

3 Redefining Lexical Semantics and Pragmatics

The aim of this chapter is to arrive at a better understanding of lexical semantics and pragmatics. The main challenge addressed in this book is that of pinning down exactly what constitutes the content of lexical items and how this content is exploited in context. While the notion of *concept* is often used, its complex nature is at the origin of debate across and within different theoretical frameworks. Construction Grammar and Relevance Theory, for instance, have been developed on the basis of opposite understandings of what concepts are and how they contribute to the interpretation of an utterance. Yet more recent developments in RT lead me to believe that the two approaches might not be incompatible. In fact, I intend to show that combining insights from the two frameworks provides an interesting view of the semantics and pragmatics of lexical elements.

The perspectives on semantics adopted in CxG and RT have long been diametrically opposed (e.g. Fodor, 1998; Levine and Bickhard, 1999). On the one hand, RT rests on the Fodorian assumption that concepts are necessarily atomic, whereas in CxG it is argued that concepts are encyclopedic in nature. When combining the two theories, this divergence need not (arguably) be a challenge, however. One could simply decide to focus on those distinctive aspects of each theory for which the other crucially lacks an explanation (for instance, lexical pragmatics in CxG, and constructional semantics in RT). Yet I am strongly convinced it could be a mistake to do so. Indeed, the more you look into each of the two frameworks, the more you realize that internal developments have been greatly influenced by their respective approaches to semantics, especially in RT. It is therefore essential to address this question of lexical semantics so as to pave the way for a genuine integration of RT and CxG.

The notion of concepts adopted in CxG has been relatively unchallenged within the theory itself, and remains rather stable (see Section 2.1.2.2). In comparison, the status of concepts in RT is more controversial. As pointed out, the picture painted in the previous chapter is a simplified version of a much more complex situation. There is a real debate in the relevance-theoretic literature as to what exactly constitutes the nature and content of concepts and, as a natural consequence, that of ad hoc concepts. I will introduce this

debate in the first sections of this chapter. Eventually, I will argue that the perspective on conceptual content adopted in CxG might provide an interesting alternative to most of the approaches developed within RT. At the same time, it will be shown that the relevance-theoretic approach to lexical pragmatics can also shed new light on the actual function of conceptual content, i.e. on the actual function of lexical semantics, and how this content is exploited in context.

In this book, the terms (lexical) *semantics* and *pragmatics* will be used very regularly. Although these terms are described in different ways in RT and CxG, there is a general agreement that semantics has to do with conventional aspects of meaning whereas pragmatics refers to inferred meanings and inferential processes. This is how I will use these terms in the rest of this book (see Leclercq (2020) for an alternative view, however).³⁹

3.1 On the Nature of Concepts and Ad Hoc Concepts in RT

The difficulty in understanding exactly what constitutes conceptual content in RT dates back (at least) to Sperber and Wilson's ([1986] 1995) use of the term. In *Relevance*, Sperber and Wilson originally treat "concepts as *triples of entries* ... spelling out its logical, lexical and encyclopaedic content" (Sperber and Wilson, 1995: 92; emphasis mine). From this perspective, a concept consists of the combination of those three entries, as shown in Figure 3.1.⁴⁰

In order to understand the implications of such a view of concepts, I will apply this model to the concept CAT, a representation of which is given in Figure 3.2. In accordance with the basic approach outlined in Sperber and Wilson ([1986] 1995), this concept consists of three entries. The lexical entry is composed of the word that is used to express the concept (*cat*), its phonological




Figure 3.1 Concepts (adapted from Sperber and Wilson, 1995: 92)

³⁹ Please note that this chapter focuses only on conceptual meaning. The type of meaning which in RT is identified as *procedural* (see Section 2.2.3.2) will be discussed in the next chapter.

⁴⁰ See Section 2.2.3.1 for a detailed description of each entry.




Figure 3.2 Concept CAT

properties and its morphosyntactic features. In the logical entry a number of inferential rules are stored that enable the use of that concept (e.g. that cats are animals, that cats are mammals, etc.). And finally, the encyclopedic entry stores all of the encyclopedic information (or world knowledge) that directly concerns cats, along with the information that individuals have gradually acquired.

With such a description, the difficulty remains to identify exactly what constitutes the semantics of the lexical item that is used to express the concept. Indeed, Sperber and Wilson argue that “the ‘meaning’ of a word is provided by the associated concept (or, in the case of an ambiguous word, concepts)” (Sperber and Wilson, 1995: 90). Yet, in Figure 3.2, it is clear that what constitutes the concept CAT is the sum of the information found in the lexical, logical and encyclopedic entries. This therefore suggests that the meaning of the lexical item *cat* partly consists of the lexical entry itself. It is unlikely that Sperber and Wilson endorse this view, however. Groefsema (2007: 138) also argues that this is “an undesirable conclusion.”

Given this scenario, either the notion of concept or of meaning has to be redefined. According to Groefsema, there are three possible views that emerge from this observation:

1. The logical and encyclopaedic entries of a concept constitute the content of the concept.
2. Conceptual addresses are simple, unanalysable concepts whose entries do not constitute their content.
3. The logical entry of a concept constitutes the content of that concept, while information in the encyclopaedic entry does not contribute to the content of the concept. The role of the encyclopaedic entry is to contribute to the context in which an utterance encoding the concept is interpreted. (Groefsema, 2007: 139)

In the next sections, I will introduce each of these alternatives, taking into account their respective advantages and limits. In particular, the aim is to identify how much each of these views can account for the notions of monosemy and polysemy, as well as provide a sound basis for the derivation of ad

hoc concepts. It is important to note that these views do not receive equal attention in RT. Concerning the first view, for instance, only Groefsema explicitly argues in favor of it (or, at least, a version of it; see Section 3.1.3). For that reason, I will introduce the first view only after I have presented View 2 and View 3. Both these views indeed receive support from a number of relevance theorists. The aim is to try and identify which of these views is most compatible with the perspective adopted in CxG.

3.1.1 *Concepts as Atoms*

The first solution to the issue identified above is to consider that concepts are not, in fact, triples of entries but rather, as represented in Figure 3.3, consist of a different object (in bold) which itself gives access to the different entries (View 2). This is the perspective that was presented in the previous chapter.

According to this view, the meaning of a lexical entry consists of the atomic concept with which it is associated. The logical and encyclopedic information that a concept also gives access to only serves during the inferential phase of comprehension, and is used to derive ad hoc concepts, explicatures and implicatures (see Section 2.2.3.1). This broadly Fodorian approach to concepts⁴¹ helps to compensate for the difficulty identified above. In this case, the lexical entry is no longer a constitutive element of the concept, and the meaning of the lexical entry can be identified explicitly: it is the atomic concept itself. As mentioned in the previous chapter, this is the view most largely adopted within RT.

Although this view solves the ambiguity inherent in the description given in the previous section, it also raises a number of issues itself, a couple of which were briefly discussed in the previous chapter and will be taken up here. One of




Figure 3.3 Concepts as atoms

⁴¹ Fodor argues that concepts are atomic, but does not give any particular status to logical and encyclopedic information, at least originally (see, for instance, Fodor, 1998). Later he adopted a similar perspective by explicitly placing encyclopedic knowledge within “mental files” (Fodor, 2008: 94).

the very first difficulties with this view is to understand exactly what is considered to constitute the content of an atomic concept. Fodor (1998: 12) argues that the content of concepts consists of their “causal-cum-nomological” relations with the world (i.e. a lawlike causal dependency between the concept and a specific referent in the real world). In Carston’s (2010) words, as mentioned before,

the content or semantics of this entity is its denotation, what it refers to in the world, and the lexical form that encodes it, in effect, inherits its denotational semantics. (Carston, 2010: 245)

From this perspective, the content of a concept is not some internal representation but consists only of the relationship between the concept and the object or property that it refers to in the mind-external world. A number of objections have been raised against this type of referential semantics, in particular by philosophers (e.g. Wittgenstein, 1958; Quine, 1960; Davidson, 1967, 1984; Chomsky, 2000, 2003; Brandom, 2000; Ludlow, 2003; Sztencel, 2012a, 2012b, 2018, *inter alia*). The aim of this section is not to look individually at each of these objections, especially since they have often been presented on the basis of various theoretical commitments to meaning, truth and reference. Rather, the aim is to show that such a view of referential semantics does not align with the more general picture of lexical pragmatics developed in RT (as opposed to what might be expected to be the case) and as a result challenges the relevance-theoretic commitment to conceptual atomism.

3.1.1.1 Relevance Theory and Referential Semantics: Compatible? The content of an atomic concept, according to Fodor, consists in the *necessary* (causal-cum-nomological) relationship that the concept has with a particular object in the mind-external world. One of the questions addressed by Sztencel (2012a) concerns the origins of this necessity, and in particular why it is that concepts refer to the things they do. Specifically, she nicely shows that Fodor’s definition of conceptual content is relatively circular and therefore fails to be fully explanatory (Sztencel, 2012a). Given the relevance-theoretic commitment to Fodor’s approach, this also becomes an issue for RT. Sztencel’s (2012a) argument runs as follows: if concepts necessarily have to refer to (or *lock onto*, to use Fodor’s terminology) a specific referent in the external world, then there must be something internal to concepts (and prior to reference) that forces us to select this specific referent:

The question is: what is it about any particular concept in and of itself that makes it lock onto the things it does lock onto and not other things? The problem is that if primitive concepts are to be in a non-arbitrary relation to what they “lock to” then they must have logically prior internal content independent of, but determinative of, the “locking” relation just like complex concepts. (Sztencel, 2012a: 491)

For instance, what is it about the concept CAT that *necessarily* makes it refer to cats? According to Fodor, however, the only element that composes the content of the concept CAT is the referential relation itself. This is why Sztencel argues that there is circularity in Fodor's view of conceptual meaning: in this account, what can explain reference to cats is the referential relation itself.

The reason this question is directly relevant to RT comes from the particular way in which Fodor tried to handle the issue (although he most probably did not consider it to be an issue). One of the ways in which Fodor could have simply objected to the circularity just mentioned is by arguing that, after all, he considers atomic concepts to be innate (Fodor, 1998: 124). In this case, although the circularity identified above remains, it cannot be used as an argument against Fodor any more. Concepts refer to the things they do in virtue of what they are programmed to refer to. Assuming Fodor is right, "since normal adults command a vocabulary of at least 60,000 words, it would seem that, at a bare minimum, they possess 60,000 *innate* concepts" (Laurence and Margolis, 2002: 28; emphasis mine). It is doubtful that we have (as many as) 60,000 innate concepts and I agree with Churchland (1986: 389) that it is "difficult to take such an idea seriously." Many arguments have been presented in the literature on concepts against the innateness hypothesis (see, among many others, Locke [1690] 1975; Berkeley, 1709; Hume, [1739] 1978; Wittgenstein, 1958; Johnson, 1987; Lakoff, 1988; Putnam, 1988; Elman et al., 1996; Cowie, 1998, 1999; Barsalou, 1999; Levine and Bickhard, 1999; Prinz, 2002; Tomasello, 2003; Sampson, 2005, and the references cited in Section 2.1). There is convincing evidence in the above papers, and I agree with what their authors write. For reasons of space, however, I will not elaborate on those arguments.

The questionable plausibility of the innateness argument is not the main issue here. More generally, to view concepts as innate has a direct consequence for the relevance-theoretic perspective on lexical pragmatics. It is strongly argued in RT that the concepts that are expressed in utterances are not lexically encoded but need to be pragmatically inferred by the hearer in order to derive the intended interpretation (see Section 2.2.3.1). If one considers concepts to be innate, however, it becomes doubtful whether any (ad hoc) atomic concepts have to be inferred and that such inferential mechanisms as described in RT are ever useful. Speakers simply have a number of concepts at their disposal.

Naturally, it could be argued that many of those innate concepts are non-lexicalized and still need to be contextually derived (or selected) in accordance with the principle of relevance. After all, it is often argued in RT that many (supposedly) ad hoc concepts are not purely the product of pragmatics, in the

sense that they are not entirely new but consist of stored concepts that are simply not lexicalized.⁴²

The pragmatic process of inferring ad hoc concepts in utterance interpretation . . . may result in a tokening of one of these stable, albeit non-lexicalised, concepts, already established in the hearer's conceptual system. (Carston, 2010: 251)

Neither Fodor nor relevance theorists have adopted exactly this view, however. On this issue, RT radically departs from Fodor's account. On the one hand, Fodor somehow argues against the possibility that there might be unlexicalized atomic concepts. According to him, there is a strict one-to-one mapping between lexical words and atomic concepts (Fodor, 1998: 55). Discussing different uses of the verb *keep* (*NP kept the money*, *NP kept the crowd happy*, etc.), he argues that the differences in meaning come from the verb's arguments but not from the verb itself: in all instances, *keep* expresses the single atomic concept *KEEP* (p. 52). As a result, as mentioned in Section 2.2.3.1, Fodor also considers that "there is no such thing as polysemy" (1998: 53). In this case, a theory of ad hoc concepts (and more generally of inference) is not required. On the other hand, RT takes a much more nuanced approach. First of all, it has long been argued within RT that there is not a one-to-one mapping between the public and the mental lexicon (i.e. between words and concepts), but that there is a one-to-many mapping (Sperber and Wilson, 1998). That is, individuals also store and use concepts that are not lexicalized. In fact, on the basis of the work of Barsalou, Sperber and Wilson explicitly argue that "the idea that there is an exhaustive, one-to-one mapping between concepts and words is quite implausible" (1998: 185). I also share this view. As a result, even if concepts were innate, a theory of concept inference/selection such as developed in RT is required. This is all the more true since relevance theorists seem to assume that concepts are not innate but are instead acquired or learned. This position is more or less explicit in the work of Carston (2002a) and Wilson and Sperber (2012) and becomes very explicit in the work of Wharton (2004, 2014), who specifically discusses the acquisition of lexical meaning, i.e. conceptual content.⁴³ It therefore also follows that many of the concepts that are communicated are considered to "be quite new" (Carston, 2010: 251) and need to be inferred pragmatically. That is, unlike what is suggested by Fodor, speakers are not innately equipped with a finite set of concepts, but there is an infinite number of concepts that they can use and gradually acquire (hence the need for the comprehension procedure developed in RT).

⁴² In this case, however, it is no longer clear in what sense these concepts are ad hoc (cf. Mioduszewska, 2015: 87).

⁴³ To my knowledge, only Carston assumes that "we are born with at least some innate concepts" (2010: 250). It is unclear whether she still holds this view, however (see Section 3.3.2).

