

# *Generativity, Relevance and the Problem of Polysemy\**

INGRID LOSSIUS FALKUM

---

## **Abstract**

In this paper, I discuss the generative account of word meaning and interpretation of words in context proposed by Pustejovsky (1995). The goal of the generative lexicon theory (GLT) is to provide a formal account of ‘logical polysemy’, involving cases where a basic word meaning gives rise to different meanings in different contexts. My main claim is that the GLT fails to account for the flexibility of the processes involved in the modulation of lexical meaning in context, and thus makes a range of wrong predictions with regard to the derivation of compositional interpretations. I argue that the range of phenomena falling under the label of ‘logical polysemy’ on Pustejovsky’s theory are better given a pragmatic account within the relevance-theoretic framework (Sperber & Wilson 1995, Carston 2002), which not only provides a simpler, more unified account, but also avoids the problems of misinterpretation associated with GLT.

## **1 Introduction: The Problem of Polysemy**

There are several different ways in which a word can have more than one meaning. There is the traditional distinction between *homonymy*, involving cases where two distinct lexical items accidentally have the same linguistic form (e.g. the English form *bank*, which may refer to a financial institution or a riverside), and *polysemy*, referring to cases where a single lexical item has several (related) senses. Also, any word, whether monosemous, homonymous or polysemous, may have its meaning further modulated on a particular occasion of use by the linguistic or non-linguistic context in which it occurs. However, the distinctions between these cases of a word being associated with more than one meaning are not clear-cut. For instance, the question of where to draw the distinction between homonymy and polysemy remains a source of theoretical discussion in the lexical semantics literature. Similarly, the distinction between polysemy and cases of pragmatic modulation in context is not always easy to draw, and there is a question concerning what kind of phenomena fall under the label of ‘polysemy’. Consider (1)-(4):

---

\*I am grateful to Robyn Carston for her helpful comments and suggestions on previous drafts of this paper.

- (1) a. *The newspaper* fired its editor.  
 b. Mary spilled coffee on *the newspaper*.  
 c. *The newspaper* revealed surprising details of the trial.
- (2) Mary *began* a book.
- (3) a. John *baked* a potato.  
 b. John *baked* a cake.
- (4) a. Every chef needs a *good* knife.  
 b. ‘The Kite Runner’ is a *good* book.

In (1), *newspaper* has different meanings in (1a), where it refers to the publisher of the newspaper, (1b), where it denotes a physical object, and (1c), where it refers to the newspaper as an information type. The verb *began* in (2a) takes an NP as its complement, thereby leaving the VP underspecified with regards to the kind of activity associated with it (e.g. reading, writing, binding, etc.). The verb *bake* has different interpretations in (3a), where it has a change of state sense, and (3b), where it has a creative sense, while the adjective *good* expresses different properties in (4a) and (4b).

These kinds of examples have been given different treatments in the literature. The noun *newspaper* in (1) has been analysed as a *densely metonymous* term (Nunberg 1996: 126) with a number of interdefined denotations (e.g. organisation, edition or publication type, physical object or token, etc.), licensing a “widespread predicate transfer from the properties of one of its denotations to another”. The verb *began* in (2) has been analysed as having a derived sense meaning ‘began reading’, where the process specified by the VP is made apparent from context or the other lexical items in the sentence (Langacker 1984: 185). The verb *bake* in (3) has been argued to be systematically ambiguous between a creative and a change of state sense (Atkins et al. 1988), while the adjective *good* in (4) has been analysed as an incomplete one-place predicate of individuals (Szabó 2001: 133-134), which is associated with a set of individuals only when additional information is provided (e.g. under one completion we get the set of individuals/objects which are knives that are good for cutting, under another we get the set of books that are good to read, etc.). Szabó postulates a variable in the lexical representation of *good* (and for many other adjectives too), which is fixed by the context in which the word is used.

Although the examples given in (1)-(4) arguably exhibit quite different characteristics, Pustejovsky (1995) has suggested treating them all as instances of a single phenomenon labelled ‘logical polysemy’, defined as the case where a word has two or more meanings, but where “there is no change in lexical category, and the multiple senses of the word have overlapping, dependent, or shared meanings” (ibid. 28). In the case of (1), Pustejovsky takes the polysemy of *newspaper* to be encoded directly into the semantics of the word. In (3), the different senses of *bake* are determined by the semantics of the complement noun, while the meaning of the

adjective *good* in (4) is assumed to be dependent on the semantics of the noun it modifies. Example (2), however, stands out as different from the previous cases, as the polysemy of the verb *begin* is viewed as related to the different complement types it may select for. Pustejovsky's notion of polysemy is thus a broad one.

With the goal of providing a formal account of examples such as (1)-(4), Pustejovsky (1991, 1995, 1998a, 1998b) develops a generative theory of word meaning structured around a representation and treatment of cases of 'logical polysemy'. The theory is designed to provide an account of how words can have different meanings in different linguistic contexts and how new senses can emerge as a result of the linguistic compositionality process, while at the same time limiting the number of entries in the lexicon. He rejects the view of lexical sense organisation assumed by so-called *SELS* (*Sense Enumeration Lexicons*), according to which polysemy can be accounted for in terms of multiple listings of senses in the lexicon (where these senses are stored either in separate lexical entries or as multiple listings under a single lexical entry), as he believes that these accounts fail to explain the nature of the phenomenon, and miss the important generalisations that can be made on the basis of the available data. Instead, he proposes that the lexicon should reflect the following two assumptions: (1) the meaning of a lexical item cannot be divorced from the structure that carries it; and (2) word meaning should mirror our non-linguistic conceptual organising principles. The result is a lexicon consisting of complex lexical entries, over which a set of generative operations may apply in order to yield compositional interpretations such as the ones discussed for (1)-(4).

My aim in this paper is to discuss and assess the generative account of word meaning and interpretation of words in context proposed by Pustejovsky (1991, 1995, 1998a, 1998b). My main point is that this theory, while providing an interesting perspective on the relation between word meaning and compositional interpretations, nevertheless fails to account for the flexibility of the processes of sense modulation shown in (1)-(4), and makes a range of wrong predictions with regard to the derivation of compositional interpretations. I will argue that it is doubtful that the complex lexical representations assumed by the generative lexicon theory actually succeed in picking out the right denotations in the world. I will also argue, along with previous authors, that the theory does not provide a proper distinction between word meaning (what exactly is encoded by a word) and world knowledge. Finally, I will try to show that the range of phenomena falling under the label of 'logical polysemy' on Pustejovsky's theory are better given a pragmatic account within the relevance-theoretic framework (Sperber & Wilson 1995, Carston 2002), without having to posit complex lexical entries or any arbitrary cut-off point between aspects of general knowledge that are part of the lexicon and those that are part of the general knowledge system.

In the next section, the central tenets of the generative lexicon theory (henceforth GLT) will be presented. Section 3 discusses some objections to the theory, while section 4 presents a reanalysis of the data from the viewpoint of relevance theory.

## 2 The Generative Lexicon Theory (GLT)

The *generative lexicon* is viewed as a system consisting of (at least) four levels of linguistic representation (Pustejovsky 1995: 61). These are (1) the *Argument structure*, which includes the specification of number and type of logical arguments; (2) the *Event structure*, which is a definition of the event type of an expression, e.g. a state, process, or transition; (3) the *Qualia structure*, involving the essential attributes of an object or event as defined by the lexical item; and (4) the *Lexical inheritance structure*, describing the way in which a lexical item is globally related to other concepts in the lexicon. The semantics of a lexical item  $\alpha$  is viewed as a structure consisting of these four components, which are connected by a set of generative devices, referred to as *type coercion*, *co-composition* and *selective binding*, that allow for the compositional interpretation of words in context.

