . 1975b. Universals and cultural specifics in human categorisation. In *Cross-Cultural Perspectives on Learning*, edited by R. Brislin, S. Bochner and W. Lonner. New York: Halsted Press. 177-206.
- . 1999 [1978]. Principles of categorization. In *Concepts: Core Readings*, edited by E. Margolis and S. Laurence. Cambridge, Massachusetts: The MIT Press. 189-206.
- Rosch, E., and C. B. Mervis. 1975. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 4:573-605.

- Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem. 1976. Basic objects in natural categories. *Cognitive Psychology* 8:382-439.
- Rubenstein, H., S. S. Lewis, and A. M. Rubenstein. 1971. Homographic entries in the internal lexicon: Effects of systematicity and relative frequency of meanings. *Journal of Verbal Learning and Verbal Behavior* 10:57-62.
- Ruhl, C. 1989. *On Monosemy: A Study in Linguistic Semantics*. Albany, New York: State University of New York Press.
- Sag, I. 1981. Formal semantics and extralinguistic context. In *Radical Pragmatics*, edited by P. Cole. New York: Academic Press.
- Sandra, D. 1998. What linguists can and can't tell you about the mind: A reply to Croft. *Cognitive Linguistics* 9 (4):361-378.
- Saussure, F. de. 1974. *Course in General Linguistics*. Translated from the French (1916) by Wade Baskin. Glasgow: Collins.
- Schiffer, S. R. 1972. *Meaning*. New York: Oxford University Press.
- Schwanenflugel, P. J., K. K. Harnishfeger, and R. W. Stowe. 1988. Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language* 27 (5):499-520.
- Searle, J. R. 1978. Literal meaning. *Erkenntnis* 13 (1):207-224.
- . 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- . 1992. *The Rediscovery of the Mind*. Cambridge, Massachusetts: The MIT Press.
- Seto, K. 2003. Metonymic polysemy and its place in meaning extension. In *Polysemy: Flexible Patterns of Meaning in Mind and Language*, edited by B. Nerlich, Z. Todd, V. Herman and D. D. Clarke. Berlin: Mouton de Gruyter. 195-214.
- Seyfarth, R. M., and D. L. Cheney. 2003. Signalers and receivers in animal communication. *Annual Review of Psychology* 54 (1):145-173.
- Seyfarth, R. M., D. L. Cheney, and P. Marler. 1980. Monkey responses to three different alarm calls: evidence of predator classification and semantic communication. *Science* 210 (4471):801.
- Shannon, C. E., and W. Weaver. 1949. *The Mathematical Theory of Communication*. Urbana, Illinois: University of Illinois Press.

- Shields, C. 1999. *Order in in Multiplicity: Homonymy in the Philosophy of Aristotle*. Oxford: Oxford University Press.
- . 2009. Aristotle. In *The Stanford Encyclopedia of Philosophy (Winter 2009 Edition)*, edited by E. N. Zalta. URL: <http://plato.stanford.edu/archives/win2009/entries/aristotle/>.
- Smith, E. E., and D. L. Medin. 1981. *Categories and Concepts*. Cambridge, Massachusetts: Harvard University Press.
- Smith, E. E., and D. N. Osherson. 1984. Conceptual combination with prototype concepts. *Cognitive Science* (8):337-361.
- Smith, E. E., D. N. Osherson, L. Rips, and M. Keane. 1999 [1988]. Combining prototypes: A selective modification model. In *Concepts: Core Readings*, edited by E. Margolis and E. Laurence. Cambridge, Massachusetts: The MIT Press. 355-390.
- Smith, E. E., E. Shoben, and L. Rips. 1974. Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review* 81 (3):214-241.
- Spelke, E. 2000. Core knowledge. *American Psychologist* 55:1233-1243.
- Spelke, E. S., G. Katz, S. E. Purcell, S. M. Ehrlich, and K. Breinlinger. 1994. Early knowledge of object motion: Continuity and inertia. *Cognition* 51 (2):131-176.
- Sperber, D. 1994a. The modularity of thought and the epidemiology of representations. In *Mapping the Mind: Domain Specificity in Cognition and Culture*, edited by L. A. Hirschfield and S. A. Gelman. Cambridge: Cambridge University Press. 39-67.
- . 1994b. Understanding verbal understanding. In *What Is Intelligence?*, edited by J. Khalfa. Cambridge: Cambridge University Press. 179-198.
- . 1997. Intuitive and reflective beliefs. *Mind & Language* 12 (1):67-83.
- . 2000. Metarepresentations in an evolutionary perspective. In *Metarepresentations: A Multidisciplinary Perspective*, edited by D. Sperber. Oxford: Oxford University Press. 117-137.
- . 2007 [1975]. Rudiments of cognitive rhetoric. *Rhetoric Society Quarterly* 37 (4):361-400.

