Meaning, Concepts, and the Lexicon*

Michael Glanzberg
Northwestern University

Abstract

This paper explores how words relate to concepts. It argues that in many cases, words get their meanings in part by associating with concepts, but only in conjunction with substantial input from language. Language packages concepts in grammatically determined ways. This structures the meanings of words, and determines which sorts of concepts map to words. The results are linguistically modulated meanings, and the extra-linguistic concepts associated with words are often not what intuitively would be expected. The paper concludes by discussing implications of this thesis for the relation of word to sentence meaning, and for issues of linguistic determinism.

Keywords: meaning, concept, lexicon, semantics, syntax-semantics interface

This paper will explore one aspect of the way our words relate to our wider conceptual abilities. It is natural to think that in many cases, that relation is quite simple: we have concepts, and our words express them. This paper will argue that this picture is too simple. In many cases, linguistic meaning stands in a far more complex relation to our concepts. As I shall put it below, language *packages* concepts in linguistically, and substantially grammatically,

*Versions of the work presented here were presented at the Conference on Philosophy of Linguistics in Dubrovnik, 2010 and the conference on Formal Semantics and Pragmatics in Riga, 2010. Thanks to all the participants at those events for valuable discussion. I am especially indebted to John Collins, Kai von Fintel, Itamar Frances, Guy Longworth, Dean Pettit, Craige Roberts, Barry Smith, and Zoltán Szabó for helpful questions and comments. The references in this paper do not adequately show my debt to work of Paul Pietroski. This paper is, at heart, an extended reflection on some ideas he has been forcefully advocating in many of his writings.

determined ways. Not only does this structure the meanings of our words, it also determines which sorts of concepts map to words.

The force of this thesis can be made more vivid if adopt a few assumptions about the mind, and the place of language in it. Let us suppose that we have a reasonably large repertoire of ordinary concepts, and that much of our thinking involves manipulating such concepts in some way. Let us also suppose that language is a special faculty of the mind, as Chomsky has long argued (e.g. Chomsky, 1986, 2000). In particular, let us suppose that the core grammatical and semantic properties of words are fixed by the language faculty. We expect that in certain cases, this is done by establishing a connection between a given 'word' (a given phonological shape) and a concept provided by our extra-linguistic cognitive abilities. Not every word needs to have its content fixed this way. Arguably some functional elements do not. For instance, whatever meaning a complementizer has (e.g. *that* or *which* in English) is unlikely to be provided by an extra-linguistic concept. But we should still expect genuine lexical elements, like nouns, verbs, or adjectives, to have their meanings provided by some association with a concept in our wider conceptual faculties. My claim in this paper is not that there is no such association; rather, I make a claim about what this association is like. I shall argue that associated concepts often provide meanings for words only in conjunction with substantial input from the language faculty proper, and that the results are linguistically modulated meanings. Moreover, the extra-linguistic concepts associated with words are often not the ones we might have intuitively expected.

I should pause to note that however natural the picture of word meaning as provided simply by a mapping of words to concepts might sound, it has been the subject of much critical discussion over the years. It is held explicitly by some philosophers. Fodor certainly holds it (e.g. Fodor, 1975, 1998). Something like it is often a starting point for psychologists (e.g. Clark, 1983; Murphy, 1991, 2002), but typically it is merely a stepping stone for exploring, say, effects of wider cognition on word meaning, or relations between conceptual development and language development. Philosophers of many sorts have challenged the picture in more fundamental ways, by positing some role for language in wider cognition (e.g. Carruthers, 2002; Dennett, 1996; Dummett, 1989). (Some exploration of these issues from the point of view of psychology and language acquisition can be found in Bloom (2000) and Carey (2009).) In light of this, arguing against an overly simplistic picture that has already been subject to extensive criticism is not itself going to be

of great interest. However, the arguments I shall present here will not focus on the role of complex language or grammar in thought, but on the nature of the lexicon. I shall conclude that the structure of the lexicon does allow points where word meanings can be formed by association with concepts; but, I shall argue, not the concepts we might have expected. Grammatical aspects of the lexicon indicate that lexical entries encode concepts in highly linguistically constrained ways, and build a great deal of the content of a word not out of the content of the concept encoded, but by distinctively linguistic means. I shall thus defend a nuanced view which has some affinities with its simplistic cousin, as well as some surprising features. I shall not reject the basic idea of associating concepts with words, as some philosophers have.