### 2.1 Complex Lexical Entries

*2.1.1 Argument and Event Structures.* In GLT, the Argument structure of a word is taken to be the minimal specification of its semantics. The arguments for a lexical item  $\alpha$  are represented as a list structure where the type of argument is directly encoded in the argument structure (ARGSTR), in the following way (Pustejovsky 1995: 67):

$$(5) \quad \left[ \begin{array}{l} \alpha \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = \dots \\ \text{ARG2} = \dots \\ \text{D-ARG1} = \dots \\ \text{S-ARG1} = \dots \end{array} \right] \\ \dots \end{array} \right]$$

As the representation in (5) shows, GLT posits different types of arguments for lexical items: (1) *True arguments* (ARG1 and ARG2) are those that must be syntactically realised; (2) *Default arguments* (D-ARG1) are not necessarily expressed syntactically, but are necessary for the logical well-formedness of the sentence; (3) *Shadow arguments* (S-ARG1) also refer to semantic content that is not necessarily expressed, but which is expressible under certain circumstances. The

different types of arguments are exemplified by the argument structures for the verbs *butter* and *build* in (6) and (7) (ibid.):

- (6) 
$$\left[ \begin{array}{l} \mathbf{butter} \\ \text{ARGSTR} = \\ \dots \end{array} \left[ \begin{array}{l} \text{ARG}_1 = \mathbf{human} \\ \text{ARG}_2 = \mathbf{phys\_object} \\ \text{S-ARG}_1 = \mathbf{butter} \end{array} \right] \right]$$
- (7) 
$$\left[ \begin{array}{l} \mathbf{build} \\ \text{ARGSTR} = \\ \dots \end{array} \left[ \begin{array}{l} \text{ARG}_1 = \mathbf{animate\_individual} \\ \text{ARG}_2 = \mathbf{artifact} \\ \text{D-ARG}_1 = \mathbf{material} \end{array} \right] \right]$$

The representations in (6) and (7) show that the verb *butter* takes a human as its subject, a physical object as its complement, as well as a shadow argument, *butter*, which is the material used in performing the buttering act, (e.g. *John*<sub>ARG1</sub> buttered *his sandwich*<sub>ARG2</sub> (with *butter*<sub>S-ARG1</sub>)). Similarly, the verb *build* takes an animate individual and an artefact as its subject and complement respectively, but also takes a ‘default argument’, which is the material involved in the building, which may be explicit or implicit (e.g. *Mary*<sub>ARG1</sub> built a *house*<sub>ARG2</sub> (out of *bricks*<sub>SD-ARG1</sub>)).

The *Event structure* of a lexical item or a phrase defines the type of event associated with that lexical item or phrase. Events are classified as *processes*, *states* or *transitions*, and may also have subevental structures associated with them. These subevents may have one of the following three restrictions on their temporal ordering: ‘<<sub>α</sub>’ describes a subevental structure involving a development process and a resulting state (e.g. *build*), ‘< o<sub>α</sub>’ describes two simultaneous subevents, where one starts before the other (e.g. *walk*, where the movement of the legs brings about the movement of the body), and ‘o<sub>α</sub>’ describes two completely simultaneous subevents (e.g. *accompany*, the person accompanying and the one being accompanied). The relative prominence of the subevents is also specified, where the *head* of an event is defined as “the most prominent subevent of the event structure of a predicate, which contributes to the “focus” of the interpretation” (Pustejovsky 1995: 72). The event structure for the verb *build* is given in (8):

$$(8) \left[ \begin{array}{l} \mathbf{build} \\ \text{EVENTSTR} = \left[ \begin{array}{l} E_1 = \mathbf{process} \\ E_2 = \mathbf{state} \\ \text{RESTR} = <_{\alpha} \\ \text{HEAD} = E_1 \end{array} \right] \\ \dots \end{array} \right]$$

This event structure includes two subevents,  $E_1$  and  $E_2$ , of which the former is the head (i.e. the event that figures as the most prominent in interpretation), and specifies that  $E_1$  must precede  $E_2$ .

*2.1.2 Qualia Structure.* The Qualia structure of a lexical item is the distinctive feature of Pustejovsky's theory, and is supposed to be "the set of properties or events associated with a lexical item which best explain what that word means" (Pustejovsky 1995:77). The qualia consist in a specification of four basic roles: the constitutive, formal, telic and agentive roles, each of which is taken to provide a distinct but essential component to uniquely determine the meaning of a word.

The *constitutive* role of a lexical item captures the relation between an object and its constituents, or proper parts (this may include a specification of material, weight, and of parts and component elements). The *formal* role specifies what distinguishes the object within a larger domain (and may include information about the orientation, magnitude, shape, dimensionality, colour and position of the object). The *telic role* defines the purpose and function of the object, if this may be specified, while the *agentive* role describes the factors involved in the origin or 'bringing about' of an object (including distinctions such as e.g. natural kind/artefact). The decompositional semantics of a lexical item thus includes a specification of these four aspects of meaning. (9) presents a schematic representation of the qualia structure for a given lexical item (Pustejovsky 1998a: 295):

$$(9) \left[ \begin{array}{l} \alpha \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = x \\ \dots \end{array} \right] \\ \text{QUALIA} = \left[ \begin{array}{l} \text{CONST} = \mathbf{what\ } x \mathbf{ is made of} \\ \text{FORMAL} = \mathbf{what\ } x \mathbf{ is} \\ \text{TELIC} = \mathbf{function\ of\ } x \\ \text{AGENTIVE} = \mathbf{how\ } x \mathbf{ came into being} \end{array} \right] \end{array} \right]$$

The representation in (10) is a minimal semantic description for the noun *novel*, including values for each of the qualia roles (adapted from Pustejovsky 1991: 427, 1995: 78):

$$(10) \left[ \begin{array}{l} \mathbf{novel} \\ \dots \\ \text{QUALIA} = \left[ \begin{array}{l} \text{CONST} = \mathbf{narrative(x)} \\ \text{FORMAL} = \mathbf{book(x)} \\ \text{TELIC} = \mathbf{read(e,y,x)} \\ \text{AGENT} = \mathbf{write(e',z,x)} \end{array} \right] \end{array} \right]$$

This representation shows how the noun *novel* is supposed to encode information about particular properties and activities associated with it (in this representation, *x* is a distinguished variable which represents the object itself). A novel is a narrative, which has the form of a book; its purpose is to be read and it comes into being by a process of writing (*e* and *e'* refer to the event types of reading and writing, which are both transitions).<sup>1</sup> Pustejovsky assumes that a qualia structure may be defined for all types of lexical items, but that not all lexical items need carry a value for each role. For instance, the lexical representation associated with the verb *break* is taken to be the following (Pustejovsky 1995: 80):

$$(11) \left[ \begin{array}{l} \mathbf{break} \\ \text{EVENTSTR} = \left[ \begin{array}{l} E_1 = \mathbf{e_1:process} \\ E_2 = \mathbf{e_2:state} \\ \text{RESTR} = <_{\alpha} \end{array} \right] \\ \text{QUALIA} = \left[ \begin{array}{l} \text{FORMAL} = \mathbf{broken(e_2,y)} \\ \text{AGENTIVE} = \mathbf{break\_act(e_1,x,y)} \\ \dots \end{array} \right] \end{array} \right]$$

The event structure encoded by the causative verb in (11) includes two subevents: a process followed by a resulting state. These subevents are assumed to map directly

---

<sup>1</sup> However, there is more than this to the qualia structure for *novel*, since this noun is taken to be a subtype of *book*, which is represented as a *Lexical Conceptual Paradigm* (on this notion, see section 2.1.2.1), being polysemous between a physical object and a textual information sense. The noun *novel* inherits this polysemy from *book*, but its lexical representation is more informative in the sense that it includes a specification of the sort of textual information associated with it (i.e. a narrative) (see Pustejovsky 1995: 156).

into the formal and agentive qualia roles, where the resulting state corresponds to the formal role and the causing event to the agentive role.

*2.1.2.1 Lexical Conceptual Paradigms.* So how does Pustejovsky use this idea of complex lexical entries to account for the examples of polysemy we set out in the introduction? Let us consider again the example in (1), repeated here as (12):

- (12) a. *The newspaper* fired its editor.  
 b. Mary spilled coffee on *the newspaper*.  
 c. *The newspaper* revealed surprising details of the trial.

In this case, where *newspaper* is used to denote an organisation (12a), a physical object (12b), and an information type (12c), GLT takes the polysemy to be encoded directly into the lexical representation of the object. The notion of a *Lexical Conceptual Paradigm* (lcp) provides a means of characterising a lexical item as a *meta-entry* (Pustejovsky 1995, 1998b). Consider the uses of the noun *book* in (13):

- (13) a. *The book* is sitting on the coffee table.  
 b. Mary found *the book* interesting.  
 c. Paul has started writing *his second book*.