It will start to become clear why the referential approach to conceptual content proposed by Fodor does not fit well with RT, unlike what relevance theorists believe. By rejecting the innateness hypothesis, RT inevitably also has to abandon the referential approach to meaning put forward by Fodor (which they do not). Indeed, without the innateness hypothesis, there is no explanation given for the necessary relationship between a concept and its referent. Naturally, the critics will argue that scientific theorizing often consists in working with hypotheses for which we may as yet lack an explanation, and RT, therefore, does not have to abandon referentialism. Nevertheless, there are (at least) two more reasons why RT might prefer to move away from referentialism.

The first reason follows from the observation that referentialism is inconsistent with one of the most central claims put forward within RT. The notion of *explicatures* was developed within RT to account for the observation that the logical form of an utterance never fully determines the speaker's intended interpretation, which has to be pragmatically derived (see Section 2.2.2). This is referred to as the *underdeterminacy* thesis. And from this perspective, logical forms cannot be defined in terms of truth-conditions:

Various terms for this are used in the literature; the linguistic expression employed is described as providing an incomplete logical form, a 'semantic' skeleton, 'semantic' scaffolding, a 'semantic' template, a proposition/assumption schema (see, for instance, Sperber and Wilson, 1986/95; Recanati, 1993; Bach, 1994b; Taylor, 2001). What all of these different locutions entail is that the linguistic contribution is not propositional, it is not a complete semantic entity, not truth evaluable. (Carston, 2002b: 134)

In RT, it is generally assumed that only the enriched propositions that are derived pragmatically, i.e. explicatures and implicatures, can be described in such truth-conditional terms (Carston and Hall, 2012: 76; see also Moeschler, 2018). Yet, when considering that concepts have referential semantics (by locking onto a specific entity or property in the mind-external world), then they "must have truth-theoretic content" (Sztencel, 2011: 379). By extension, the logical forms in which concepts occur should also be defined in terms of these truth-conditions. Even if the concepts that occur in the logical forms do not correspond exactly to what the speaker intends to communicate, they still have truth-conditions.⁴⁴ It is therefore inconsistent for RT to argue both that concepts have referential content and that logical forms are not truth-evaluable. For this reason, "it has become unclear in what form Relevance Theory still holds the underdeterminacy thesis" (p. 376). As a result, either RT needs to

⁴⁴ Note that, outside RT, a number of researchers have precisely provided descriptions along these lines (e.g. Borg, 2004, 2016; Cappelen and Lepore, 2005, inter alia). According to them, although sentences fail to provide the speaker's intended meaning, the words that a speaker uses do have truth-theoretic content.

change its view concerning the nature of conceptual content, or it needs to revise (or abandon, even) the underdeterminacy thesis (cf. Burton-Roberts, 2005). Some might want to stick with the Fodorian assumption that concepts are necessarily referential. In this case, the difficulty is to determine exactly in what sense logical forms underdetermine the speakers' intended interpretation. It is most likely that relevance theorists will prefer to keep the underdeterminacy thesis as it is currently formulated, and which provides a strong basis on which the rest of the theory has been developed (see, in particular, Section 2.2.2), and rather to reconsider the nature of lexical concepts. I also believe this is the most preferable option. From this perspective, concepts do not have truth-theoretic content, but truth values are only accessible once the explicatures have been derived (and which relevance theorists call *real* (as opposed to *linguistic*) semantics, cf. Clark, 2013a: 299).⁴⁵ Naturally, this requires the explicit spelling out of what constitutes the content of lexical concepts. As we will see in the next sections, there have been a few suggestions in RT.

There is a second important argument against a purely referential account of conceptual content: Fodor himself seems to be aware of the difficulties that such an account of meaning faces. Sztencel (2018) aptly captures the reason why concepts must have some internal content:

If concepts do actually lock onto things in the world, we want to say that they do so non-arbitrarily – in other words, that there is something about the concept itself (some property of the concept, which I am calling its internal content) that determines that it locks onto the things it does lock onto and not anything else. The question is then: should we align 'semantics' with (internal) content or with (external) reference? Having so distinguished between content and reference, it seems reasonable to say that content is metaphysically prior to and a precondition for reference. Insofar as 'semantics' is referential at all, such semanticity derives from, is parasitic on, internal conceptual content. It is arguable, then, that it is internal content that is fundamentally 'semantic'. (Sztencel, 2018: 14)

It appears that Fodor is not a complete stranger to this line of argumentation. After all, he once argued "(a) that the *content* of a linguistic expression should be distinguished from such of its semantic properties as its truth conditions; and (b) that content *is* – though truth conditions are not – a construct out of the communicative intentions of speaker/hearers" (Fodor, 1982: 105–106; original

⁴⁵ Although 'meaning' in CxG (as in cognitive linguistics more generally) is not defined in truth-conditional terms, which is a perspective I largely endorse myself, this second view (i.e. an approach in terms of 'real' semantics) is closer to the approach generally adopted in CxG. Indeed, the way RT relocates truth-values at another, pragmatic, level does not radically differ from the perspective (most probably) adopted in CxG (which, as mentioned earlier, cannot entirely get away without truth-conditions), according to which truth-values are secondary to meaning and are evaluated in context.

emphasis). From this perspective, meaning is in the head and not purely referential.⁴⁶ As mentioned in the previous chapter, the idea that meaning is *in the head* is a view that cognitive linguists in general and construction grammarians in particular have largely adopted. In fact, Chomsky himself also subscribes to this view. He explicitly argues that “the semantic properties of words are used to think and talk about the world in terms of the perspectives made available by the resources of the mind” (Chomsky, 2000: 16). Moreover, Fodor also understands that reference alone (and, with it, truth-conditions) cannot distinguish between the meaning of different lexical items, especially those that are co-referential (e.g. *Superman* and *Clark Kent*; Fodor, 2008: 86). He specifically argues that to distinguish between the two, one also has to associate with the lexical items different *modes of presentation* (Fodor, 1998: 15).⁴⁷

Unfortunately, in spite of these observations, Fodor still clings to the idea that concepts must be atomic, referential items. According to him, as far as conceptual *content* (i.e. meaning) is concerned, neither the mode of presentation nor the mental files (cf. footnote 41) that he introduces are actually relevant: “all that matters is the extension” (Fodor, 2008: 87). Relevance theorists do not have to comply with Fodor, however. It will not be the first time that they part company from him (e.g. on how modular the mind is, on the acquisition of concepts, on inference rules). RT could be stronger as a theory if it were to change its view of conceptual content and adopt a non-referential approach to lexical meaning. The challenge in this case, of course, is to identify exactly what constitutes this content. Is it the inference rules located in the logical entry? This is a possibility that some relevance theorists entertain and which will be introduced in Section 3.1.2. Or is the content of a concept determined by the information stored in the encyclopedic entry? In Section 3.1.3, we will see that Groefsema (2007) favors this position. As mentioned in the previous chapter, this is also a view that is strongly defended in cognitive linguistics and, therefore, in CxG. In fact, it is also worth noting that some philosophers have argued that the information stored within Fodor’s mental files (which more or less correspond to RT’s encyclopedic knowledge) should be considered to be the content of those concepts (e.g. Lee, 2017). In Section 3.4, I will also argue in favor of such a perspective on conceptual content.

3.1.1.2 Ad Hoc Concepts and Atomism: Problems It is important to understand the immediate relevance of the previous section to the question addressed

⁴⁶ Arguing against Hilary Putnam, Fodor (1982: 110) says that he finds “unattractive” a position whereby “meaning is not in the head.”

⁴⁷ Fodor’s modes of presentation closely resembles the notion of construal/vantage point put forward (in different ways) in cognitive linguistics.

in this chapter. In the previous chapter, it was conceptual atomism, and not referentialism, that was considered to be the main issue for RT. Conceptual atomism does not (arguably) presuppose referentialism, or vice versa. It could therefore be argued that, having rejected referentialism, one could still maintain atomism. According to Fodor, however, referentialism and atomism are necessarily related. (As we will see in this section, this is also the case for Carston.) Fodor actually calls himself a “referential atomist” (Fodor, 2008: 99). It has been shown in the previous section, however, that referentialism is generally incompatible with the relevance-theoretic enterprise. One is therefore also entitled to question the necessity (and, even, the possibility) for RT to argue that concepts are atomic. The aim of this section is precisely to show that atomism is also incompatible with the relevance-theoretic approach to lexical pragmatics and that alternative perspectives have to be considered.⁴⁸

Conceptual atomism and the relevance-theoretic approach to pragmatics are incompatible for a number of reasons, all of which are closely related. One of them, however, particularly stands out from the others and seems to be the most problematic case for RT. This issue was briefly addressed in the previous chapter and will be elaborated on here. The difficulty for RT is to explain, on the basis of atomic concepts, how hearers manage to derive ad hoc concepts.⁴⁹ Take the following examples:

- (64) It's clear that *stress* can contribute to chronic disease, but fixing stress is not as simple as taking a deep breath or an occasional yoga class. (NOW)
- (65) During the adoption process in 2015, Avinash told the court that he loves his Taya and Tayi (uncle and aunt) who are now his *parents* and they too love him. (NOW)

⁴⁸ As mentioned several times already, it is important to note that the focus of this book is ‘meaning’. Therefore, when arguing against atomism, I only mean to argue that conceptual content (i.e. the meaning of lexical items) is not atomic. I make absolutely no claim regarding the *form* that concepts take when occurring in thoughts, however. They may well be atomic (say, a conceptual address, or Fodor’s ‘file name’), but I challenge the view that their content consists of this atom.

⁴⁹ The systematic derivation of ad hoc concepts in itself is not at issue here. Barsalou indeed provides solid evidence that individuals systematically create ad hoc categories (cf. footnote 31 for references). This view is also very much in line with most approaches developed in cognitive linguistics, such as Langacker’s notion of conceptualization which he uses to refer to the dynamic nature of conceptual content:

Rather than being fixed, the values of linguistic elements are actively negotiated; and rather than being static, the meanings of complex expressions emerge and develop in discourse. ... Meaning is not identified with concepts but with **conceptualization**, the term being chosen precisely to highlight its dynamic nature. (Langacker, 2008: 30; original emphasis)

The only issue being tackled in this section is the apparent incompatibility between the particular way ad hoc concepts are said to be derived and the atomic commitment to the nature of concepts and ad hoc concepts.

There are various ways in which the lexemes *stress* and *parent* can be understood depending on how they are used by the speaker. It is theoretically possible (in the required contexts, of course, not here) that their interpretation will lead to the recovery of the encoded concepts STRESS and PARENT (whatever these are). This scenario is relatively unproblematic, since in this case the intended interpretation (and therefore the content of the communicated concept) equals that of the lexical (encoded) concept. Difficulties emerge, however, if the hearer has to infer an ad hoc concept. This is the case for the interpretation of the sentences in (64) and (65). In the sentence in (64), the word *stress* can be understood as communicating the more specific (i.e. narrower) concept STRESS*: ‘high/undue levels of stress’. In the sentence in (65), hearers have to infer the more general (i.e. broader) concept PARENT* in which the biological aspect of parenthood is dropped. From the relevance-theoretic standpoint, these two examples illustrate cases of conceptual narrowing and broadening, respectively (see Section 2.2.3.1).

As mentioned in the previous chapter, the difficulty here is to understand how RT can argue in favor of both conceptual atomism and the systematic narrowing and/or broadening of conceptual content. Considering that lexical concepts and ad hoc concepts are atomic, it is unclear in what sense exactly ad hoc concepts can be narrower/broader than lexical concepts. Somehow, these notions necessarily require some form of internal structure that our minds can exploit in different ways (in accordance, of course, with the comprehension procedure). The notion of ad hoc concepts is often introduced in the relevance-theoretic literature in direct reference to the work of Barsalou on ‘ad hoc categories’. Yet, looking carefully at what Barsalou himself argues, it is relatively clear that concepts are not atomic, and that the process of conceptual adjustment is possible only because concepts are considered to be “bodies of knowledge in long-term memory” (Barsalou et al., 1993: 57). Interestingly, in some specific accounts of RT, one could also be led to believe that a similar view is adopted by relevance theorists themselves. After all, they often argue that the derivation of ad hoc concepts crucially depends on the information stored in the encyclopedic entry:

On this approach, *bank* . . . might be understood as conveying not the encoded concept BANK but the ad hoc concept BANK*, with **a more restricted encyclopedic entry** and a narrower denotation. (Wilson and Sperber, 2004: 618; emphasis mine)

In another utterance situation, **different items of encyclopaedic information** about children might be more highly activated making most accessible such implications as that Boris doesn’t earn his keep, expects others to look after him, is irresponsible, etc., resulting in a distinct ad hoc concept CHILD** in the explication. (Carston, 2012: 613; emphasis mine)

In this case, the information stored in the encyclopedic entry provides the structure on the basis of which ad hoc concepts can be derived. In other words, the content of ad hoc concepts more or less resembles that of lexical concepts depending on the degree of overlap between the information stored in their respective encyclopedic entries. Naturally, this strongly suggests that encyclopedic information must be content-constitutive. Yet, this is incompatible with RT's commitment to atomism. This incompatibility has already been pointed out by a number of relevance theorists, among whom is Anne Reboul:

Relevance Theory rests on a Fodorian account of concepts according to which concepts are atomic, hence not definitions. Ad hoc concepts, however, are supposed to be formed by modifying the definition of the original concept by deleting features or introducing them in the definition. This directly contradicts a Fodorian view of concepts. (Reboul, 2014: 20)

In addition to Reboul (2014), this contradiction has also been discussed, among others, by Vicente (2005: 190), Burton-Roberts (2007: 106), Groefsema (2007: 146), Vicente and Martínez-Manrique (2010: 49), Allott and Textor (2012: 198), Assimakopoulos (2012: 23), and Mioduszevska (2015: 83). As a result, either RT has to change its approach to the formation and nature of ad hoc concepts, or it needs to rethink exactly what it considers conceptual content to consist of. From the perspective outlined above, the information stored in the encyclopedic entry (and the logical entry) is a good candidate. In the rest of this chapter, I will strongly argue in favor of this option. Before doing so, it is worth investigating a bit further this incompatibility between atomism and the relevance-theoretic notion of ad hoc concepts.

