

2 Understanding Construction Grammar and Relevance Theory

Construction Grammar and Relevance Theory will now be presented in turn. In the case of Construction Grammar, it will be shown that its main strength resides in its capacity to provide a thorough understanding of linguistic knowledge. Its usage-based take on language provides profound insights into the forms and functions of the linguistic units that individuals can use. At the same time, the discussion will show that its focus on conventions makes for only a partial understanding of linguistic communication. Concerning Relevance Theory, the opposite observation will be made. I will show that while it provides a very elaborate analysis of the pragmatic processes that make verbal communication successful, the argumentation is sometimes weakened by theory-internal assumptions about linguistic knowledge.

2.1 Construction Grammar

Construction Grammar is a cognitively oriented theory of language whose central aim is to account for the entirety of linguistic knowledge. The term *construction grammar* was first used by Charles Fillmore and Paul Kay (Fillmore, 1985a, 1988, 1989; Fillmore, Kay and O'Connor, 1988; Fillmore and Kay, 1995), who were concerned about the lack of attention given in derivational generative grammars to allegedly more peripheral linguistic phenomena (e.g. idiomatic expressions, 'irregular' clausal structures). From a constructionist perspective, these phenomena are considered as much a part of an individual's linguistic knowledge as any general grammatical rules, and not merely by-products of some combinatorial or transformational operations (Hoffmann and Trousdale, 2013b: 3). That is, instead of a core-periphery view, constructionists adopt a more holistic approach to language. In this approach, knowing a language only (or mostly) consists in knowing *constructions*, hence the name of the theory. Like the Saussurean sign (de Saussure, 1916), constructions are defined as arbitrary form-function mappings (Goldberg, 1995: 4). However, whereas the Saussurean sign only applies to lexemes (and morphemes), the notion of construction extends to all aspects of grammar, including idioms as well as abstract phrasal patterns. To use

Goldberg's (2003: 223, 2006: 18) much-cited phrase, it is "constructions all the way down."

Adopting such a symbolic view of language is of course not distinctive of Construction Grammar. This idea is largely shared by functional/cognitive linguists (e.g. Lakoff, 1987; Langacker, 1987, 1991a, 2008; Talmy, 1988, 2000a, 2000b; Wierzbicka, 1988; Halliday, 1994; Givón, 1995, *inter alia*). Yet Construction Grammar stands out from other functional/cognitive-oriented frameworks in terms of how these symbols (i.e. *constructions*) are said to be acquired and mentally represented, as well as how they interact with one another. There are (naturally) different points of contention between constructionists themselves as well. The term 'Construction Grammar' in fact covers a range of different constructionist approaches (cf. Croft and Cruse, 2004: 257–290; Goldberg, 2006: 213–214, 2013; Hoffmann, 2022: 256–271; Ungerer and Hartmann, 2023). In this book, I will mostly work with the ideas developed within (Goldbergian) Cognitive Construction Grammar (Goldberg, 2006: 213; Boas, 2013: 233). The absence of formalism in this approach seems particularly well suited for its integration with Relevance Theory. Nevertheless, since I will not de facto ignore other constructionist approaches,¹ I will continue to use the umbrella term Construction Grammar and its conventional acronym CxG to refer to the theory.

2.1.1 *Fundamental Principles*

The use of the term *construction* in CxG can sometimes be unsettling when you are not familiar with the theory, for it does not only refer to complex combinations or grammatical structures such as is usually the case elsewhere in linguistics. Rather, all objects of linguistic knowledge are argued to be constructions: morphemes, lexemes, idioms as well as larger phrasal patterns (Goldberg, 2003: 219). In CxG, what defines a construction is not its internal complexity but its symbolic nature: constructions are conventional pairings of a specific form and a particular semantic or pragmatic function (Goldberg, 1995: 4, 2006: 5; Langacker, 2008: 5). In order to be conventional, i.e. in order to be part of the speaker's knowledge and obtain construction status, these pairings should exhibit at least one of two properties: (i) non-predictability, and/or (ii) sufficient frequency of occurrence. Goldberg (2006) puts it as follows:

Any linguistic pattern is recognized as a construction as long as some aspect of its form or function is not strictly predictable from its component parts or from other constructions recognized to exist. In addition, patterns are stored as constructions even if they are fully predictable as long as they occur with sufficient frequency. (Goldberg, 2006: 5)

¹ I also largely embrace Langacker's approach, Cognitive Grammar, which comes very close to Goldberg's Cognitive Construction Grammar (see Langacker, 2009).

Originally, Goldberg (1995: 4) defined non-predictability as the only defining criterion for construction status. From this perspective, constructions were all assumed to be either semantically or formally non-predictable, the paradigm case of which are idioms. The semantics of *piece of cake* and *kick the bucket*, for instance, are non-predictable because they are non-compositional, i.e. they cannot be understood solely on the basis of the individual lexemes that compose them. In the sentence in (2), the form of the construction *many a day* is non-predictable, given that *many* usually selects for a plural noun.

(2) I have waited many a day for this to happen. (Hilpert, 2019: 10)

Linguistic expressions that show non-predictability are naturally good candidates for construction status, for they require language users to store them independently of the canonical patterns from which they cannot be derived (Hilpert, 2019: 12). This explains why morphemes and words are constructions, since their forms and functions are non-predictable and language users have to learn them individually (Goldberg, 2002a: 1). However, not all linguistic patterns are non-predictable. The phrases *Make a wish* and *I miss you*, for instance, are neither semantically nor formally deviant. The same is true of the multi-word patterns *legal action* and *in exchange* as well as the inflected forms *smaller* and *students*. What gives these patterns construction status is not non-predictability but frequency of occurrence. That is, these patterns are used frequently enough to be stored by the speaker as distinct constructions (cf. Langacker, 1988; Stemberger and MacWhinney, 1988; Arnon and Snider, 2010; Hanna and Pulvermüller, 2014, inter alia).²

In other words, in CxG, knowing a language consists in knowing patterns that combine a form and a meaning either non-predictively or that occur with sufficient frequency. According to this definition, all of the patterns in (3) to (9) are constructions.³

- (3) ADJ-*ish*
- a. Part of it is yellowish. (COCA, spoken)
 - b. But, generally, I think of myself as a youngish person in an oldish body. (NOW)
- (4) *Roof*
- a. The roof is leaking in a lot of places. (COCA, spoken)
 - b. Smoke rises through a hole in the roof. (COCA, spoken)

² See Ungerer and Hartmann (2023) for a recent discussion on the definition of constructions.

³ Most of the examples used in this book were extracted from corpora available on Mark Davies' interface (english-corpora.org): BNC (British National Corpus), COCA (Corpus of Contemporary American English) and NOW (News on the Web) corpora (Davies, 2004, 2008-, 2016-).

- (5) *Private property*
- a. For example, there's no private property in the Soviet Union. (COCA, spoken)
 - b. Trespassing upon private property is unlawful in all States. (COCA, written)
- (6) *Break the ice*
- a. What can I say to break the ice with a guy? (COCA, written)
 - b. Did you try to do anything to break the ice with them? (COCA, spoken)
- (7) *X is the new Y* construction
- a. So that's why I say, you know, land is the new gold. (COCA, spoken)
 - b. Strong is the new skinny according to the *New York Post*. (COCA, spoken)
- (8) WAY construction (form: [SUBJ V *one's way* OBL])
- a. Mickey Mouse tootled his way across the screen. (NOW)
 - b. You can't buy your way into someone's heart and mind. (NOW)
- (9) CAUSED-MOTION construction (form: [SUBJ V OBJ OBL])
- a. Henry's friend moved the bookcases in Mr Emerson's study. (COCA, spoken)
 - b. My constituents will vote me out of office. (COCA, spoken)

In the sentences in (3), the morphological schema *ADJ-ish* is used to indicate approximation or vagueness: vaguely yellow in (3a), and relatively young/old in (3b). The noun *roof* in (4) refers to the cover of a building. Regardless of their frequency of occurrence, these two constructions are neither semantically nor syntactically predictable from their component parts. This is not the case of the multi-lexeme construction *private property* in (5). This construction, used to indicate an individual's land or building, can be predicted both syntactically (as a particular instance of the [Adj N] pattern) and semantically (i.e. it is semantically transparent). Still, it certainly has construction status for most (native) English speakers due to its high frequency of occurrence. In (6), the idiomatic expression *break the ice* is syntactically predictable. It can be seen as an instantiation of the more general [V NP] and [Det N] constructions. However, it is not semantically predictable. Nothing in the individual meanings of *break*, *the* and *ice* can predict the interpretation of the idiom in terms of a particular social behavior between individuals who are meeting for the first time. In (7), the *X is the new Y* construction is neither entirely syntactically nor semantically predictable. This pattern is not syntactically predictable since Y can be realized by an adjective, as in (7b), although the string *the new* would normally select a nominal head (to instantiate the more regular [Det Adj N] pattern). Neither is it semantically predictable, because none of the elements that occur in this construction suggest that the X and Y items should be interpreted not literally but in a metonymic relationship to a bigger category




Figure 2.1 Lexicon–syntax continuum in CxG

that hearers need to infer in context, e.g. in (7b) *strong* and *skinny* have to be understood in relation to the category of what body type currently seems to be more attractive (cf. Dancygier and Sweetser, 2014: 154). The WAY construction in (8), used to convey manner (or means) of motion, is arguably syntactically predictable. It is not, however, semantically predictable. The meanings of the different items in (8b), for instance, and in particular that of the verb *buy*, do not themselves convey (metaphorical) manner of motion interpretations (cf. Jackendoff, 1990: 218; Israel, 1996: 218). Finally, a similar analysis can be given to the CAUSED-MOTION construction identified in (9). While the form of the construction can be derived from more canonical patterns, examples like (9b) show that the caused-motion meaning with which it is associated is not always predictable from its component parts (cf. Goldberg, 1995: 152–179).

CxG therefore establishes no principled distinction between elements of the lexicon and larger phrasal (or ‘syntactic’) patterns. Instead of a dichotomy between the two, it is assumed that there is a continuum of constructions from more lexical to more syntactic. This continuum is often referred to as the *lexicon–syntax continuum* (Langacker, 2005: 102; Croft and Cruse, 2004: 255; Goldberg, 2006: 220). One way of representing this continuum is to locate constructions on a gradient of lexical fixedness, i.e. from lexically fixed to lexically open (or schematic) constructions, as in Figure 2.1 (inspired by Kay and Michaelis, 2012: 4; Michaelis, 2017, 2019).

There are different reasons why no strict distinction is made between lexicon and syntax in CxG, all of which are closely related. The main reason has to do with the general aim of the theory. Although CxG directly takes its name from arguing that all levels of linguistic knowledge can be described in terms of constructions, it is primarily concerned with how linguistic knowledge relates to cognition in order to provide a “psychologically plausible account of language” (Boas, 2013: 233). A central assumption within CxG is that language does not require a specific cognitive mechanism but is the product of general cognitive abilities (Lakoff, 1987: 58, 1991: 62; Langacker, 1991b: 1; Tomasello, 2003: 3; Goldberg, 2006: 12, 2019: 52;

Bybee, 2010: 6–8, 2013: 49).⁴ Like other models in functional/cognitive linguistics, CxG therefore rejects a modular view of language and in particular the autonomy of syntax (Croft and Cruse, 2004: 1; Fried and Östman, 2004: 24). That is, grammatical constructions are not separated from the rest of our linguistic knowledge and abilities. In addition, constructionists consider that the “primary function of language is to convey information” (Goldberg, 2013: 2).⁵ From this perspective, all components of language are considered to be meaningful. Hence, like lexical items, grammatical constructions are assumed to have a specific meaning that contributes to the understanding of the sentences in which they occur. This is the case for the WAY construction and the CAUSED-MOTION construction in (8) and (9) discussed above. It is also true of the DITRANSITIVE construction, different instantiations of which are found in (10).

- (10) a. The United Nations was giving them food. (COCA, spoken)
 b. Heloise passed me the wooden bowl. (COCA, written)
 c. He told his wife the same thing. (NOW)

Although the interpretations of these sentences differ, they are composed of similar constructions, one of which is called the DITRANSITIVE construction (Goldberg, 1995: 141–151). In terms of semantics, it is this construction that conveys the notion of transfer, or more specifically *X CAUSES Y TO RECEIVE Z* (Goldberg, 1995: 141). And this meaning is said to be associated with the abstract phrasal form [SUBJ V OBJ1 OBJ2], which all sentences in (10) instantiate. Specifically, all the slots of this pattern are associated with a specific function:

- (11) DITRANSITIVE: Syn: SUBJ V OBJ1 OBJ2
 () () () ()
 Sem: Agent CAUSE-RECEIVE Recipient Theme

As the representation in (11) indicates, each of the open slots of the construction is associated with a particular function which, in context, is inherited by the

⁴ The psychological and neurological reality of this assumption is supported by empirical and experimental evidence (cf., for instance, overviews by Bates and Goodman, 1997, 1999; Tomasello and Slobin, 2005; Behrens, 2009, and references cited therein). See Cappelle (to appear) for a critical discussion, though.

⁵ This quote is interesting as it could suggest that Construction Grammar falls prey to the “descriptive fallacy” (Austin, 1962: 3), namely the assumption that words and utterances are only used to convey (truth-evaluable) propositions while language is also used to perform a number of other functions (i.e. speech acts). Yet this is clearly not the view adopted in CxG. Not only does it reject truth-conditionality, but it also acknowledges the other types of (pragmatic) functions that language is used to perform (see next section). Nevertheless, as will be shown in Section 2.1.2.3, this quote is a good example of the symptomatic vagueness with which issues that relate to the semantics–pragmatics interface are treated.

lexical items that occur in that slot.⁶ In (10a), for instance, *them* and *food* are respectively interpreted as ‘recipient’ and ‘theme’ because of their occurrence in the OBJ1 and OBJ2 slots of the DITRANSITIVE construction. Of course, it could also be argued that these interpretations of the lexemes are not due to their being used in a distinct DITRANSITIVE construction but to the subcategorization frame (i.e. valence)⁷ of the main verb *give* of which they are the arguments. Although this might sometimes be the case, the perspective developed in CxG nonetheless seems to provide better insights into an individual’s linguistic knowledge and about their use of the language. First of all, experimental data reveal that these constructions are psychologically real and that language users do store grammatical patterns in association with a specific function independently of the lexical items that occur inside them (cf. Hare and Goldberg, 1999; Bencini and Goldberg, 2000; Kaschak and Glenberg, 2000; Chang, Bock and Goldberg, 2003; Goldberg and Bencini, 2005; Ye, Zhan and Zhou, 2007; Bencini and Valian, 2008; Boyd, Gottschalk and Goldberg, 2009; Johnson and Goldberg, 2013; Shin and Kim, 2021; Li et al. 2022, inter alia).

More importantly for us, the observation that grammatical constructions, like lexical items, are meaningful necessarily shifts the semanticist’s focus of attention. In (12), for instance, *kick* is interpreted in terms of transfer and the expressions *Bob* and *the football* respectively receive the roles of ‘recipient’ and ‘theme’ not because of the subcategorization frame of the verb *kick*, which usually only selects one object (e.g. *Pat kicked the ball*), but because of their occurrence in the DITRANSITIVE construction.

(12) Pat kicked Bob the football. (Goldberg, 1995: 11)

(13) Lyn crutched Tom her apple. (Kaschak and Glenberg, 2000: 512)

Similarly, the DITRANSITIVE construction is responsible for the transfer interpretation in (13) of the denominal verb *crutch*, whereby Lyn (SUBJ/agent) is understood to have used a crutch in order for Tom (OBJ1/recipient) to receive her apple (OBJ2/theme). In this case, not only are the respective roles of *Tom* and *her apple* inherited from the DITRANSITIVE construction, but also the CAUSE-RECEIVE interpretation of *crutch*. The particular interaction between a lexeme

⁶ The representation in (11) is adapted from Goldberg (1995: 50, 2006: 138). Please note that although specific thematic roles and (Jackendovian) primitives (CAUSE and RECEIVE) have been used to describe the construction’s semantics, they have no “theoretical significance” here (Goldberg, 1995: 49). That is, these labels are only used because they facilitate the semantic description of the argument-structure construction. However, the actual semantics of the DITRANSITIVE construction are more complex and require more than a list of rudimentary components (see Section 2.1.2.2).

⁷ Subcategorization frames have been “an essential part of linguistic theorizing since Chomsky (1965)” (Goldberg, 2006: 65).

and a construction such as in (13) is often referred to as *coercion* and will be addressed more fully in Section 2.1.3.

Finally, CxG assumes no a priori distinction between the lexicon and syntax because of its usage-based approach to language. That is, no syntactic structures or linguistic items of any sort are considered to be innate. Rather, a central tenet within CxG consists in viewing all aspects of linguistic knowledge as resulting from language use (Langacker, 1991b: 264; Croft, 2001: 59; Goldberg, 2006: 44; Diessel, 2013: 347). From this perspective, one's linguistic knowledge consists in "the cognitive organization of one's experience with language" (Bybee, 2006: 711). In particular, regardless of their internal complexity, it appears that linguistic patterns emerge from a process of categorization (and generalization) over exemplars, i.e. concrete realizations (Kemmer and Barlow, 2000; Tomasello, 2003, 2006; Bybee, 2010). In CxG, these concrete realizations – which are found in utterances – are called *constructs*, while the generalizations that emerge from them are what form *constructions* (Fried, 2015: 980). Consider the sentences in (14) to (16).