This noun is analysed as being polysemous between a physical object sense (13a), a textual information sense (13b), as well as the combination of the two senses (13c), all of which are captured in the meta-entry for *book*. The lcp acts as a type constructor, which creates a complex type for a term  $\alpha$ , carrying the polysemous senses  $\sigma_1$  and  $\sigma_2$ , on the basis of the following rule (Pustejovsky 1995: 93):

- (14)  $\frac{\alpha: \sigma_1 \quad \alpha: \sigma_2}{\text{lcp}(\alpha): \sigma_1 \cdot \sigma_2}$

The lcp of a polysemous lexical item may then be represented as in (15), consisting of the complex type (a Cartesian product), as well as the two base senses (ibid.):

- (15)  $\text{lcp} = \{\sigma_1 \cdot \sigma_2, \sigma_1, \sigma_2\}$

Pustejovsky thus suggests that three senses of a lexical item, clustered in one single lexical entry, may be constructed from two base types. The lexical representation for *book* may thus look like (16) (Pustejovsky 1995: 101):



$$(16) \left[ \begin{array}{l} \mathbf{book} \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = \mathbf{x:information} \\ \text{ARG2} = \mathbf{y:phys\_obj} \end{array} \right] \\ \text{QUALIA} = \left[ \begin{array}{l} \mathbf{information.phys\_obj\_lcp} \\ \text{FORMAL} = \mathbf{hold(y,x)} \\ \text{TELIC} = \mathbf{read(e,w,x,y)} \\ \text{AGENT} = \mathbf{write(e',v,x,y)} \end{array} \right] \end{array} \right]$$

The same rule can be used to give a lexical representation for the senses of *newspaper* (Pustejovsky 1995: 156):

$$(17) \left[ \begin{array}{l} \mathbf{newspaper} \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = \mathbf{x:org} \\ \text{ARG2} = \mathbf{y:info . physobj} \end{array} \right] \\ \text{QUALIA} = \left[ \begin{array}{l} \mathbf{org.info.physobj\_lcp} \\ \text{FORMAL} = \mathbf{y} \\ \text{TELIC} = \mathbf{read(e_2,w,y)} \\ \text{AGENT} = \mathbf{publish(e_1,x,y)} \end{array} \right] \end{array} \right]$$

In this way, the different senses of *newspaper* will be encoded in its lexical representation as part of a lexical conceptual paradigm in the qualia structure (where the ‘info.physobj’ sense is itself an lcp). However, as Pustejovsky himself notes, the way in which the noun *newspaper* denotes is different from that of *book*, since the former cannot be used to denote the combination of the senses stored in the lcp, while the latter can, cf. (13c). I will return to this issue in section 3.2.

Nunberg (1996: 126), who analyses *newspaper* as a densely metonymous term, notes that a problem with these kinds of metonymies is that it is often hard to decide which use of the word is ‘basic’ and which uses are ‘derived’, and, as a corollary, to know in which direction the transfer operates. The GLT seems to solve this problem by letting all senses be encoded as a lexical conceptual paradigm in the qualia structure for the word; hence there is no need to assume that one sense is more ‘basic’ than the other senses.

## 2.2 Generative Mechanisms

Now let us consider the remaining cases of ‘logical polysemy’ that we looked at in the introduction. The examples in (2)-(4) are different from the previous case in

that their relevant interpretations are assumed to be conditioned by a set of generative operations, described as *type coercion*, *co-composition*, and *selective binding* (Pustejovsky 1995), whose application is determined by linguistic context.

*2.2.1 Type Coercion.* The semantic transformation of type coercion is defined as a “semantic operation that converts an argument to the type that is expected by a function, where it would otherwise result in a type error” (Pustejovsky 1995: 111). This generative mechanism accounts for examples of the sort in (2), repeated here as (18a). Compare (18a-c):

- (18) a. Mary began *a book*.  
 b. Mary began *reading a book*.  
 c. Mary began *to read a book*.

In (18a) the verb *begin* selects an NP as its complement (specifying an object), while in (18b) and (18c) its verbal complements are both of the semantic type ‘event’. Instead of assuming that there are different lexical entries for *begin*, that is, one for each complement type the verb may select for, this verb is treated as having a single lexical entry, with the following event and argument structures (cf. Pustejovsky 1995: 116):

- (19) 
$$\left[ \begin{array}{l} \mathbf{begin} \\ \text{EVENTSTR} = \left[ \begin{array}{l} E_1 = \mathbf{transition} \\ E_2 = \mathbf{transition} \\ \text{RESTR} = \langle o_\alpha \end{array} \right] \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = \mathbf{x: human} \\ \text{ARG2} = \mathbf{e}_2 \end{array} \right] \end{array} \right]$$

The lexical representation in (19) shows that *begin* selects for an event as its complement type (cf. ARG<sub>2</sub>). In cases where this requirement is not directly satisfied by the surface syntactic structure, as in (18a) above, coercion applies in order to change the type of the complement NP into an event. We saw that the qualia structure for the noun *book* in (16) above included a specification of two events in the agentive and telic roles (namely those of reading and writing), hence the denotation of the NP *a book* is coerced into an event denotation consistent with the interpretations ‘Mary began to write a book’ and ‘Mary began to read a book’ respectively.

2.2.2 *Co-Composition*. The generative operation referred to as *co-composition* (Pustejovsky 1991, 1995: 122), accounts for the example in (3), repeated below as (20). This mechanism operates over structures that allow for more than one function application.

- (20) a. John *baked* the potato.  
 b. John *baked* the cake.

As we have seen, *bake* has different meanings in (20a), where it has a change of state sense, and (20b), where it has a creation sense. Rather than treating these two senses of *bake* as separate lexical entries (or, alternatively, as different senses stored under a single lexical entry), Pustejovsky assumes only one sense for *bake*, a change-of-state reading, and that any other readings are derived through the generative mechanism of co-composition, which applies at the level of the entire verb phrase. This mechanism allows for the information carried by a complement to act on the governing verb, by taking the verb as argument and shifting its event type. Consider the lexical representation for the verb *bake* (Pustejovsky 1995: 123):

- (21) 
$$\left[ \begin{array}{l} \mathbf{bake} \\ \text{EVENTSTR} = \left[ \begin{array}{l} E_1 = \mathbf{e_1:process} \\ \text{HEAD} = \mathbf{e_1} \end{array} \right] \\ \text{ARGSTR} = \left[ \begin{array}{l} \text{ARG1} = \mathbf{1} \left[ \begin{array}{l} \mathbf{animate\_ind} \\ \text{FORMAL} = \mathbf{physobj} \end{array} \right] \\ \text{ARG2} = \mathbf{2} \left[ \begin{array}{l} \mathbf{mass} \\ \text{FORMAL} = \mathbf{physobj} \end{array} \right] \end{array} \right] \\ \text{QUALIA} = \left[ \begin{array}{l} \mathbf{state\_change\_lcp} \\ \text{AGENTIVE} = \mathbf{bake\_act(e_1,1,2)} \end{array} \right] \end{array} \right]$$

Here *bake* is specified as having a change of state denotation. When combined with the complement *potato*, whose agentive role is specified as ‘natural kind’, the change of state interpretation of *bake* remains unchanged. Pustejovsky sees this as a case where the noun does not, on its own, contain a specification of an event structure and hence the process of co-composition does not apply. The semantics of the complement *cake* in (20b) does, however, shift the core meaning of *bake* from a change-of-state to a creation sense. This is a result of the lexical representation associated with *cake*, given in (22) (Pustejovsky 1995: 123):

(22)	<table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;"><b>cake</b></td> <td></td> </tr> <tr> <td style="padding-right: 10px;">ARGSTR =</td> <td style="border-left: 1px solid black; border-right: 1px solid black; padding: 5px;"> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">ARG1 = <b>x:food_ind</b></td> </tr> <tr> <td style="padding-right: 10px;">D-ARG1 = <b>y:mass</b></td> </tr> </table> </td> </tr> <tr> <td style="padding-right: 10px;">QUALIA =</td> <td style="border-left: 1px solid black; border-right: 1px solid black; padding: 5px;"> <table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">CONST = <b>y</b></td> </tr> <tr> <td style="padding-right: 10px;">FORMAL = <b>x</b></td> </tr> <tr> <td style="padding-right: 10px;">TELIC = <b>eat(e<sub>2</sub>,z,x)</b></td> </tr> <tr> <td style="padding-right: 10px;">AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b></td> </tr> </table> </td> </tr> </table>	<b>cake</b>		ARGSTR =	<table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">ARG1 = <b>x:food_ind</b></td> </tr> <tr> <td style="padding-right: 10px;">D-ARG1 = <b>y:mass</b></td> </tr> </table>	ARG1 = <b>x:food_ind</b>	D-ARG1 = <b>y:mass</b>	QUALIA =	<table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">CONST = <b>y</b></td> </tr> <tr> <td style="padding-right: 10px;">FORMAL = <b>x</b></td> </tr> <tr> <td style="padding-right: 10px;">TELIC = <b>eat(e<sub>2</sub>,z,x)</b></td> </tr> <tr> <td style="padding-right: 10px;">AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b></td> </tr> </table>	CONST = <b>y</b>	FORMAL = <b>x</b>	TELIC = <b>eat(e<sub>2</sub>,z,x)</b>	AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b>
<b>cake</b>													
ARGSTR =	<table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">ARG1 = <b>x:food_ind</b></td> </tr> <tr> <td style="padding-right: 10px;">D-ARG1 = <b>y:mass</b></td> </tr> </table>	ARG1 = <b>x:food_ind</b>	D-ARG1 = <b>y:mass</b>										
ARG1 = <b>x:food_ind</b>													
D-ARG1 = <b>y:mass</b>													
QUALIA =	<table style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding-right: 10px;">CONST = <b>y</b></td> </tr> <tr> <td style="padding-right: 10px;">FORMAL = <b>x</b></td> </tr> <tr> <td style="padding-right: 10px;">TELIC = <b>eat(e<sub>2</sub>,z,x)</b></td> </tr> <tr> <td style="padding-right: 10px;">AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b></td> </tr> </table>	CONST = <b>y</b>	FORMAL = <b>x</b>	TELIC = <b>eat(e<sub>2</sub>,z,x)</b>	AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b>								
CONST = <b>y</b>													
FORMAL = <b>x</b>													
TELIC = <b>eat(e<sub>2</sub>,z,x)</b>													
AGENTIVE = <b>bake_act(e<sub>1</sub>,w,y)</b>													

The qualia structure for *cake* specifies that there is an event associated with this object's coming into being, namely the process of baking. In this case, the agentive role of the complement noun makes reference to the very process denoted by the verb phrase, a relation referred to as *co-specification* (Pustejovsky 1991, 1995). The 'creation' sense of *bake* thus arises from the semantics of *cake* by virtue of its being an artefact (i.e. originating out of the act of baking). On this account, therefore, it is not the verb *bake* itself that is polysemous, but the difference in meaning between *bake a cake* and *bake a potato* is the result of principles of semantic composition.<sup>2</sup> A verb such as *use*, which is said to be underspecified with respect to the activity being performed, is analysed along similar lines (examples from Pustejovsky 1995):

- (23) a. John used *the new knife* on the turkey.  
 b. Mary has used *soft contact lenses* since college.  
 c. This car uses *unleaded gasoline*.  
 d. My wife uses *the subway* every day.