In spite of the demonstration just made, Carston maintains that she has been “unable to find any arguments supporting the alleged incompatibility” (Carston, 2010: 247) and argues that the underlying thinking that points towards this incompatibility “is quite wide of the mark since the account of ad hoc concept formation is not semantic and not internal to the linguistic system” (p. 247). This last statement has long puzzled me for a number of reasons. First of all, it is doubtful whether anyone (among relevance theorists or not) has ever considered the process of conceptual adjustment to be a purely semantic process. But more importantly, it is unclear why the fact that this is indeed a pragmatic process can be used as an argument against the incompatibility discussed by many. If it were so obvious, then presumably the apparent incompatibility would have been raised less often. Furthermore, regardless of whether one considers conceptual adjustment to be a semantic or a pragmatic process, the type of conceptual narrowing/broadening discussed in RT requires some internal structure to be exploited which the atomic approach does not offer. For this reason, although the incompatibility between the atomic view of concepts and the relevance-theoretic approach to ad hoc concepts formation

strikes many as blatant, it is important to take a closer look at the reasons why Carston assumes that there is no such incompatibility.

In order to understand Carston's comment, one needs to take into account all of her theoretical commitments. Of all relevance theorists, Carston is probably the most Fodorian in her approach to meaning. She is in particular very faithful to the referential approach that he advocates. This is particularly clear from the way in which she uses the term *denotation* (see, for instance, Carston, 1997b, 2002a, 2010, 2012). It is most probably this commitment to referential semantics that can explain Carston's surprise (and rightly so). It is clear to Carston that the information stored in the encyclopedic entry is not content-constitutive but only provides contextual information about concepts, i.e. about their denotation. Although this stored encyclopedic knowledge is indeed exploited during the comprehension phase to recover the intended concept, it does not eventually form the content of this ad hoc concept either. The content of ad hoc concepts is determined, like lexical concepts, solely by their denotation. The encyclopedic and logical information that a lexical concept provides only serves in context to recover the denotation of the communicated ad hoc concept. According to Carston, the level at which conceptual narrowing/broadening matters is not the encyclopedic level (which is not content-constitutive) but the level of the denotation itself. Ad hoc concepts can be said to be narrower/broader than the lexical concepts from which they are derived depending on the degree of overlap between the set of items that fall within their denotation (Carston, 2002a: 353). This is why, according to Carston, it is not incompatible to argue both that ad hoc concepts are atomic (since they consist of a referential relation) and that they can be narrower/broader than the lexical concept from which they are formed (depending on their denotational overlap).⁵⁰

Carston's perspective seems internally coherent and therefore appears to solve the incompatibility mentioned above. The question is whether this view is theoretically and descriptively accurate, however. There are two crucial points that need to be addressed here. First, Carston's account is flawed since it ignores a basic constraint imposed by referential atomism. According to Hall (2011), the atomic account requires that neither the logical nor the encyclopedic types of information made accessible by a concept be "constitutive or reference-determining" (Hall, 2011: 7). Yet, as was just explained, it is clear in Carston's account that the logical and encyclopedic entries directly contribute to establishing reference. As a result, her view suggests (in spite of what she argues) that these two types of information *are* content-constitutive (since content determines reference).⁵¹ This is why her approach to concepts and ad

⁵⁰ That is why Carston (2010: 247) argues that the derivation of ad hoc concepts takes place outside the linguistic system.

⁵¹ That content determines reference seems to be implicitly assumed by those who discuss the incompatibility identified above (see, for instance, Reboul, 2000: 60).

hoc concepts runs counter to the atomic view of conceptual content, which compromises the coherence of her account. Second, Carston's approach is all the more problematic in that her defense of conceptual atomism rests primarily on the assumption that concepts are referential. Besides the issue of determining how reference is established, it was shown in the previous section that referential semantics is incompatible with some of the core notions developed within RT (e.g. the one-to-many mapping between words and concepts, the difference in truth-conditions between the logical form and explicatures, etc.). So Carston advocates for a referential account, but there are a number of critical questions that she fails to address (see Section 3.1.1.1), which further threatens the coherence of the account sketched in the previous paragraph. It is therefore highly questionable whether conceptual atomism can be maintained in RT. Rather, it is preferable to drop conceptual atomism, together with referentialism. It was shown in the previous section that in order to be able to establish reference, one must have some internal content (prior to reference) which itself can be considered to constitute the content of that concept. Carston assumes that the recovery of an ad hoc concept's denotation is largely made possible by the encyclopedic entry made accessible by the lexical concept with which it is associated. From this perspective, the information stored in the encyclopedic entry once more seems to be the best candidate for what constitutes conceptual content. This view, which will be adopted later in this chapter, is further supported by the observation that the account of ad hoc concepts presented in RT was developed on the basis of Barsalou's work on (ad hoc) categories which, as mentioned above, he conceives of as bodies of information (and not as atoms).

It is worth noting that Carston finds Fodor's arguments for conceptual atomism "unassailable" (Carston, 2010: 245) and argues that "the most compelling of these, perhaps, is that no-one has been able, despite centuries of trying, to give adequate definitions for any but a tiny group of words" (p. 245). This viewpoint will be discussed in more detail in Section 3.5. Suffice it to say that it remains dubious why Fodor's arguments are unassailable, in particular the one Carston mentions and which actually seem to be the most inconsequential argument Fodor puts forward.⁵² That no one has ever been able to provide specific definitions for most of the words we use is not bulletproof evidence that conceptual content is necessarily atomic. It might, however, tell us something about how the mind works, as I will suggest in Section 3.5. But Fodor's own skepticism about definitions cannot be used as an argument for atomism.

⁵² As far as the content of concepts is concerned, there is quite a large amount of evidence (see Section 2.1.2.2), starting from the work of Barsalou himself, that shows that conceptual content consists of a structured network of information rather than an atom. Therefore, Fodor's arguments are, in fact, assailable.

The difficulties that an atomic view of conceptual content involves go beyond the challenge represented by the derivation of ad hoc concepts. One further difficulty concerns how conceptual content can be argued to be acquired. If, like Fodor, one assumes concepts to be innate and to be in a one-to-one mapping with the lexicon, then there is no such issue. As mentioned in the previous section, however, relevance theorists tend not to view concepts as innate objects and consider that there are more concepts than just those that are lexicalized. In this case, it is unclear exactly, given the variety of atomic concepts that a lexeme could lock onto, which of those it should actually be associated with.⁵³

Directly related to this issue is the question of monosemy and polysemy. If one considers concepts to be atomic, then it is also not straightforward how polysemy can possibly be represented in one's mind. As mentioned in the previous chapter, indeed, Fodor himself assumes that there is no such thing as polysemy since he considers there to be a one-to-one mapping between words and concepts. Relevance theorists, however, do not adopt such isomorphism. They admit that some words might be (conventionally) polysemous. Yet, considering that concepts are atoms, only two outcomes for the representation of meaning are possible, polysemy not being one of them. In the first case, if one assumes that conventional polysemy exists (although it is obvious to construction grammarians, it is much less so in the pragmatics literature), then the identified meanings cannot be distinguished from cases of homonymy. Indeed, in this case it is opaque how one mentally represents the assumed relationship between the different atoms. It is doubtful, however, that the senses of the lexeme *wood* (e.g. 'material' and 'geographical area') are unrelated in the same way as the senses of the lexeme *bank* are (e.g. 'financial institution' and 'land alongside a river'). Considering that conventional polysemy exists, it is preferable to keep the notion distinct from cases of homonymy. Yet, viewing concepts as atoms does not make this possible. The alternative option is to consider that although lexemes can be used to convey different concepts in different contexts, only one of those concepts is actually encoded by the lexeme which is used to express them. This is what Carston (2013: 187) calls *pragmatic* polysemy, which is a view also defended by Falkum (2011, 2015). Here, although lexemes are considered to be pragmatically

⁵³ Given the aim of this book, I will not go into the detail of this issue here. Nevertheless, it is worth noting that this question has led to quite diverging opinions among the relevance theorists who try to answer the question. On the one hand, for instance, Tim Wharton is "very much into nativist accounts of concept acquisition, in the sense of Bloom 1996, Dehaene 1997 and Margolis 1998" (Tim Wharton, p.c.). In comparison, Sandrine Zufferey argues that frequency must have a key role in the acquisition of concepts (Zufferey, 2010: 146, 2015).

polysemous, they are assumed to be semantically monosemous. I have expressed strong doubts about such a perspective in the previous chapter. In Section 3.3, I will elaborate on these doubts.

3.1.1.3 Concepts and Atomism: Conclusions Conceptual atomism provides an interesting solution to the challenge identified above in terms of what exactly constitutes the meaning of a lexical item. The meaning comes from the referential relation made accessible by the atomic concept. I have argued that this perspective – which is the most adopted in RT – faces a number of problems. The first difficulty concerns the referential nature of this concept. Referential semantics, beyond its own limits, is not compatible with some of the most central tenets of RT. In addition, conceptual atomism does not seem to be compatible with the account of ad hoc concept formation largely discussed within RT. Instead, it has been suggested that an account whereby concepts have internal content is better equipped to face the different challenges that referential atomism faces. Before introducing such an account, it is worth looking at the other accounts of conceptual meaning that have been argued for in RT.

3.1.2 Conceptual Content and the Logical Entry

Another way of dealing with the issue identified at the beginning of this chapter is to consider that the content of concepts, regardless of whether or not they are atomic, (also) consists of the information stored in the logical entry that the concept gives access to. This view is the third possibility discussed by Groefsema (2007) and is illustrated in Figure 3.4.

This possibility can be more or less explicitly found in the work of a number of relevance theorists. Sperber and Wilson (1987), for instance, specifically argue that they “see [the semantic properties of a word] as provided by the logical entry filed at the same address” (Sperber and Wilson, 1987: 741). Later, Carston (2002a) adopts a similar view. She argues that logical information is a defining property of concepts, i.e. it is




Figure 3.4 Logical entry as conceptual content

content-constitutive, and that conceptual narrowing and broadening is located at the level of the logical entry:

An ad hoc concept formed by strengthening a lexical concept seems to involve elevating an encyclopaedic property of the latter to a *logical (or content-constitutive) status* . . . an ad hoc concept formed by the loosening of a lexical concept seems to involve dropping one or more of the *logical or defining properties* of the latter. (Carston, 2002a: 339, emphasis mine)

This view quite explicitly suggests that the logical entry made accessible by a concept provides the meaning of a lexical entry. Yet this comes in complete opposition to (and again seems incompatible with, see Section 3.1.1.2) Carston's strong atomic commitment to conceptual content. More recently, she argues that:

A decompositional view might also seem to have been implied by my talk (Carston 2002a: 339) of the dropping of logical properties (in the case of loose uses) and the promoting of encyclopaedic properties (in the case of narrowing), although this does not strictly follow, since these properties are clearly not internal components of the lexical concepts themselves and need not be taken that way for ad hoc concepts either. In fact, it was my aim then, as now, to maintain a consistently atomic view of concepts if at all possible. (Carston, 2010: 249)

The opposition between these two quotes quite clearly shows the tension that can be found in Carston's own work and more generally within RT. Carston is aware of this tension and explicitly points out (in Carston, 2010) that she should not be understood as arguing that the logical properties are content-constitutive. Nevertheless, this post-hoc clarification may not be entirely convincing for it generally seems that she is struggling to explain the different facets of her approach. Indeed, a few pages before denying that this is the perspective she adopts, Carston (2010) repeats that the inference rules found in the logical entry of a concept "are, crucially, taken to be content constitutive" (Carston, 2010: 246).

Elsewhere in the relevance-theoretic literature, Falkum (2011) writes that logical properties "are thought to be content-constitutive of a concept" (Falkum, 2011: 118). This view is in particular strongly defended by Horsey (2006).

This approach to conceptual content provides, like the first view discussed previously, an interesting solution to the issue identified at the beginning of this chapter when trying to pin down exactly what Sperber and Wilson consider to be the content of a concept. In this case, the meaning of a lexical item consists of the information stored in the logical entry.⁵⁴ Like in the previous scenario,

⁵⁴ After the first publication of *Relevance*, Macnamara (1987: 724) is probably the first to have noticed the difficulty of identifying exactly what Sperber and Wilson consider the content of a concept to consist of. It is interesting to note that it is when responding to Macnamara that Sperber and Wilson (1987) argue that the information stored in the logical entry constitutes the content of a concept.

the information stored in the encyclopedic entry only serves during the inferential phase of comprehension to derive explicatures and implicatures. However, although this perspective does not share the limits of the atomic view, it also faces a number of issues.

Unlike the atomic account of conceptual content, this view can capture more easily in what sense exactly conceptual narrowing and/or broadening is possible. As described by Carston above, it is the set of inference rules that one stores in the logical entry that provides (together with encyclopedic information) the necessary structure on the basis of which narrowing/broadening can occur. Somehow, this view is also equipped to answer the challenge that referentialism represents. In this case, one has internal content from which a specific reference can be established. And yet, this view does not entirely solve the issue of how to determine reference. First of all, the content of the logical entry of some concepts is not sufficient to justify exactly why they refer to the entity/object they do and not to another. Consider, for instance, the concepts *ROTTWEILER* and *DOBERMANN*. In order to be able to establish reference, the internal content must be sufficiently detailed to enable us to pick the right referent. When using the two concepts *ROTTWEILER* and *DOBERMANN*, one refers to two distinct kinds of dogs even though they (arguably) look quite similar. Yet it is not clear in RT whether the logical entry associated with each of these two concepts provides enough (in the sense of *distinguishing*) information to determine the right reference, i.e. reference to rottweilers and to dobermanns respectively (and not the other way round). As mentioned in the previous chapter, it is often argued in RT that the information stored in the logical entry never fully defines a concept. Sperber and Wilson (1995) argue that:

Our framework allows for empty logical entries, logical entries which amount to a proper definition of the concept, and logical entries which fall anywhere between these two extremes: that is, which provide some logical specification of the concept without fully defining it. (p. 92)

Such a view is problematic if we assume that the information stored in the logical entry is content-constitutive and should, therefore, be reference-determining. It is unclear how one can accurately establish reference if the concepts one uses are not fully defined (and sufficiently distinct from other concepts). As a result, either RT needs to argue that the logical entry always fully defines a concept, or it needs to consider that the content of a concept is not (only) determined by the information stored in the logical entry, for otherwise reference cannot be established. The latter option is preferable, and as we will see in the next section, the information stored in the encyclopedic entry once again seems to be a good candidate.