- (14) a. It's about a cat who stole a dog's bed. (NOW)
 b. Why don't you have a cat? (COCA, spoken)
 c. The cat wanted a little air time. (COCA, spoken)
- (15) a. She was as calm as a pond on a windless day. (COCA, written)
 b. I felt as proud as a president. (COCA, written)
 c. Clare acted as serious as a nun. (COCA, written)
- (16) a. It was you who begged for those loans in the past. (NOW)
 b. In some cases, it is their wives who are the chief wage earners. (COCA, written)
 c. It is my son who made it. (COCA, written)

Constructionists believe that just like the form and meaning of the lexeme *cat* are acquired by generalizing over different usage events such as in (14), the *as ADJ as a N* construction is itself acquired by generalizing over examples like those in (15), and the *It-CLEFT* construction by generalizing over examples such as in (16). That is, all aspects of linguistic knowledge are acquired by a gradual process of categorization and generalization across usage events, and no grammatical pattern is therefore considered innate. As a result, linguistic knowledge is "viewed as emergent and constantly changing" (Bybee, 2013: 49). Indeed, new constructs have a systematic impact on the representation of a construction. The lexicon–syntax continuum represented in Figure 2.1 can therefore be seen as a consequence of this usage-based acquisition process, with different constructions being more or less abstract depending on the degree of generalization made possible by the input received by an individual. It also follows from this perspective that all constructions (lexical to grammatical) are not stored separately but are located in the same repository of linguistic

knowledge. This repository is referred to in CxG as the *construct-i-con* (Jurafsky, 1992: 28; Goldberg, 2003: 219).

It is important to note that the construct-i-con does not contain only constructions, i.e. generalizations (cf. Section 2.1.2.2, footnote 14.), but these are stored alongside the individual constructs from which they emerge (Abbot-Smith and Tomasello, 2006; Bybee, 2010, 2013; Goldberg, 2006). According to Langacker (1987), arguing that linguistic knowledge is either composed of broad generalizations or specific instantiations amounts to committing to what he calls the “rule/list fallacy” (Langacker, 1987: 29), i.e. an either/or idealization that may not correspond to a speaker’s cognitive reality. This assumption is in particular supported by the observation that frequency plays a major role in the mental representation of constructions (cf. Ellis, 2002; Diessel, 2007). The effects of frequency are often discussed in terms of a construction’s degree of entrenchment (Langacker, 1987: 59, 2008: 16; Schmid, 2020: 205ff.). The more frequently a linguistic expression is used, the more cognitively entrenched it is. Among other characteristics, a high degree of entrenchment correlates with higher cognitive salience (i.e. accessibility) and faster processing (Harris, 1998; Schmid, 2007, 2017; Blumenthal-Dramé, 2012). How this particular view has an impact on the representation of meaning, which is the focus of this book, will be fully discussed in Section 2.1.2.2.

2.1.2 *Constructions: What They Are*

Constructions are considered the basic building blocks on the basis of which complex structures and sentences can be constructed. Given that constructions combine a form with a meaning, the interpretation of an utterance therefore depends on which constructions are being used in a given context and how they are being assembled. In order to understand the individual contribution of these constructions to the interpretation process, it is necessary to look more closely at how CxG defines the notions of *form* and especially that of *meaning*.

2.1.2.1 *The Forms of Constructions* The previous section already referred to the possible forms that constructions can have. It remains to be established exactly what constitutes the form of a construction. CxG considers that the formal pole of a construction includes phonological and morphosyntactic properties (cf. Boas, 2013: 234). To give one example, knowing the construction *admire* consists in knowing that it is pronounced /əd'maɪə/ for instance,⁸

⁸ I say ‘for instance’ since it follows logically from CxG’s usage-based approach that not all English speakers have the pronunciation shown here (e.g. /əd'maɪə/ with a rhotic American accent).

and that it shares the morphosyntactic properties of verbs (e.g. subject–verb agreement, tense inflection, etc.). Not all constructions are phonologically specific, however. Because they are gradually acquired in context, it was shown earlier that constructions may be more or less schematic depending on the degree of abstraction involved (cf. the *lexicon–syntax continuum*, Figure 2.1). Like the verb *admire*, the lexeme *audience* and the idiom *by and large*, for instance, are lexically (and phonologically) fixed constructions. On the other hand, constructions like the *as* ADJ *as a* N construction identified in (15) are only partially specific. Some parts of this construction, the *as* and *as a* elements, are lexically (and phonologically) fixed. The two open slots ADJ and N, however, only specify the morphosyntactic properties that the (phonologically specific) items that fill them should have. Other constructions, such as the DITRANSITIVE construction, are, however, entirely schematic and only specify for morphosyntactic properties. As described in (11), for instance, the DITRANSITIVE construction takes the schematic form [SUBJ V OBJ1 OBJ2]. Only the items that fill the different slots of this construction are phonologically specific, e.g. *me*, *ball*, *threw*, *the* and *Jake* in (17).

(17) Jake threw me the ball. (COCA, written)

The forms of constructions can therefore vary from fully lexically (and phonologically) specific to more schematic. This is not, however, the only way in which constructions have been approached and described in CxG. Constructions are also often discussed in terms of another continuum from atomic to complex constructions (see Croft and Cruse, 2004: 255; Langacker, 2005: 108). That is, as illustrated in Figure 2.2, in addition to being lexically specific or schematic, constructions can also vary in size. From this perspective, increased complexity does not correlate with increased schematicity. Rather, lexically specific constructions can also be very complex. In Figure 2.2, the idiom *break the ice* and the phrase *as soon as possible* are good examples of lexically specific and complex constructions. Partially specific constructions can also be relatively simple (e.g. the *How* ADJ! construction, as in *How adorable!* or *How confusing!*) as well as more complex (e.g. the *X is the new Y* construction, as in (7)). Finally, fully schematic constructions need not always be complex, such as the DITRANSITIVE construction, but can also be simpler (e.g. the AUX V construction as in *have slept* or *should write*).

Representations such as that in Figure 2.2 perfectly illustrate the position adopted in CxG that linguistic knowledge is not strictly divided between words on the one hand and syntactic rules on the other, but that it is composed of a network of more or less complex and schematic constructions. As such, it also nicely captures the perspective adopted in CxG that all of these forms gradually emerge from language use.




Figure 2.2 Fixity and complexity of constructions (adapted from Leclercq, 2023: 67)

There is, however, a central implication of the usage-based approach adopted in CxG that I have not yet discussed, and which directly concerns the form of constructions and in particular that of argument structure constructions (e.g. the DITRANSITIVE and CAUSED-MOTION constructions). The different surface forms in (18) to (19) illustrate what is commonly referred to as the “dative alternation” (cf. Pinker, 1989: 82; Rappaport Hovav and Levin, 2008: 129; Perek, 2015: 154).

- (18) [SUBJ V OBJ1 OBJ2]
 a. We’ll give them a voucher. (COCA, spoken)
 b. They’ll send you the tune beforehand. (COCA, spoken)
- (19) [SUBJ V OBJ2 to OBJ1]
 a. She gave the money to the suspect. (COCA, spoken)
 b. You can send a postcard to us. (COCA, spoken)

In a Chomskyan transformational account of grammar, it has been argued that these different surface forms are derived from a single (deep) underlying syntactic structure (cf. Akmajian and Heny, 1975: 185). In CxG, however, these forms are not treated as variants of the same structure but as two distinct constructions (cf. Perek, 2015: 148, and references cited therein). The pattern in (18), as mentioned before, is referred to as the DITRANSITIVE construction, and the pattern in (19) is referred to as the *To-DATIVE* construction. This distinction is argued to follow logically from the usage-based nature of linguistic knowledge, with generalizations emerging from surface structures (Goldberg, 2002b: 329).⁹ In CxG, what is true for the dative

⁹ Perek (2015: 149) argues that a speaker’s knowledge most probably also includes knowledge of the commonalities between the two constructions, which he discusses, following Cappelle (2006), in terms of an *allostruction*. See Perek (2015: 151–166) for more details.

alternation is (of course) also true for other alternations, such as the causative alternation (cf. Romain, 2017, 2022) and the locative alternation (cf. Perek, 2015: 158). Here, each pattern in the alternation is considered a construction in its own right since each can be identified with its own set of idiosyncratic properties.

Note that the focus on form here is relevant to the semantics–pragmatics interface for a simple reason. Constructions are defined as form–meaning pairings. The DITRANSITIVE and the *To-DATIVE* constructions identified in (18) and (19) should therefore each be associated with a specific meaning. The main question has to do with what meaning is expressed exactly. The DITRANSITIVE construction was described in (11) in terms of a notion of transfer, whereby *X CAUSES Y TO RECEIVE Z*. At first sight, the *To-DATIVE* construction in (19) seems to convey a similar meaning. It is assumed in CxG, however, that differences in form should systematically correspond to differences in meaning. This has been discussed in terms of the *principle of no synonymy* (Goldberg, 1995: 67), recently reframed as the *principle of no equivalence* (Leclercq and Morin, 2023). According to this principle, the *To-DATIVE* construction should therefore serve a different function from the DITRANSITIVE construction. Thompson and Koide (1987: 400) argue that the iconic distance between the SUBJ and OBJ positions in fact reflects a conceptual distance, whereby the *To-DATIVE* construction conveys greater physical distance between the referents of SUBJ and OBJ than the DITRANSITIVE construction. Similarly, Goldberg (1995: 90) argues that the sentences in (18) are better interpreted as conveying *X CAUSES Y TO MOVE TO Z*.¹⁰ That is, both analyses consider the *To-DATIVE* construction to convey greater motion than the DITRANSITIVE construction. As Diessel (2015) points out:

This explains why the verbs *bring* and *take* are particularly frequent in the *to-dative* construction, whereas verbs such as *give* and *tell* are proportionally more frequent in the ditransitive (cf. Gries and Stefanowitsch, 2004). (p. 313)

This observation is meant to show that even seemingly similar patterns can convey slightly different meanings.¹¹ Therefore, it is important for the semanticist, and in particular the pragmaticist, to pay careful attention to the forms of the constructions that speakers use, as they provide rich clues as to the intended interpretation. (In the next sections, it will be shown that this is not systematically the case in Relevance Theory.) It will have become clear that, like in the

¹⁰ Note that Goldberg (1995: 89–97) considers the *To-DATIVE* alternative to be a metaphorical extension of the CAUSED-MOTION construction (SUBJ V OBJ OBL), which she calls the TRANSFER-CAUSED-MOTION construction and which explains her semantic description.

¹¹ That there is a semantic difference between the two constructions has received support from experiments in neuroscience (cf. Allen et al., 2012).

Chomskyan tradition, CxG also tries to account for the generativity of language, i.e. the ability to produce novel sentences (Fried and Östman, 2004: 24). Unlike in the Chomskyan tradition, however, this generativity is not attributed to transformational syntactic rules. Rather, generativity originates from the possibility for meaningful constructions to combine with (and be embedded in) other meaningful constructions. Therefore, as mentioned before, complex sentences are not only syntactically complex but also semantically complex, given that both their form and meaning have to combine. Some of the results behind this combination process will be discussed in Section 2.1.3.

2.1.2.2 Semantics in Construction Grammar The previous section illustrates the challenge that describing the form of a construction in isolation from its meaning can represent. The next step therefore naturally consists in spelling out more explicitly how meaning is defined in CxG. The reader will already have noticed that in spite of this section's title, I have just used the term *meaning* (twice) instead of the term *semantics*. This might appear as a confusing terminological *laissez-faire* to those working on the semantics–pragmatics interface. However, this is a deliberate choice that, as will become clear in this section and the next, actually reflects much of the CxG viewpoint with regard to the functional pole of constructions. For this reason, I will continue using the term *meaning* here and gradually elucidate the reasons why it is preferred – together with the term *function* – to the notion of *semantics*.

The perspective on meaning adopted in CxG can be attributed in particular to Charles Fillmore (1975, 1976, 1982, 1985b), George Lakoff (1987, 1988, 1989) and Ronald Langacker (1987, 1991a, 1991b), whose work has largely contributed to the development of CxG. It is important to understand, however, that CxG also generally embraces most of the ideas on meaning developed in the wider context of cognitive linguistics (see Geeraerts and Cuyckens (2007: 25–418), Geeraerts (2010: 182–272, 2017, 2021) and Lemmens (2016) for detailed overviews). To put it simply, the meaning of a construction is often discussed in terms of a concept, a conceptual structure or a conceptualization (cf. Langacker, 2008: 46). This view is *prima facie* similar to the one adopted in Relevance Theory, which, as we will see, also discusses meaning in terms of concepts (cf. Section 2.2.3.1). However, the two frameworks have a radically different understanding of the nature of concepts. In CxG, as in cognitive linguistics more generally, concepts are understood not in terms of atomic primitives but as more or less complex units of our conceptual system that are internally structured (cf. Lakoff, 1987). This approach was developed in direct opposition to atomic accounts of conceptual content such as the one developed by Jerry Fodor (cf. Fodor, Fodor and Garrett, 1975; Fodor et al., 1980; Fodor, 1998: 40–87, *inter alia*). That is, concepts are considered to be

complex structures. The aim of this section therefore is to understand what type of information concepts make accessible and how this information is organized.


In order to discuss the nature of these conceptual structures, different theoretical constructs have been developed, such as *frames* (Fillmore, 1985b), *idealized cognitive models* (Lakoff, 1987) and *domains* (Langacker, 1987). Although these terms reflect slightly different standpoints, they “are often interchangeable” (Langacker, 2008: 46). For this reason, I will not delve into the particularities of each proposal but will discuss more generally the core assumptions that they all share.¹² A central assumption is that concepts are cognitive objects: “meanings are in our head” (Gärdenfors, 1999: 21). Meaning is therefore not understood in CxG as a bearer of truth-conditions in relation to the external (or some possible) world. Rather, meaning is understood in terms of the way speakers themselves construe and conceptualize the world and particular situations. This has been discussed in cognitive linguistics in terms of the notion of *construal* (Langacker, 1991b: 61, 2019). Consider the sentences in (20) to (21).

- (20) a. The rock is in front of the tree. (Langacker, 2008: 76)
 b. The tree is behind the rock. (Langacker, 2008: 76)
- (21) a. [This type of bird] spends its life on the ground. (Fillmore, 1982: 121)
 b. [This type of bird] spends its life on land. (Fillmore, 1982: 121)

In the sentences in (20), the same situation is being depicted (i.e. both sentences would have the same truth-conditions). They differ, however, in terms of their *vantage point* (cf. Langacker, 2008: 73), i.e. the perspective adopted by the speaker to describe the situation. Similarly, in the sentences in (21), the nouns *ground* and *land* can be used to refer to the same “dry surface of the earth” (Fillmore, 1982: 121). Choosing one or the other, however, depends on whether you construe this surface in relation to the air (22a), or in relation to the sea (22b).¹³ That is, it is argued that their meanings consists of these particular construals, in which some content is understood in relation to a particular background.

¹² Note that there is a tendency among construction grammarians to discuss meaning in terms of *frames* (see Goldberg, 1995: 25; Boas, 2021). This is particularly true when describing the semantics of verbs and argument-structure constructions, for which *frames* offer a relatively more adapted perspective. As Goldberg (1995: 7) herself points out, however, all three approaches have been crucial to the development of CxG.

¹³ Specific terminology has been used to describe the particular construal depicted here: *figure* and *ground* (Langacker, 2008: 58) or *profile* and *base* (p. 66). From this perspective, the nouns *ground* and *land* are used to denote a similar *figure/profile* but in relation to a different *ground/base*.



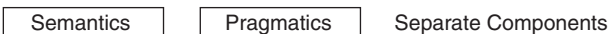
You will notice the particular schematic imagery that underlies the two representations in (22). This captures another central assumption in cognitive linguistics with respect to what meanings are actually composed of. It is assumed that much of a construction's meaning is made of a number of pre-conceptual *image schemas* (Johnson, 1987: xix; Langacker, 2008: 32; Hampe, 2005). Image schemas, as Evans and Green (2006: 184) point out, are not exactly the type of symbolic mental images such as the ones depicted in (22). Still, the notion of *image schema* is meant to capture the observation that much of our conceptual system is shaped by our perceptual and physical experiences, from which conceptual patterns can be abstracted. It was mentioned earlier that a central tenet of CxG is to view language as drawing on general cognitive mechanisms and emerging from language use. This does not only hold for linguistic forms but is also true at the level of meaning. Meaning also gradually emerges from renewed experiences and language use, and it is clear in cognitive linguistics that this experience is not purely mentalistic or intellectual but involves all of our perceptual and physical senses, as well as social and cultural practices:

“Experience,” then, is to be understood in a very rich, broad sense as including basic perceptual, motor-program, emotional, historical, social, and linguistic dimensions. I am rejecting the classical empiricist notion of experience as reducible to passively received sense impressions, which are combined to form atomic experiences. By contrast, experience involves everything that makes us human – our bodily, social, linguistic, and intellectual being combined in complex interactions that make up our understanding of our world. (Johnson, 1987: xvi)

In other words, a central assumption within cognitive linguistics is that meaning is not a purely linguistic notion (and therefore not autonomous) but is encyclopedic in nature, i.e. concepts include knowledge about the world and how we experience it (Croft and Cruse, 2004: 30; Geeraerts and Cuyckens, 2007: 5; Langacker, 2008: 39; Lemmens, 2016: 92; Diessel, 2019: 93; Goldberg, 2019: 12). The meaning (i.e. semantics) of the noun *strawberry*, for instance, includes a whole set of knowledge ranging from its particular shape and color, that it is a (summer) fruit, as well as facts about how they grow and how they are usually sold (i.e. in a punnet), etc. Similarly, as Lemmens (2016: 92) points out, the meaning of the construction *school night* necessarily

includes cultural knowledge of how weeks are divided and organized as well as social practices that are related to *school nights* with regard to the rest of this cultural/social organization (e.g. weekends). This analysis also applies to more phrasal patterns such as the DITRANSITIVE construction discussed earlier. Each of the open slots of this construction, [SUBJ V OBJ1 OBJ2], was described via conceptual primitives: *Agent CAUSE-RECEIVE Recipient Theme*. As mentioned in footnote 6, however, the exact meaning of this construction is actually more complex and cannot be reduced to these primitives (see Goldberg, 1995: 49). The meaning of the DITRANSITIVE construction more largely includes knowledge of how humans engage in acts of transfer (Goldberg, 1995: 39), i.e. knowledge of what a transfer actually involves, of the respective roles of agents, recipients and themes and the relation between them, as well as who/what can usually perform these roles (Goldberg, 1995: 142–151).