In these cases, it is the object NPs, by virtue of their qualia structures, which contribute towards specifying the interpretation of *use*. In (23a) *use* refers to the action of *cutting*, in (23b) to *wearing* the contact lenses, and similarly in (23c) and (23d) the actions expressed by *use* are determined by the kind of event associated with the object NPs. The telic roles of the objects thus play an important role in determining which action is expressed by the verb; it is assumed that our

---

<sup>2</sup> If the lexical entry for *potato* is taken to be the same as that for *cake* except for the agentive role, which could be specified as **grow(e,w,y)**, then *potato* would, just like *cake*, have an event structure associated with it. Pustejovsky (1991: 422) acknowledges that "relative to the process of *growing*, the noun *potato* does denote an event". He does not, however, make it clear whether he thinks that *potato* acquires this event structure when combined with the verb *grow* (e.g. *Mary grew the potato*), in which case it would be the verb that is modifying the meaning of its complement (and not vice versa), or if this event structure is somehow cancelled when *potato* is combined with a verb such as *bake*.

understanding of sentences such as the ones in (23) is facilitated by default interpretations of properties and activities associated with the objects. In syntactic terms this means that the semantics of a governing expression (in this case the head of VP) is determined by the semantic content of the expression it governs (in (23) by the object NPs).

*2.2.3 Selective Binding.* The third generative mechanism presented in Pustejovsky (1991, 1995, 1998a, 1998b) is that of selective binding, which is described as a process whereby an adjective “is able to make available a selective interpretation of an event expression contained in the qualia for the head noun” (Pustejovsky 1995: 128). This is the process that is assumed to take place for an evaluative adjective such as *good* in (4), repeated here as (24):

- (24) a. Every chef needs a *good* knife.  
 b. ‘The Kite Runner’ is a *good* book.

In (24a), it is assumed that *good* selectively modifies the event description given by the telic role of the noun *knife* (i.e. the action of cutting), so as to yield the interpretation ‘a knife that cuts well’. Similarly, in (24b), the adjective acts as an event predicate for the function of the object to yield the interpretation ‘a good read’, rather than modifying the entire denotation of the noun.

### 3 Problems with the Generative Lexicon Theory

Pustejovsky’s theory of word meaning takes the semantics of a lexical item to involve a specification of argument, event, qualia and lexical inheritance structure.<sup>3</sup> It is assumed that verbs, nouns and adjectives may all function actively in determining the denotation of a phrase, as a result of the application of the generative operations of type-coercion, co-composition and selective binding. The theory thus seems to fulfil its goal of accounting for polysemy without having to assume multiple listings of word meanings stored in the lexicon, save in the case of

---

<sup>3</sup> The notion of lexical inheritance structure is not given much attention in Pustejovsky (1995). Pustejovsky (1991: 433) distinguishes between two inheritance mechanisms for representing the conceptual relations in the lexicon: *fixed* inheritance, involving e.g. the hypernym/hyponym relation (e.g. *book/novel*), and *projective* inheritance, which is assumed to operate “*generatively* from the qualia structure of a lexical item to create a relational structure for ad hoc categories.” (ibid.). This mechanism is supposed to account for a relation such as that between e.g. *escape* and *prisoner* in a sentence like *The prisoner escaped last night*, and for the information conveyed by this sentence having a more prototypical character than that of e.g. *The prisoner ate dinner last night*.

words encoding a lexical conceptual paradigm, where two of the meanings will be listed in the qualia structure. The theory also makes some interesting observations with regard to the interaction between word meanings in linguistic context. However, I believe that, ultimately, the theory is unworkable; first of all because it makes a range of wrong predictions, and, second, because the correct predictions that it makes are better treated pragmatically. The details of such a pragmatic account will be described in the next section. In this section, I will consider three central objections to the GLT, concerning (1) the way certain inferences seem to come out as strictly necessary entailments as a result of the generative operations over lexical entries; (2) the (in)ability of the complex lexical representations to pick out the right denotations in the world, and (3) the extent to which world knowledge is assumed to be built into the lexicon. The two latter issues are in fact closely related, as they both concern the kind of information stored as the semantic component of the lexical representation of a word in the GLT.

### 3.1 Necessary Inferences?

An important objection against the GLT, which has been advanced by a number of authors (e.g. Blutner 2002, Fodor & Lepore 1998), is the way certain inferences seem to come out as strictly necessary entailments as a result of the application of the generative operations over lexical entries. Consider the use of *bake* in (25), which is Pustejovsky's prime example of the process of co-composition:

(25) Mary *baked* the pizza.

Here Pustejovsky's account would predict a creative reading of *bake* (*pizza* being an artefact in much the same way as *cake* is), but a non-creative reading, on which Mary is simply heating up a frozen pizza, would be just as likely to be the preferred reading as the creative one. The inference via co-composition from *Mary baked the pizza* to 'Mary baked the pizza (creatively)' is therefore defeasible, hence not strictly necessary. The same point applies to inferences involving type coercion and selective binding. For example, *want a cigarette* is type coerced into 'want to smoke a cigarette' (Fodor & Lepore 1998: 274), but it could also be interpreted as e.g. 'want to hold a cigarette' in a context where someone has recently quit smoking but wants to hold a cigarette to keep the craving away (or as 'want to smell a cigarette'). By selective binding, *a fast book* yields 'a book that can be read in a short time', but in a context of book-sellers, *fast books* could be interpreted as the ones that sell quickly, and *slow books* as the ones that stay on the shelves for a long time (and there are numerous other interpretive possibilities).

Blutner (2002), following Fodor and Lepore (1998), points out that there is also a problem related to the restrictiveness of the coercion mechanism, in that the

account incorrectly predicts that *begin a car* should mean ‘begin to drive a car’ and *begin a thermometer* should mean ‘begin to measure the temperature’. This problem also extends to cases of selective binding, where the account would predict that e.g. *a fast cake* should be interpreted as ‘a cake that is fast to eat’ (cf. the telic role for *cake* given above). Copestake and Briscoe (1996), who combine an HPSG-like approach to syntax with Pustejovsky’s notion of qualia structure in their account of polysemy, try to solve this problem by treating compositional interpretations from the qualia as defeasible. An NP such as e.g. *fast typist* is normally interpreted as someone who types rapidly, but by allowing for specific (linguistic) contexts to suggest other interpretations that are not inherent to the qualia of the lexical item, the NP *fast typist* might also be used to refer to a typist who is e.g. running quickly. In such cases of non-default interpretations, the linguistic context is expected to explicitly give the exceptional component used in the interpretation.

However, the move to (defeasible) default interpretations does not seem to solve the problem. Consider the type coercion mechanism proposed by Pustejovsky again. Even if we assume that the interpretations generated by this mechanism are default but defeasible, we still cannot prevent it from computing the incorrect readings described above. Also, the coercion mechanism fails to indicate which is the default interpretation in cases where there is more than one possible interpretation compatible with the selection requirement of a verb, such as in e.g. example (18) above, where *Mary began a book* is compatible with both the interpretation ‘Mary began to write a book’ and ‘Mary began to read a book’. If we assume that there are default (but defeasible) compositional interpretations associated with a qualia structure, which interpretation of *Mary began a book* should we take to be the default one, and why? Both possible interpretations may be generated by the process of type coercion, acting upon information specified by the qualia structure for *book*. Assuming that, say, ‘Mary began to read a book’ is the default interpretation, does this mean that the information specified by the telic role (which involves ‘read’) of a noun is more ‘default’ than that specified by the agentive role (which involves ‘write’)? In these cases, therefore, it seems that the coercion mechanism just generates another underspecified structure (i.e. two possible event interpretations), and it would be a task for pragmatics to choose between them.