Now, in spite of the latter observation, let's assume for a moment that the logical entry fully determines the content of a concept. In this case, the meaning of a concept is clearly identifiable, it is prior to reference, enables conceptual

narrowing/broadening, and makes reference assignment possible. There are two reasons why this option still remains relatively problematic, however. The first reason comes from the observation that this perspective faces exactly the same limits as “classical” (Aristotelian) definitions in terms of necessary and sufficient conditions.⁵⁵ Indeed, if the logical entry of a concept fully defines that concept, then it will inevitably consist of a set of necessary and sufficient inferential rules which can be used to compute (Sperber and Wilson, 1995: 89) the logical forms in which these concepts occur. Yet there are a number of issues with this view. One of them follows from the observation that concepts are not mentally represented in such rigid terms, but are more flexibly organized (around a prototype) and have fuzzy boundaries (see, for instance, Rosch 1975).⁵⁶ This explains why, for instance, both eagles and ostriches can be described as birds even though the feature ‘fly’ (which could be argued to be central to the concept BIRD, i.e. common to all birds) applies to eagles but not to ostriches. As a result, if one still wants to argue, on the one hand, that (like eagles) ostriches are birds and, on the other, that concepts are defined in terms of inferential rules, then the property $\{\text{BIRD} \rightarrow \text{FLY}\}$ needs to be dropped from the list. The difficulty for a theory of concepts is that this might result in concepts that possess very few inferential rules and that, as a consequence, fail to be fully defined since they are not sufficiently distinguishable from one another. That is, whether we like it or not, the logical entry cannot fully define a concept. In that case, this leads us back to the challenge identified above. Directly related to this issue is the observation that for some concepts, in particular abstract ones, it is quite hard to understand exactly what inferential rules can possibly make up their logical entry. For instance, it is unclear what rules are attached to the concept FREEDOM, and in particular in what sense these rules enable us to distinguish it from another abstract concept such as LIBERTY.⁵⁷ Generally, then, inferential rules alone cannot possibly define the content of concepts, which strongly suggests that other elements associated with concepts (e.g. encyclopedic information) must constitute their content.

There is a second reason why it is problematic to view the logical entry of a concept as being its content. If one is to define concepts in terms of inferential rules, then one must at least have some idea of what these rules actually consist of. Yet the relevance-theoretic perspective faces a number of limits which

⁵⁵ As Lewandowska-Tomaszczyk (2007: 146) points out, issues with classical definitions have in particular been addressed by Quine (1953, 1960), Wittgenstein (1958), and by Rosch (1973, 1975).

⁵⁶ It is worth noting that Barsalou himself considers that concepts are not defined by strict common features but are composed of a graded structure (Barsalou, 1985, 1987).

⁵⁷ This observation receives support from experimental evidence. On the basis of property-generation tasks, Wiemer-Hastings and Xu (2005) show that “abstract concepts tend to be more schematic in nature, involving a larger proportion of unspecific features than concrete concepts” (p. 733).

greatly weaken the role of the logical entry as a distinct entity. The approach developed in RT will be discussed in the rest of this section.

When Sperber and Wilson (1995) establish the distinction between the information stored in the logical and encyclopedic entries, they argue that the main difference is one of mental representation:

The information in encyclopaedic entries is representational: it consists of a set of assumptions which may undergo deductive rules. The information in logical entries, by contrast, is computational: it consists of a set of deductive rules which apply to assumptions in which the associated concept appears. (Sperber and Wilson, 1995: 89)

From this perspective, logical information consists of deductive rules of the type {CAT → MAMMAL}. These rules help to compute (i.e. draw inferences from) the logical forms in which a lexical concept occurs. This view is problematic, however. First, when they distinguish between encyclopedic and logical information, Sperber and Wilson explicitly argue that the distinction between the two generally corresponds to the philosophical analytic–synthetic dichotomy (Sperber and Wilson, 1995: 88). That is, by virtue of defining a concept, the logical information must necessarily be true of that concept (i.e. analytic), while the (non-defining) information stored in the encyclopedic entry is only true in virtue of one’s experience in the world of that concept (i.e. synthetic). Sperber and Wilson explicitly argue, however, that the analogy only means to capture the observation that knowledge can be stored in different ways (cf. quote above) but not that it necessarily entails different types of truth (p. 88). This is convenient since the analytic/synthetic distinction is rather controversial. Quine (1953, 1960) in particular believes that it is not possible to distinguish between different types of truth and specifically argues against analyticity (i.e. necessary truth). More recently in the relevance-theoretic literature, a similar view is adopted by Horsey, who argues that the information stored in the logical entry is not analytic in the traditional philosophical sense (Horsey, 2006: 74). Rather, he argues that the truth of information is subjective and may be different across individuals (p. 25), and that whether an individual chooses to place a specific piece of information in the logical or encyclopedic entry crucially depends on whether that person takes this piece of information to be content-constitutive (p. 75). Sperber and Wilson adopt a similar perspective when they argue that the same piece of information can function “now as part of the content of an assumption [i.e. logical entry], now as part of the context in which it is processed [i.e. encyclopedic entry]” (Sperber and Wilson, 1995: 89). In this case, however, it means that the distinction between the logical and the encyclopedic entry is only a “psychological distinction” (Carston, 2010: 275), i.e. a perceptual difference that may not translate into different cognitive processes. As a consequence, one could question the necessity both to distinguish between the two entries and, in particular, to argue that

only the information stored in the logical entry constitutes the content of a concept. Carston (2002a) explicitly doubts that there is “really a clear logical/encyclopaedic distinction” (Carston, 2002a: 322). The necessity to distinguish between the two kinds of information seems to originate from a particular intuition. As Horsey (2006) points out, for Quine the distinction is meant to capture “intuitions of centrality” (Horsey, 2006: 13), i.e. intuitions about which pieces of information are more central to a concept than others. A similar intuition can be found in the work of Sperber and Wilson when they say:

Intuitively, there are clear-enough differences between encyclopaedic and logical entries. Encyclopaedic entries typically vary across speakers and times Logical entries, by contrast, are small, finite and relatively constant across speakers and times. (Sperber and Wilson, 1995: 88)

In other words, it is assumed that individuals classify as content-constitutive (i.e. logical) those pieces of information that are more central (and therefore stable) to a concept and as contextual (i.e. encyclopedic) those that are less central (and therefore less stable). This perspective is highly problematic, however. From a theoretical standpoint, apart from the philosopher’s intuition, there is nothing that justifies the view according to which less stable, more peripheral aspects of concepts do not directly contribute to their content too. From a more psychological viewpoint, it is unclear why individuals should so categorically treat less central elements as necessarily not being content-constitutive, especially since it is not clear in this case how individuals even manage to decide when a given piece of information is central to a concept or not. Horsey (2006: 75) himself admits this is a challenge. That is, the distinction is not straightforward. As a consequence, there is no reason to distinguish between logical and encyclopedic information: there is simply conceptual knowledge. Of course, this does not mean that individuals do not categorize this knowledge in different ways, with some aspects of it being considered more central to a concept than others. The different networks discussed in cognitive linguistics, for instance, introduce the notion of a prototype precisely so as to account for this intuition (see also Goldberg, 2019: 16). Yet the rest of the conceptual network in which a prototype occurs is also considered to be content-constitutive.

This last observation provides a transition to the last issue concerning logical information. If one considers that only the information stored in the logical entry is content-constitutive, then it is once more unclear how exactly (conventional, encoded) polysemy can be mentally represented, if it is at all possible. In this case, the logical entry simply consists of a hodgepodge of inferential rules that can be computed once a concept appears in the logical form of an utterance. Yet, considering that polysemy is possible, then one would want to be able to

distinguish between different, organized sets of logical information that can be exploited independently of one another. It is unclear how that is possible in RT.

In this section, I considered the possibility for conceptual content to be composed of the information stored in the logical entry that a concept is argued to give access to in RT. Although this view solves the ambiguity identified at the beginning of the chapter, namely that of understanding exactly what constitutes the meaning of a lexical item, it faces a number of issues itself. In particular, the difficulty is to know whether this entry alone can ever fully define a concept. This is notably supported by the observation that pinning down exactly what constitutes the nature of this entry and in what sense it differs from the information stored in the encyclopedic entry represents a real challenge. In the following sections, I will argue that there is no need to distinguish between two types of entries. Before doing so, I will introduce the last possibility identified by Groefsema (2007: 139).

3.1.3 *Conceptual Content, Logical and Encyclopedic Knowledge*

The third possibility to deal with the ambiguity left by Sperber and Wilson is to consider that the content of a concept is determined by the information stored in both the logical and the encyclopedic entries (see Figure 3.5). This is the first possibility that Groefsema (2007) mentions, but I have treated it as the last possibility since there are only very few relevance theorists who adopt this perspective. This is unfortunate, as will become clear, since it is the best solution for RT (or at least, a version of it).

Like the first two options discussed in the previous sections, this view can avoid a number of challenges. In this case, it is clear what the meaning of a concept is: the combination of logical and encyclopedic information. Furthermore, it is also clear how reference assignment can be established and how ad hoc concepts can be derived.

This view, although not largely adopted within RT, may seem to follow from a particular assumption in RT. Although the atomic account presented earlier is strongly defended in RT, the concept that a lexical entry is associated with is often




Figure 3.5 Logical and encyclopedic entries as conceptual content




Figure 3.6 Logical and encyclopedic entries as conceptual content (2)

described simply as an address in memory, or a point of access, that enables us to retrieve the information stored in the different entries (represented in Figure 3.6).

[A concept] appears as an address in memory, a heading under which various types of information can be stored and retrieved. (Sperber and Wilson, 1995: 86)

In RT, concepts are psychological objects and each consists of a label or address. (Romero and Soria, 2010: 22)

The assumption is that a concept is a kind of ‘address’ in memory which provides access to three kinds of ‘entry’. (Clark, 2013a: 244)

The conceptual address corresponds to the form that a concept takes in thought, while the information provided by the different entries constitutes the content of this very concept: “the distinction between address and entry is a distinction between form and content” (Sperber and Wilson, 1995: 92). As mentioned already, the lexical entry of a concept provides the linguistic counterpart used to express the concept (i.e. a specific word/sign), while the logical and encyclopedic entries can be understood here as specifying the actual content of the concept, some aspects of it being more central and stable than others.

Few relevance theorists adhere to this view, yet some have put forward very similar hypotheses. Before looking at these proposals, it is worth briefly discussing the limits of such an account of conceptual content. The main limit of this view corresponds very closely to the limit identified for the previous view with respect to the distinction between logical and encyclopedic entries. It is unclear exactly in what sense it is necessary to distinguish between logical and encyclopedic information, since this distinction may not actually reflect a cognitive reality. I will not repeat the arguments here (see page 86), but it is preferable simply to get rid of this distinction and to argue instead that there is simply conceptual content. As a result, a conceptual address does not give access to three but only to two types of entry: its linguistic form and its content. Of course, the difficulty in this case is to determine the nature of this content and whether it is more of the logical type or the encyclopedic type, as illustrated in Figure 3.7.




Figure 3.7 Conceptual content?

The answer to this question has to be encyclopedic knowledge. In the previous section, a number of arguments were presented against the use of inference rules to explain the nature of conceptual content. Inferential rules are hard to distinguish from other (non-deductive) inferential processes, they cannot easily account for the flexible nature of conceptual structures and therefore fail to fully define a concept. Encyclopedic knowledge does not face the same limits. First, it is not inferential in nature and thus can be distinguished from particular processes of inference involved during the comprehension phase. In addition, provided it has the right representational format (e.g. prototype theory), it can also account for the flexible nature of concepts and, therefore, fully define them (see Section 3.2). Furthermore, that conceptual content is encyclopedic in nature also seems to follow from its usage-based origin. It is rather clear when one adopts a usage-based approach to language and communication (which is the case in RT) that one primarily conceptualizes the world through one's experience with it, and in particular through the repeated exposure to particular pieces of encyclopedic information which are then organized in one's mind. This is the reason why in CxG, as in cognitive linguistics more generally, concepts are primarily described in encyclopedic terms.

It is interesting to note that the perspective on concepts just described, solely in encyclopedic terms, can sometimes be found in the relevance-theoretic literature. Having carefully considered the different options she introduces, Groefsema (2007) concludes that the only solution to the challenge she identifies is to consider that it is the encyclopedic entry of a concept that makes up its content (Groefsema, 2007: 155). This is also a view which can be found in the work of Anne Reboul. She explicitly argues, for instance, that “the distinction between *word* and *concept* is presumably nearer to that between *lexical* and *encyclopedic knowledge* than any other distinction” (Reboul, 2000: 60).⁵⁸ It also seems to be the underlying assumptions of Wilson and Sperber (2004) when they argue that “the encoded conceptual address is merely a point of access to an ordered array of encyclopedic assumptions from which the hearer is expected to select in constructing a satisfactory overall interpretation”

⁵⁸ See also Assimakopoulos (2008).




Figure 3.8 Concepts and encyclopedic content

(Wilson and Sperber, 2004: 619). This enables me to represent concepts as shown in Figure 3.8.

Already, it is worth noting that this perspective is quite reminiscent of the way in which CxG defines constructions: form–meaning pairings. Beyond this resemblance, it is interesting to note that this view faces none of the issues that the previous account does.⁵⁹ In the next section, I will try to show that this approach, which is very similar to that adopted in CxG, provides the best alternative to define lexical semantics. Before doing so, I would like to make a final observation concerning the nature of this encyclopedic entry.

As mentioned several times already, RT provides a very convincing, explicit and tangible explanation of how meaning is determined in context and, concurrently, is able to explain the origin of polysemy. Unfortunately, as pointed out by Lemmens (2017), given that the encyclopedic entry of a concept is only considered to be a “grab bag” of knowledge, it is unclear exactly how new interpretations can affect these conceptual structures in the long term and, therefore, how polysemy is assumed to be represented in one’s mind. Specifically, he convincingly argues that it is “unclear how *different* modulations of one and the same lexical item will be

⁵⁹ Note that, in this representation, the lexical entry is still considered not to be radically distinct from the concept to which it is attached, but is also part of the knowledge that one has about a concept. That is, here, lexical knowledge is conceptual knowledge. There is a debate whether or not lexical knowledge and conceptual knowledge should be distinguished (e.g. Murphy (2000), and references cited therein). This issue will not be addressed here. Nevertheless, it is worth noting that the view represented in Figure 3.8 is not unfamiliar both in RT and in CxG. Relevance Theory, in spite of its Chomskyan approach to language, assumes that concepts are made up of three entries, one of which is the lexical entry. In CxG, as in cognitive linguistics more generally, it also follows from the non-modular, usage-based approach to language that lexical knowledge is necessarily part of one’s conceptual knowledge. Paradis (2003: 265) very explicitly says that:

The conventionalized routines of linguistic expression may be referred to as the lexicon. However, such a lexicon is not an encapsulated component. On the contrary, it is inseparable from conceptual knowledge and cognitive abilities in general. Lexical items evoke and are evoked by concepts, which involve all kinds of meaning specification that we use in various usage-events. I see no reason for postulating a separate non-conceptual type of lexical knowledge.

represented” (Lemmens, 2017: 104). Lemmens naturally recognizes that, when the encyclopedic entry is viewed as being content-constitutive, then the relevance-theoretic approach “is but one step away from being fully compatible with a cognitive view” (2017: 102). Lemmens’ comment here is particularly relevant since it shows that arguing that conceptual knowledge is encyclopedic in nature does not suffice; one also needs to specify the particular way in which this information is structured. In the next section, we will see that the perspective adopted in CxG provides such a structure.