Already, it should be clear why the more general terms *meaning* and *function* are therefore preferred to the term *semantics*. The perspective adopted here indeed rejects the traditional division between purely linguistic content on the one hand (usually referred to as *semantics*) and encyclopedic knowledge on the other (usually attributed to *pragmatics*). That is, what is often attributed to pragmatics – as is the case in Relevance Theory – is considered to directly contribute to a construction's *semantics*. There is therefore no strict division between the two (as in (23)), but rather a gradation from semantics to pragmatics (as in (24)), both adapted from Langacker (2008: 40, Fig. 2.4):

(23)  Separate Components

(24)  Gradation

What the representation in (24) is meant to capture is that it is not necessarily clear to what extent, during the interpretation of an utterance, some particular piece of encyclopedic knowledge is already part of a given conceptual structure or is pragmatically derived from the context. Rather, because of the constantly changing nature of conceptual structures, some pieces of knowledge are already well established in the speaker's conceptual structure (i.e. semantic) while others are only in the process of conventionalizing (i.e. partially semantic), and yet others are wholly contextual (i.e. pragmatic). This is not the only reason why the term *semantics* is not often used in cognitive frameworks. One of the reasons comes from another central assumption that meanings are not seen as (context-free) disposable packages that speakers and hearers simply access when using a particular construction. Rather, it is assumed that using a construction only provides a point of access to all of its associated knowledge,

and that meaning is constructed in context (see Evans, Bergen and Zinken, 2007: 9; Radden et al., 2007: 1; Langacker, 2008: 41; Taylor, 2017: 261). That is, the actual meaning of a construction largely depends, in context, on “which portions of this encyclopedic knowledge are activated, and to what degree” (Langacker, 2008: 42). Some parts of this knowledge are so central to the understanding of a particular construction that they systematically get activated across usages, but other (more ‘peripheral’) aspects of knowledge will only get activated in some contexts and not others, i.e. will be more salient in some contexts and not others. For this reason, Langacker (2008: 30) prefers to talk about meaning in terms of *conceptualizations* rather than *concepts*, the former term conveying greater dynamicity than the latter notion, which conveys more stativity. It will become clear in the next chapter that some of these assumptions are also central to Relevance Theory, which I will present in Section 2.2.3.1.

Adopting an encyclopedic view of meaning (or semantics) necessarily requires some further explanation in terms of how this knowledge is organized and represented in the speaker’s and hearer’s minds. It is generally understood in cognitive linguistics that the conceptual structure associated with a particular construction does not simply represent an unstructured “grab bag” of encyclopedic knowledge (Lemmens, 2017: 107). Rather, this knowledge is well structured and organized. Conceptual structures are usually described in terms of categories. There are, however, various ways in which these categories can be described. In CxG, as in cognitive linguistics more generally, categories are often discussed in terms of either a *radial network* (Lakoff, 1987) or a *schematic network* (Langacker, 1987) of encyclopedic knowledge. In either case, it is assumed that our knowledge is organized in a number of interconnected bundles (or clusters) of knowledge, one of which is more central to a given category than others. This more central cluster of knowledge is usually referred to as the *prototype* (Rosch, 1975, 1978, 1983; Rosch and Mervis, 1975; Taylor, 1995).¹⁴ Via an analogical process, the encyclopedic information associated with a given exemplar (i.e. construct) is located within the category in relation to the prototype, either as a more or less specific instance of that prototype or as an extension depending on its resemblance to previously encountered

¹⁴ Note that two perspectives have developed concerning these networks: a prototype model and an exemplar-based model (see Bybee, 2006, 2010, 2013). The difference is whether the individual clusters are made up of all the individual traces (and uses) of encyclopedic knowledge (exemplar-based), or whether there is some form of abstraction involved (prototype). These two approaches are often seen as being in opposition. Like Barsalou (1990) and Hampton (2016), however, I think these are not necessarily opposite analyses but simply two sides of the same coin: while individuals do categorize on the basis of the individual usage events (i.e. exemplars), there is also most probably some abstraction of the knowledge involved (prototype). The possibility for abstractions to be derived does not preclude the possibility for individual exemplars to be stored, and vice versa. In this book, I will continue using the term *prototype*, but without the theoretical assumption that this necessarily precludes the storage by individuals of particular exemplars.




Figure 2.3 Radial network of *play*

exemplars. For instance, the verb *play* gives access to the radial category represented in Figure 2.3.¹⁵

This network represents the category of encyclopedic knowledge that the verb *play* gives access to, and which is organized in different clusters of various resemblance. These different dots (or clusters) constitute the different senses of the verb (which can then be understood as different but related concepts). In this network, one of the different clusters of encyclopedic knowledge (i.e. senses) – shown here in the bold circle – is more central than others and all other senses develop as extensions from this central sense. This representation captures the conventional polysemy of the verb *play*, each of the different clusters representing one of the senses of the verb.

Radial networks such as in Figure 2.3 nicely enable the identification and understanding of the various senses of a given construction by identifying the relation between the different clusters of encyclopedic information that a construction is associated with. Now, independently of whether any kind of abstraction or generalization occurs within the clusters themselves (see footnote 14), Langacker (1991b: 266) suggests that there is schematization (i.e. abstraction) across clusters, and that some of these senses may be more schematic than others, but also that there may be a “superschema” (1991b: 267) that accounts for all of the senses that compose the conceptual network. He calls such a network a *schematic network*. He discusses, for instance, the schematic network of the verb *run*, as shown in Figure 2.4.

Like in the radial network, this network is composed of different clusters or senses, one of which is more central than the others (box in bold). Other senses are seen as semantic extensions from this more central sense (broken arrows). In addition to the radial network, however, some senses are seen as schematic

¹⁵ Note that this figure is a relatively simple representation of the radial category that is associated with the verb *play*, which in reality is probably more complex. This figure was drawn on the basis of the definitions provided by the online *Cambridge Dictionary*, available at: <http://dictionary.cambridge.org/fr/dictionnaire/anglais/play> (last accessed: May 31, 2023).

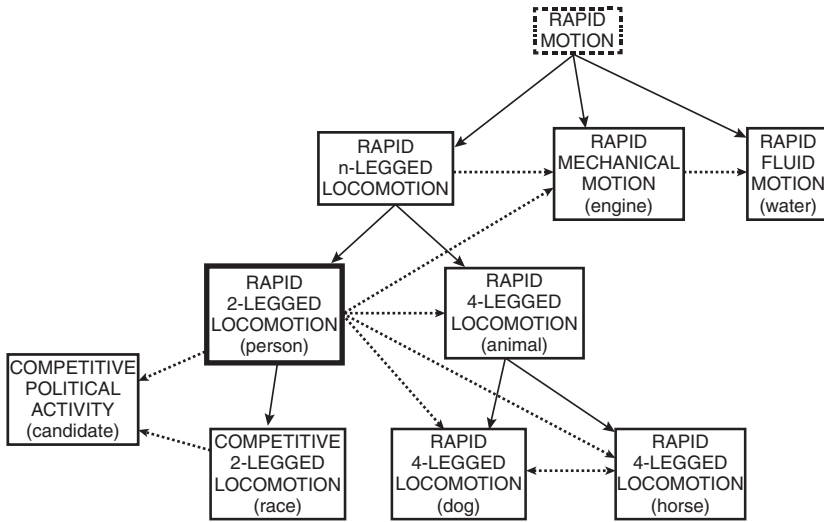


Figure 2.4 Schematic network of *run* (from Langacker, 1991b: 267)

relative to other senses (plain arrows), and one of them is schematic to all these senses (broken box). In the next chapter, I will try to show how this perspective can help to shed some light on different issues that concern the semantics–pragmatics interface, and in particular in relation to the ideas developed in Relevance Theory (see Section 3.4). For now, in order to avoid possible misunderstandings, a few points concerning these representations are worth mentioning. It is true that the different clusters identified correspond to different senses of the construction to which the network is associated, e.g. the verb *run* in Figure 2.4. In CxG, conventional polysemy (a network of interrelated senses) is the norm rather than the exception (Goldberg, 1995: 31, 2019: 20). However, it is necessary to understand, as mentioned before, that this polysemy is neither predetermined nor fixed, and that these senses are not context-free packages. First, these conceptual networks are gradually acquired via exposure to actual exemplars and emerge from this experience. As a consequence, within and across languages, not every individual will share exactly the same conceptual structure (although speakers of the same speech community most certainly have very similar ones). As Langacker (1991b: 267) himself points out, English speakers may not all have within their conceptual structure of *run* clusters of knowledge as specific as some that can be found in Figure 2.4 (e.g. the “bottom” *dog* and *horse* type of running senses). Similarly, not everyone will necessarily abstract the (same) superschema, in this case *rapid motion*. In other words, the categories that individuals possess are relatively flexible, and

constantly change depending on their experience. Second, and directly related to this last observation, it is interesting to note that this usage-based approach easily explains language change and grammaticalization (Bybee, 2010; Traugott and Trousdale, 2013), since it is usage that determines the shape of these categories and how they gradually develop.

2.1.2.3 The Pragmatics of Constructions It will now have become clear that meaning in CxG is to be understood in terms of rich conceptual structures that emerge via our experience of the world and which are constantly evolving and changing. In the next section, I will discuss the interaction between different constructions and, therefore, between different conceptual structures. Before doing so, I will address one last point, which can also explain why the terms *meaning* and *function* are preferred to the term *semantics* in CxG.¹⁶ In this section, I will briefly look at how the notion of *pragmatics* has been discussed in relation to constructions in CxG. It is essential to understand that constructionists, although primarily focusing on what constitutes our linguistic knowledge, do not ignore the role of pragmatics. Ideally, CxG tries to account for all of “the rich semantic, pragmatic, and complex formal constraints” that somehow regulate the use of individual constructions (Goldberg, 2003: 220). However, there has so far been little attention paid to pragmatics in CxG (cf. Cappelle, 2017; Finkbeiner, 2019a; Leclercq, 2020). Furthermore, pragmatics in this framework is understood and approached in a way that differs from how it is generally discussed in the literature on the semantics–pragmatics interface. This is a potential source of confusion.

This particular approach is best illustrated in Goldberg (2004), who distinguishes between *non-conventional* pragmatics and *conventional* pragmatics (Goldberg, 2004: 428). The former kind of pragmatics has to do with online computations of contextual effects such as are usually discussed in (post-/neo-) Gricean pragmatics. The latter, *conventional* type of pragmatics is concerned with “the conventional association of certain formal properties of language with certain constraints on pragmatic contexts” (p. 428). It is this latter type of pragmatics that constructionists are mostly interested in (cf. Kay, 2004; Nikiforidou, 2009; Lee-Goldman, 2011; Cappelle, 2017; Kuzai, 2020). Yet, by virtue of being conventional, one may wonder to what extent this type of ‘pragmatics’ really is ‘pragmatics’ rather than semantics, and what exactly the term is meant to capture. This question is the focus of this section.

It is generally considered in CxG that “some constructions have pragmatic content built into them” (Cappelle, 2017: 116). Some of this pragmatic content follows from the usage-based nature of meaning representation. As Bybee

¹⁶ See Leclercq (2020) for a critical discussion of the use of the terms ‘semantics’ and ‘pragmatics’ in CxG.

(2010) points out, because semantic structures are gradually acquired via repetition, “frequently made inferences from the context can become part of the meaning of an expression or construction” (Bybee, 2010: 52). An often-discussed example in CxG is the *What’s X doing X?* (or WXDY) construction, illustrated in (25).

- (25) a. What’s it doing in the box? (COCA, written)
 b. What’s THAT book doing in the library? (NOW)
 c. And what’s he doing in my kitchen? (COCA, written)

All of the sentences in (25) express a notion of incongruity (or disapproval) regarding a specific situation (Kay and Fillmore, 1999: 4). In (25b), for instance, the speaker seems to disapprove of a given book being available in a specific library. Although Kay and Fillmore recognize that this meaning most probably originated as a conversational implicature, “the semantics of incongruity is now CONVENTIONALLY associated with the special morphosyntax of WXDY constructs” (p. 5, original emphasis). That is, this part of the communicated meaning is not (re)calculated each time the hearer comes across the WXDY construction but is accessed as part of their knowledge of the construction. In this case, a previously pragmatic aspect has become conventional and is now part of the meaning of the construction itself. It is for this reason that Kay and Fillmore refer to it in terms of *semantics* rather than in terms of *pragmatics*.

Cappelle (2017: 118) points out, however – and this seems to have been the underlying reason for Goldberg to discuss the notion of *conventional pragmatics* – that it is not necessarily clear to what extent a conventionalized pragmatic aspect necessarily becomes a semantic aspect of a construction. The functional pole of a construction, it is argued, may actually be composed of both semantic and conventional pragmatic aspects of meaning (and hence the preference for using the words *meaning/function*). I will discuss the case of the *let alone* construction to explain this view. In their oft-cited paper, Fillmore, Kay and O’Connor (1988: 514) discuss the use of *let alone* and its communicative function in sentences such as in (26) to (28) (original emphasis).

- (26) I don’t even want to read an article ABOUT, let alone a book written BY, that swine.
 (27) Max won’t eat SHRIMP, let alone SQUID.
 (28) He wouldn’t give A NICKEL to his MOTHER, let alone TEN DOLLARS to a COMPLETE STRANGER.

In these sentences, the *let alone* construction (or rather, the X *let alone* Y construction) is argued to introduce each of the two conjoined elements in terms of “contrasted points on an implicational scale” (Cappelle, Dugas and Tobin, 2015: 72). In the sentence in (27), for instance, Max’s eating squid is

understood as being less probable than his eating shrimp (which is itself already very unlikely). This aspect is argued to be the semantic contribution of the construction. In addition, it is argued that the construction is also used to indicate to the hearer that the second conjunct (e.g. that there is no chance that Max is going to eat squid) provides the most relevant – in the Gricean sense – piece of information to the context at hand (Cappelle, Dugas and Tobin, 2015: 72). In this case, the construction therefore conventionally provides the hearer with the tools that they would otherwise have had to work out in context. Yet according to Fillmore, Kay and O'Connor (1988: 532), in spite of being conventionally associated with the *let alone* construction, this later contribution is argued not to be purely semantic but is instead pragmatic (see also Cappelle, Dugas and Tobin, 2015: 73).

Note that in addition to the ‘pragmatics’ of examples such as the WXDY construction and the X *let alone* Y construction, a number of other pragmatic functions associated with particular constructions have been identified, such as in (29) to (30).

(29) Can you pass the salt? (Stefanowitsch, 2003: 108)

(30) It's not pretty, it's gorgeous. (Kay and Michaelis, 2012: 2286)

The construct in (29), for instance, is argued to instantiate the *Can you X?* construction, which is conventionally associated with a *request* indirect speech act (Stefanowitsch, 2003: 109). That is, this indirect speech act is not calculated online, but is accessed by the hearer as part of their knowledge of the construction. From the perspective of CxG, this part of the function of the construction is not purely semantic (since it is non-propositional) but is instead a pragmatic convention. Similarly, in the construct in (30), it is argued that the adverb *not* is used here not to negate the proposition itself but as a metalinguistic device to cancel a specific quantity implicature (cf. Kay and Michaelis (2012: 17), and references cited therein). That is, in constructional terms, this function is not semantic but pragmatic.

The previous paragraphs show that there is no clear agreement as to what counts as purely ‘semantic’ or ‘pragmatic’. As discussed in Leclercq (2020), though CxG initially considers the semantics–pragmatics distinction “more or less obsolete” (auf der Straße, 2017: 61), two opposite views can be distinguished: one that considers that semantics is the domain of conventional meaning while pragmatics pertains to contextual inference (e.g. Kay and Fillmore's 1999 analysis of the WXDY construction), and one that views semantics as contributing to propositional (i.e. truth-conditional) content and pragmatics to non-truth-conditional content (e.g. Fillmore, Kay and O'Connor's 1988 analysis of the X *let alone* Y construction). It has to be understood that as much as CxG does not define meanings in terms of truth-

conditions, and sees meaning as a very dynamic object (see previous section), constructionists still primarily seem to think of the notion of *semantics* as being propositional in nature (although it is not clear exactly what their position on the matter really is, e.g. Kay and Michaelis, 2012: 2277). This assumption is in particular defended by Cappelle (2017: 122), who argues that it is “useful to make a distinction between lexical or propositional semantics . . . and pragmatic information.” From this perspective, the information provided for example by the *let alone* construction about the second conjunct in terms of its particular relevance to a given discourse is indeed not semantic since it is non-propositional. This assumption is also central to Goldberg’s (2004) distinction between conventional and non-conventional pragmatics, both of which concern non-truth-conditional aspects of meaning. As mentioned in Leclercq (2020), one might argue that this view comes into direct contradiction with the approach presented in Section 2.1.2.2, in which it was argued that truth-conditions do not define the meaning of a construction. As I see it, there is no necessary contradiction, however: “[o]ne can maintain the view that meaning is not restricted to truth conditions and at the same time argue that there is a level at which some aspects of meaning (more than others) will eventually contribute to establishing the truth value of the proposition expressed” (Leclercq, 2020: 232). Indeed, as Gärdenfors (1999) puts it, “the truth of expressions is considered to be secondary, since truth concerns the relation between the mental structure and the world. [. . .] Meaning comes before truth” (p. 21). So it is possible for CxG to preserve a clear semantic/pragmatic distinction and to treat as purely *semantic* those (conventional) aspects of meaning that are eligible to contextual truth values and as *pragmatic* those that are not. We will see in Chapter 4 how this question relates to the conceptual/procedural distinction established in Relevance Theory.