Also, the move to defeasible default interpretations raises another kind of question related to whether we want our semantics to include defeasibility when we have already got it as part of our pragmatic system. No reason is given for not simply leaving it to pragmatics, using information from the immediate context as well as general world knowledge, to construct the intended interpretation.

### 3.2 Adequate Lexical Semantic Representations?

In the GLT, the notion of ‘qualia structure’, and the four types of roles (constitutive, formal, telic, and agentive roles) seem to apply mainly to the description of the semantics of nouns, and to nouns denoting concrete objects in particular. In fact, with only a few exceptions, the nouns discussed in Pustejovsky (1995) all denote physical objects. Pustejovsky maintains that a qualia structure can be defined for all types of lexical items, but that not all lexical items need carry a value for each role. It is, however, hard to see how the four types of roles associated with a qualia structure are supposed to apply to e.g. prepositions, connectives, etc. Even for lexical categories such as verbs, adjectives, and adverbs, it seems like a category mistake to try to describe their semantic content in terms of a qualia structure, as defined in (9). Does it make sense to speak of e.g. what an action/process is made of (i.e. its constitutive role) or what the function of an adjective is (i.e. its telic role)? Furthermore, there seems to be a problem with regard to Pustejovsky’s specifications of argument structure for a number of lexical items. The lexical representation he gives for e.g. *build* in (7) specifies that this verb takes an animate individual as its subject. However, this makes it unclear how he would accommodate examples such as e.g. *The army built the wall in a day*, *The cognitive system builds and stores representations in memory* into this analysis. Similarly, the subject of the verb *begin* in (19) is specified as ‘human’, which seems to rule out examples such as e.g. *The engine began to make a strange noise*, etc. There seems, in fact, to be exceptions to all of Pustejovsky’s uses of ‘animate’, ‘human’, etc., as they figure as part of his lexical entries.

One may also question GLT’s notion of so-called ‘default’ and ‘shadow’ arguments, as arguments which are not necessarily syntactically expressed, but nevertheless encoded as part of the semantic representation for a lexical item. The question is whether the information contained by such ‘arguments’ should be regarded as linguistic. In the case of *John buttered his sandwich* (involving the shadow argument *butter*) and *Mary built a house* (involving a default argument specified as ‘material’) it is clearly a matter of metaphysical necessity that the bread has to have been buttered with some kind of material (e.g. butter, margarine, peanut butter, chocolate butter, etc.) and that the house has to have been built out of some material or other (e.g. bricks, wood, stone, corrugated iron, etc.), but this fact alone does not entail that they are semantically necessary. Consider an alternative view based on the concept of an ‘unarticulated constituent’, that is, a constituent of the statement made (by uttering a sentence in a given context) which does not correspond to any linguistic entity in the sentence. Recanati (2002: 305-306) makes a distinction between *metaphysical* and *communicational* unarticulated constituents. For something to count as a communicational unarticulated



constituent, it must be part of “what the speaker means by his or her utterance” (ibid. 306), while a metaphysical unarticulated constituent is subject to no such requirement. In the event of a speaker uttering *I’ve had a large breakfast* as a response to the question *Are you hungry?*, it is assumed that the hearer has to identify the temporal location of the breakfast event (as the day of the utterance) in order to be able to derive the intended implicated meaning (implicature) that the speaker is not hungry, hence this would count as an unarticulated constituent in the communicational sense. A case of a metaphysical unarticulated constituent would be the location of the dancing in an utterance of *Juliet danced*; obviously there is, as a matter of metaphysical necessity, a place at which Juliet danced but comprehension of the proposition expressed by the utterance does not require the recovery of a location constituent. The GLT ‘shadow’ and ‘default’ arguments posited in the case of *John buttered his sandwich* and *Mary built the house* would, on this account, be cases of metaphysical unarticulated constituents, since (in most instances, at least) the hearer will be able to understand these utterances without having to identify the material with which the sandwich was buttered or out of which the house was built. While a case could be made for there being some covert linguistic entity corresponding to the communicational constituents (see Stanley 2000), there is no such case for the metaphysical unarticulated constituents, which suggests that Pustejovsky’s shadow and default arguments should not be regarded as linguistic, but rather as part of our general knowledge about the world.

The idea of lexical conceptual paradigms as a way to account for metonymous senses of a word also involves some problems. In cases like *book* and *newspaper*, it seems that Pustejovsky’s account is not as different from the sense enumeration lexicon approach as he makes out, in that the two basic senses forming the input to the lcp operator are both listed in the lexical representations for these words. Furthermore, the fact that the type constructor generates a sense for *newspaper* for which there is no corresponding denotation (i.e. the ‘org.info.physobj’ sense) does not seem to favour Pustejovsky’s analysis. We would not be happy with a generative grammar that generated ungrammatical sentences; similarly, a generative lexicon that generates impossible senses is unsatisfactory.

Finally, consider the lexical representation for the noun *cake* given in (22) above. If we were to provide a representation of the noun *pizza* discussed in 3.1 along the same lines, it is actually hard to see how this semantic representation would differ from the one given for *cake*, as all the properties and activities associated with *cake* in (22) may also be said of *pizza* (and possibly of other artefacts of similar type). How, then, are these lexical representations supposed to pick out the right denotations in the world? It seems that Pustejovsky’s lexical entries do not specify sufficient conditions for the application of a concept. Furthermore, as pointed out above in the discussion of *begin a book*, some of the features generated are not necessary ones. Thus, there are three kinds of failure with regard to denotations: (1)

sometimes the lexical representations do not distinguish between what are obviously distinct denotations (e.g. *pizza, cake*); (2) sometimes they predict a narrower denotation than is in fact the case (e.g. *begin a book, enjoy a book*); and (3) sometimes they simply predict a wrong denotation (e.g. *begin a car, begin a thermometer*).

The problem of giving necessary and sufficient conditions for the application of a concept is well-known within philosophy and cognitive science, and Fodor (1970, 1981, 1998) and Fodor et al. (1980) have provided convincing arguments against the possibility of a definitional analysis of word meaning. Given this ‘undefinability’ of most words (including nouns), it seems likely that, in many cases, it will be difficult to give any precise values for e.g. the qualia roles (or for any of the other levels of linguistic representation proposed by Pustejovsky) associated with a lexical item. Try, for instance, to provide qualia structures for abstract nouns such as e.g. *semantics, freedom, justice*, etc.

### 3.3 World Knowledge and the Lexicon

Another problematic issue regarding the lexical representations assumed by the GLT is the way it presupposes world knowledge being built into the lexicon. This is one of Fodor and Lepore’s (1998) main objections to the theory; that it lacks a distinction between what the language tells us and what we know about the world. Fodor and Lepore argue that even though it may seem at first glance that Pustejovsky’s theory is able to account for the differences in interpretation of *bake* in *bake a cake* and *bake a potato* solely by distinguishing artefacts from natural kinds, this is clearly not sufficient. Although e.g. knives and trolley cars are artefacts, *bake a knife* and *bake a trolley car* resist a creative reading just as much as *bake a potato* does. They further claim that even *bake a cake* is ambiguous, since it is perfectly possible to bake a (pre-existing) cake by putting it in the oven and non-creatively baking it. Instead, they argue, the reason we have an intuition that *bake* is lexically ambiguous in the case of (20) above, is because the verb is genuinely polysemous between the two readings: The verb *bake* has a single lexical entry, but denotes two different processes (the creative and the change of state sense), and it is our knowledge about the world that tells us which reading of *bake* to pick out in the case of *bake a potato*, and that makes *bake a knife* sound funny. On their view, it is not the case that the lexicon encodes or generates this information.

Fodor and Lepore (1998: 281) further claim that Pustejovsky’s account offers no convincing cases of the meaning of a governing expression being modulated by the lexical content of the expression that it governs. In the case of the verb *use*, they note that it is perfectly possible to use things that do not have uses (e.g. a rock to break a window, snow to make a snowball, etc.) and that this presents a serious

problem for Pustejovsky's theory, as there may, in fact, be many cases in which a verb makes a demand on an argument that the lexical entry for that argument does not satisfy (and we do not want to say that the telic role for e.g. *rock* includes the function *break*, i.e. that this function is part of the semantics of the noun). This problem also extends to cases of selective binding, where the lexical representation of a governing noun may, in many cases, not contain a specification of a telic role at all, in which case there would be no event description for the adjective to selectively modify. Szabó (2001: 132) points out that the meaning of the sentence *This is a good pebble* depends on the context of its use (the pebble could be good e.g. for playing marbles, breaking a window, etc.), and thus concludes that the meaning of a noun *N* in a phrase '*a* is a good *N*' is often not sufficient in order to decide in which way *a* is said to be good. Moreover, Fodor and Lepore argue that it is perfectly possible to understand a sentence such as e.g. *Sarah likes to use the subway* and to wonder at the same time what she likes using it for (she might like to use the subway for sightseeing, shelter, etc.). The inference from *likes to use the subway* to 'likes to travel on the subway' is made possible by what we know about subways but this information should not be regarded as contained as part of the lexical entry for *subway*. On their view, therefore, Pustejovsky's theory conflates linguistic knowledge and world knowledge, and furthermore offers no way of constraining the amount of world knowledge entering into the lexical representation of a given lexical item.