3.2 Lexical Semantics: A Structured Body of Encyclopedic Knowledge

The previous section shows the problem involved in pinning down the content of a concept (and, as a consequence, the meaning of lexical concepts) in RT. Following Groefsema (2007), it was shown that there are at least three views that emerge from the relevance-theoretic literature, one of which stands out particularly with respect to the others (the atomic view). The strengths and weaknesses of each of those views were presented in turn. Eventually, it was suggested that the information stored in the encyclopedic entry is the best candidate for what constitutes lexical ‘meaning’. At the theoretical level it is the view that best fits the other underlying assumptions of the theory (namely, how the underdeterminacy thesis, the notion of explicatures, and the derivation of ad hoc concepts have been formulated). From a more descriptive level, and as far as meaning is concerned, this perspective does not face many of the challenges that the other views encounter. As mentioned several times already, a relatively similar view is adopted in CxG which, given its usage-based approach, also views meaning in terms of encyclopedic knowledge. In this section, I will briefly point out some of the advantages of adopting such a perspective generally, as well as discuss the ways in which the view adopted by constructionists can provide further insights into the difficult question that defining lexical semantics represents. It is on this basis that the notion of lexical pragmatics will be discussed in Section 3.4.

In order to discuss the advantages of the CxG approach, I will look in turn at all of the issues and difficulties that were identified in the previous section and show that it can handle most of them. In particular, I will focus on the way that encyclopedic knowledge is assumed to be structured in CxG and I will argue that this approach provides the required, solid basis to explain polysemy (and language change more generally).

It is traditional when discussing questions of lexical meaning, at least in philosophy, also to address the question of *reference*. In Section 3.1.1.1, it was shown that this question can sometimes represent a challenge depending on the

way one defines meaning. In order to be able to establish the right reference, one necessarily needs to possess some internal conceptual content prior to reference. The perspective adopted here, according to which lexical meaning is to be defined in encyclopedic terms, precisely enables reference. This knowledge is internally stored in the minds of speakers and can be used to establish reference to a specific item/person in context. This content does not constitute the reference itself but forms the basis from which reference is possible. In this case, for instance, the reason why the concept CAT is used to refer to *cats* simply follows from what we know about cats and which enables us to refer to cats in the real world.

At the same time, unlike the different views presented earlier, this content is not considered to constitute the necessary and sufficient conditions that a concept gives access to and that are systematically used when establishing reference. That is, unlike an atomic concept or inferential rules, encyclopedic content is not taken to be necessarily and systematically true of that concept in all contexts and, therefore, is not used only to refer to items that share exactly the same properties (e.g. the concept BIRD, Section 3.1.2, can be used to refer both to eagles and ostriches, although the property ‘fly’ applies only to eagles and not to ostriches). This is due to the dynamic nature and graded structure of encyclopedic knowledge which individuals gradually acquire from the different contexts in which a concept occurs. As mentioned in Section 2.1.2.2, from a usage-based approach such as CxG, the conceptual structures that one has in mind emerge from one’s experience with these concepts in the world. This experience involves a categorization process whereby new uses of a concept systematically affect the mental representation of that concept and new information is stored alongside old information. This process does not result in a grab bag of information, but new information is systematically placed within a conceptual network. This network is organized around a prototype, which contains the most salient features of a concept, and forms different bundles of knowledge (i.e. different *senses*) which are organized via analogy on the basis of a judgment of similarity. (The main difference between the radial network and the schematic network introduced in Section 2.1.2.2 mostly concerns the extent to which individuals actually abstract away from their experience.) A result of this process of categorization is that the different “features” that a concept makes accessible need not be necessarily activated across contexts but only get activated in the relevant contexts. That is, the conceptual network is a relatively flexible mental object that speakers and hearers can exploit in different ways. In this scenario, concepts do not make particular reference *necessary* (contra Fodor), but a given reference (and with it its truth values) is solely determined in a specific context by a particular speaker (and therefore has to be retrieved by the hearer).

This last observation provides a nice transition to the next advantage of holding an encyclopedic view of conceptual content. As mentioned above, if one considers concepts to be atomic, to consist of inferential rules or simply to be grab bags of knowledge, then it is unclear how different modulations can affect the original concept from which they are derived (and therefore, it is unclear how RT can possibly explain language change).⁶⁰ Considering instead that concepts are structured networks of acquired knowledge, this is no longer an issue. The particular way in which (old and new) information is exploited within a particular context directly impacts the conceptual network to which those pieces of information relate (either by entrenching an already existing bundle of knowledge, i.e. sense, or by creating a new one). And this perspective, together with strong pragmatic principles (such as those proposed in RT), can explain how language change actually works. I will come back to lexical pragmatics later.

In addition, and in direct relation to the previous point, this approach can also make sense of the (much-discussed) process of ad hoc concept creation. First of all, once we assume that individuals store conceptual networks, it then becomes clear in what sense we manage to derive narrower and/or broader concepts in different contexts. (Although, as we will see in Section 3.4, the terms *narrowing* and *broadening* might be slightly inappropriate.) Indeed, the conceptual network provides the necessary structure on the basis of which narrowing/broadening can occur.⁶¹ There is conceptual narrowing when the information conveyed by a particular concept is more specific than that originally provided by the stored conceptual network; and there is conceptual broadening when the information provided by a concept is less specific than the information found in the original network. This is particularly interesting since, as mentioned above, this process also directly leaves a trace in the conceptual network, and the repeated derivation of a given ad hoc concept will lead to its entrenchment (and conventionalization) in the conceptual network from which it is originally derived. In Section 3.4, I will come back to the pragmatic process which is involved in the derivation of ad hoc concepts and how it fits in with the picture of lexical semantics presented here. As mentioned in the previous chapter, the pragmatic principles involved during meaning adjustments are often omitted in the cognitive linguistics literature, and it is not clear exactly what is meant by

⁶⁰ Note that in RT, although it is assumed that assumptions may be strengthened or contradicted in context (p. 37), the structure of encyclopedic knowledge is not addressed since they generally assume that concepts are atomic and that the encyclopedic entry plays only a secondary role.

⁶¹ It is interesting to note, as mentioned several times already, that Barsalou himself (a standard reference in RT) actually develops his notion of “ad hoc categories” on the assumption that they are structured bodies of knowledge (and not atoms/inferential rules). So this perspective is also consistent with the original view from which the notion of ad hoc concepts was introduced in RT. (This is true to the extent that Barsalou only talks about narrowing processes. This will be discussed in Section 3.4.)

the words *context* and *pragmatics*. Nonetheless, it is important to point out that, regardless of how one defines pragmatics, it is largely accepted and argued within cognitive frameworks such as CxG that meaning (and therefore the conceptual networks that one stores) is not conceived of as fixed items that one simply invokes each time a specific linguistic element is used. It is a central assumption in cognitive linguistics that the meaning of a word is constantly negotiated and it is argued, like in RT, to “emerge and develop in discourse” (Langacker, 2008: 30). That is, as mentioned in Section 2.1.2.2, a crucial commitment in cognitive frameworks is that understanding an utterance does not simply consist in unpacking information but rather involves a systematic process of meaning construction (see, for instance, Taylor, Cuyckens and Dirven, 2003; Croft and Cruse, 2004; Evans, 2006, 2009; Evans and Green, 2006; Evans, Bergen and Zinken, 2007; Radden et al., 2007; Langacker, 2008; Geeraerts, 2016; Taylor, 2017; Schmid, 2020, and references cited therein). The next difficulty is of course to understand exactly how this meaning-construction process is actually carried out, since cognitive linguists often fail to make explicit the pragmatic principles involved (see Section 2.1.2.2). We will see that RT provides very interesting insights into the matter. Nevertheless, it is not because one argues that concepts are encyclopedic in nature that one therefore abandons the idea that meaning is primarily contextually derived. Indeed, quite the opposite is true. As Lemmens (2017: 106) points out, “no one will deny the importance of contextual modulation, but this does not provide evidence that meaning should not, or cannot, be encyclopaedic.” Although this type of meaning is quite rich, it remains to be exploited and negotiated by individuals in context. How exactly this is carried out will be explained more fully later in this chapter, the aim of which is to show that although neither of the two theories provides a full account of lexical semantics–pragmatics, their integration precisely enables one to achieve greater descriptive accuracy.

As a matter of fact, the richness of conceptual content constitutes the last point that I want to address in this section. When one views concepts as being primarily encyclopedic in nature given that they result from one’s experience, then one faces the necessary conclusion that concepts must be internally quite rich (cf. Lemmens, 2017: 106; Hogeweg and Vicente, 2020). That is, concepts are not decontextualized, abstract objects but are filled with contextual information. Two consequences directly follow from this observation. First of all, this richness can explain why encoded polysemy is considered to be the norm in frameworks such as CxG (and why it is possible in the first place). In this case, one necessarily has to organize one’s knowledge, and the different nodes of the mental network that one derives represent the different senses of the lexical item associated with that concept. This view is naturally compatible with the perspective adopted in RT. Carston (2016a), for instance, very explicitly argues

that “polysemy very often has its basis in pragmatics. . . . Lexical meaning evolves and very often it is a (recurrent) pragmatic inference that lies at the root of new meanings” (Carston, 2016a: 621). To the best of my knowledge, no (or few) cognitive linguists would dispute Carston’s claim. However, unlike the type of atomic semantics advocated by Carston, it is interesting to notice that only an encyclopedic view of conceptual content enables us to capture exactly how pragmatics can impact on the mental representation, as discussed above, and in particular how one can mentally represent the link between different modulations of the same concept and therefore make (semantic) polysemy possible. Relevance theorists are fully aware of (semantic) polysemy. Carston (2002a: 219), however, rightly suggests that postulating polysemy is not enough; one must also explain how the meaning of polysemous items is actually used and exploited in different contexts. This is an important question which will be addressed in the next sections. Before doing so, I want to discuss very briefly another consequence of assuming rich conceptual knowledge. The account of ad hoc concepts discussed in RT primarily rests on the assumption that the words we use largely fail to convey the speaker’s intended interpretation (i.e. the underdeterminacy thesis). Assuming that concepts are essentially encyclopedic and, therefore, are rich representations, one could dispute the necessity to postulate the underdeterminacy thesis exactly as it is presented in RT (and strongly defended by Robyn Carston). Indeed, do the words we use always underdetermine what we want to communicate? Instead, it might also be appropriate to view the systematic derivation of ad hoc concepts (i.e. meaning construction) as resulting from some form of *indeterminacy*. Indeed, the assumption that pragmatics functions exclusively to complete the meaning of utterances that semantics fails to provide is not satisfactory. (Note that relevance theorists have never suggested that this is the case; quite the contrary since they argue for the systematic derivation of ad hoc concepts, see next section.) Rather, given the picture presented above in terms of meaning construction, semantics and pragmatics are inextricably intertwined. Yet the notion of underdeterminacy fails to capture this aspect of utterance interpretation. First, as mentioned before, it strongly suggests that pragmatics only serves to compensate for defective semantics. Yet it does not. Furthermore, it also suggests that the content provided by the words we use is quite poor. Again, as mentioned above, it is not. Rather, pragmatics and semantics are simultaneously exploited in context to derive a relevant interpretation. And in that sense, as far as polysemy is concerned, it seems preferable to argue that semantics does not (necessarily) underdetermine what a speaker wants to communicate but rather is indeterminate with respect to their intentions. This move is, of course, not meant to diminish the role of pragmatics during the process of utterance interpretation. It simply consists in giving back to

semantics more room in a theory of comprehension than is often given in RT. This will be discussed more fully in the next two sections.

The encyclopedic view of conceptual knowledge such as adopted in CxG and defended here faces a major challenge: that of understanding exactly how this content is exploited in context, i.e. understanding in what sense it can fit with an account of lexical pragmatics and whether semantics and pragmatics are necessarily to be distinguished. This issue will be addressed in the next two sections. The conclusion to this section is that the encyclopedic view of conceptual content nicely resolves a number of issues that the other perspectives (discussed in the previous sections) fail to answer. These issues are primarily theoretical in nature. Moreover, the encyclopedic view nicely ties in with the accounts developed in RT and CxG. Beyond its theoretical adequacy, the encyclopedic view of concept also provides a psychologically sound assumption about concepts and is also consistent with most of the experimental research carried out in cognitive science (see, for instance, Barsalou, 2012, 2016, and references cited therein).

3.3 Concepts and Literalness: Issues of Representation or Computation?

The aim of this chapter is to provide a better understanding of lexical semantics and pragmatics. The first part of this chapter was primarily concerned with questions relating to lexical semantics. It was eventually suggested that the content of lexical items is best described in usage-based, encyclopedic terms. This view seems to be not only descriptively accurate, but it is also compatible with the views on meaning developed in both RT and CxG. The aim now is to try and position this perspective on meaning in the larger context of utterance comprehension and to understand more specifically how lexical pragmatics operates. This might seem a relatively straightforward and easy task, especially given the somewhat shared usage-based approach to ‘meaning construction’ adopted within both RT and CxG. However, as we will see below, this is not straightforward.

Understanding the nature of concepts and describing the manner in which they are used are, of course, two closely related issues. Precisely at the interface between lexical semantics and pragmatics, however, comes another issue that I have been careful not to mention in the first part of this chapter. It is only once this question has been addressed that I will be able to detail exactly in what sense lexical pragmatics is understood to operate.

As mentioned in the previous chapter, the account of ad hoc concepts presented in RT was originally developed on the basis of two observations. There is, of course, the underlying assumption captured by the underdeterminacy thesis that the content of the words we use often fails to fully determine

what we are actually trying to communicate. But more importantly, as Assimakopoulos (2012: 17) points out, the account of ad hoc concepts developed in RT primarily rests on the rejection of the “encoded first” hypothesis (see Section 2.2.3.1). That is, in accordance with the work of Barsalou (and psycholinguists more generally), relevance theorists assume that individuals do not first test for relevance the encoded concept (or concepts) that a lexical item gives access to and then modulate this concept in context if it is not relevant enough; instead they systematically derive ad hoc concepts (i.e. systematically reconstruct a context-specific concept) across contexts. The same view is also largely adopted within CxG, in which the systematic process of meaning construction is often discussed. Within RT, however, it has recently been argued that this assumption somehow raises a dilemma concerning the nature of concepts. As we will see in Section 3.4, this does not have to be an issue, and it is not considered to be one in cognitive linguistics. The aim of this section, however, is to try and understand exactly what this issue consists of and how it has been dealt with within RT.