Now, regardless of how semantics and pragmatics are understood in CxG, it is relatively explicit that non-conventional pragmatics (Goldberg, 2004: 428), i.e. conversational pragmatics as discussed in the (post-/neo-)Gricean tradition, is not the focus of interest in CxG. Rather, the interest remains to a large extent centered on knowledge itself. Nevertheless, because of CxG’s usage-based approach to language, as mentioned before, it does recognize the role of pragmatics during verbal communication. Unfortunately, constructionists sometimes leave some room for ambiguity concerning their exact stance on the question. For instance, the title *What’s pragmatics doing outside constructions?* in Cappelle (2017) is a bit surprising. It might seem as though the author is rejecting the possibility for pragmatics to exist outside of constructions. The *What’s X doing Y?* construction used here, which Cappelle himself discusses, indeed introduces a notion of incongruity and disapproval for the relation between X and Y. Does this mean that Cappelle (and constructionists more generally) rejects the role of conversational, non-conventional pragmatics?

No. Although this is what his title could suggest, this is not the position of CxG. Cappelle himself acknowledges the major role of (non-conventional) pragmatics during the interpretation of an utterance (Cappelle, 2017: 117). The aim of Cappelle (2017) is only to show that pragmatics can also be part of constructions, and that it does not only take place outside of constructions.¹⁷ Of course, the aim of CxG is only to describe an individual's linguistic knowledge, and therefore it could be argued that it does not need to explain the ins and outs of conversational pragmatics. In this case, however, CxG on its own is not enough to fully explain how linguistic communication works and succeeds. In order to do so, one needs an account of how non-conventional pragmatics operates. Goldberg (2006) argues that

[a] focus on form to the neglect of function is like investigating a human organ such as the liver, without attending to what the liver does: while this is not impossible, it is certain to fail to be explanatory. (p. 168)

Similarly, one might argue that focusing on semantics without looking at pragmatics, although not impossible, has little explanatory power. As González-García (2020: 112) puts it, “the treatment of semantic and/or pragmatic facts in [Cognitive Construction Grammar] is at best somewhat inconsistent with the theoretical premises invoked.” This is why an approach along the lines of Relevance Theory is necessary.

2.1.3 *Constructions in Interaction: Coercion*

There is one more aspect of the theory that I will now discuss and which will prove particularly relevant when comparing CxG with Relevance Theory. This has to do with the interaction between various constructions that are used in an utterance. As mentioned before, utterances result from the complex combination of a number of different constructions, all of which contribute to this utterance in terms of both form and meaning. Goldberg (2003: 221) nicely captures this complexity in Figure 2.5, for instance.

Figure 2.5 illustrates the constructionist view according to which one utterance (in (a), *What did Liza buy the child?*) results from the combination of various constructions, identified in (b). Interestingly, Figure 2.5 is but one illustration of the generativity of language which can be explained by the infinite combinatorial possibilities that constructions offer (Goldberg, 1995: 7; Fried and Östman, 2004: 14). The aim of this section is to understand more specifically the scope of these possibilities and the ways in which

¹⁷ A similar observation can be made about Bergs and Diewald's (2009) *Contexts and Constructions* volume. Although some chapters do address non-conventional pragmatics, most are concerned with how pragmatics can be part of constructions.

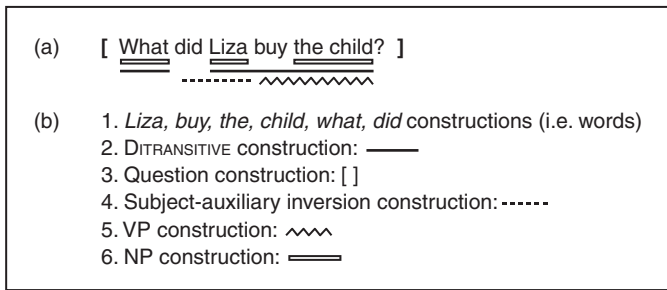


Figure 2.5 Constructing an utterance (adapted from Goldberg, 2003: 221)

constructions, and in particular lexemes, may (or may not) interact with other constructions. Consider for instance the use of *fly* in the following examples.

- (31) The pilot announced that geese were *flying* in the sky. (COCA, written)
- (32) Our son was just three months old when we first *flew* him across the Atlantic. (COCA, written)
- (33) It was a breathtaking experience to fly over the beautiful valley of Palampur. Why walk, when you can *fly* your way down! (NOW)

Example (31) contains a rather prototypical use of the verb, in both form and meaning. The uses of *fly* in (32) and (33), however, are comparatively more unusual. In (32), *fly* is interpreted in terms of a caused motion, whereby the parents have taken their son on a flight across the Atlantic at three months of age. In (33), there is a particular focus placed on the specific *manner* of motion which is encoded by the verb *fly*. In CxG, it is considered that these interpretations actually result from their being used inside another construction, the CAUSED-MOTION construction and the WAY construction respectively. We can see indeed that these sentences differ from (31) not only in meaning but also in form. In (32), the form [SUBJ V OBJ OBL] can be easily identified. As mentioned in Section 2.1.1, this form is associated with the meaning *X causes Y to go Z* to form the CAUSED-MOTION construction. It is understood in CxG that the interpretation of *fly* in (32) follows from its being used in that construction. Similarly in (33), the form [SUBJ V *one's way* OBL] is argued to be associated with a particular manner of motion interpretation to form the WAY construction, and it is the use of *fly* in this construction that provides its particular interpretation in (33). So the interpretation of the verb *fly* here depends as much on the semantics of the constructions in which it is used as on its own lexical meaning.

In the example just discussed, it could be argued that the verb *fly* readily combines with the semantics of the two constructions in which it occurs.¹⁸ This naturally follows because there is *semantic coherence* and *semantic correspondence* (cf. Goldberg, 1995: 50) between the lexical item and the two constructions, i.e. the semantic features of the lexical item that is used closely correspond to those of the constructions in which it occurs and therefore can combine (or *fuse*) with each of them. The verb *fly* indeed refers to a specific motion event which is a central aspect of both the CAUSED-MOTION and the WAY constructions. Sometimes, however, a lexeme and the construction in which it occurs are not semantically (and morphosyntactically) coherent and compatible. Consider the sentences in (34) to (36).

- (34) He drank three *beers* hoping that would help. It did not. (COCA, written)
- (35) I just *Google Mapped* my way to an exam because I didn't know where Engineering South was. #senioryear (Twitter, @Brittany_N_Lee, 21 apr. 2016)
- (36) The doc (. . .) had *happied* himself to death on his own laudanum two months before. (COCA, written)

In (34), the noun *beer* is a mass noun that should not be specified for number, i.e. it should neither be inflected with the plural *-s* suffix nor with the number *three*, and which is prima facie incompatible with the morphosyntactic context in which it occurs.¹⁹ Nevertheless, instead of the mass reading (i.e. the liquid), *beer* is here interpreted in accordance with its morphosyntactic context in terms of a countable portion of beer (given our world knowledge of how beer is usually served, most probably three bottles or pints). That is, its interpretation is somehow inherited from the semantics of the construction in which (and with which) it occurs. Similarly in (35), the noun *Google Maps*, which usually refers to a GPS application, is here used as a denominal verb to refer to a particular means of motion. That is, the speaker communicates that they used Google Maps in order to arrive at the location of their exam. CxG has a specific explanation for the origin of this interpretation. This interpretation is argued to be inherited from the WAY construction (SUBJ *V one's way* OBL) in which it occurs. The lexeme *Google Maps* is originally neither a verb nor does it encode means of motion (which is expected in that position of the WAY construction). At first sight, there is a semantic (and morphosyntactic) mismatch between the noun and the WAY construction. Nevertheless, it is interpreted as a means of motion verb in accordance with the position it occupies in that construction. Finally, in (36), the adjective *happy* is also used creatively as a de-adjectival

¹⁸ This will be discussed more fully in Chapter 4.

¹⁹ Hilpert (2019: 17) discusses a similar example.

verb. In this case, it is not used to refer to the doc's mental state, but to the (metaphorical) act of leading himself to his own death by using drugs. This can once again be explained in terms of the larger construction in which it occurs: in this case the CAUSED-MOTION construction (SUBJ V OBJ OBL). Although *happy* is neither a verb nor encodes caused-motion (as required by the CAUSED-MOTION construction), it is so interpreted in accordance with the larger construction in which it occurs.

The examples just discussed show that sometimes a lexeme can be used in a construction with which it is seemingly incompatible, i.e. there can be a mismatch between the semantic and morphosyntactic properties of the lexeme and those of the construction in which it occurs. In these cases, the lexeme is systematically reinterpreted in accordance with the semantics of that construction. This resolution process has been discussed in CxG in terms of *coercion* (Goldberg, 1995: 159; Michaelis, 2003a, 2003b, 2004; Lauwers and Willems, 2011; Hilpert, 2019: 17; Leclercq, 2019). The term *coercion* was first used outside the framework of Construction Grammar. The terms 'coerce' and 'coercion' were initially used in programming languages (Ait-Kaci, 1984) and artificial intelligence (Hobbs, Walker and Amsler, 1982; Hobbs and Martin, 1987; Hobbs et al., 1993) and were soon adopted and developed by formal semanticists interested in aspectual meaning (Moens and Steedman, 1988: 17; Pustejovsky, 1991: 425, 1995: 106, 2011: 1401; de Swart, 2000: 7, 2011: 580). It is from this work on aspect that Construction Grammar has borrowed the term *coercion*. In particular, Michaelis (2004) took up the notion and adapted it to the needs of the theory. The term has been used to describe a variety of phenomena in the different frameworks just mentioned. Nevertheless, they all share the view that coercion is concerned with the resolution of an incompatibility between a selector (e.g. argument structure construction) and a selected (e.g. lexeme) whereby the latter adapts to the former. This has been referred to by Michaelis (2004) as the "override principle":

The override principle. If a lexical item is semantically incompatible with its morphosyntactic context, the meaning of the lexical item conforms to the meaning of the structure in which it is embedded. (p. 25)

The interpretations of the lexemes *beer*, *Google Maps* and *happy* as they are used in (34) to (36) precisely follow from such coercion effects: they are interpreted in accordance with the semantics of the different constructions in which they occur. As mentioned before, the same coercion effect is involved in example (13), repeated here:

(37) Lyn crutched Tom her apple. (Kaschak and Glenberg, 2000: 512)

The interpretation of the denominal verb *crutch* in terms of transfer (whereby Lyn used the crutch to give Tom her apple) is inherited from the DITRANSITIVE construction in which it occurs (SUBJ V OBJ1 OBJ2). Although there is originally a semantic (and morphosyntactic) mismatch between the noun *crutch* and the verb position it occupies in the DITRANSITIVE construction, the noun is reinterpreted in accordance with the semantics of the construction.

There are three further points concerning coercion that I wish to address. First, it has to be understood that the notion of coercion has been widely discussed in a number of different frameworks and that it is perceived and described in slightly different ways in each of them (see Audring and Booij (2016) for an interesting discussion). For instance, in the tradition of formal semantics (cf. Pustejovsky, 1991, 1995, 2011; de Swart, 2000, 2011; Jackendoff, 1997), on the basis of which Michaelis has elaborated her own account, coercion seems to be understood as an autonomous linguistic process whereby language itself coerces (or shifts) the meaning of a particular lexeme. Lauwers and Willems (2011: 1224) point out that these approaches have indeed given very little attention to the role of language users and context during the interpretation process. Such a position somehow seems to resonate in Michaelis' own work in any case, and in particular in the way the override principle has been stated. Because of CxG's usage-based approach to language, however, I believe that most construction grammarians view coercion as involving a process whereby the language users themselves have to solve the mismatch between the lexeme and the morphosyntactic context in which they occur. Such a perspective has been nicely captured by Langacker (1987):

Putting together novel expressions is something that speakers do, not grammars. It is a problem-solving activity that demands a constructive effort and occurs when linguistic convention is put to use in specific circumstances. (p. 65)

This particular point of view will be addressed more fully in Chapter 4 and I will take into account other arguments when relating the notion of coercion to some of the work developed in Relevance Theory, for, as we will see, this notion raises a lot of questions. In particular, and this is my second point, one of these questions has to do with how the "problem-solving activity," i.e. the mismatch resolution, is actually accounted for. After all, as Yoon (2012) points out, although coercion in CxG is seen as a process carried out by language users, "the psychological process toward the resolution [is] not dealt with" (Yoon, 2012: 7). This question will also be addressed in Chapter 4.

Finally, it is also worth mentioning that there is a limit to coercion. Coercion is possible because of the productivity of the constructions that speakers use (Lauwers and Willems, 2011: 1224), i.e. the possibility for a construction to produce novel forms and combine with new lexemes. However, this

productivity is constrained (cf. Cappelle, 2014), that is, it is not possible for any construction to combine with any new lexeme and for coercion (i.e. mismatch resolution) to take place. This restriction is often discussed in terms of *partial productivity*, *coverage*, *competition* and *statistical preemption* (Goldberg, 1995: 120, 2019; Boyd and Goldberg, 2011; Goldberg, 2011; Suttle and Goldberg, 2011). That is, as productive as constructions may be, there are systematic constraints that limit the range of possible coinages. Given the focus of this book on lexical meaning, however, I will not delve into the specific constraints on constructional productivity but only focus on those cases where novel forms (such as in (34) to (37)) are possible and give rise to coercion effects (see also Bergs, 2018; Hoffmann, 2018; and references cited therein).

2.1.4 Construction Grammar: Summing Up

Construction Grammar is a cognitively grounded theory of language that mostly focuses on knowledge. As a functionalist approach, it assumes that all forms that a language is composed of are essentially meaningful. From this perspective, meaning should be at the very heart of linguistic analysis. In particular, construction grammarians often discuss the semantics (or rather *function*) that is associated with any given construction. Due to its particular appeal to usage, we have seen that CxG has a particular understanding of the notion of semantics. As mentioned in Section 2.1.2.2, meaning is discussed in terms of encyclopedic knowledge. It therefore rejects the traditional division between semantics (as purely linguistic knowledge) and pragmatics (in terms of encyclopedic knowledge), but rather believes in a more gradual distinction from more to less conventional aspects of encyclopedic knowledge. As a result, it is assumed that the functional pole of constructions is rather rich and that polysemy is almost systematic. In the next chapter, this view will be compared to the relevance-theoretic approach, and I will discuss whether or not the two are compatible. What is particularly interesting in CxG is that it considers all different forms of a given language to be systematically associated with different functions. As discussed in the case of the dative alternation, even small differences in form are related to differences in function. This therefore provides the analyst – whether it be a semanticist or a pragmaticist – with a relatively clear agenda. Constructions indeed provide a solid source of information in order to identify the speaker's intended interpretation. Hence, trying to identify a speaker's meaning should always involve looking carefully at the particular form of an utterance, from which a number of specific functions can be recovered and to which pragmatic principles can apply. As we will see in Chapter 3, however, this multifaceted strategy is not always adopted in Relevance Theory.

The difficulty with Construction Grammar is to pin down exactly its view on pragmatics. On the one hand, it has been shown that much pragmatic content is considered to be part of a construction's function. This was referred to as conventional pragmatics. On the other hand, it is less clear how non-conventional (i.e. conversational) pragmatics generally fits in with the general enterprise of CxG, which gives rather little space to these facets of 'meaning'. Of course, constructionists may argue that CxG aims only at providing a framework for linguistic knowledge and that (non-conventional) pragmatics falls outside its scope. However, it is not clear why this should necessarily be the case. CxG assumes that linguistic knowledge results from one's experience with language. From this perspective, the experience itself – which involves non-conventional pragmatics – should be as much an aim of study as the resulting knowledge. It has been shown, however, that construction grammarians tend to focus on the result itself (i.e. on knowledge) more than they do on the experience. This is the reason why, for instance, speakers are sometimes credited with too much knowledge (e.g. Sandra and Rice (1995) on prepositions, see also Sandra (1998: 368)) that instead should probably be attributed to pragmatics. In order to arrive at a cognitively more accurate description of linguistic knowledge and language more generally (which was earlier stated as being one of CxG's main goals, CxG may therefore have to integrate principles of pragmatics more explicitly in its framework.

2.2 Relevance Theory

Relevance Theory is a theory of cognition and cognitive processes which has mostly been applied to verbal communication. In this framework, much focus is placed on the semantics–pragmatics interface and the processes involved during the interpretation of an utterance. It was originally developed by Dan Sperber and Deirdre Wilson and was first fully spelled out in their seminal book *Relevance* ([1986] 1995). Relevance Theory grew out of Sperber and Wilson's desire to make sense of our capacity to understand the world, and in particular our capacity to communicate effectively. As such, it had a direct impact on the field of pragmatics. Indeed, Paul Grice – whose work provided a major incentive for the development of pragmatics – had already addressed some of the most central issues discussed in Relevance Theory (Grice, 1989). Yet, although sharing a number of Gricean assumptions, Sperber and Wilson developed Relevance Theory as an alternative account to that of Grice (and other post- and neo-Gricean theories) and thus challenged traditional perspectives on pragmatics. In Relevance Theory, the success of verbal communication is not attributed to a number of maxims, or rules, that interlocutors follow, but to a single cognitive mechanism referred to as the *principle of relevance* (see Section 2.2.1).

Since the publication of *Relevance* in 1986, the theory has been extensively revised and extended to address many of the issues discussed in the pragmatics literature (cf. Carston, 2002a; Wilson and Sperber, 2012; Clark, 2013a; Wilson, 2017; Allott, 2020). The length of Francisco Yus' up-to-date online bibliography (<http://personal.ua.es/francisco.yus/rt.html>), which gathers almost all of the research embedded in a relevance-theoretic perspective, bears witness to this. The variety of contributors to the development of the theory has naturally led to diverging points of view within the framework itself. Nevertheless, relevance theorists have remained relatively united and, for that reason, I will continue using the general term Relevance Theory (and its acronym RT) in spite of individual differences across some of its advocates.

2.2.1 *Principle(s) of Relevance*

The reason Relevance Theory has remained a stable framework for so many years is most probably the fact that in spite of internal differences among relevance theorists, the underlying assumption responsible for the development of the theory has never changed and still inspires many researchers. This assumption I referred to earlier as the *principle of relevance*. This principle was first introduced in *Relevance* (Sperber and Wilson, [1986] 1995) and remains today the most central element of the theory around which other ideas have been developed. It is already worth noting that although a number of relevance-theoretic notions will be challenged in this book, the *principle of relevance* will not be one of them. This principle offers key answers to some of the questions that Construction Grammar fails to address, and this is why it deserves a section of its own.