The same point can be made with regard to the analysis of *bake the pizza* in terms of co-composition: we certainly do not want to say that the lexical representation for *pizza* actually contains some kind of distinction between frozen and home-made pizzas as part of its qualia structure, which is one way in which Pustejovsky could account for the two interpretations of *bake the pizza* if he wants to maintain his analysis of *bake* as dependent on the semantics of the complement noun. However, this distinction is clearly part of our world knowledge, and contextual factors (including extra-linguistic ones) will indicate which interpretation to select or construct in each case.

#### 4 A Relevance-Theoretic Approach

The problems discussed in the previous section illustrate the limitations of relying on linguistic context alone in the derivation of compositional interpretations. An important motivation for a re-analysis of Pustejovsky's examples of 'logical polysemy' in pragmatic terms is the fact that the GLT, although incorporating a considerable amount of world knowledge into the lexicon, still leaves a lot of work for pragmatics to do in finding the interpretation intended by a speaker. First of all, pragmatics is needed to correct the cases of wrong interpretations predicted by the

theory. Second, when there is more than one possible default interpretation associated with a qualia structure, it would be entirely a matter of pragmatics to choose which is the correct one (if any of them) in a given context. Žegarac (2006: 1707) argues that a pragmatic theory which is capable of doing the considerable amount of work that a GLT-based account of lexical meaning leaves for it to do can also do that part of the interpretive work that a GLT does adequately. I follow him in this and would add that a wholly pragmatic account is to be preferred, not only because it would give a simpler, more unified account overall (with a much leaner semantics), but also if it can be shown to avoid the problems of misinterpretation discussed in the previous section. I believe that relevance theory (Sperber & Wilson 1986/1995, Carston 2002) provides a pragmatic theory that does meet these requirements.

Relevance-theoretic pragmatics is based on a definition of relevance and two principles: the Cognitive and Communicative Principles of Relevance. Relevance is a potential property of any input to cognitive processes, and may be assessed in terms of processing effort and cognitive effects. An input is relevant to an individual when it yields a positive cognitive effect as a result of being processed in a context of available assumptions (Wilson & Sperber 2004). Positive cognitive effects include contextual implications (that is, conclusions that are deducible from input and context together), strengthening and/or elimination of existing assumptions. Other things being equal, the greater the cognitive effects of an input to an individual who processes it, and the smaller the processing effort required to derive these effects, the greater the relevance of that input to that individual at that time. The two principles, given in (25) and (26), are proposed on the basis of this definition:

(26) **The Cognitive Principle of Relevance**

Human cognition tends to be geared to the maximisation of relevance. (Sperber & Wilson 1995: 260)

(27) **The Communicative Principle of Relevance**

Every act of ostensive communication communicates a presumption of its own optimal relevance. (ibid.)

Following the definition of relevance, the Cognitive Principle of Relevance predicts that hearers will aim at deriving as many cognitive effects as possible for as little processing effort as possible, while it follows from the Communicative Principle of Relevance that the speaker, by the very act of addressing the hearer, is communicating that her utterance is the most relevant one she could have produced given her own preferences and abilities, and is at least relevant enough to be worth the hearer's processing effort (i.e. it is 'optimally' relevant). The hearer's task,

then, is to find an interpretation of the speaker's utterance that is consistent with this presumption, i.e. the first one he accesses which yields sufficient cognitive effects for no gratuitous processing effort. (For more detail, see Wilson & Sperber 2004.)

#### 4.1 Linguistic Meaning and Explicature

The relevance-theoretic account assumes that communication is essentially an inferential process. Inferential communication is taken to involve the application of general-purpose inference rules, applying to any type of information which is represented conceptually (Sperber & Wilson 1986/1995: 176). On this view, linguistic communication is but one way in which communication may obtain, but it is special in the sense that it can achieve a degree of precision and complexity which is rarely achieved in non-linguistic communication. Verbal communication is taken to involve two kinds of communication/comprehension process: one is coding and decoding, the other ostension and (non-demonstrative) inference. The first one provides input to the pragmatic inferential process, which is seen as the main part of the comprehension process.

Now let us focus in a bit more detail on the relevance-theoretic notion of the input to the inferential process, that is, the view of linguistic meaning assumed by the theory. The general idea is that the linguistic meaning of an utterance greatly underdetermines its interpretation: it not only underdetermines the implicit meaning of the utterance (i.e. its implicature(s)), but also the proposition expressed by it, i.e. its *explicit* content (Carston 2002: 29). On this view, the semantic representation encoded by a sentence is taken to be an incomplete logical form, which is, at best, a fragmentary representation of a thought. In order to recover the complete (truth-evaluable) proposition the speaker has expressed from the linguistically decoded semantic representation, it must be inferentially enriched (by disambiguation, assignment of referents, free enrichment) in a context of available assumptions. The resulting proposition is an *explicature*, which is defined as follows (Carston 2002: 377):

[An explicature is] an ostensively communicated assumption which is inferentially developed from one of the incomplete conceptual representations (logical forms) encoded by the utterance.

It is assumed that, at every stage of the process of developing the linguistic meaning into a complete proposition, the hearer will choose the solution involving the least processing effort, and this solution will be abandoned only if it does not yield an interpretation consistent with his expectations of relevance (Sperber & Wilson 1986/1995: 185). This view of linguistic meaning, on which semantic

representations are incomplete logical forms that must be contextually developed, departs radically from the one assumed by GLT, where a lot of the information which is taken to be pragmatically inferred on the relevance-theoretic account, is encoded in the lexicon.

## 4.2 A Re-analysis

*4.2.1 Free Enrichment.* Now let us return to our examples of ‘logical polysemy’. I will leave aside the metonymy example in (1) for a moment, and focus on the one in (2) (repeated here as (28)):

(28) Mary *began* a book.

Recall that Pustejovsky analyses this as a case of type coercion, where the type of the complement is changed into an event consistent with the interpretations ‘Mary began reading a book’ and ‘Mary began writing a book’, as a result of the events of reading and writing being encoded as part of the qualia structure for *book*. On the relevance-theoretic account, of course, no such information is assumed to be encoded by this word, and the relevant event associated with *begin a book* will be supplied entirely on pragmatic grounds. This would count as a case of ‘free’ pragmatic enrichment (e.g. Carston 2002: 323), where a conceptual constituent which is not articulated in the linguistic form of the utterance is contextually derived, as a result of the hearer’s trying to figure out which proposition the speaker has explicitly expressed.<sup>4</sup> Let us imagine a context where John, Mary and Sue have been given the job of cleaning their late grandmother’s house before it is advertised for sale. The bookshelves haven’t been dusted for years, and, because they are completely covered in dust, each book has to be dusted individually. After they have been cleaning for a while, John asks if they should all take a break and go for coffee, and Sue replies (29):

(29) Let’s wait for a couple of minutes. Mary just began a book.

In this context, the interpretation according to which the explicit content of the last part of Sue’s utterance is ‘Mary just began [DUSTING] a book’ would clearly be the most relevant one. It is the most accessible interpretation, that is, the one requiring

---

<sup>4</sup> However, if it turns out that the semantics of the verb *begin* comes with some kind of parameter indicating that its complement is an event or activity, then this will be a case of saturation (rather than free enrichment). Either way, though, the recovery of the ‘reading’, ‘writing’ or other specific book-involving activity is a matter for context and pragmatic inference.

the least processing effort on the part of the hearer, and it is also one that offers an adequate explanation for why Sue thinks they should wait before leaving for coffee (it functions as a premise for the conclusion given in the first part of Sue's utterance). John derives this interpretation according to the relevance-theoretic comprehension heuristic (Wilson & Sperber 2004), which is applied automatically to verbal input. According to this procedure, the hearer (1) takes the decoded linguistic meaning, follows a path of least effort in mutually adjusting explicit content, contextual assumptions and contextual implications, and (2) stops when the interpretation he arrives at satisfies his expectations of relevance. At no stage in the interpretation process, therefore, is it assumed that a 'default' interpretation (according to which Mary began e.g. reading a book) is computed and then cancelled by context.<sup>5</sup> It doesn't take much imagination to think of contexts in which the correct (intended and easily retrieved) interpretation is 'began binding a book', 'began mending a book', 'began ripping up a book', and so on.