The issue with the “non-encoded-first” hypothesis defended in RT has in particular been discussed by Deirdre Wilson and Robyn Carston:

Why should a hearer using the relevance-theoretic comprehension heuristic not simply test the encoded (‘literal’) meaning first? What could be easier than plugging the encoded concept into the proposition expressed, and adjusting it only if the resulting interpretation fails to satisfy expectations of relevance? In other words, what is there to *prevent* the encoded concept being not only activated, but also deployed? (Wilson, 2011: 12)

However, the worry is that, given that the relevance-based comprehension heuristic explicitly licenses hearers to follow a path of least effort in accessing and testing interpretations for relevance, it seems natural to suppose that the encoded concept, which is made instantly available by the word form, would be tried first and only pragmatically adjusted if it didn’t meet the required standards of relevance. (Carston, 2013: 195)

In other words, it seems more relevant (in the technical sense) to test the encoded concept for relevance before trying to derive an ad hoc concept. The question here is to know whether arguing both for the relevance-guided comprehension heuristics and against the “encoded-first” hypothesis does not lead to a theoretical contradiction. This is particularly true when one adopts, as I do in this book, a relatively rich type of lexical semantics. In Section 3.4, I will argue that the answer to this question does not concern lexical semantics but lexical pragmatics and that there is no necessary contradiction. Both Wilson and Carston, however, have treated this problem assuming that it concerns lexical semantics directly rather than lexical pragmatics, and that one therefore

needs to deal with this contradiction. I will examine their views in the rest of this section.⁶²

3.3.1 *Deirdre Wilson's (2011, 2016) Procedural Account*

Wilson (2011) is perhaps the first to have directly discussed the contradiction in adopting the relevance heuristics (*follow a path of least effort*) and rejecting the “encoded first” hypothesis. She puts forward the following solution. According to her, the reason concepts are not directly accessed but ad hoc concepts are systematically derived is to be found at the level of lexical semantics (i.e. the level of the encoded meaning of a word). She argues that the systematicity involved in the derivation of ad hoc concepts might reflect much more complex semantics than previously assumed. Specifically, she argues that, in addition to being associated with a particular concept, lexemes might automatically “trigger a procedure for constructing an ad hoc concept on the basis of the encoded [one]” (Wilson, 2011: 17). In order to explain the paradox, Wilson thus suggests that lexical words are semantic hybrids that both activate a concept and trigger a procedure to construct an ad hoc concept. (Procedures, as mentioned in the previous chapter, consist of specific instructions for the processing of conceptual information which are meant to guide the hearer towards (optimal) relevance.) In this case, it is clear how Wilson gets rid of the issue she identifies in the first place. By virtue of encoding an instruction to construct an ad hoc concept, lexemes can never simply give access to the encoded concept. We observe the instruction and do so by following a path of least effort. Paradox resolved.

According to Wilson (2011, 2016), an account in procedural terms provides an elegant explanation both for the theoretical contradiction identified earlier in this chapter and for the underpinnings of lexical pragmatics more generally. For a number of reasons, however, I share Carston's (2013) skepticism about this proposal. First, an account in procedural terms makes the derivation of ad hoc concepts not only systematic but also compulsory. Yet it is sometimes argued in RT, as Carston (2013: 196) points out, that “the encoded concept can, on occasion, be the concept communicated (Sperber and Wilson, 1998, 2008).” If the derivation of an ad hoc concept is viewed as obligatory, however, it is unclear whether or not it is ever possible to reconstruct the encoded concept (i.e. whether the procedure enables the recovery of the encoded concept). Assuming it is possible, then Wilson needs to account for the observation that reconstructing the original concept (arguably) takes more effort than simply testing it as it is, which makes the overall interpretation less relevant than it could have been (since the more processing effort, the less relevance).

⁶² Sections 3.3.1 and 3.3.2 were published in Leclercq (2022).

Assuming it is not possible to reconstruct the original concept, the challenge is to understand how that concept (and the associated procedure) was acquired in the first place and what exactly the function of that concept is, as well as what is the relevance of storing a concept that is never actually entertained and communicated by individuals. Second, this view also suggests that words that encode a concept therefore all encode exactly the same procedure, namely that of constructing an ad hoc concept. Yet, as Carston (2013) points out, this tremendously weakens the approach to procedural meaning developed in RT. Just like no two words encode exactly the same concept, it is implicitly assumed in RT that no two words encode exactly the same procedure. Yet this assumption is seriously challenged here. (One may argue, however, that this need not directly be an issue for Wilson.) In fact, third, Wilson's proposal is all the more surprising since it assumes that all words are thus (at least partly) procedural. Yet, there is growing consensus that procedural encoding is a property of grammatical units of the language and not of lexical items (see Section 4.2.2). Finally, the challenge with Wilson's proposal also comes from the observation that the task she attributes to a particular procedure is in RT originally supposed to be taken care of by the relevance-guided comprehension heuristics (Carston, 2013: 196; Escandell-Vidal, 2017: 88). That is, individuals are said to adjust concepts in RT because of their expectations of relevance. Adding a specific procedure is quite unnecessary since it is redundant with respect to one of the central claims of the theory. Carston (2013: 193) in fact argues that this move "seems like overkill." For all these reasons, a different solution to the paradox might be preferable.⁶³

3.3.2 Robyn Carston's (2013) *Underspecific Content*

In spite of disagreeing with Wilson's proposal, Carston shares the concern that rejecting the "encoded first" hypothesis is inconsistent with arguing for the relevance-guided comprehension heuristics. Therefore, she puts forward an alternative solution. Carston (2013: 196) suggests that, maybe, the reason encoded concepts are never tested first (and then adjusted only when they do not meet one's expectations of relevance) simply follows from the fact that words never actually encode full concepts but only conceptual schemas or templates (i.e. underspecific schematic meanings).⁶⁴ In order to recover the

⁶³ Of course, this does not mean that the interpretation of lexical items never involves procedures. Meaning is largely compositional, and I will in fact argue in Chapter 4 that the derivation of ad hoc concepts is directly guided by the procedural function of the grammatical constructions in which lexemes occur, thus sometimes giving rise to coercion effects. However, I reject the idea that *all* lexical items themselves encode the same procedure of having to create an ad hoc concept.

⁶⁴ This is a suggestion that, for other reasons, she already made in Carston (2002a: 360). Her solution thus comes across as slightly less post hoc than Wilson's.

full-fledged concepts intended by the speaker, these conceptual schemas thus have to be contextually enriched. As in Wilson's proposal, this perspective makes the construction of an ad hoc concept necessary and hence explains why, while following a path of least effort, encoded concepts are not tested first (since there is no concept to start with; see below). Unlike Wilson's proposal, however, it has the advantage of not putting any additional burden on the lexicon. Nonetheless, Carston's proposal also faces a number of crucial issues.

Carston argues that her account is as explanatory as Wilson's without sharing any of its limitations. She argues, for instance, that, unlike Wilson's account, hers "does not entail an obligatory process that is sometimes unnecessary (as when the encoded concept is the concept communicated)" (Carston, 2013: 197). Two comments can be made about this observation. First, it is not clear in what sense her account does not require an obligatory process of concept construction. By virtue of being underspecific, concept schemas necessarily have to be enriched in context in order to arrive at a specific interpretation (i.e. to derive a specific proposition). This process is therefore precisely required by the type of semantics that Carston argues for. Second, she suggests that the reason why the construction of an ad hoc concept in this account is not necessary follows from the observation that the communicated concept might be the one which is encoded. It is difficult to reconcile what seem like two opposite hypotheses. On the one hand, she argues that words do not encode concepts but concept schemas, while on the other, she argues that the communicated concept might be the encoded one. Yet either words encode full concepts or concept schemas, but the advantage of concept schemas cannot possibly be that they provide a full concept. In spite of what she might argue, Carston's account thus suffers from limitations similar to Wilson's.

The proposal that Carston develops here once more quite strikingly illustrates the tension that there is in her own work in terms of how to define concepts. If one assumes that words encode concept schemas, and not full-fledged concepts, one necessarily drops the idea according to which concepts are referential, atomic objects (a position, as mentioned before, Carston has quite staunchly defended until very recently). This is not the only issue with Carston's proposal, however. There is at least one other critical theoretical implication that needs to be discussed. The relevance-theoretic approach to the semantics–pragmatics interface was developed on the assumption, called the underdeterminacy thesis (Carston, 2002a: 19), that words alone do not suffice to recover the speaker's intended meaning and that, besides implicatures, much inferential work is also needed at the explicit level of communication. Sperber and Wilson (1995: 182) coined the term *explicature* precisely to capture the hybrid nature (semantic and pragmatic) of explicit propositions. As I understand it, though, the standard argument within relevance theory has always consisted in highlighting some form of pragmatic underdeterminacy.

That is, the sentences we use do carry a specific meaning (which occurs in the logical form of an utterance), and this meaning only has to be pragmatically enriched (e.g. disambiguation, reference assignment, conceptual adjustment) in order to derive the explicature. If one now assumes that words merely encode concept schemas, however, then one necessarily has to postulate some form of semantic underdeterminacy whereby language does not simply fail to provide the speaker's intended interpretation but altogether fails to provide any meaning at all. This seems to be Carston's underlying assumption when she says that "while sentences encode thought/proposition templates, words encode concept templates; *it's linguistic underdeterminacy all the way down*" (Carston, 2002a: 360; emphasis mine). However, this perspective is hardly plausible. For one, such a view generally seems to undermine the relevance-theoretic approach to the semantics–pragmatics interface and in particular the notion of explicatures. Indeed, from this perspective, explicatures are essentially pragmatic in nature, which means that they can never truly be explicit (cf. Sperber and Wilson, 1995: 182), which therefore adds confusion as to their role and status in utterance comprehension (cf. discussion in Borg, 2016). As will become clear in Section 3.4, I contend that individuals do have rather rich conceptual knowledge. Within relevance theory, Wilson (2011) also questions the plausibility of such an underspecification account. The idea that some words might not encode full-fledged concepts but simply act as pointers for the recovery of conceptual content can be found in Sperber and Wilson's (1998) discussion of *pro-concepts*. This notion (which is more of an assumption) only applies to a specific set of words, however (e.g. pronouns, gradable adjectives, etc.), and it is not Sperber and Wilson's intention to argue that all words encode such pro-concepts. Wilson specifically points out that "while the assumption that some words encode pro-concepts is quite plausible, the idea that all of them do is unlikely" (Wilson, 2011: 16; see also Carston, 2012: 619). Carston in fact identifies some of the limitations of her proposal herself:

Even if these abstract non-semantic lexical meanings could be elucidated, it is entirely unclear what role they would play in the account of language meaning and use. On the relevance-based pragmatic account of how ad hoc concepts/senses are contextually constructed in the process of utterance interpretation, the real work is done by the encyclopaedic information associated with a concept (a semantic entity) and there is no further constraining or guiding role to be played by a schematic (non-semantic) meaning. Nor does the schema appear to play any role in a child's acquisition of word meaning; in fact, the child's first "meanings" for a word are the (fully semantic) concepts/senses grasped in communication, so the abstract (non-semantic) meaning could only be acquired subsequently by some process of induction. Even supposing we could give an account of how this is done, what would be missing is an explanation of why it would be done, what purpose it would serve. (Carston, 2016b: 158)

There are at least two points in this quote that are worth commenting upon. First, Carston argues that the reason concepts are probably not schematic comes from the observation that these schemas would have no particular role in the comprehension procedure since it is the information stored in the encyclopedic entry that constrains the derivation of ad hoc concepts. When saying this, it is interesting to note that Carston once more gives encyclopedic information a central role in (linguistic) communication. Although it is not her intention, this view is fully compatible with the encyclopedic approach to lexical semantics introduced previously. From this perspective, it is indeed unclear what could possibly motivate the necessity of storing a single (and independent, here) schematic meaning as well as how this meaning might be used (see next section). The most convincing argument against meaning schematism, however, comes from the second part of the quote. Carston rightly points out that the main difficulty is to understand exactly how these schemas might be acquired. These schemas can only be acquired via a gradual process of abstraction on the basis of the full-fledged concepts that one directly accesses in context. Yet the necessity for such a level of abstraction is unclear and seems rather counterintuitive (in the sense of less *relevant*, in the technical sense of the term). Abstracting away such a schematic meaning forces us to derive systematically a specific ad hoc concept in context that we might otherwise store as such and access directly. Intuitively, it could be more economical to store and organize these concepts directly in one's mind, even if some abstraction is involved (see, for instance, footnote 14 (Chapter 2) on exemplar-based and prototype models), rather than to abstract away from these concepts to such an extent that one may not even need this schematic meaning during comprehension.

Carston thus concludes that the underspecification hypothesis needs to be dropped (see also Carston, 2019, 2021). While I fully support this move, it nonetheless raises the question of whether and how Carston still intends to explain the theoretical paradox that her underspecification account was meant to resolve in the first place: if words do have specific meanings attached to them, then why aren't these tested first for relevance? Carston sketches an alternative approach:

This requires making a distinction between the kind of lexicon that features in a narrowly construed I-language, with its focus on syntactic computations and constraints, and the lexicon of the broader public language system, which is a repository of communicative devices whose conceptual contents are what the inferential pragmatic system operates on. In the narrow I-lexicon, the words (or roots) listed have no meaning, conceptual or schematic, while in the C-lexicon of the broader communicational language system, words are stored with their polysemy complexes (bundles of senses/concepts that have become conventionally associated with a word and perhaps others that are not yet fully established as stable senses). (Carston, 2016b: 159)

Carston, however, does not develop this account any further; the information in the quote only contains a basic hypothesis and is not yet developed into a full-fledged theory.⁶⁵ Unfortunately, it is not clear exactly in what sense distinguishing between I- and C-lexicon might help us deal with the issue identified above. First of all, it is unclear what is meant by *polysemy complexes* and *bundles of senses/concepts*. As mentioned several times already, there is quite a lot of tension in Carston's work as to exactly what concepts are. The terminology used, e.g. "bundles of senses," is often found in the literature on prototypes, yet this is most likely not the perspective endorsed by Carston. Importantly, placing the conceptual network at a different level of representation simply pushes the issue to a different level of analysis but does not necessarily solve it. This is particularly true because Carston argues that it is the C-lexicon that "provides input to the pragmatic processes of relevance-based comprehension" (Carston, 2019: 157). That is, it remains a challenge to understand why we should still systematically build an ad hoc concept and not try and test first for relevance any of the stored senses of the C-lexicon.