2.2.1.1 *Defining Relevance* Understanding the term *principle of relevance* in RT can be a challenge for at least two reasons. First, part of the difficulty lies in the ambiguity of the term used, for there are actually not one but two principles of relevance: the first (or cognitive) principle of relevance and the second (or communicative) principle of relevance (Sperber and Wilson, 1995: 261). In the literature, however, it is common to refer to the second principle as *the principle of relevance* (see Section 2.2.1.2).²⁰ Second, another difficulty concerns the meaning of the term *relevance* itself. This notion is used in RT in a very technical sense, which differs both from the everyday perception of 'relevance' and from Grice's understanding of the notion. Yet, in order to understand the two principles of relevance (i.e. how and to what phenomena

²⁰ This is most probably because the second/communicative principle of relevance was originally the only one discussed in the 1986 edition of *Relevance*. Claims about cognition were only turned into a principle in the second edition of the book in 1995.

the notion of relevance applies), it is essential to define exactly what is meant by *relevance* in the first place.

As mentioned before, RT is first and foremost a theory about cognition. In particular, it aims at explaining, in cognitively realistic terms, information processing, and notably how inferential processes are constrained (Wilson and Sperber, 1991: 586; Sinclair and Winckler, 1991: 13; Sperber and Wilson, 1995: 32, 66; Clark, 2013a: xv). The notion of relevance therefore does not apply only to language but to all possible types of cognitive stimuli: visual, auditory, kinesthetic, etc. The perspective developed in RT was originally based on Jerry Fodor's *language of thought* and modularity of mind hypotheses (Fodor, 1975, 1983).²¹ The term *cognition* is understood in RT as having "to do with 'thinking'" (Clark, 2013a: 91). Like Fodor, RT assumes that thoughts are language-like mental (or conceptual) representations and that thinking (i.e. cognition) is the computation of these mental representations (Sperber and Wilson, 1995: 71). External stimuli are taken as input by specialized modules (or input systems) which transform the type of information they receive (visual, linguistic, etc.) into mental representations of the same format: *logical forms* (p. 72). These mental representations are then used as input information by central cognitive systems which perform computations over them. The information provided by the various input systems is integrated with information stored in memory and various inferential tasks are then performed (p. 71). These inferential tasks form the basis of comprehension processes and belief fixation. It is such processes, for instance, that enrich the often-incomplete logical forms into full-fledged assumptions (i.e. fully propositional conceptual structures). As we will see in the next section, most linguistic logical forms provided by a given utterance need to be enriched. More generally, the inferential tasks performed by the central cognitive systems enable an individual to keep their representation of the world, their belief system, as accurate as possible by comparing the newly formed assumptions with those already stored in memory (Clark, 2013a: 96).

According to Fodor, the computations that take place within the central systems are primarily inferential and require different types of pragmatic abilities (Fodor, 1983: 110). He argues, however, that it is not possible to describe exactly how these inferential tasks are carried out and how all of the information that comes into somebody's central systems is actually processed in order to keep their belief system up to date (p. 112). Sperber and Wilson disagree with Fodor, however, and they introduced the notion of *relevance* as an attempt to provide such an explanation (Sperber and Wilson, 1995: 66).

²¹ Note that RT has since moved away from Fodor's (1983) original ideas and now adopts a slightly different view of the modularity of the mind (see, for instance, Sperber, 1994, 2001, 2005; Carston, 1997a, 2006; Wilson and Sperber, 2004; Wilson, 2005). Issues related to this debate will not be addressed in this book.

According to Sperber and Wilson, not all information is equally worth processing. Information worth processing is *relevant* information. First introduced as a lawlike generalization, the notion of relevance has laid the foundations of RT.

Relevance is not an either/or property of a given input but, rather, is a matter of degree (Sperber, 2005: 63). The relevance of an input is defined in terms of a balance between cognitive effects and processing effort (Wilson and Sperber, 2004: 609):

- (38) **Relevance of an input to an individual**
- a. Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.
 - b. Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time.

The first part of this definition makes it clear that the more positive cognitive effects a stimulus provides to an individual in a particular context, the more relevant it is. A stimulus has cognitive effects whenever, once integrated with information already stored in memory, it causes a change in an individual's set of assumptions (their *cognitive environment*, Sperber and Wilson, 1995: 38). According to Sperber and Wilson (1995: 109), there are in particular three ways in which a given stimulus can have such cognitive effects. First, a new stimulus achieves cognitive effects if it strengthens an already existing assumption. This happens, for instance, if you suspect two of your friends might be dating, yet do not have much proof. Then one day, you see them walking in the street holding hands. This new information achieves cognitive effects by strengthening your assumption that they are dating. Second, a new stimulus also achieves cognitive effects if it contradicts an already existing assumption. For instance, you might think that your husband is out working in the garden, but then you suddenly see him in front of the computer checking his email. This new information achieves cognitive effects by contradicting a previously held assumption. Finally, a new stimulus also has cognitive effects if, when integrated with older assumptions, it leads to the derivation of a new assumption (called a *contextual implication*, Sperber and Wilson, 1995: 107). For instance, you have been informed that the concert will be canceled if it rains. A few hours before the event, as you walk back home, a storm hits town. This new information achieves cognitive effects since you derive the contextual implication that the concert will be canceled (and therefore you can stay at home, etc.). According to Wilson and Sperber (2004: 608), the derivation of contextual implications provides the most important type of cognitive effects.

The relevance of a given input is not a specific quantitative notion that can be easily measured but is, rather, comparative (Sperber and Wilson, 1987: 742,

1995: 132). The relevance of an input indeed varies depending on the number of cognitive effects it provides in a specific context. The same input may be more or less relevant depending on the context or the strength²² of the assumptions that are manifest to an individual at a given time. More importantly, as the definition in (38b) clearly indicates, the relevance of an input is also a function of the mental effort spent on processing it (Sperber and Wilson, 1995: 124). The amount of effort required depends on a number of different factors such as “recency of use, frequency of use, perceptual salience, ease of retrieval from memory, linguistic or logical complexity and size of the context” (Clark, 2013a: 104; see also Allott (2013: 66) and references cited therein). The more processing effort is needed by a given stimulus, the less relevant it is.

2.2.1.2 Principles of Relevance It will have become clear from the preceding section that various sources of information can be more or less relevant depending on the mental processing effort they require as well as the number of cognitive effects they achieve. A central claim within RT is that we do not pay equal attention to these different inputs. Rather, it is claimed that we tend to allocate cognitive resources only to those inputs that are maximally relevant. This is referred to as the first (or cognitive) principle of relevance (Wilson and Sperber, 2004: 610):

(39) **Cognitive Principle of Relevance**

Human cognition tends to be geared to the maximization of relevance.

This means that both our specialized input systems (which compute external stimuli) and the central systems (which compute internal stimuli) tend to devote cognitive resources to inputs that provide as many cognitive effects as possible for as little processing effort as possible (Sperber and Wilson, 1995: 261). This tendency, it is argued, results from a biological evolution caused by the systematic pressure towards relevance (Sperber, 1996: 114, 2005: 67; Sperber and Wilson, 2002: 13; Wilson and Sperber, 2004: 110). Of course, it is important to understand that, in and of themselves, most natural stimuli do not provide any indication of their potential relevance. As a consequence, the selection of maximally relevant information is said to follow from a specific heuristic that involves “local arbitrations, aimed at incremental gains, between simultaneously available inputs competing for immediately available resources” (Sperber and Wilson, 1995: 261; cf. also Sperber, 2005: 63). How exactly such heuristics apply to language will be

²² The strength of an assumption refers to the extent to which an assumption is accepted as true by an individual (Sperber and Wilson, 1995: 75; see also Wilson, 2022 on the strength of communicated propositions).

discussed in the next section. According to Sperber and Wilson (1995: 156), however, some stimuli do provide an indication of their own potential relevance to an individual: *ostensive* stimuli. This hypothesis forms the basis of the second (or communicative) principle of relevance (Wilson and Sperber, 2004: 612):

(40) **Communicative Principle of Relevance**

Every ostensive stimulus conveys a presumption of its own optimal relevance.

A stimulus is ostensive whenever it indicates an individual's intention to communicate and to be informative (Sperber and Wilson, 1995: 54–64). Utterances are probably the paradigm case of ostensive stimuli (hence the focus on language in RT), but other types of stimuli can also be ostensive (e.g. gesture). In RT, it is argued that ostensive stimuli – such as a speaker's utterance – systematically communicate a presumption of their own optimal relevance, as described in (41) (from Wilson and Sperber, 2004: 612).

(41) **Presumption of optimal relevance**

- a. The ostensive stimulus is relevant enough to be worth the audience's processing effort.
- b. It is the most relevant one compatible with the communicator's abilities and preferences.

From this perspective, ostensive stimuli are systematically expected to provide an individual with enough cognitive effects to justify the amount of effort required to process them. The consequences of this presumption are twofold. First, because of this expectation, an individual systematically processes ostensive stimuli in such a way as to optimize their relevance. In addition, it also follows from this presumption that individuals who intend to communicate a given assumption need to make sure the ostensive stimulus they use can be optimally processed by the addressee (Sperber and Wilson, 1995: 157).²³ Both these corollaries have been tested and have received support from experimental evidence (cf. Sperber, Cara and Girotto, 1995; Girotto et al., 2001; van der Henst, Carles and Sperber, 2002; van der Henst, Sperber and Politzer, 2002; van der Henst and Sperber, 2004; van der Henst, 2006; Gibbs and Bryant, 2008). More generally, it has to be understood that the communicative principle of relevance is not meant as a specific rule that individuals need to follow (and could therefore violate) but is introduced in RT as a lawlike generalization about what the human mind does whenever it is faced with ostensive stimuli (Sperber and Wilson, 1995: 162). In the next section, we will see how this principle applies to linguistic stimuli and in particular how it is relevant to the field of lexical pragmatics. Note

²³ This second consequence has in fact recently led Park and Clark (2022) to aptly argue that there is a *relevance-focused production heuristics*.

that in order to facilitate the discussion, I will generally refer to the second (communicative) principle of relevance as *the* principle of relevance. Given the focus on language and lexical pragmatics, this is the most relevant of the two principles for our discussion.

2.2.2 *Meaning, the Underdeterminacy Thesis and Comprehension Heuristics*

The notion of relevance has been primarily applied to linguistic communication. As mentioned above, the use of language is indeed a paradigm case of an ostensive stimulus (see discussion in Assimakopoulos, 2022) and therefore raises expectations of optimal relevance. Speakers are thus expected to make their contribution worth the hearer's processing effort and, concurrently, hearers tend to look for an interpretation that provides enough cognitive effects to justify the amount of effort invested in the process. This observation has been used in RT in particular to explain the pragmatics of linguistic communication.²⁴ In RT, the term *ostensive-inferential communication* is in fact generally preferred to the term linguistic communication (Sperber and Wilson, 1995: 50). It is a central assumption in RT that linguistic communication does not only consist in the ostensive use of linguistic conventions. Rather, in order for linguistic communication to be effective, and for the speaker's meaning to be fully recovered, much inferencing also needs to take place. Inference is needed, for instance, for the derivation of *implicatures*, this latter term having been introduced by Paul Grice (1989: 24). Consider the following dialogue:

- (42) LAURA: Would you like some more chicken?
 PETER: I'm full, thanks.

Here, in order to understand Peter's answer and act in accordance with the information it conveys, Laura needs to derive the implicature that Peter does not want any more chicken. Implicatures are textbook examples of what an individual needs to infer and what requires good pragmatic competence (Zufferey, Moeschler and Reboul, 2019). It is clear for relevance theorists, however, that much more than just implicatures actually need to be inferred in order for communication to be successful. It was noted earlier that linguistic logical forms often need to be enriched into full-fledged (i.e. fully propositional) assumptions. It is indeed argued in RT that a linguistic input often fails to provide all of the information that is being explicitly communicated by an

²⁴ In RT, the term *pragmatics* refers to inferential processes and inferred meanings and the term *semantics* usually refers to the encoded meaning of an expression (Wilson, 2011: 9).

individual. This is usually discussed in terms of the *underdeterminacy thesis* (Carston, 2002a: 19):

(43) **The underdeterminacy thesis**

Linguistic meaning underdetermines what is said.

That is, it is assumed in RT that there is a gap between the meaning of the words that we use (i.e. *linguistic* meaning) and the content of the proposition that is explicitly expressed (i.e. *what is said* in Gricean terms (Grice, 1989: 25)). The communicated propositions are systematically richer than the meaning of the linguistic conventions used to convey those propositions. According to this view, much inferencing is therefore also needed at the explicit level of communication in order to derive a speaker's intended interpretation. There are various sources of linguistic underdeterminacy. Carston (2002a: 28) identifies six such sources. The examples given in (44) to (49), and in particular the italicized items, illustrate each of these sources.

- (44) Multiple encodings (i.e. ambiguities)
- a. Give me my *bow*. (COCA, written)
 - b. Both of them really get to every *ball*. (COCA, spoken)
- (45) Indexical references
- a. *She* pointed out some consequences of not wearing the correct shoe. (NOW)
 - b. Some of the wealthiest people in the world live *there*. (NOW)
- (46) Missing constituents
- a. I was just wondering if you were good *enough*. [for what?] (BNC, written)
 - b. Chelsea are a *better* bet for trophies. [than whom?] (BNC, written)
- (47) Unspecified scope of elements
- a. Everyone *isn't* perfect. (COCA, written)
 - b. We need 103 Canberrans to bake *a* cake. (NOW)
- (48) Underspecificity or weakness of encoded conceptual content
- a. What it comes down to is trying to give *children* a better Christmas. (BNC, spoken)
 - b. Let's be more efficient and make the tax payers' *money* be used wisely. (BNC, spoken)
- (49) Overspecificity or narrowness of encoded conceptual content
- a. It's so *empty* I can hear the tick of a wristwatch from three rows away. (NOW)
 - b. That's exactly what it was. I feel that I loved a – a *teddy bear* for 15 years, and suddenly I've met this young man who – who has everything I wanted my son to have. (COCA, spoken)

Prior to Carston, Grice had already acknowledged the types of underdeterminacy identified in (44) and (45), for which disambiguation and reference assignment are required (Grice, 1989: 25). He does not explain, however, how

these processes are carried out (Carston, 2002a: 21). Utterances such as in (46) are said to be missing a constituent (and therefore underdetermine what is said) since no truth value can be attributed to them without this constituent.²⁵ In (46a), for instance, one needs to know exactly what quality or skill is being referred to in order to be able to evaluate whether ‘you’ is good enough. The sentences in (47) are typical examples of scopal ambiguities, whereby the scope of a given linguistic item needs to be contextually worked out in order to understand the speaker’s intended interpretation. In (47a), for instance, it is not clear whether the speaker is communicating that not everyone is perfect or that no one is perfect. In (48) and (49), the different examples used more directly fall within the field of lexical pragmatics, the focus of this book. In (48a), the encoded content of the noun *children* underspecifies what is actually being communicated since not just any children are being referred to but only children in need. In (49a), the encoded content of the adjective *empty* over-specifies what is explicitly expressed given that the room was not literally empty, but only sufficiently so that the speaker could hear the ticking of someone else’s watch. In both cases the content of the word used is either too specific or not specific enough (with regard to the speaker’s intended interpretation) and needs to be pragmatically adjusted. More examples will be discussed in the next section.

It is argued in RT that the interpretation of the utterances in (44) to (49) involves a single inferential process. As mentioned before, the aim of RT is to explain how this inferential process is constrained on the basis of the principle of relevance. This approach will be introduced in the rest of this section. First, however, it is important to note that it is clear from the relevance-theoretic perspective that inferential processes are not only required for the derivation of implicatures and implicated content. Rather, pragmatic inference is also involved at the explicit level of communication, whereby the conceptual material provided by the linguistic logical form of an utterance also needs to be pragmatically enriched into a fully propositional assumption (Sperber and Wilson, 1995: 181; Carston, 1988: 41, 2002a: 107, 2016a: 614). For this reason, Sperber and Wilson (1995) coined the term *explicature* (by analogy with the term implicature) to refer to the pragmatically enriched explicit content of an utterance (Sperber and Wilson, 1995: 182; see also Carston, 1999, 2004, 2009, 2010).²⁶ It follows from this perspective that RT draws the line between semantics and pragmatics somewhere at the explicit level of communication (cf. Carston and Hall, 2012). While an implicature is purely the product of

²⁵ Perry (1993: 206) uses the term *unarticulated constituents* (see also Bach, 1994a: 269).

²⁶ Similar terms have been used elsewhere in the pragmatics literature (see, for instance, Recanati’s (1989: 297–302) use of the phrase *what is said* and Bach’s (1994b: 125) notion of *implicature*).

pragmatic inference, an explicature is a “semantic-pragmatic hybrid” (Carston, 2004: 819) since it is a pragmatic development of a linguistic logical form.

One of the goals of RT is to explain exactly how the inferential derivation of explicit and implicit content is constrained. It will have become clear from the previous sections that in RT the principle of relevance is considered as being the main driving force behind interpretation processes. Ostensive stimuli come with a presumption of their own optimal relevance. As a consequence, speakers need to optimize the relevance of their utterances, and hearers, whether for the recovery of explicatures or implicatures, need to look for an interpretation that provides them with enough cognitive effects to justify the amount of effort they put into the interpretation process. In this sense, inferential processes are constrained by the search for relevance, which is systematically expected given the ostensive nature of utterances. Still, it remains to be spelled out exactly how hearers actually proceed to recover the speaker’s intended interpretation. The principle of relevance specifically indicates that the more cognitive effort, the less relevance. As a result, hearers do not consider all possible interpretations and then choose the most relevant one. This would require too much processing effort and therefore be self-defeating (see Sperber, 2005: 64). Rather, it is argued in RT that the principle of relevance naturally lays the foundations for the following comprehension procedure (Sperber and Wilson, 2002: 18; Wilson and Sperber, 2004: 613):

(50) **Relevance-theoretic comprehension procedure**

- a. Follow a path of least effort in computing cognitive effects. In particular, test interpretive hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.
- b. Stop when your expectations of relevance are satisfied.