However, there is not much doubt that Pustejovsky is right in assuming that certain interpretations come more readily to mind than others in a 'null' context. In the absence of real-life contextual constraints, the utterance in (29) would probably most often be interpreted as meaning that Mary began reading a book. Sperber and Wilson (1986/1995: 185) see this as a case of hearers favouring the least effort-consuming conceivable interpretation. A person reading a book may be regarded as a stereotypical event, which may be stored in encyclopaedic memory as a chunk and accessed as a single unit of information. Retrieving this information from encyclopaedic memory during the interpretation of (29) would require little processing effort, while deriving the interpretation according to which Mary began e.g. dusting a book would involve accessing several units of information and hence be more costly in processing terms (see discussion in Sperber & Wilson 1986/1995: 186). In this way, relevance theory provides a natural explanation for the fact that certain interpretations are often favoured over others, without being committed to the view that these interpretations are always computed first as a result of default inferences associated with lexical items.

*4.2.2 Ad Hoc Concept Construction.* In order to account for Pustejovsky's remaining cases of 'logical polysemy', we need to take a closer look at the relevance-theoretic account of concepts and lexical-pragmatic processes. Not surprisingly, relevance-theorists take a radically different stance on the question of

---

<sup>5</sup> The different predictions made by GLT and relevance theory with regard to the derivation of compositional interpretations in cases like (29) should be experimentally testable. Lexical priming experiments (cf. e.g. Rubio Fernández 2007) may provide evidence either for the view that a so-called default interpretation is accessed first and then cancelled by context (GLT), or for the view that only the relevant interpretation is derived (RT).

word meaning from the decompositional view assumed by the GLT, in that they do not believe that it is generally possible to provide definitional word meanings (i.e. components of meaning which are individually necessary and jointly sufficient to capture the encoded content of a word). On the relevance-theoretic account, therefore, most words encode simple, unanalysable, atomic concepts in the sense of Fodor (1981, 1998). The word *cat*, for instance, encodes the atomic concept CAT. Such atomic concepts are represented in our ‘language of thought’, and consist of an address, or entry, in memory that may give access to three kinds of information: (1) the logical properties of the concept (e.g. a one-way meaning postulate, such as  $CAT \Rightarrow ANIMAL\ OF\ A\ CERTAIN\ KIND$ );<sup>6</sup> (2) a set of assumptions, or encyclopaedic information, about the denotation of the concept (e.g. ‘cats have four legs’, ‘cats purr’, various kinds of culture-specific information related to the concept, etc.), and (3) the lexical (i.e. phonological and syntactic) information related to the linguistic form that encodes the concept (Sperber & Wilson 1995: 85-93). On this view, lexical interpretation typically involves constructing an ad hoc concept, or an occasion-specific sense, which may be narrower (e.g. *bird* used to express ‘bird of a specific sort, e.g. woodland bird’) or broader (e.g. *wizard* used metaphorically to express ‘a person who is very skilled in a particular field or activity’) than the linguistically-specified denotation. The hearer takes the encoded concept and its associated logical and encyclopaedic entries as input, and uses the relevance-theoretic comprehension heuristic to derive warranted conclusions about the speaker’s meaning (Wilson 2003, Wilson & Carston 2007), a process which includes a mutual adjustment of explicit content, contextual assumptions and contextual implications, guided and constrained by expectations of relevance. In many cases, an ad hoc concept is constructed as a result of a process of ‘backward inference’, where there is an inference from expected contextual implications to the explicitly expressed proposition containing the ad hoc concept (For more detail, see Carston & Powell 2006, Wilson & Sperber 2002, Wilson & Carston 2006). Concept narrowing and broadening are thus seen as the outcome of a single interpretive process, which adjusts the interpretation of individual words (Carston 1997, 2002, Wilson & Carston 2007).

Now let us consider again the example involving the verb *bake*, and see how the difference in meaning between *bake a cake* and *bake a potato* can be analysed in relevance-theoretic terms. The example is repeated below as (30):

---

<sup>6</sup> In espousing logical entries for concepts, relevance theory stands with Fodor in his earlier manifestation when he supported atomism (no definitions) together with meaning postulates (e.g. Fodor et al. 1980, Fodor 1981), rather than with his more recent position which eschews any logical/ encyclopaedic distinction (Fodor 1998).



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

As we have seen, Pustejovsky analyses the verb *bake* as having a single change of state denotation, which is changed into the creation sense as a result of the process of co-composition, whose application is dependent on the semantics of the complement noun (e.g. natural kind/artefact). Etymologically, the basic sense of *bake* is assumed to be ‘cook by dry heat’,<sup>7</sup> and I will therefore follow Pustejovsky in considering the change of state denotation to be the one that is indeed linguistically specified by this word. However, in what follows I will try to show that, contrary to Pustejovsky’s view, the change in meaning of the verb in (30b) is a pragmatic matter rather than a lexical semantic operation internal to the language system.

If we assume that the linguistically-specified denotation of the concept BAKE is the change of state meaning, then the concept communicated by use of *bake* in (30b) may be analysed as a case of pragmatic narrowing, since ‘creation’ in fact entails a ‘change of state’ (but not vice versa). In this case, the hearer will construct the ad hoc concept BAKE\*, which will have a more restricted denotation than the encoded concept, i.e. picking out a subset consisting of those instances of the process of cooking by dry heat which bring something into existence (cake, pizza, bread, etc). The hearer arrives at this interpretation as a result of applying the relevance-theoretic comprehension heuristic. The decoded meaning of the sentence in (30b) will contain the concepts BAKE and CAKE, both of which will give access to a range of encyclopaedic information associated with their denotations. The encyclopaedic entry for CAKE will contain assumptions about the way in which cakes come into being, i.e. as a result of a ‘creative’ process of baking. The interpretation according to which BAKE communicates the creation sense BAKE\*, would in this case be the one that is the most accessible to the hearer. Similarly, in (30a), where *bake* is used to communicate its encoded concept (BAKE), an important step in the hearer’s accepting the encoded change of state meaning will be his accessing of the encyclopaedic entry for the concept POTATO, where he would find information to warrant this interpretation of *bake*.<sup>8</sup>

Some may argue that the difference between this analysis of *bake* and the one given by Pustejovsky is really just a matter of whether one takes the information

---

<sup>7</sup> *The Oxford Dictionary of English Etymology*. 1966/1985. Oxford: Clarendon Press.

<sup>8</sup> It would be interesting to do more research on the etymology of the verb *bake* in Germanic languages: in a language like Norwegian, for instance, the creative sense of the word is more likely to be the basic one, while the change of state sense is only accessible when the verb is combined with specific complements (such as e.g. *potato*). Such cases would induce a broadening of the linguistically-specified denotation of *bake*.

associated with the complement nouns to be linguistically encoded or part of the encyclopaedic information associated with the encoded concept. However, the pragmatic analysis has a great advantage over the semantic one, in that it allows for the necessary flexibility in lexical interpretation (in a given context, *bake* in *bake a cake* could just as well be used to convey the change of state meaning), while on the semantic approach one is committed to assuming that there is a single interpretation for *bake a cake*, which can be overridden only in the case of explicit contextual evidence pointing to a different interpretation.

However, a possible objection to the analysis of the creative sense of *bake* as an instance of on-line ad hoc concept construction could be that this sense may in fact be lexicalised. On this view, *bake* would be a case of genuine polysemy, as argued by Fodor and Lepore (1998), and there would be a single lexical entry for *bake*, with two senses, each denoting different but semantically related processes. This may well be so (at least for some people), in which case the pragmatic analysis just given can be seen as providing an account of the semantic relation between these two senses: assuming that the ‘change of state’ sense is the basic one, the ‘creation’ sense may have become lexicalised as a result of frequent narrowing. (For further discussion of the polysemy issue, see Wilson & Carston 2006.)

Turning to Pustejovsky’s examples of selective binding, these can be straightforwardly analysed as cases of lexical narrowing on the relevance-theoretic account. Consider again the case of *good*, repeated below as (31):

- (31) a. Every chef needs a *good* knife.  
b. ‘The Kite Runner’ is a *good* book.

This adjective may be seen as encoding a very broad concept, which, on most occasions of use, will have to be pragmatically adjusted in order to arrive at the interpretation that was intended by a speaker. In (31a), the concept communicated by *good* would be the ad hoc concept GOOD\*, conveying ‘good for cutting’, which would have a much narrower denotation than the one linguistically-encoded by the word. In (31b), a different ad hoc concept, GOOD\*\*, would be communicated by this word, conveying that ‘The Kite Runner’ is a ‘good read’.

There are, however, many other ways in which a book can be good. A *good book* can be one that is easily read, one that is entertaining, one of a high academic or intellectual quality, one that is good to kill flies with, etc. Similarly, it is not difficult to imagine contexts in which *good knife* would be given entirely different interpretations than the one in (31a). Consider (32)-(33):

- (32) To become a member of Billy's exclusive gang you had to have a *good* knife.<sup>9</sup>
- (33) This is a *good* knife for people with wrist pain.