- Sperber, D., and D. Wilson. 1982. Mutual knowledge and relevance in theories of comprehension. In *Mutual Knowledge*, edited by N. V. Smith. London: Academic Press. 61- 87.
- . 1985/1986. Loose talk. *Proceedings of the Aristotelian Society LXXXVI*. In *Pragmatics: A Reader*, edited by S. Davis. 1991. Oxford: Oxford University Press. 153-171.
- . 1986/1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- . 1987. Précis of *Relevance: Communication and Cognition*. *Behavioural and Brain Sciences* 10:697-710.
- . 1998. The mapping between the mental and the public lexicon. In *Language and Thought*, edited by P. Carruthers and J. Boucher. Cambridge: Cambridge University Press. 184-200.
- . 2002a. Pragmatics, modularity and mind-reading. *Mind & Language* 17:3-23.
- . 2002b. Pragmatics, modularity and mind-reading. *Mind & Language* 17 (1/2):3-23.
- . 2008. A deflationary account of metaphor. In *The Cambridge Handbook of Metaphor and Thought*, edited by R. W. Gibbs. Cambridge: Cambridge University Press. 84-108.
- Stainton, R. J. 1994. Using non-sentences: An application of Relevance Theory. *Pragmatics and Cognition* 2:269-284.
- Stallard, D. 1993. Two kinds of metonymy. In *Proceedings of the 31st Annual Meeting of the Association for Computational Linguistics*. Columbus, Ohio: Association for Computational Linguistics. 87-94.
- Stanley, J. 2000. Context and logical form. *Linguistics and Philosophy* 23 (4):391-434.
- . 2005. Semantics in context. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by G. Preyer and G. Peter. Oxford: Oxford University Press. 221-254.
- Stern, J. 2000. *Metaphor in Context*. Cambridge, Massachusetts: The MIT Press.
- . 2006. Metaphor, literal, literalism. *Mind and Language* 21 (3):243-279.
- Struhsaker, T. T. 1967. Auditory communication among vervet monkeys (*Cercopithecus aethiops*). *Social Communication Among Primates*:281-324.

- Sweetser, E. 1990. *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.
- Szabó, Z. G. 2001. Adjectives in context. In *Perspectives on Semantics, Pragmatics and Discourse: A Festschrift for Ferenc Kiefer*, edited by I. Kenesei and R. Harnish. Amsterdam: John Benjamins. 119-146.
- Taylor, J. R. 1989/2003. *Linguistic Categorization. Third Edition*. Oxford: Oxford University Press.
- . 2003. Polysemy's paradoxes. *Language Sciences* 25 (6):637-655.
- . 2006. Polysemy and the lexicon. In *Cognitive Linguistics: Current Applications and Future Perspectives*, edited by K. G., M. Achard, R. Dirven and F. J. Ruiz de Mendoza Ibáñez. Berlin: Mouton de Gruyter. 51-80.
- Tendahl, M., and R. W. Gibbs. 2008. Complementary perspectives on metaphor: Cognitive linguistics and relevance theory. *Journal of Pragmatics* 40:1823-1864.
- Tomasello, M. 2008. *Origins of Human Communication*. Cambridge, Massachusetts: The MIT Press.
- Traugott, E. C., and R. B. Dasher. 2002. *Regularity in Semantic Change*. Cambridge: Cambridge University Press.
- Travis, C. 1981. *The True and the False: the Domain of the Pragmatic*. Amsterdam: John Benjamins.
- . 1985. On what is strictly speaking true. *Canadian Journal of Philosophy Edmonton* 15 (2):187-229.
- . 1997. Pragmatics. In *A Companion to the Philosophy of Language*, edited by B. Hale and C. Wright. Oxford: Blackwell. 87-107.
- Traxler, M. J., M. J. Pickering, and B. McElree. 2002. Coercion in sentence processing: evidence from eye-movements and self-paced reading. *Journal of Memory and Language* 47:530-547.
- Tuggy, D. 1993. Ambiguity, polysemy, and vagueness. *Cognitive linguistics* 4 (3):273-290.
- Tyler, A., and V. Evans. 2001. Reconsidering prepositional polysemy networks: The case of *over*. *Language* 77 (4):724-765.
- . 2003. *The Semantics of English Prepositions*. Cambridge: Cambridge University Press.