That is, for a given utterance, hearers do not process all possible interpretations but only focus on those that are most salient and which they test (for optimal relevance) in order of accessibility. Once an interpretation provides them with enough cognitive effects to justify the amount of processing effort involved, they stop searching and consider this interpretation to be the one intended by the speaker.²⁷ From a relevance-theoretic perspective, this comprehension procedure does not involve complex conscious computations. Rather, it is understood as an unconscious and intuitive process that involves “fast and frugal heuristics” (Wilson and Sperber, 2004: 624; see also Gigerenzer et al., 1999). In other

²⁷ As Wilson and Sperber (2004) point out, this naturally requires that “a speaker who wants her utterance to be as easy as possible to understand should formulate it (within the limits of her abilities and preferences) so that the first interpretation to satisfy the hearer’s expectation of relevance is the one she intended to convey” (Wilson and Sperber, 2004: 614).

words, it is computationally the simplest procedure possible (cf. Allott, 2002: 80). Consider the following exchange:

- (51) PETER: I just thought we could put the sofa in Tom's car.
 RICHARD: It's too big!

In order to understand Richard's answer, Peter first needs to recognize the ostensive nature of Richard's behavior. The recognition of Richard's ostensive behavior creates an expectation of optimal relevance, whereby Peter is to allocate cognitive resources in such a way as to optimize the relevance of the logical form provided by Richard's utterance. Peter therefore derives the explicature by assigning a referent to the pronoun *it* (i.e. the sofa) as well as enriching the adjectival phrase *too big* (i.e. too big to fit in Tom's car). This is the most relevant explicature that Peter can derive, and he stops at this interpretation, since it enables him to infer the implicatures that they cannot put the sofa in Tom's car and that they need to find another solution. These implicatures are directly relevant to Peter since they contradict (at least) one of his assumptions and thus modify his cognitive environment. As such, this interpretation provides Peter with enough cognitive effects to justify the effort invested in processing Richard's utterance, and therefore he stops looking for other interpretations. Note that, from the relevance-theoretic approach, the derivation of explicatures and implicatures is not treated as sequential. In the example just discussed, for instance, Peter does not first derive the explicature *The sofa is too big to fit in Tom's car* and then the implicature *We cannot put the sofa in Tom's car*. Rather, the derivation of explicatures and implicatures is coordinated and both are gradually derived on the basis of contextual assumptions in order to optimize relevance. This is referred to by Wilson and Sperber (2004: 617) as a process of *mutual parallel adjustment*.

In the next section, I will focus particularly on how this comprehension heuristics has been applied to lexical meaning and will therefore mostly discuss the derivation of explicatures. It is first important to understand why exactly this comprehension procedure is possible in the first place. Following Grice (1989: 213–223), it is argued in RT that inferential communication is primarily possible because of our ability to attribute particular mental states (such as beliefs, desires and intentions) to our interlocutors and vice versa (Sperber and Wilson, 1995: 24, 2002: 3; Carston, 1999: 103). This mind-reading ability is usually discussed in terms of the *theory of mind* hypothesis (see Carruthers and Smith, 1996; Goldman, 2012). According to Sperber and Wilson (1995), ostensive-inferential communication is possible notably because of our capacity to recognize someone's communicative and informative intentions (Sperber and Wilson, 1995: 50–64). From this recognition follows a presumption of relevance that systematically constrains the inferential processes that are required to derive a speaker's communicated content and to build mental representations about an individual's

thoughts and desires as well as specific attitudes.²⁸ The relationship between mind-reading abilities on the one hand (and in particular intention recognition) and pragmatic abilities on the other has been widely discussed in the pragmatics literature (cf. Haugh, 2008; Haugh and Jaszczolt, 2012). Unlike other pragmatic theories, however, RT has more systematically integrated theory of mind into its framework, which naturally leads to the development of new ideas within the theory. One such development concerns the type of information that hearers actually recover when interpreting a given utterance. On the basis of their mind-reading abilities, it is argued in RT that hearers do not only recover a speaker's communicated assumptions (be they *explicatures* or *implicatures*), but also recover a speaker's commitment and attitude towards those. That is, hearers also recover speakers' meta-representations of their communicated assumptions (Wilson, 2000). One particular kind of meta-representation that hearers are said to recover is called a *higher-level explicature* (Wilson and Sperber, 1993: 4; Sperber, 2000b: 6; Ifantidou, 2001: 80; Carston, 2004: 825). Consider the following example (from Wilson and Sperber, 1993: 4):

- (52) a. PETER: Can you help?
 MARY (SADLY): I can't.
 b. Mary says she can't help Peter to find a job.
 c. Mary believes she can't help Peter to find a job.
 d. Mary regrets that she can't help Peter to find a job.

In order to understand Mary's answer in (52a), Peter will derive the explicature *Mary cannot help me to find a job* on the basis of his expectation of optimal relevance. In addition, it is argued in RT that Peter will also embed this explicature within a higher-level representation (i.e. a meta-representation) such as in (52b) to (52d). This higher-level explicature includes either the representation of a particular speech act (as in (52b)) or the representation of a propositional attitude (as in (52c) and (52d)) (cf. Carston, 2004: 825). Exactly which of these higher-level explicatures are derived by Peter naturally depends on the context and therefore on which assumptions are manifest to him when interpreting Mary's utterance. Note that from the relevance-theoretic standpoint, however, "hearers always infer at least one higher level of embedding for

²⁸ In cognitive psychology, it is argued that our mind is equipped with a specific "theory-of-mind module" that enables mind-reading inferential processes (see Leslie, 1992, 1994; Baron-Cohen, 1994, 1995; Scholl and Leslie, 1999, *inter alia*). Within RT, Sperber and Wilson argue that pragmatics (and in particular the search for optimally relevant interpretations) is a specific type of mind-reading ability and even suggest that it constitutes a specialized submodule: the *comprehension* (or *Relevance*) module (Sperber, 2000a: 129; Wilson, 2000: 42; Sperber and Wilson, 2002: 5). Whether or not this is actually the case will not be discussed here. However, it is worth noting that according to Zufferey (2010), the relevance-theoretic approach "seems to offer the best model to account for the role of theory of mind in verbal communication" (Zufferey, 2010: 25; see also Zufferey, 2015: 91).

any proposition we express” (Clark, 2013a: 209). The systematic derivation of higher-level explicatures therefore makes them an integral part of the interpretation of an utterance. Carston (2004: 825) in fact argues that sometimes the relevance of an utterance is to be found precisely in higher-level explicatures.²⁹

Note that the notion of higher-level explicatures is directly relevant to our discussion because of the particular way in which they can be derived. In the example in (52) above, their derivation seems to be entirely pragmatic. Yet it has been suggested in RT that, in precisely the same way as some words provide rich clues about the speaker’s intended interpretations, there must be linguistic items whose sole function is to help the hearer recover higher-level explicatures (cf. Wilson and Sperber, 2012: 166). This will be discussed more fully in Section 2.2.3.2. For now, it is worth noting from the relevance-theoretic perspective that questions of semantics and pragmatics pervade linguistic communication. For this reason, RT prefers to talk about ostensive-inferential communication.

In this section, the focus was placed on the pragmatics of linguistic communication. Any good theory of pragmatics, however, necessarily rests upon a particular theory of semantics, and vice versa. Exactly what approach to semantics is adopted in RT will be addressed in the next section. This will enable us to identify explicitly the similarities and differences between RT and CxG and how compatible they are in terms of theoretical description. Before doing so, I wish to make a small observation. It was shown previously that CxG mostly focuses on the semantic (i.e. conventional) side of linguistic communication without much consideration for (non-conventional) pragmatics. As we have seen in this section, however, much of what is actually communicated by an individual is inferred in context and not provided by the linguistic items that they use. From this perspective, CxG alone can therefore not explain exactly how linguistic communication succeeds. This is the reason why CxG needs to be combined with a theory of pragmatics such as RT to be fully explanatory. At the same time, integrating RT and CxG can prove to be a real challenge for the following reason: the relevance-theoretic approach to pragmatics is actually based on a view of semantics that radically differs from that adopted in CxG.

²⁹ As Reboul (2001) rightly points out, this notion “also enables [relevance theorists] to propose a solution to *Moore’s paradox*” (Reboul, 2001: 46; emphasis mine), which has to do with the absurd nature of sentences like “I called you and I don’t believe I did.” In RT, when a speaker asserts a proposition P, her interlocutor usually derives the higher-level implicature that she believes P to be true.

2.2.3 *Semantics in Relevance Theory*

RT is particularly well known as a theory of pragmatics that tackles issues related to non-conventional aspects of linguistic communication. Yet RT also offers a specific understanding of the nature of linguistically encoded content, i.e. of semantics.

As mentioned in Section 2.1.2.2, in a purely terminological sense, Relevance Theory and Construction Grammar share a similar view since both frameworks discuss the notion of semantics in terms of *concepts*. It is argued in RT that the (optimally relevant) assumption communicated by a speaker forms “a structured set of concepts” (Sperber and Wilson, 1995: 85). From a theoretical standpoint, however, the perspectives developed in RT and CxG are fundamentally different since they are based on two opposite understandings of the nature of concepts. Indeed, relevance theorists follow Jerry Fodor’s hypothesis (Fodor et al., 1980; Fodor, 1998), which postulates that concepts are *atomic* (Sperber and Wilson, 1995: 91, 1998: 187; Carston, 2002a: 321). This specific approach is categorically rejected in the constructionist approach. In addition, concepts in RT are only discussed in relation to lexical items: “the ‘meaning’ of a word is provided by the associated concept” (Sperber and Wilson, 1995: 90). Unlike in CxG, however, comparatively little attention is paid to the semantics of larger linguistic patterns (see below).

The RT-specific approach to ‘concepts’ is explained in the next section. In particular, I will try to show how RT’s commitment to atomism has heavily influenced its analysis of lexical pragmatics, which results in theoretical (in)compatibility with CxG. However, it is worth first mentioning that RT discusses the notion of semantics not only in terms of concepts but also in terms of *procedures*. On the basis of the work of Blakemore (1987, 2002), more recent developments of RT consist in arguing that the encoded content of some linguistic expressions might best be described in terms of *procedural* meaning. What the term procedural meaning exactly captures (and why it is relevant to our discussion) will be explained in Section 2.2.3.2.

2.2.3.1 *Concepts and Ad Hoc Concepts* In their seminal book *Relevance*, Sperber and Wilson (1995) discuss the encoded content of linguistic items in terms of their associated concepts (Sperber and Wilson, 1995: 90). In particular, following Fodor, Sperber and Wilson adopt an atomic, non-decompositional approach: “the meaning of a word such as ‘yellow’, ‘giraffe’ or ‘salt’ is an *irreducible* concept” (p. 91; emphasis mine). RT’s commitment to conceptual atomism becomes explicitly clear in Carston (2002a):

I follow Jerry Fodor in assuming that concepts encoded by (monomorphemic) lexical items are atomic and so not decompositional; they don’t have definitions (sets of necessary and sufficient component features) and they are not structured around

prototypes or bundles of stereotypical features (for the arguments, see Fodor et al., 1980; Fodor, 1998; Laurence and Margolis, 1999). (p. 321)

It will have become clear that this approach to concepts is in direct opposition to that adopted in CxG, in which concepts are precisely understood in terms of encyclopedic knowledge organized in a network of related bundles around a prototype (see Section 2.1.2.2). The aim of this section is to spell out explicitly what the atomic account adopted in RT consists of. In addition, I will introduce the relevance-theoretic approach to lexical pragmatics and show how this perspective is directly influenced by the atomic approach to lexical semantics (and as a result cannot be easily integrated into a framework such as CxG). It is important to note here that the notion of *concept* represents a real issue in RT and that, within the theory itself, there are at least “three different possible views of what constitutes the content of concepts” (Groefsema, 2007: 139). In this section, I will introduce what I consider to be the more traditional approach to concepts in RT, as well as some of the issues that it raises for the semantics–pragmatics interface. A more comprehensive analysis of the relevance-theoretic perspective on concepts is given in the next chapter.

From a Fodorian point of view, the information provided by an atomic concept does not consist of a set of specific features, and even less so of a particular definition. Rather, conceptual information is said to consist of a nomological mind–world relation, i.e. a lawlike mind–world dependency (Fodor, 1998: 12).³⁰ The information provided by the concept *CAT*, for instance, consists of a necessary relationship between that concept and a specific instantiation of a cat in the real world. From this perspective, the argument goes, it is impossible to define what the word *cat* actually means: *cat* simply means *CAT* (p. 67). This lays the foundations for Fodor’s (1998: 55) *disquotational lexicon* hypothesis, which RT largely adopts. In order to discuss a given lexical concept, relevance theorists disquote the word and put it in capitals: *love* means *LOVE*, *happy* means *HAPPY*, etc. This formalization serves to highlight the hypothesis that words and concepts belong to two different types of lexicon (see Sperber and Wilson, 1998). When the linguistic item directly contributes to a particular sentence (public lexicon), the conceptual counterpart systematically feeds the language of thought (mental lexicon). This is one of the two functions of lexical concepts: they directly appear in the linguistic logical forms that form the basis of the thoughts that we communicate (Sperber and Wilson, 1995: 86).

In addition, atomic concepts also perform another function which appears to be essential within RT. In addition to contributing to the language of thought, atomic concepts also serve as an address (i.e. a point of access) to a variety of

³⁰ More simply, Carston (2010) says that the content of an atomic concept consists in “its denotation – what it refers to in the world” (Carston, 2010: 265).

information stored in memory. Specifically, there are three types of information that a concept gives access to: lexical, logical and encyclopedic (Sperber and Wilson, 1995: 86). The lexical entry of a given concept provides details about the linguistic item used to express that concept. This ranges from the phonological and morphosyntactic properties as well as co-occurrence possibilities. The logical entry of a concept provides information about the logical implications of that concept and consists of “a set of deductive rules which apply to logical forms of which that concept is a constituent” (Sperber and Wilson, 1995: 86). That is, for instance, the logical entry of the concept *CAT* includes inference rules such as $\{CAT \rightarrow ANIMAL\}$ and $\{CAT \rightarrow MAMMAL\}$ which enable an individual to compute the logical form in which the concept occurs and to derive new assumptions. From this perspective, the logical entry does not provide representational information about the concept but only computational information about how to use that concept (p. 89). In opposition, the encyclopedic entry precisely consists of representational information. The encyclopedic entry includes all of the real-world knowledge and assumptions that an individual stores in association with a specific concept and which provides a rich contextual background for the derivation of new assumptions. In the case of the concept *CAT*, this includes assumptions such as:

General knowledge about the appearance and behaviour of cats, including, perhaps, visual images of cats, and, for some people, scientific knowledge about cats, such as their anatomy, their genetic make-up, or their relation to other feline species, etc., and, for most people, personal experiences of, and attitudes towards, particular cats. (Carston, 2002a: 321)

Sperber and Wilson (1995: 88) observe that “such notions as *schema*, *frame*, *prototype* or *script*” are often used to discuss encyclopedic knowledge and in particular how the encyclopedic entry is internally structured and organized. Sperber and Wilson commit to none of those views, however (Sperber and Wilson, 1995: 88). In fact, unlike the different theoretical models that introduce these notions, RT is generally not inclined to explain how encyclopedic information is structured. To be precise, the encyclopedic entry is actually viewed as a “‘grab bag’ of encyclopaedic information” (Hall, 2017: 94). The only type of structure which is discussed is “in terms of the degree of accessibility of the items of information it contains, which implies that the internal structure of this entry is in constant flux” (Wałaszewska, 2011: 316; see also Carston, 2002a: 321). In comparison to the different networks discussed in Section 2.1.2.2, RT’s notion of encyclopedic knowledge is therefore relatively structureless. This can be easily explained due to the status attributed to encyclopedic information in RT.

From the relevance-theoretic standpoint, the types of information stored in the logical and encyclopedic entries of a given concept only serve to compute

the logical form in which that concept occurs. It is often argued, however, that they do not directly contribute to the semantics of the lexical item that is associated with that concept. Here, RT and CxG therefore provide opposite analyses. While in CxG encyclopedic knowledge is considered a central element of conceptual content, and is particularly structured, in RT it is only perceived as “contextual information” (Ribeiro, 2013: 104; see also Sperber and Wilson, 1995: 89; Carston, 1997b: 119). Note that it is not always clear, though, what the status of logical and encyclopedic information actually is in RT. Groefsema (2007: 139) convincingly shows how the relevance-theoretic perspective on concepts developed in Sperber and Wilson (1995) leaves room for various interpretations, some of which would actually consist in viewing encyclopedic and/or logical information as being content-constitutive. This will be discussed at length in the next chapter. Nevertheless, there is a general tendency to consider that the logical and encyclopedic information that a concept gives access to does not constitute its content:

Neither the encyclopaedic nor the logical information associated with a concept can be thought of as constitutive of the concept or as being its content. (Reboul, 2008: 522–523)

[Logical and encyclopaedic] properties are clearly not internal components of the lexical concepts themselves. (Carston, 2010: 249)

None of the information – logical or encyclopaedic – is constitutive of the concept. (Hall, 2011: 4)

The fact that the logical and encyclopedic entries are not content-constitutive is often supported, in addition to Fodor’s own arguments (see Fodor et al., 1980; Fodor, 1998), by a number of observations within RT. First of all, some concepts may not have both of these entries. Sperber and Wilson (1995: 92) argue, for instance, that the concept to which the lexical item *and* is associated “may lack an encyclopaedic entry,” i.e. it is not itself associated with real-world knowledge. Similarly, it is argued that concepts associated with proper names may not trigger inferential rules and therefore lack a logical entry (p. 92). In addition, relevance theorists appear to share the assumption that only a limited set of inferential rules can occur in the logical entry of a concept. (They never discuss more than one or two inference rules for each concept they look at.) Why this should be the case is not necessarily clear, but this motivates relevance theorists to assume that the logical entry “generally [falls] short of anything definitional” (Carston, 2002a: 321). This is further supported by the observation that different concepts may share similar inferential rules, which therefore cannot be used to distinguish between them. The concepts *CAT* and *DOG*, for instance, both contain the inferential rule {ANIMAL OF A CERTAIN KIND} (p. 322). Ultimately, it is argued that – unlike logical information – encyclopedic knowledge varies a lot across individuals and time

(Sperber and Wilson, 1995: 88) and is therefore too unstable to possibly be content-constitutive.