3.3.3 *Concepts and Literalness: Issues of Computation*

I have argued in the previous sections that adopting a schematic view of meaning is as undesirable as Wilson's procedural account, and Carston herself recognizes that this perspective is somewhat problematic. However, this means that we are left with no specific explanation for the apparent contradiction identified earlier on (namely, that of arguing for the relevance-based comprehension heuristics while at the same time arguing against the "encoded first" hypothesis). Although Carston's (2013, 2016b) proposals raise a number of issues, she asks important questions. In order to account for the dilemma identified by Wilson (2011), Carston (2013) brings into the discussion experimental work by psycholinguists so as to provide an explanation which is not only theoretically plausible within RT but generally psychologically plausible and descriptively accurate. In particular, Carston refers to Steven Frisson (and colleagues), a psycholinguist whose work precisely consists in looking at the processing of lexical items. The findings of Frisson's experiments seem to corroborate Carston's claim that meaning is underspecific (which then explains why ad hoc concepts are systematically derived). In Carston (2016b), the argument is different. She looks at the same data but this time she takes a different view and argues that meanings are not underspecific. (In this case, however, we saw above that it is unclear how the systematic derivation of ad hoc concepts is explained.) The results of these

⁶⁵ In more recent work, Carston (2019, 2022) shows how the distinction between I- and C-lexicon can be applied to issues in morphology (in particular in order to account for the use of cross-categorical words). However, in those papers she does not address the issue discussed here.

experiments will be reported on below. Then, in the next section, I will explain how both the results of these experiments and the dilemma identified by Wilson can be explained when adopting an encyclopedic view of lexical meaning.

The particular experiments that Carston refers to aim at a better understanding of the processing of polysemous lexemes. (She explicitly refers to the work of Frazier and Rayner, 1990; Frisson and Pickering, 2001; Pickering and Frisson, 2001; Frisson, 2009.)⁶⁶ Using eye-tracking methods, Frisson and his colleagues have tried to pin down any differences between the processing of polysemous words (which give access to distinct but related meanings) and that of homonymous terms (which give access to unrelated meanings). As polysemy and homonymy give access to more than one interpretation, one might expect that the same type of selection procedure may be involved in both cases. The results of their experiments do not confirm this hypothesis, however. Indeed, homonyms require significantly more processing time than polysemous terms (which are processed much more like monosemous items). In particular, it is shown that the different senses of polysemous items do not compete in the way that the different meanings of homonyms do. While the competing meanings of homonyms seem to be directly accessed, and therefore need to be processed, this is not the case for the different senses of polysemous terms. When interpreting the data, Frazier and Rayner (1990) argue that polysemous terms provide only an *immediate partial interpretation*, i.e. some form of common ground which can provide access to more specific senses in context. In a similar way, Frisson and Pickering argue that these results provide support for what they call the “underspecified account” (Pickering and Frisson, 2001: 567). In this case, polysemous lexemes are argued to give access not to the different senses they can be used to express but to an underspecific meaning which forms the basis from which the different senses can be arrived at in context (via some “homing in” process). This is the reason why Carston (2013) naturally sees these results as providing evidence for her claim that words might only encode underspecific meanings.

The following observations are particularly relevant given the aim of this chapter. First of all, the results of these experiments provide yet further evidence that the meanings of lexical items are not simply accessed by individuals but are instead systematically built (or constructed) in context. This is consistent with both the relevance-theoretic and the constructionist enterprises. These experiments are also particularly interesting since they directly challenge the notion of lexical semantics. On the face of it, it could seem as though individuals only store some underspecific meaning and not the rich type of

⁶⁶ In her Ph.D. thesis *The semantics and pragmatics of polysemy: A relevance theoretic account*, Ingrid Falkum (2011) also refers to a large set of experimental work on the processing of polysemous items, most of which provides similar results to those discussed by Carston. See Falkum (2011: 63) for specific references.

conceptual structures defended in Section 3.2. Carston (2013) specifically follows this line of argumentation, which not only provides evidence for her previous claim that words encode concept schemas (Carston, 2002a) but can also explain why the literal interpretation of a lexeme is never tested first. Carston (2016b), however, expresses strong doubts that individuals do indeed only store such underspecific meanings. I share Carston's latest skepticism, and in the next section I will show that the results of these experiments may not be incompatible with the view adopted here in terms of rich conceptual structures. It is worth pointing out that this possibility is actually mentioned by those who developed the experiments in the first place:

The underspecification model is in principle compatible with both [the radical monosemy and the radical polysemy] views, at least as long as underspecified meanings are also part of a polysemous lexicon. (Frisson and Pickering, 2001: 166)

The idea of underspecification is perfectly compatible with a representation of all individual senses at some level, though the claim made here is that these individual senses do not play a role in the earliest stages of processing. (Frisson, 2009: 119)

The different experiments discussed by Carston therefore mostly provide evidence not against rich types of lexical semantics but in favor of relatively complex processes of lexical pragmatics. Exactly how the type of semantics adopted in this book easily accommodates the different questions addressed in this chapter is the focus of the next section.

3.4 Lexical Pragmatics: Lexically Regulated Saturation

The previous sections have highlighted the difficulty of identifying exactly what constitutes lexical semantics and how this knowledge is actually put to use in context. The aim of this section is to try and develop an approach to lexical pragmatics which addresses each of the issues identified earlier. It will have become clear that in this book I am largely arguing in favor of the type of semantics adopted in CxG in terms of rich conceptual networks of encyclopedic knowledge (see Sections 3.1 and 3.2). The main challenge now is to understand exactly how this type of semantics can be integrated into a larger framework of lexical pragmatics and answer the different questions raised so far.

First, it is essential for me to remind the reader that, in spite of being rich, the type of semantics adopted in cognitive linguistics is not considered to provide context-free packages that hearers systematically access and necessarily take to be the speaker's intended interpretation. Rather, the conceptual material that a lexical item gives access to is by definition highly context-sensitive. As in RT, the idea that interpreting an utterance does not simply consist in the selection of a particular sense within a conceptual network but rather involves the

systematic construction of meaning (or conceptualization, as Langacker puts it) forms one of the central tenets of cognitive linguistics. Adopting such a rich type of semantics, therefore, should not be perceived as a rejection of pragmatics. Quite the opposite, since cognitive linguists generally fail to see the relevance of decontextualized semantics. By comparison, as Leclercq (2022) points out, it is precisely because they adopt a rigid ('dictionary') view of meaning that proponents of RT face issues such as those discussed in Section 3.3. Yet numerous arguments have been given against such an approach to meaning (Reddy, 1979; Haiman, 1980; Fillmore, 1982; Lakoff, 1987; Langacker, 1987; Murphy, 1991; Pustejovsky, 1995). What is true, however, is that the pragmatic principles that govern the process of meaning construction, and exactly how this rich type of semantics is actually exploited in context, are largely lacking within the cognitive framework. There is, of course, a considerable body of work on metaphors. Outside this area of research, however, the domain of lexical pragmatics has generally been given little attention in cognitive linguistics and, therefore, in CxG.⁶⁷

Lexical semantics, within both RT and CxG, is therefore the starting point on the basis of which lexical pragmatics can operate, and here is how it actually happens. The conceptual network that a lexical item provides access to is organized as a structure in which one has stored related bundles of knowledge (i.e. different senses) around a specific prototype via an analogical process of family resemblance. This network provides the basis for a process of *lexically regulated saturation*.

The term *lexically regulated saturation* was introduced by Depraetere (2010, 2014) when discussing the interpretation of modal expressions in English. In particular, she develops this notion to reconcile monosemous and polysemous approaches to modal meaning. She herself argues in favor of a polysemous analysis of modal verbs. Yet she also believes that understanding modal verbs is more complex than simply selecting one of the stored senses. Rather, she claims that the specific senses that modal verbs encode are entirely context-dependent and that they are systematically reconstructed by individuals on the basis of some context-independent layer of semantics. This independent layer of semantics forms the "semantic core" (Depraetere, 2010: 83) of modal verbs which needs to be contextually saturated by hearers in order to arrive at one of the (contextually-dependent) encoded senses. In this sense, understanding modal verbs is a *saturation* process, since the semantic cores they give access to need to be contextually enriched to provide the hearer with a specific interpretation.⁶⁸ This saturation process is, however, *lexically regulated* since

⁶⁷ I am of course talking about the *non-conventional* type of lexical pragmatics.

⁶⁸ The term *saturation* was first introduced by François Recanati (1989: 304). He defines saturation as a linguistically mandated pragmatic process, whereby an open variable needs to be contextually determined (e.g. reference assignment for pronouns). In Recanati's original

it is constrained not only by pragmatic principles but also by the contextually-dependent layer of semantics, i.e. by the specific senses that belong to the conceptual networks attached to modal verbs.

The aim of this section is to extend the notion of lexically regulated saturation beyond the field of modality and to argue that this process is central to lexical pragmatics more generally. In keeping with Leclercq (2023), I contend that Depraetere's original explanation remains too mechanistic, however, especially as it also seems to (implicitly) rely on a dictionary view of meaning. Not only do I reject the idea of a stable "semantic core" that needs to be enriched, I also condemn the view that speakers merely need to enrich this core into one of a set of already established senses. This leaves too little – if any – room for novel interpretations and for language variation and change. So what exactly is involved in lexically regulated saturation? One key ingredient is given to us in the cognitive linguistics literature. Metaphor (and metonymy) aside, there is a tendency to discuss the notion of "meaning construction" (or conceptualization) mostly in terms of activation (emphases mine):

An expression's meaning presupposes an extensive, multifaceted conceptual substrate that supports it, shapes it, and renders it coherent. Among the facets of this substrate are (i) the conceptions evoked or created through the previous discourse; (ii) engagement in the speech event itself, as part of the interlocutors' social interaction; (iii) apprehension of the physical, social, and cultural context; and (iv) any domains of knowledge that might prove relevant. . . . Precisely what it means on a given occasion – which portions of this encyclopedic knowledge **are activated**, and to what degree – depends on all the factors cited. (Langacker, 2008: 42)

Any given word will provide a unique **activation** of part of its semantic potential on every occasion of use. This follows as every utterance, and thus the resulting conception, is unique. (Evans, 2006: 501)

Making meaning for a word like *antelope* involves **activating** conceptual knowledge about what antelopes are like. (Bergen, 2016: 143)

It is relatively clear from these quotes that (in cognitive linguistics) the conceptualization process involved during the interpretation of a lexeme mostly has to do with activating (to different degrees) parts of the conceptual knowledge associated with that lexeme.⁶⁹ Langacker mentions a few factors that are meant to explain how this activation happens. To put it simply, the

account, we are dealing with a relatively "open-ended valuation" (Depraetere, 2014: 170), since there is no predetermined set of values that one is expected to aim for. Although linguistically mandated, this type of pragmatic enrichment may therefore go in any direction. Depraetere's notion of *lexically regulated* saturation differs precisely in this regard: the stored senses provide such predetermined values that one should aim for when enriching the semantic core.

⁶⁹ In relation to the notion of meaning construction, the words *activate*, *activated* and *activation* are also highly frequent, for instance, both in Nick Riemer's (2016) *Routledge handbook of semantics* and in Geeraerts and Cuykens' (2007) *Oxford handbook of cognitive linguistics*.

underlying idea is that different contexts (linguistic and non-linguistic) will activate slightly different parts of our conceptual knowledge, and this motivates the claim that each (contextual) conceptualization is therefore unique. However, there are a number of issues with this approach. First, Langacker does not really elaborate on the different contextual factors that he mentions, and it is not clear in what sense these directly affect the activation process which lies at the root of conceptualization. More importantly, it is unlikely that the process of conceptualization (or, in relevance-theoretic terms, ad hoc concept creation) can be reduced merely to some activation process. Suffice it to look at language change to understand that the interpretation of a lexeme cannot simply be reduced to activating parts of the conceptual knowledge which it gives access to, for otherwise meaning would never actually change (different parts of the same conceptual network systematically getting activated). In order for language change to be possible, more than conceptual activation is required in the first place. Like relevance theorists, I assume that communication is primarily intentional, and that interpreting an utterance precisely involves taking into account the speaker's intentions. An activation account of conceptualization, however, cannot account for this important factor. Instead, one needs pragmatic (i.e. non-logical, non-deductive) inference to account for this observation (see Mazzarella (2013, 2014) and references cited therein). Within cognitive linguistics (and CxG), however, pragmatic inferences are seldom referred to in relation to lexical meaning.⁷⁰ There are basically two contexts in which the term *inference* is used in cognitive linguistics. It is often used as an umbrella term for all kinds of implicated content, i.e. for the type of content which occurs in implicatures. This is, for instance, the case in Traugott and Dasher's (2002) work on semantic change (and in particular grammaticalization), where the term *inference* seems to be equated with the notion of implicatures only. Yet it is clear in frameworks like RT that pragmatic inferences do not only concern implicatures. In addition, the notion of inference is often referred to in the literature on metaphors and metonymy (e.g. Lakoff and Johnson, 1980; Lakoff, 1987), where it is argued that lexemes that are used metaphorically inherit most of the inferences that are associated with the conceptual domain in which they are used. Yet the inferences referred to here are more of the logical type (i.e. entailments, presuppositions) rather than purely pragmatic ones. Generally speaking, the notion of pragmatic inference is barely referred to in discussion on lexical meaning in cognitive linguistics. In spite of this observation, cognitive linguists are undoubtedly sensitive to the primarily communicative function of language and, therefore, of meaning. This

⁷⁰ One only needs to look at how the word *inference* is used in Geeraerts and Cuykens' (2007) *Oxford handbook of cognitive linguistics*. (Talmy's contribution remains a noticeable exception.)

is very clear in Traugott and Dasher (2002), for instance, who mention the following two quotes:

As pointed out by Bartsch: “semantic change is possible because the specific linguistic norms, including semantic norms, are hypothetical norms, subordinated to the highest norms of communication (the pragmatic aspect of change)” (1984: 393). (p. 25)

We agree wholeheartedly with [Lewandowska-Tomaszczyk’s (1985)] claim that meanings have “a starting point in the conventional *given*, but in the course of ongoing interaction meaning is negotiated, i.e. jointly and collaboratively constructed . . . This is the setting of semantic variability and change” (Lewandowska-Tomaszczyk, 1985: 300). (p. 26)

It is clear for Traugott and Dasher that meaning construction is primarily a collaborative communicative activity rather than the simple recovery (or activation) of conventional aspects of meaning. At least, this is what those quotes strongly suggest. And this is exactly the view I am defending here: meaning construction involves more than activating part of our conceptual knowledge; it involves the recovery of the speaker’s intentions and, therefore, it requires much pragmatic inference (see Rubio-Fernández (2008) for experimental evidence).⁷¹

Very much in the spirit of Relevance Theory, I want to argue that lexically regulated saturation consists of an inferential process. This process is *lexically regulated* in the sense that, naturally, in different contexts, different parts of the conceptual network associated with a lexeme will be activated (and some features of a concept may be so central that they systematically get activated) and will serve as the basis for the interpretation process. These most salient features, unlike what cognitive linguists believe, only provide evidence about what particular interpretation might have been intended by the speaker, however. That is, they only raise awareness of the type of meaning that might have been intended by the speaker. (In this regard, I strongly endorse Bartsch’s (1984: 393) idea that semantic norms are only “hypothetical norms,” see the quote above.) It is then on the basis of those activated conceptual features that the hearer will be able to construct a relevant interpretation.⁷² In that sense, it is a *saturation* (i.e. mandatory, inferential) process, since there is no meaning

⁷¹ By no means do I intend to diminish the role of contextual conceptual activation. George Lakoff, for instance, has done considerable work on the role of conceptual activation (and spreading activation) in relation to metaphorical interpretations (see, for instance, Lakoff, 2014). The aim here is only to argue that activation is but the starting point of the interpretation process and not the only operation involved. Recently, Pritchard (2019) and Glynn (2022) argued that the interpretation process involves analogical (i.e. similarity-based) reasoning. My account includes, though is not limited to, analogical thinking.