- Ullmann, S. 1962. *Semantics: An Introduction to the Science of Meaning*. Oxford: Blackwell.
- Vega-Moreno, R. 2004. Metaphor interpretation and emergence. *UCL Working Papers in Linguistics* 16:297-322.
- . 2007. *Creativity and Convention: The Pragmatics of Everyday Figurative Speech*. Amsterdam: John Benjamins.
- Vicente, A., and F. Martínez Manrique. 2008. Lexical meaning and concepts: A defence of decompositionalism. Unpublished ms.
- . 2010. On relevance theory's atomistic commitments. In *Explicit Communication: Robyn Carston's Pragmatics*, edited by B. Soria and C. Romero. Basingstoke: Palgrave.
- von Fintel, K., and L. Matthewson. 2008. Universals in semantics. *Linguistic Review* 25 (1/2):139-201.
- von Frisch, K. 1976. *The Dance Language and Orientation of Bees*. Cambridge, Massachusetts: Belknap Press of Harvard University Press.
- Vygotsky, L. 1986. *Thought and Language*. Translated from the Russian (1934) by Alex Kozulin. Cambridge, Massachusetts: The MIT Press.
- Ward, G. 2004. Equatives and deferred reference. *Language* 80 (2):262-289.
- Ware, R. X. 1975. Some bits and pieces. *Synthese* 31 (3):379-393.
- Weinreich, U. 1964. *Webster's Third: A critique of its semantics*. *International Journal of American Linguistics* 30:405-409.
- . 1966. Explorations in semantic theory. In *Current Trends in Linguistics, Vol. 3*, edited by T. A. Sebeok. The Hague: Mouton de Gruyter. 395-477.
- Wettstein, HK. 1979. Indexical reference and propositional content. *Philosophical Studies* 36 (1):91-100.
- Wharton, T. 2009. *Pragmatics and Non-Verbal Communication*. Cambridge: Cambridge University Press.
- Wierzbicka, A. 1996. *Semantics: Primes and universals*. Oxford: Oxford University Press.
- Willems, K. 2006. Logical polysemy and variable verb valency. *Language Sciences* 28:580-603.

- Williams, J. N. 1992. Processing polysemous words in context: Evidence for interrelated meanings. *Journal of Psycholinguistic Research* 21 (3):193-218.
- Williamson, T. 1994. *Vagueness*. New York: Routledge.
- Wilson, D. 2003. Relevance and lexical pragmatics. *Italian Journal of Linguistics* 15:273-292.
- Wilson, D. 1998. Linguistic structure and inferential communication. In *Proceedings of the 16th International Congress of Linguistics (Paris, 20-25 July 1997)*, edited by B. Caron. Amsterdam: Elsevier Sciences.
- . 2000. Metarepresentation in linguistic communication. In *Metarepresentations: A Multidisciplinary Perspective*, edited by D. Sperber. Oxford: Oxford University Press. 411-48.
- . 2005. New directions for research on pragmatics and modularity. *Lingua* 115:1129-1146.
- . 2008. Relevance theory and pragmatics. Course in relevance theory and pragmatics at the CSMN - Centre for the Study of Mind in Nature, University of Oslo, October 2008.
- . 2009. Parallels and differences in the treatment of metaphor and Relevance theory and Cognitive Linguistics. *Studies in Pragmatics* 11:42-60.
- Wilson, D., and R. Carston. 2006. Metaphor, relevance and the 'emergent property' issue. *Mind & Language* 21 (3):404-433.
- . 2007. A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts. In *Pragmatics*, edited by N. Burton-Roberts. London: Palgrave. 230-259.
- Wilson, D., and D. Sperber. 1981. On Grice's theory of conversation. In *Conversation and Discourse*, edited by P. Werth. London: Croom Helm. 155-179.
- . 1993. Linguistic form and relevance. *Lingua* 90:1-25.
- . 2002. Truthfulness and relevance. *Mind* 111 (443):583-632.
- . 2004. Relevance theory. In *The Handbook of Pragmatics*, edited by L. R. Horn and G. L. Ward. Oxford: Blackwell. 607-632.
- Wimmer, H., and J. Perner. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13 (1):103-128.

- Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.
- Žegarac, V. 2006. Believing in: A pragmatic account. *Lingua* 116 (10):1703-1721.
- Zlatev, J. 2003. Polysemy or generality? Mu. In *Cognitive Approaches to Lexical Semantics*, edited by H. Cuyckens, R. Dirven and J. R. Taylor. Berlin: Mouton de Gruyter. 447-494.
- Zwicky, A., and J. Sadock. 1975. Ambiguity tests and how to fail them. In *Syntax and Semantics, Volume 4*, edited by J. P. Kimball. New York: Academic Press. 1-36.