In (32), the communicated meaning of *good* could be 'good for stabbing people' or 'good for threatening people with', while in (33), it could be 'good for holding'. Again, the relevance-theoretic approach predicts that *good* may communicate different occasion-specific senses in these cases, by assuming that lexical interpretation is a matter of adjusting the interpretation of individual words in accordance with one's context-specific expectations of relevance. The variety of possible interpretations of *good* in the above examples also strongly suggests that Pustejovsky's analysis of the meaning of *good* as dependent on the function of the object it modifies is inadequate.<sup>10</sup>

4.2.3 'The Problem of Metonymy'. Finally, let us consider our metonymy example in (1), repeated below as (34). The word *newspaper* has different denotations in (34a-c):

- (34) a. *The newspaper* fired its editor. (organisation)  
 b. Mary spilled coffee on *the newspaper*. (physical object)  
 c. *The newspaper* revealed surprising details about the trial. (information type)

On Pustejovsky's account, the different senses of *newspaper* were taken to be encoded directly into its lexical representation as a lexical conceptual paradigm, an analysis which involved the generation of a sense for *newspaper* for which there was no corresponding denotation (as *newspaper* cannot simultaneously denote an organisation, an information type and a physical object), a clearly unsatisfactory outcome. However, as Wilson and Carston (2007: 253) concede, the question of how to analyse cases of metonymy within the relevance-theoretic framework (in terms of ad hoc concept construction) is not straightforwardly answerable.<sup>11</sup> They suggest, following Nunberg's (1996, 2004) account, that at least some cases of metonymy could be analysed as involving lexical broadening. Nunberg argues that

---

<sup>9</sup> I owe this example to Robyn Carston.

<sup>10</sup> It is worth mentioning that the flexibility lacking in Pustejovsky's account is allowed for in an approach like Szabó's (2001), which postulates a variable in the semantic representation of *good*, inducing a pragmatic operation of saturation.

<sup>11</sup> Papafragou (1996) gives a preliminary relevance-theoretic account of metonymy, where metonymy is seen as a variety of interpretive use of language as defined by Sperber and Wilson (1986/1995: 228-229).

in many metonymical utterances (such as e.g. an utterance of *I'm parked out back*) the meaning transfer applies to the conventional meaning of the predicate (i.e. to *parked out back* and not to the subject *I*), and such cases can be analysed in terms of lexical broadening on the relevance-theoretic account. However, in (34) above, it is clearly the noun *newspaper* to which the alternation in meaning applies, and this case does not seem straightforwardly analysable in terms of lexical narrowing or broadening. Wilson and Carston (2007: 253-254) note that the notion of metonymy is intuitively harder to grasp than that of e.g. metaphor, and there is not always agreement on which cases are instances of metonymy and which are not. I believe it is a genuine question whether the different senses of *newspaper* should be regarded as semantic (that is, as encoded in the lexicon), in which case the noun could be regarded as being polysemous in the traditional sense (and associated with a set of conventionalised metonymic senses), or pragmatic, in which case one of the senses should be taken to be the encoded one and the other senses pragmatically derived. In such a case, however, it is not obvious which of the meanings in (34a-c) should be taken to be the basic one, although the physical object sense in (34b) would be a likely candidate. (For further discussion of the semantics/pragmatics issue, see Nunberg 2004: 351-354).<sup>12</sup> I will not choose between these possible analyses here, but will simply acknowledge that this is an issue which remains to be adequately explained within the relevance-theoretic framework.

## 5. Conclusion

In this paper I have discussed Pustejovsky's generative theory of word meaning and compositional interpretation, and its treatment of cases of 'logical polysemy'. I hope to have shown that the range of phenomena falling under this label in Pustejovsky's theory are better given a pragmatic account, and that the data can be reanalysed in terms of the independently-motivated pragmatic processes of free enrichment and ad hoc concept construction within the relevance-theoretic framework. This analysis allows for, in fact predicts, the flexibility that we find in the derivation of compositional interpretations.

---

<sup>12</sup> However, whether the metonymy of *newspaper* is regarded as a semantic or as a pragmatic issue, an adequate account of metonymy should have something to say about the distinction between clearly pragmatic cases, such as e.g. *The ham sandwich left without paying*, and other more predictable/systematic cases, such as e.g. *newspaper, book*, etc.

## References

- Atkins, B. T. & Kegl, J. & Levin, B. 1988. Anatomy of a verb entry: from linguistic theory to lexicographic practice. *International Journal of Lexicography* 1 (2), 84-126.
- Blutner, R. 2002. Lexical semantics and pragmatics. *Linguistische Berichte* 10, 27-58.
- Carston, R. 1997. Enrichment and loosening: complementary processes in deriving the proposition expressed? *Linguistische Berichte* 8, 103-127.
- Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. & Powell, G. 2006. Relevance theory – new directions and developments. In *The Oxford Handbook of Philosophy of Language*, Lepore, E. & Smith, B. (eds.), 341-360. Oxford: Oxford University Press.
- Copestake, A. & Briscoe, T. 1996. Semi-productive polysemy and sense extension. In *Lexical Semantics: The Problem of Polysemy*. Pustejovsky, J. & Boguraev, B. (eds.), 15-67. Oxford: Clarendon Press.
- Fodor, J. A. 1970. Three reasons for not deriving “kill” from “cause to die”. *Linguistic Inquiry* 1, 429-438.
- Fodor, J. A. 1981. The present status of the innateness controversy. In *Representations: Philosophical Essays on the Foundations of Cognitive Science*, 257-316. Brighton, Sussex: Harvester Press.
- Fodor, J. A. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
- Fodor, J. A., Garrett, M. F., Walker, E. C. T. & Parkes, C. H. 1980. Against definitions. *Cognition* 8, 263-367.
- Fodor, J. A. & Lepore, E. 1998. The emptiness of the lexicon: reflections on James Pustejovsky’s ‘The generative lexicon’ *Linguistic Inquiry* 29, 269-288.
- Langacker, R. 1984. Active zones. *Proceedings of the Tenth Annual Meeting of the Berkeley Linguistics Society*, 172-188.
- Nunberg, G. 1996. Transfers of meaning. In *Lexical Semantics: The Problem of Polysemy*. Pustejovsky, J. & Boguraev, B. (eds.), 109-132. Oxford: Clarendon Press.
- Nunberg, G. 2004. The pragmatics of deferred interpretation. In *The Handbook of Pragmatics*, Horn, L. & Ward, G. (eds.), 344-364. Oxford: Blackwell.
- Papafragou, A. 1996. On metonymy. *Lingua* 99, 169-195.
- Pustejovsky, J. 1991. The generative lexicon. *Computational Linguistics* 17 (4), 409-441.
- Pustejovsky, J. 1995. *The Generative Lexicon*. Cambridge, Massachusetts: The MIT Press.
- Pustejovsky, J. 1998a. Generativity and explanation in semantics. Reply to Fodor and Lepore. *Linguistic Inquiry* 29, 289-311.
- Pustejovsky, J. 1998b. The semantics of lexical underspecification. *Folia Linguistica* 32 (3-4), 323-347.
- Recanati, F. 2002. Unarticulated constituents. *Linguistics and Philosophy* 25, 299-345.
- Rubio Fernández, P. June, 16 2007. Suppression in metaphor interpretation: differences between meaning selection and meaning construction. *Journal of Semantics* doi:10.1093/jos/ffm006.
- Sperber, D. & Wilson, D. 1995 [1986]. *Relevance: Communication and Cognition*. Oxford: Blackwell. [2<sup>nd</sup> edition].
- Stanley, J. 2000. Context and logical form. *Linguistics and Philosophy* 23, 391-434.

- Szabó, Z. G. 2001. Adjectives in context. In *Perspectives on Semantics, Pragmatics and Discourse: A Festschrift for Ferenc Kiefer*, Kenesei, I. & Harnish, R. M. (eds.), 119-146. Amsterdam: John Benjamins.
- Wilson, D. 2003. Relevance theory and lexical pragmatics. *Italia Journal of Linguistics/Rivista di Linguistica* 15 (2): 273-291. [Special Issue on Pragmatics and the Lexicon].
- Wilson, D. & Carston, R. 2006. Metaphor, relevance and the 'emergent property' issue. *Mind & Language* 21 (3), 404-433.
- Wilson, D. & Carston, R. 2007. A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In *Pragmatics*, Burton-Roberts, N. (ed), 230-259. Palgrave.
- Wilson, D. & Sperber, D. 2002. Truthfulness and relevance. *Mind* 111, 538-632.
- Wilson, D. & Sperber, D. 2004. Relevance theory. In *The Handbook of Pragmatics*, Horn, L. R. & Ward, G. (eds.), 607-632. Oxford: Blackwell.
- Žegarac, V. 2006. 'Believing in': a pragmatic account. *Lingua* 116 (10), 1703-1721. [*Language in Mind: A Tribute to Neil Smith on the Occasion of his Retirement*].