⁷² In a sense, this is not radically different from the relevance-theoretic perspective in which the information stored in the encyclopedic entry enables the recovery of ad hoc concepts. The main difference here is that the encyclopedic information is considered to be content-constitutive.

available to the hearer as long as a specific interpretation has not been inferred. Exactly how hearers manage to derive the speaker's intended interpretation is a question for which RT provides a very good answer: the relevance-theoretic comprehension heuristics. That is, in accordance with their expectations of relevance, individuals follow a step-by-step inferential procedure and test various interpretations for *relevance* in order of accessibility. They do so by taking into account many factors, such as the speaker's intentions, extra-linguistic factors, previous discourse contexts and stored assumptions. Once an interpretation provides them with enough cognitive effects to justify the effort put into the interpretation process, they can stop searching. The result of this saturation process may then consist in an ad hoc concept/a conceptualization that has already been derived previously in similar contexts. This will directly lead to the entrenchment (and, potentially, conventionalization) of this particular sense within a conceptual network. More importantly, this process may also lead to the derivation of relatively new ad hoc concepts which lay the foundation for semantic change. In this regard, the type of process discussed here does not radically differ from the type of acquisition process and strategies that children use when hearing a particular word for the first time. The main difference is that whereas children rely a lot (and sometimes exclusively) on extra-linguistic factors to derive the speaker's intended interpretation, adults who already possess large conceptual networks can rely on this knowledge much more and therefore (probably) more easily derive the intended meaning. Like children, however, adults also need to infer in context what the speaker actually intends to communicate, and which interpretation seems to be the most relevant. In other words, conceptual networks are never taken as given, but only provide solid evidence for the type of interpretation that the speaker may intend to communicate (and storing conceptual networks might be relevant precisely in that sense, one of their functions being that of facilitating the saturation process). Exactly how we manage to recover (or try to recover) the speaker's actual interpretation is, as argued in RT, provided by the relevance-based comprehension heuristics.

The process of lexically regulated saturation can answer many of the issues discussed previously in this chapter. First, it can account for the observation that in spite of storing rich conceptual networks, the encoded sense(s) are not tested first for relevance. Indeed, the different senses that a concept gives access to are not context-free packages that one directly has access to and from which one needs to select the most relevant sense. First of all, different parts of this conceptual network will get activated in different contexts. In addition, depending on which part of the conceptual network actually gets activated, the hearer will also have to construct a specific interpretation in accordance with their expectations of relevance. That is, the contextual activation of part of this network does not suffice to provide the speaker's interpretation (although it is most probably the case, here, that the most salient features that have been

activated will be tested first for relevance). This is consistent with most of the work carried out in psycholinguistics, such as that of Barsalou (see previous sections), according to whom the construction of meaning is a complex context-sensitive process.⁷³ Finally, the saturation process also nicely accounts for the type of experimental evidence discussed by Carston (see Section 3.3.3), according to which “individual senses do not play a role in the earliest stages of processing” (Frisson, 2009: 119). Indeed, it will have become clear that words do not directly provide any specific sense to the interpretation process *per se*.

In Section 3.5, I will make final observations concerning the nature of lexical concepts and the type of pragmatic process involved during the interpretation of lexemes. Before doing so, I would like to point out one last consequence that follows from arguing for lexically regulated saturation. As mentioned in Section 3.2, this view challenges the appropriateness of using both the terms *broadening* and *narrowing* in relation to the creation of ad hoc concepts (see Bardzokas, 2023 for a similar observation). Indeed, by virtue of inferentially deriving a specific meaning on the basis of the activated conceptual features, only the term conceptual *NARROWING* seems appropriate. In fact, it is interesting to note that the particular way in which Barsalou himself discusses the creation of “ad hoc categories” mostly supports a *NARROWING* perspective (e.g. Barsalou, 1987). Yet this crucially depends on what constitutes the focus of description, and whether one is discussing the saturation process itself or the resulting ad hoc concept. It is true that only the term *NARROWING* seems appropriate to describe lexically regulated saturation, since the eventual conceptualization will (most often) be more specific than the set of activated conceptual features on the basis of which it has been constructed. Looking at ad hoc concepts directly, however, and comparing ad hoc concepts with the conceptual networks from which they are derived, it seems that both the terms *narrower* and *broader* can be used depending on how much their content actually overlaps. The use of these terms simply depends on whether one is focusing on the saturation process itself (a *NARROWING* process), or on the resulting conceptualization (which can be narrower or broader than the encoded concept).

3.5 Lexical Semantics and Pragmatics: Setting the Story Straight

The aim of this chapter was to try and define lexical semantics and lexical pragmatics. In the first part, I strongly argued in favor of a usage-based, encyclopedic approach to lexical semantics according to which individuals

⁷³ The saturation account presented here therefore provides a nice alternative to Wilson’s procedural approach, since it avoids positing additional semantic constraints. While I concur with Wilson that the meaning of a lexical item is systematically (re)constructed in context, I do not attribute this mechanism to some encoded instruction but assume that it follows logically from the context-sensitive nature of the conceptual material associated with lexical items.

store complex conceptual networks. The challenge is to determine in what sense such a framework integrates lexical pragmatics. In the previous section, I have tried to show that the perspective on lexical semantics adopted here easily combines with the view of lexical pragmatics developed within RT, given that the conceptual networks are not seen as context-independent units but are instead highly context-sensitive. The aim of this section is twofold. First, I will discuss a number of assumptions about semantics which might account for the limits identified in both RT and CxG. Then, the aim is to show that, in spite of rejecting the type of lexical semantics RT generally adopts, its view of lexical pragmatics so far provides the best account of how people manage to communicate effectively.

I have shown in Sections 3.1 and 3.3 that it is difficult to determine exactly how concepts are defined within RT. Recently, the challenge of understanding why stored concepts are not necessarily tested first for relevance has in particular led to some rather peculiar hypotheses concerning the nature of lexical semantics. Generally speaking, the more RT is developed, the more room is given to pragmatics (at the cost of semantics). As explained in Section 3.1.1.2, the commitment to referential atomism has pushed relevance theorists into arguing for relatively poor semantics as opposed to increasingly more pervasive pragmatics. By contrast, in spite of recognizing the central role of usage and pragmatics in communication, there is a tendency in CxG to view the rich (semantic) networks of conceptual knowledge associated with a particular linguistic unit as sufficient: they provide most of the information communicated by an individual. In this case, much room is given to lexical semantics, and pragmatics is often marginalized to the level of implicatures.⁷⁴ Although these two analyses are in direct opposition, their respective limits originate from a relatively common assumption about the mental status of semantics. There is a tendency in the two frameworks to assume (more or less implicitly) that once a particular interpretation is entrenched and conventionalized, and becomes part of our “semantic knowledge,” then this knowledge is almost necessarily consciously available to us. In RT, this seems to be one of the underlying reasons why Carston so strongly defends Fodor’s atomic account (Carston, 2010: 245). She also argues, for instance, that concepts are available to consciousness and introspection (Carston, 2016b: 156). In CxG, and in cognitive linguistics more generally, we have seen that the “meaning construction” process involved during the interpretation of an utterance simply consists in the activation of parts of the network, and that inference is only involved on the implicit side of communication. As a result, semantics and (non-conventional) pragmatics are often put in opposition, with pragmatics simply bridging the gap left by semantics during the interpretation of an

⁷⁴ In that regard, Hoffmann (2022) is a nice exception.

utterance. This explains why, given the respective aims of the two theories, more or less room is given to lexical semantics. Relevance theorists explicitly focus on pragmatics (and hence reduce the role of semantics), whereas construction grammarians are primarily interested in linguistic knowledge (and hence reduce the role of pragmatics). In this section, I briefly want to argue that this (implicit) assumption is ill-founded. As mentioned in the previous section, semantics and pragmatics need not be put on opposite ends of some comprehension scale. Rather, they are two tightly intertwined aspects of the comprehension procedure. Therefore, it is possible to argue both for a rich type of semantics as well as for a ubiquitous type of pragmatics.

Many issues have been discussed in this chapter. If it has taught us anything, it is no doubt the observation that what is traditionally referred to as the semantics (or function) of a construction is not easily brought into consciousness and is not readily available for introspection. This observation explains why it is difficult to define exactly what lexical semantics is. For this very reason, I want to argue that it is not appropriate to discuss the function of constructions (i.e. lexemes, or larger patterns) in terms of *knowledge*. That is, the term *knowledge* can too easily be interpreted as though it is clear to individuals what it is that they have stored. The semantics of a particular construction, however, is (often) not consciously learned but unconsciously acquired, and the actual content is only manifest to us. Of course, in different contexts, different aspects of this content are particularly salient and accessible to an individual. But everything that is stored and composes the semantics of a particular construction can never be consciously accessed as a whole at a given time. Rather, the semantics of a particular lexeme only functions as meaning potential which is exploited in context to derive the speaker's intended interpretation.⁷⁵ (Note that the notion of meaning/semantic potential has also been used and discussed, though in different ways, by Halliday, 1973; Bezuidenhout, 2002; Allwood, 2003; Fauconnier and Turner, 2003; Croft and Cruse, 2004; Recanati, 2004; Evans and Green, 2006; Norén and Linell, 2007; La Mantia, 2018; Verschueren, 2018; Leclercq, 2022, *inter alia*.) That is, to put it slightly differently, we do not know concepts, but our minds make concepts

⁷⁵ This more or less relates to the distinction in philosophy between “knowing-that” and “knowing-how” (see Ryle, 1946; Fantl, 2017). In simple terms, knowledge-that consists in our propositional/theoretical understandings, while knowledge-how refers to the actual use of a specific capacity (mental or physical). I want to argue that (most of) the concepts we acquire originally constitute part of our knowledge-how. (I will come back to this in Chapter 4.) That is, we do not simply understand our concepts (know-that), but we actually know how to use these concepts in order to understand the world around us and engage in communication (know-how). Having and using concepts is therefore primarily a capacity of the mind. In philosophy, for instance, this perspective is defended by Robert Brandom (1998: 135). In linguistics, I believe it is a similar (although implicit) assumption that has led some linguists to describe semantics in terms of inferential rules (for instance). After all, the terms *norms* and *conventions* precisely describe practices (know-how) and not merely theoretical understandings (knowing-that).

available to us. This observation is what explains why it is possible to argue for both a rich type of semantics as well as a rich type of pragmatics. Our minds make available complex semantic structures which in different contexts will be exploited differently (see previous section).⁷⁶

I have argued quite strongly in favor of the rich type of semantics adopted in CxG, the nature of which can be easily accommodated in a theory of pragmatics such as developed in RT. In the meantime, I have given little credit to the actual insights provided by RT on lexical pragmatics. In spite of storing complex conceptual networks, individuals still need to work out in context exactly what interpretation was intended by the speaker. Here, RT provides a very specific and detailed account of how we actually manage to do so: the relevance-based comprehension heuristics. It is important to point out that beyond making clear predictions about how we manage to communicate, these predictions have often been supported by empirical and experimental evidence (see Chapter 2). The development of experimental pragmatics is in fact largely due to the research carried out in RT when trying to test and provide evidence for the different claims of the theory (see Clark (2018) for a discussion). Of course, it will have become clear that I am not inclined to argue that inference is the main provider of meaning during the interpretation of an utterance. Nevertheless, it has been my aim to show that the underlying mechanism that RT presents concerning the interpretation process is very convincing.

3.6 Conclusion

Understanding exactly what lexical semantics and pragmatics are as well as determining how each of them actually contributes to the interpretation of an utterance is no simple task. In this chapter, I have tried to combine insights from CxG and RT to answer this difficult question. First, I tried to show the challenge involved in identifying what relevance theorists assume lexical semantics to consist of. It was shown that the commitment to referential atomism often made within RT is incompatible with some of the most central developments of the theory. I have suggested that the type of semantics adopted in CxG in terms of rich conceptual networks seems to be best suited at both the descriptive and theoretical levels. The difficulty with this perspective is to understand exactly

⁷⁶ More generally, one could ask whether we actually *know* a language or whether we simply *have* a language that our minds make available to us. As far as our native language is concerned, I am inclined to argue in favor of the latter option. The fallacy that we *know* a language comes from the simple fact that any use of language (i.e. linguistic performance) is directly accessible to us and often serves as the basis for introspection. Yet we do not directly have access to the linguistic system (i.e. linguistic competence). It is rather clear in most introductory textbooks to linguistics that linguists look at linguistic performance to draw conclusions about linguistic competence. This is evidence that we only *have* a language. This also seems an assumption shared by Croft (to appear) who recently argued against a focus on mental representations.

how much pragmatics is involved during the interpretation process and whether it still has a place in a theory of communication. I have strongly argued that, in spite of storing rich conceptual networks, individuals still have to reconstruct in context the intended interpretation in accordance with their expectations of relevance. From this perspective, the interpretation of a lexical construction systematically involves the complex combination of both semantics and pragmatics. On the basis of the work of Depraetere, I have suggested that interpreting a particular construction systematically requires a *lexically regulated saturation* process, whereby the rich conceptual networks that constitute its function provide the underlying structure against which the interpretation process operates. This process, however, is primarily pragmatic, and is carried out in accordance with our expectations of relevance following the comprehension procedure spelled out in RT.