In spite of not being content-constitutive, the information stored in the logical and encyclopedic entries is argued to play a significant role during the interpretation process of an utterance. As will become clear in the rest of this section, they are key elements in the relevance-theoretic account of lexical pragmatics.

A central assumption in RT is that the words we use underdetermine the actual content of the thoughts that we communicate. In Section 2.2.2, this was referred to as the underdeterminacy thesis. In the sentence in (48a), repeated here in (53), the word *children*, for instance is used to express not the concept CHILDREN with which it is originally associated but the unlexicalized (atomic) concept CHILDREN-IN-NEED.

(53) What it comes down to is trying to give *children* a better Christmas.

It is argued in RT that, as the example in (53) illustrates, most of the concepts that we actually communicate are not lexicalized, i.e. they lack a lexical entry (Sperber and Wilson, 1998: 189). In order to convey these concepts, speakers therefore use the lexical entry of the most resembling concept, on the basis of which hearers recover the communicated concept in accordance with their expectations of relevance and following the comprehension procedure discussed in Section 2.2.2. Consider the sentences in (54) and (55). These examples nicely show that, in different contexts, the same lexical item (here *man*) may be used to express a variety of different concepts and not necessarily the one to which it is originally associated.

(54) A: I need a man to love me. (COCA, spoken)

B: Your dad loves you.

A: Dad, come on, you know what I mean.

(55) A: I need a man to love me. (COCA, spoken)

B: Tom loves you.

A: I said a man.

In neither of the examples does the speaker use the lexical item *man* to refer to the atomic concept MAN, say ‘a male human being’. In (54), assuming that the speaker is a heterosexual woman, she probably means to communicate a concept such as ‘heterosexual bachelor ready to commit to a long-lasting relationship’. In (55), she could intend the concept ‘heterosexual bachelor with prototypically masculine features’. These two concepts are not lexicalized in English. In order to communicate these concepts, the speaker therefore uses the lexical entry of another, similar (enough) concept on the basis of which the hearer should be able to infer the intended ones in accordance with their expectations of relevance and,

therefore, by following the relevance comprehension procedure. How exactly the intended concepts are recovered by the hearer is a major concern to relevance theorists. Naturally, as will have become clear, it is argued in RT that the recovery of these unlexicalized concepts is largely constrained by the search for optimal relevance which is triggered by the ostensive nature of the speaker's utterance. Precisely how the content of these concepts is actually established still calls for specification, however. Relevance theorists propose a specific account to explain the underlying mechanisms of lexical pragmatics, i.e. the meaning-construction process of lexical items.

The relevance-theoretic account of lexical pragmatics is generally based on the work of Lawrence Barsalou and his notion of *ad hoc categories* (Barsalou, 1983, 1987, 1993). According to Barsalou, conceptual categories (i.e. concepts) are never just retrieved from memory. Rather, we systematically construct *ad hoc categories*, i.e. occasion-specific categorizations (or conceptualizations), that are tailored to the specifics of each situation in which they occur.³¹ (See references cited in footnote 31 for empirical and experimental evidence.) Following Barsalou, RT argues that utterance comprehension systematically requires the creation of *ad hoc concepts*. From this perspective, in spite of their being associated with a specific atomic concept, it is argued in RT that "all words behave as if they encoded pro-concepts: that is . . . the concept it is used to convey in a given utterance has to be contextually worked out" (Sperber and Wilson, 1998: 185). The interpretation of the sentences in (53) to (55), for instance, does not consist first in testing (for relevance) the concepts CHILDREN and MAN associated with the lexical items *children* and *man* and then in deriving the intended concept. Rather, their interpretation directly requires the contextual construction of the *ad hoc* concepts CHILDREN* in (53), MAN* in (54) and MAN**

³¹ Already, it is important to note that this theoretical choice is unfortunate since it is inconsistent with the atomic approach to concepts adopted in RT. Barsalou et al. (1993: 57) note that:

[Barsalou] has been fairly schizophrenic in his definition of *concept*. Barsalou (1987, 1989, 1993) argues that concepts are temporary representations in working memory; Barsalou (1992a) and this chapter argue that concepts are bodies of knowledge in long-term memory; Barsalou (1992b) argues that concepts determine categorization. From this paper on, the following terminology will be followed: *Concepts* are the underlying knowledge in long-term memory from which temporary *conceptualizations* in working memory are constructed.

This quote shows that it is quite a challenge to pin down exactly what Barsalou considers to be a concept. I will come back to his view in the next chapter. For now, suffice it to say that his view of concepts is indisputably not atomic and neither is his view of *ad hoc* categorization (cf. Barsalou, 2000: 247, 2012: 239, 2016: 11; Yeh and Barsalou, 2006: 352; *inter alia*). Therefore, although relevance theorists generally discuss lexical semantics (i.e. concepts) in atomic terms, their approach to lexical pragmatics is based on a non-atomic approach. As I will try to show in the rest of this section and in particular in the next chapter, this naturally leaves a number of questions unanswered and results in many inconsistencies within RT.

in (55).³² As I will show in the next chapter, this view naturally raises a number of fundamental questions. For instance, it is no longer clear what the role of the lexically encoded concept actually is. As Recanati (2004: 97) points out, lexical concepts therefore seem to be “communicationally irrelevant.” This is rather inconsistent with the general claim that human cognition is geared towards relevance. In addition, the challenge is to know exactly what the nature of these ad hoc concepts is and how they are derived. These questions are at the origin of much debate within RT (see Chapter 3). Concerning the nature of ad hoc concepts, the traditional approach in RT considers that, like lexical concepts, they are atomic (e.g. Carston, 2010: 250). As for the way they are derived, the rest of this section will describe the underlying assumption developed in RT.

In RT, the derivation of ad hoc concepts is argued to result from a single inferential process often called ‘free’ *pragmatic enrichment* (Carston, 2004: 830, 2010: 218). This pragmatic process is *free* in the sense that it is not directly required by the linguistic item which is used to express that concept. Nevertheless, it will have become clear that “free” pragmatic enrichment remains constrained (or guided) by the search for optimal relevance in order to develop the logical form into a full-fledged explicature. In RT, this process of pragmatic enrichment is said to result in three possible outcomes: a narrower concept, a broader concept, or both a narrower and broader concept (Carston, 1997b: 121, 2002a: 334; Wilson and Carston, 2006: 409, 2007: 231; Sperber and Wilson, 2008: 92).

There is conceptual narrowing whenever the sense (or denotation) of the ad hoc concept is more specific than that of the lexical concept from which it is derived (Wilson and Carston, 2007: 232). The interpretations of *children* and *man* above involve such conceptual narrowing. Consider also the following examples:

- (56)
- a. I have a *temperature*. (Sperber and Wilson, 2008: 91)
 - b. Either you become a *human being* or you leave the group. (Wilson and Carston, 2007: 240)
 - c. I’m not *drinking* tonight. (Wilson and Carston, 2007: 232)

In (56a), the noun *temperature* is not used to communicate the context-free concept TEMPERATURE (i.e. ‘some degree of heat’) which it originally encodes. That someone has a particular temperature is a simple truism that achieves no relevance to an individual since it provides them with no cognitive effects. Rather, what is communicated in (56a) is the more specific, narrower ad hoc concept TEMPERATURE*, ‘abnormally high temperature’, which is argued to be

³² It is conventional in the relevance-theoretic literature to use the * symbol to differentiate ad hoc concepts (e.g. CHILDREN*) from context-independent lexical concepts (e.g. CHILDREN). When different (but related) ad hoc concepts are being referred to, such as in (54) and (55), a number of stars may be added to distinguish between them (e.g. MAN*, MAN**, MAN***, etc.).

inferentially derived following the comprehension procedure. A similar truism can be found in (56b). It can but only be mutually manifest to the interlocutors that the hearer (already) is a human being, and the latter must therefore infer a more specific ad hoc concept HUMAN-BEING*, e.g. a ‘well-mannered person’. A similar narrowing process is also said to occur when interpreting (56c). Depending on the context, the speaker can be taken to communicate either that they will not be drinking any alcoholic drinks at all (DRINK*), or that they will not drink themselves drunk (DRINK**). Both these interpretations are more specific than the encoded concept DRINK (i.e. ‘absorption of liquids’) and have to be inferentially derived by the hearer.

The derivation of ad hoc concepts, as mentioned above, may also provide an interpretation which is broader than that of the lexical concept from which it is derived. In this case, it is argued that the sense (or denotation) of the ad hoc concept is more general than that of the lexical concept (Wilson and Carston, 2007: 234). Consider the following sentences:

- (57) a. Holland is *flat*. (Sperber and Wilson, 2008: 91)
 b. This policy will *bankrupt* the farmers. (Wilson and Carston, 2007: 234)
 c. This steak is *raw*. (Carston, 2002a: 328)

In (57a), the adjective *flat* is not used to convey the concept FLAT, i.e. that Holland is literally even. Rather, the word is loosely used to communicate the broader ad hoc concept FLAT* whereby Holland is simply understood as not being mountainous. In (57b), the verb *bankrupt* can be understood literally. But there might also be some contexts in which it is loosely used to communicate that farmers will in fact lose a great amount of money (but yet not go bankrupt). This interpretation therefore requires the derivation of the broader ad hoc concept BANKRUPT*. Finally, the sentence in (57c) may be used literally to communicate that the steak is not cooked at all. It can also be used more loosely, however. For instance, you might use (57c) when you are not pleased with the cooking of your steak in a restaurant, in which case interpreting *raw* requires the derivation of the less specific ad hoc concept RAW*, i.e. not cooked enough.

Finally, ad hoc concepts can also be both narrower and broader than the encoded lexical concept from which they are derived. In this case, the sense of the ad hoc concept merely overlaps with that of the lexical concept (Carston, 1997b: 114).

- (58) a. Robert is a *computer*. (Wilson, 2009: 44)
 b. Caroline is a *princess*. (Wilson and Carston, 2006: 406)
 c. Sally is an *angel*. (Wilson, 2009: 44)

In (58a), *computer* is used to communicate the ad hoc concept COMPUTER*. This concept is narrower in the sense that it only refers to the quality of fast

computation, and broader in that the category is extended to include individuals other than physical objects. Interpreting (58b) requires the derivation of the ad hoc concept *PRINCESS** which is also both narrower and broader than the lexical concept *PRINCESS*. It is broader since it is extended to non-royal individuals and narrower since it only selects (for instance) the particularly good physical properties often attributed to princesses. Similarly, the interpretation of *angel* in (58c) is argued to be both narrower and broader than the lexical concept *ANGEL*. It is narrower since it only includes good angels, and broader given that it extends beyond celestial creatures.³³

It will have become clear that all of these cases of conceptual narrowing and/or broadening are said to be derived following the relevance comprehension procedure discussed in Section 2.2.2. The challenge here is to understand exactly what constitutes the content of ad hoc concepts. It was mentioned earlier that, like lexical concepts, ad hoc concepts are also considered to be atomic. If such is the case, it is very unclear in what way ad hoc concepts can be narrower/broader than lexical concepts. The notions of narrowing and broadening necessarily require some internal structure that can be exploited in different ways (cf. Assimakopoulos (2012: 23), and references cited therein). By virtue of being atomic, however, concepts in RT do not provide such structure. The reason why the terms *narrower* and *broader* are used follows from the way the content of these concepts is said to be determined. Ad hoc concepts are derived not solely on the basis of the atomic (lexical) concept itself, but primarily on the basis of the information stored in the logical and encyclopedic entries that the concept gives access to. Carston (1997b) explicitly says that, in order to derive ad hoc concepts

the hearer decodes the lexically encoded concept, thereby gaining access to certain logical and encyclopedic properties; he treats the utterance as a rough guide to what the speaker intends to communicate, and, in effect, sorts through the available properties, rejecting those that are not relevant in the particular context and accepting those that are, as reflections of the speaker's view. (p. 107)

³³ As the sentences in (58) illustrate, ad hoc concepts that are both narrower and broader than the lexical concepts often include examples of metaphorical use. Within RT, the interpretation of (lexical) metaphors is argued to follow from exactly the same process of conceptual development as cases of loose and hyperbolic use. The only difference concerns the extent to which the denotation of the ad hoc concept actually overlaps with that of the lexical concept from which it originates. For this reason, the notion of a “literal–loose–metaphorical continuum” has been introduced within RT (Sperber and Wilson, 2008: 93, see also Wilson and Carston, 2007). In comparison, cognitive linguists (and, therefore, construction grammarians) usually give a different treatment to metaphors (cf. Grady (2007) for a detailed overview). I will not address this issue here. For some work on the compatibility between the relevance-theoretic approach to metaphors and that elaborated in cognitive linguistics, see Tendahl and Gibbs (2008), Tendahl (2009), Wilson (2009, 2022), Stöver (2010) and Gibbs and Tendahl (2011).

That is, the content of ad hoc concepts is argued to be determined on the basis of the information stored in the encyclopedic and/or logical entries of the context-free lexical concept (Carston, 2002a: 347). From this perspective, it is easier to understand how ad hoc concepts can be narrower/broader than the lexical concepts from which they are inferentially derived. It is the set of information stored in the logical and encyclopedic entries that is narrower/broader (p. 347). As will become clear in the next chapter, this approach to the construction of (lexical) meaning provides a solid basis for the understanding of lexical semantics and pragmatics that I will explore in more detail and put in relation to the (not so different) constructionist perspective. Nevertheless, as far as meaning is concerned, there still seems to be a contradiction within RT (cf. Mioduszewska, 2015). On the one hand, it is argued that both lexical and ad hoc concepts are atomic and that it is this atom that constitutes the content of the lexical item used. At the same time, the difference between lexical concepts and ad hoc concepts is located at the level of the logical/encyclopedic entries, in particular in terms of how the information stored in those entries actually overlaps. In a sense, this suggests that the information stored in the logical/encyclopedic entries is therefore content-constitutive, unlike what is often argued in RT. Given this dichotomy, there can be only one of two outcomes. One possibility is to keep arguing that the content of lexical items (i.e. lexical semantics) must be the atomic concept itself. In this case, the challenge is to know exactly what the nature of ad hoc concepts actually is. This is explicitly what Carston (2010) points out:

The questions in the domain of relevance-theoretic lexical pragmatics that strike me as most interesting and most in need of some long hard thought concern the nature of ad hoc concepts. Are ad hoc concepts the same kind of entity as lexical concepts (apart from not being lexicalised)? Are they atomic or decompositional (perhaps even definitional)? . . . This is a research programme with most of the work yet to be done and I do not have much to offer here but a few hunches, hopes, and intuitions. (p. 249)

Although Carston (2010: 250) argues that these questions still “remain to be answered,” however, she considers that there is no reason to think that ad hoc concepts are not atomic. Alternatively, it has been suggested that ad hoc concepts, unlike lexical concepts, are not atomic but decompositional/definitional (e.g. Vicente and Martínez Manrique, 2010; Allott and Textor, 2012). The other possibility consists in abandoning the general atomic commitment to concepts and in considering that both lexical and ad hoc concepts are decompositional/definitional (e.g. Assimakopoulos, 2012). As might have become clear from the previous discussion, I will suggest an analysis along these lines in the next chapter.³⁴ Please note, of course, that the two possibilities

³⁴ It is important to note that the focus of this book is on linguistic meaning and not conceptual format. I will therefore not argue in favor of or against conceptual atomism. This debate requires more space than is possible here and has already resulted in a considerable body of research.

just discussed both rest on the initial assumption that lexical items must encode concepts. In the next chapter, I will show that more recent developments of RT also consider an alternative possibility whereby (context-free) lexical meaning might not be conceptual at all (see Section 3.3.2).

It is crucial to understand that this discussion of the nature of lexical semantics is indispensable. The perspective on concepts adopted in RT has serious theoretical consequences. I will discuss two of them here.

One of the consequences of the relevance-theoretic approach to concepts/lexical semantics concerns its understanding of the notions of monosemy and encoded polysemy. There is a general tendency in RT to assume that lexical items are monosemous, i.e. that they only encode one concept. The different analyses of the sentences in (56) to (58) have already pointed in that direction. It is argued that the interpretation of the items *temperature* in (56a), *drink* in (56c) and *angel* in (58c), for instance, are all pragmatically inferred from context following the relevance-theoretic comprehension procedure. From a constructionist perspective, however, these senses are already encoded by the lexical items which are used to communicate them.³⁵ Of course, it can be argued that the analyses of the sentences in (56) to (58) are purely rhetorical and are only meant to explain the relevance-theoretic approach to lexical pragmatics, and that relevance theorists are well aware that these particular senses might already be part of the speaker's knowledge. In fact, relevance theorists readily recognize that some of the senses they discuss in inferential terms might be stored by speakers of English (e.g. Wilson, 2003: 277; Wilson and Carston, 2007: 238, Carston, 2016a: 618, 2019: 152, 2021: 122). Strictly speaking, encoded polysemy is therefore not rejected in RT. Yet, monosemy is still preferred to polysemy. As I have shown in Leclercq (2023), for instance, the various relevance-theoretic accounts of modal meaning in English all adopt (and strongly argue for) a monosemous account (cf. Walton, 1988; Haegeman, 1989; Groefsema, 1992, 1995; Klinge, 1993; Berbeira Gardón, 1996, 1998, 2006; Nicolle, 1996, 1997a, 1998a; Papafragou, 2000; Kisiielewska-Krysiuk, 2008). More generally, Carston (2002a) says:

I am uneasy with the assumption that a monosemous analysis is always to be preferred to a polysemous one, though the "if at all possible, go pragmatic" strategy that it entails is one that I generally follow myself, as it makes for much more elegant analyses and because, for the time being, we lack any other strong guiding principle. (p. 219)

Rather, my aim is to argue that, regardless of whether concepts are atomic or not, the encoded meaning (i.e. semantics) of lexemes must contain the encyclopedic information that a concept gives access to.

³⁵ It is worth noting that the senses discussed earlier for the item *temperature*, *drink* and *angel* actually all have an entry in different dictionaries (e.g. Cambridge Dictionary, last accessed November 29, 2017, <https://dictionary.cambridge.org/dictionary/english/temperature>).

It is not clear why the notion of *elegance* is used as a (defining) criterion when conducting a scientific investigation, especially if one is to be descriptively accurate. (That is, one should be careful not to aim for elegance at any cost.) Beyond concerns of elegance, the relevance-theoretic appeal to monosemy is primarily grounded in a number of assumptions on which the theory was developed.

First, encoded polysemy might be dispreferred simply because it does not fit well with Fodor's atomic view of conceptual content and in particular with the disquotational lexicon hypothesis (cf. Fodor and Lepore, 2002: 116–117; see also Carston, 2010: 276). Fodor (1998: 53) explicitly argues that “there is no such thing as polysemy.” Of course, not all relevance theorists share Fodor's view. Falkum (2011), for instance, very explicitly argues that “contrary to Fodor, I believe that there is such a thing as polysemy” (p. 61). Yet the kind of polysemy that Falkum has in mind is not *encoded* polysemy but some sort of *pragmatic* polysemy: words can be used to communicate different concepts in different contexts, the actual content of which has to be systematically inferred (p. 61). In other words, there is still no room given to encoded polysemy. I believe this is largely due to another theoretical commitment within RT.

Encoded polysemy might be eschewed by relevance theorists simply because it is at odds, in spite of what could be argued, with the relevance-theoretic approach to ad hoc concepts. (This will be discussed at more length in the next chapter.) Even though relevance theorists claim that they are not particularly interested in encoded polysemy, they do argue that their pragmatic approach to lexical meaning can explain its origin (see Falkum, 2011: 147, 2015: 96; Carston, 2013: 187, 2016a: 619, 2019: 152, 2021: 123, *inter alia*). The repeated derivation of an ad hoc concept will lead to its conventionalization alongside the original lexical concept. This is a point of view I share. As Assimakopoulos (2012: 19) points out, however, the notion of ad hoc concepts was originally developed within RT as a rejection of the “encoded first” hypothesis.³⁶ It is argued in RT that the lexically encoded concept is not simply retrieved from memory and tested first for optimal relevance but that an ad hoc concept is systematically constructed following the relevance-theoretic comprehension procedure. In this case, as mentioned before, the question is to know what the role of the encoded concept actually is. Intuitively, it seems more *relevant* (since relatively effortless) to test this concept first for optimal relevance and then to try and derive an ad hoc concept. In the next chapter, I will discuss some of the issues that this approach raises and some of the solutions that have been suggested. The point to be made here is that monosemy might be

³⁶ Assimakopoulos (2012: 19) actually refers to the ‘*literal first*’ hypothesis, in keeping with the relevance theoretic tradition. However, this hypothesis does not relate to literal versus non-literal (e.g. metaphorical) types of meaning but to encoded versus inferred types of meaning. This is why I prefer to use the term *encoded* here.

generally preferred within RT because it is hard to see how relevance theorists could explain the relevance (in the technical sense) of having several encoded senses when they cannot explain the relevance of having even just one. That is, in a sense, since the encoded senses are argued not to be tested first, then why bother storing them in the first place? This kind of thinking seems to be underlying the latest development in Wilson (2011, 2016) and Carston (2013, 2016b), which I will discuss at length in the next chapter.

This perspective comes in direct opposition to the view adopted in CxG, according to which polysemy is the norm and monosemy the exception. I tend to sympathize with the constructionist approach, especially since I doubt that the ‘if at all possible, go pragmatic’ strategy can achieve descriptive accuracy. In the next chapter, one of my aims will be to show that there is no necessary incompatibility between arguing both for polysemy as well as against the “encoded first” hypothesis. One only needs, as Carston herself points out, specific guiding principles. I will attempt to provide such principles.

There is another, less direct, consequence of the relevance-theoretic approach to concepts that also needs mentioning. The main aim of relevance theorists is not so much to describe linguistic competence as it is to explain the pragmatics of linguistic communication. Yet in order to provide an accurate account of linguistic pragmatics, it is essential to know exactly what constitutes an individual’s linguistic knowledge and how this knowledge actually contributes to different communicative acts. In RT, considerable attention is given to lexical concepts, their nature and how they are used in context. Unlike in CxG, however, comparatively little attention is given to other (larger) elements of the language. Only a few relevance-theoretic studies mention the possibility for larger (non-lexical) patterns to be meaningful and to contribute to the understanding of an utterance. This includes, for instance, work on sentence types (e.g. Clark, 1991), prosody (cf. Scott, 2017, 2021, and references cited therein for a detailed overview) and clefts (e.g. Jucker, 1997). More generally, however, most of the work conducted from a relevance-theoretic standpoint focuses strictly on aspects of lexical semantics/pragmatics.³⁷ As a result, it is argued, for instance, that interpreting the creative use of *flick-knife* and *wrist* in (59) and (60) consists of a purely pragmatic process, whereby the ad hoc concepts FLICK-KNIFE* and WRIST* are derived in accordance with the principle of relevance.

(59) Handguns are the new *flick-knives*. (Wilson and Carston, 2007: 236)

(60) She *wristed* the ball over the net. (Wilson and Carston, 2007: 237)

The interpretation of *flick-knives* as metonymically referring to a bigger category (e.g. favorite weapon of choice) is said to result from “a single pragmatic

³⁷ This is clear from the topics addressed in references in Francisco Yus’ detailed online bibliography (<http://personal.ua.es/francisco.yus/rt.html>).

process of lexical adjustment” (Wilson and Carston, 2007: 236). Similarly, interpreting *wristed* as a particular caused-motion verb solely depends on one’s background knowledge “about the various arm movements of competent tennis players” (p. 237). That is, in order to explain the interpretation of these lexemes, no linguistic elements other than the lexemes themselves are discussed. As a consequence, relevance theorists once more have to turn to pragmatics to explain their interpretation. From the constructionist perspective, however, this strategy can sometimes be avoided. It is clear in CxG that it is not only lexemes that are meaningful but also larger (syntactic) units of the language. The interpretation of the sentences in (59) and (60), for instance, does not depend only on the lexemes that are used but also on the meaning of the larger constructions in which they occur: in (59) the *X is the new Y* construction and in (60) the CAUSED-MOTION construction. In (59), the metonymic interpretation of *flick-knives* is required by the slot it occupies in the *X is the new Y* construction. Similarly, in (60), the caused-motion interpretation of *wristed* is already part of the CAUSED-MOTION construction in which it occurs. Exactly how the lexeme inherits its meaning from the construction will be discussed in greater detail in Chapter 4.

Naturally, it must be understood that relevance theorists do not necessarily reject constructionist ideas. Yet there remains a general tendency not to view larger patterns of language as meaningful units. (After all, RT generally adopts a Chomskyan formal approach to language, against which CxG was precisely developed. See Carston, 2000: 87; Clark, 2013a: 346.) Because of the relative absence of such insights into linguistic structures, however, relevance theorists once more have to play the “all pragmatics” strategy, an option which – although (arguably) theoretically appealing – may not always be descriptively accurate. For this reason, I believe instead that RT can benefit from the constructionist insights.

Exactly how the various lexemes inherit their meanings from the construction in which they occur will be discussed in more detail in Chapter 4, where I will show that combining insights from RT and CxG is necessary because neither account on its own can explain the interpretations of (59) and (60).

2.2.3.2 Procedural Meaning It is a central assumption in RT that communication depends on inference more than it does on language. From this perspective, language is only “subservient to the inferential process” and in particular to *relevance* (Sperber and Wilson, 1995: 176). Language indeed provides the most cost-efficient way to achieve optimal relevance. In the previous section, we saw that lexical concepts for instance provide rich clues for the recovery of the intended interpretation which is inferentially constructed. That is, lexical concepts provide a solid basis for where relevance is to be sought and found. Given this particular approach to linguistic

communication, referred to in RT as ostensive-inferential communication, the challenge is to know exactly how language contributes to the comprehension process. It is argued in RT that in addition to specific mental representations (i.e. concepts), language might also be used to provide information about how to compute these mental representations and directly constrain the inferential process involved during the search for optimal relevance. This type of information is often referred to as *procedural*, and was first introduced in RT in the work of Diane Blakemore (1987, 1990, 2002). As a result, it is argued that lexical items may have either conceptual or procedural semantics, each contributing differently to the interpretation of the utterance in which they are used.

What exactly constitutes the nature of procedural expressions is not clear (cf. Groefsema, 1992; Bezuidenhout, 2004; Curc6, 2011; Wilson, 2011, 2012, *inter alia*). This issue will be taken up in Chapter 4. It is generally argued, however, that unlike lexical concepts, procedural expressions do not map onto mental representations but are instead used to convey specific *instructions* for the processing of conceptual elements. Blakemore (1992: 151) explicitly describes procedural expressions as items that “encode instructions for processing propositional representations.” The aim of these procedures is to guide the hearer towards optimal relevance by constraining the inferential process and by limiting the range of possible inferences. That is, they act as semantic constraints on inferential processes. In order to illustrate this observation, Blakemore originally discussed the use of discourse connectives in English in such procedural terms:

- (61) Tom can open Ben’s safe. *So* he knows the combination. (Blakemore, 2002: 79)
- (62) Tom can open Ben’s safe. *After all*, he knows the combination. (Blakemore, 2002: 79)

According to Blakemore, the use of *so* and *after all* in (61) and (62), respectively, is only meant to provide the hearer with an instruction about how to connect the two propositions and guide the hearer towards optimal relevance. The markers do not contribute to these propositions but only to their computation: their use directs the hearer towards the particular consequential/causal relationship intended by the speaker and which they have to infer. Without these markers, hearers may either fail to recover the speaker’s intended interpretation or simply spend too much effort deriving it (hence making it less *relevant*). Hence, procedural expressions are primarily used to facilitate the optimization of relevance.

Since the work of Blakemore, the conceptual/procedural distinction has been extended and applied to various expressions within RT (cf. Escandell-Vidal, Leonetti and Ahern (2011) for a detailed overview). Although this has raised a

number of issues within RT (see Carston, 2016b), there are two observations that need to be mentioned with respect to procedural meaning.

The first observation is that, given that inference occurs both at the explicit and implicit level of communication in RT (see Section 2.2.2), procedural meaning can be used to constrain the derivation of both implicatures and explicatures. In the case of discourse markers, it is often argued that they put a constraint on the derivation of implicatures. In Grice's (1989: 25) famous example *He is an Englishman; he is, therefore, brave*, the discourse marker *therefore* is used to constrain the derivation of the implicated premise *Englishmen are brave*. Expressions with procedural meaning can also constrain the derivation of explicatures and higher-level explicatures, however. At the level of explicatures, it is argued, for instance, that pronouns and demonstratives provide semantic constraints for the recovery of a specific referent and therefore have procedural semantics (Wilson and Sperber, 1993; Scott, 2011, 2013, 2016). Procedural meaning can also constrain the derivation of higher-level explicatures. Clark (1991) argues, for instance, that sentence types (i.e. imperatives, exclamatives and interrogatives) provide the hearer with an instruction for how to reconstruct the speaker's attitude towards the communicated proposition, i.e. to reconstruct how the proposition is mentally represented by the speaker.

The second observation concerns the actual nature of procedural meaning. Although it is not necessarily clear what constitutes the semantics of procedural expressions, a number of criteria are often used to distinguish concepts from procedures. Carston (2016b: 160–161) lists five properties that define procedural encoding: (i) introspective inaccessibility, (ii) non-compositionality, (iii) rigidity, (iv) not susceptible to nonliteral use, and (v) not polysemous. These properties will be discussed in Chapter 4. As we will see, they may not be completely unerring. However, it is worth mentioning that rigidity is often perceived as being the best defining feature of procedural expressions (cf. Escandell-Vidal and Leonetti, 2011). The rigidity of procedural meaning has to be understood in comparison with the relative flexibility of conceptual meaning. Whenever there is an incompatibility between a lexical concept and a procedural expression, the semantics of the procedural expression always wins out over the lexical concepts, which has to be adjusted to fit with the procedural semantics. Consider the following example:

(63) John is being silly. (Escandell-Vidal and Leonetti, 2011: 93)

The authors argue that, although the stative feature inherent to the concept *SILLY* is incompatible with the dynamic nature of progressive aspect in English, it is the procedural nature of progressive marking (i.e. “procedural information about how to construct the internal representation of the state of affairs,” p. 92) that

forces a dynamic representation of the situation. This type of description, whereby the meaning of an expression adjusts to that of another linguistic item, is described in Construction Grammar in terms of coercion. Interestingly, it is exactly in those terms that Escandell-Vidal and Leonetti describe the difference between conceptual and procedural meaning.³⁸ They argue that lexical concepts are coercible but procedural expressions are not: they only have a coercive force (2011: 86). This observation will prove very useful in Chapter 4 when discussing the notion of coercion. I will show that the notion of coercion as discussed in CxG might help shed new light on the actual nature of procedural encoding.

2.2.4 *Relevance Theory: Summing Up*

Relevance Theory is a cognitive theory of communication. It is based on the assumption that our mind is equipped with the capacity to treat only those pieces of information that provide us with enough cognitive effects to justify the amount of effort involved to process them, i.e. *relevant* information. This ‘principle of relevance’ in particular makes it possible to explain exactly how linguistic communication can succeed in spite of the fact that most of the words we use often fail to determine exactly the thoughts we intend to communicate. Ostensive acts of communication raise a specific expectation of relevance. As a result, understanding an utterance simply consists in optimizing the relevance of its interpretation. This involves a particular comprehension procedure which leads to the derivation of both explicatures (enriched explicit content) and implicatures (inferred implicit content). In particular, we have seen that RT puts forward a specific account of how conceptual content is adjusted in context and how ad hoc concepts – context-specific concepts – are derived in order to meet this expectation. This approach to lexical pragmatics will prove particularly useful in the next chapter, for, as mentioned previously, CxG lacks a detailed account of pragmatics and RT provides one of the most developed accounts of pragmatics in the literature.

The challenge with Relevance Theory is not to understand how pragmatics works but to pin down the extent to which it is actually involved in linguistic communication. Relevance theorists largely adopt a broad Fodorian perspective on semantics according to which lexical concepts are atomic. Yet this perspective is not completely compatible with how they account for the derivation of ad hoc concepts. As a result, the question is left open as to what actually constitutes the content of both lexical concepts and ad hoc concepts. This

³⁸ In order to explain example (63), Escandell-Vidal and Leonetti (2011) explicitly refer to the literature on *aspectual coercion*. It is worth noting that they are, to my knowledge, the only ones in RT to use the term *coercion*.

ambiguity about the nature of concepts, however, has led relevance theorists to play the “all pragmatics” strategy. As a result, words are almost systematically considered to be semantically monosemous, and ad hoc concepts systematically have to be inferred. But also, the meanings that words are used to convey are often analyzed independently of the particular constructions in which they occur. Yet, from the perspective of CxG, a number of aspects that are considered to belong to pragmatics in RT are often viewed as semantic properties. In the next chapter, the aim will be to try and combine these two approaches to conceptual content and show that a combination of them might be more psychologically real as well as more descriptively accurate.

2.3 Conclusion

Construction Grammar and Relevance Theory are currently two of the most discussed frameworks in their respective domains. Construction Grammar has put forward a specific account of linguistic knowledge, and Relevance Theory presents a detailed perspective on communication more generally. It will have become clear, however, that the strength of one of these two frameworks often corresponds to the weakness of the other. Where CxG provides a detailed account of linguistic forms and semantics, Relevance Theory still seems to be looking for specific guidelines on what constitutes the content of linguistic expressions. At the same time, where Relevance Theory proposes a very thorough understanding of pragmatic inference, Construction Grammar fails to integrate such principles into its framework. The aim of this book thus consists in drawing a theoretical bridge between these two frameworks and to show that the two frameworks nicely complement each other. Integrating the two frameworks is more easily said than done, however, as they are based on radically opposite ways of understanding what constitutes not only meaning but even language more generally. Clark (2013a) explicitly says that Relevance Theory “is based on a broadly Chomskyan approach to language and on Fodorian assumptions about modularity” (Clark, 2013a: 95). Construction Grammar, however, is one of the early functional approaches to grammar that was developed in opposition to these two traditions. Nevertheless, the aim of the next chapters is to show that in spite of these differences, Construction Grammar and Relevance Theory are not de facto incompatible.

Most recently, the need for CxG and RT to interact and be integrated has become more and more apparent. As mentioned in footnote 33, some recent work in the field of lexical metaphors aims to combine the two approaches. In a recent (concluding) chapter, Billy Clark (2017) points out that:

One example which has often occurred to me and which has not been much considered, if at all, is the possibility of adopting ideas about the pragmatic principles which

constrain interpretation from one approach and connecting them with ideas about the nature of semantics and pragmatics from other approaches. . . . It might be possible, for example, for construction grammarians to adopt only the central relevance-theoretic principles and consider how they might constrain interpretations within a construction grammar approach. Once again, it seems that there are significant benefits from bringing together researchers from different backgrounds. (p. 352)

More recently, Finkbeiner (2019b) edited a special issue that precisely focused on how ideas from Construction Grammar can be combined with perspectives from post-/neo-Gricean approaches to pragmatics. So it generally appears that there is a desire to bring together these two frameworks (see also Xue and Lin, 2022). The aim of this book, which builds on ideas developed in previous research (cf. Leclercq, 2019, 2020, 2022, 2023), is precisely to spell out some of the directions in which this integration can be operated.