With this conclusion in hand, I shall turn to reconsider a long-standing issue surrounding Frege's context principle, about the priority of word meaning and sentence meaning. I shall suggest that the aspects of the lexicon I shall highlight here offer a modified form of the context principle that is immune to some of the common objections to the standard one. Finally, I shall gesture, briefly, to some difficult questions about the relations between language and thought that my conclusions raise.

This paper will proceed as follows. Section I will present the idea that linguistic meaning involves substantial packaging of content in grammatically determined configurations. This section will involve a review of some important ideas in lexical semantics, and the syntax-semantics interface. With those ideas in hand, section II will present the main thesis of this paper, that word meanings lexicalize concepts, but only together with linguistic packaging, leading to different concepts being lexicalized than we might have expected. The final two sections defend this thesis and explore some of its consequences. Section III will further explore the nature and range of concepts available for lexicalization, and the structural elements that package them. Section IV will consider consequences of the thesis for the relation of word to sentence meaning, and Frege's context principle, and it will raise some questions about how the thesis relates to issues of linguistic determinism.

I Roots and Packaging in the Lexicon

Much of the discussion to follow will rely on some ideas about the lexicon, and its relation to grammar. In this section, I shall begin by reviewing some

leading ideas about the lexicon in linguistics. I shall extract from these ideas a picture about how meaning is encoded in the lexicon, which I shall call the *packaging picture*. It is this picture which will lay the groundwork for the discussion of how concepts relate to word meanings in subsequent sections.

As any linguistics textbook will tell you, the lexicon is where the properties of basic elements of language are stored, including syntactic or morphological properties, phonological properties, and semantic properties. We will focus on semantic properties here, and in particular, those semantic properties that interact with grammar. There is not a lot of agreement about just what that part of a lexical entry should look like, but here is an example:

(1) a. open
    b. [[x ACT] CAUSE [BECOME [y $\langle OPEN \rangle$]]]

Here I follow the influential work of Levin and Rappaport Hovav (e.g. Levin and Rappaport Hovav, 1995; Rappaport Hovav and Levin, 1998; Levin and Rappaport Hovav, 2005). This part of the lexical entry provides what is sometimes called a 'lexical-conceptual structure' for the word. It gives the word's meaning, in terms of a decomposition of events it describes into components. The components include a subevent of an act, with an agent x, a subevent of becoming, with a theme y and result state of being open, and a causal relation between the subevents. Of course, this is not the entirety of the lexical entry, as we have not mentioned phonological information, but this is an example of a widely adopted approach to the semantic component of the lexicon.[1]

There are two main components to any lexical entry along these lines. First, there is a *structural frame*. This frame describes a type of event, in terms of specific sub-event types, and via elements like CAUSE and BECOME. The other main component of this sort of lexical entry is a *root*, here represented as $\langle OPEN \rangle$. Roots provide specific, and as we shall explore in a moment, idiosyncratic, aspects of meaning, while the structural

[1]I shall typically follow Levin and Rappaport Hovav in presenting this sort of highly structured lexical entry. However, their view is one (very influential) representative of a wider research project. Other work along related lines includes Bierwisch and Schreuder (1992), Grimshaw (2005), Jackendoff (1990), Pinker (1989), Wunderlich (1997), and many others. Though in broad outlines these authors agree about the presence of structured lexical entries, they differ on many details, and on many wider conceptual issues. Jackendoff, in particular, has been pursuing at length of project with some different conceptual underpinnings than I am advancing here. I shall return to this point in section IV, when some of these conceptual issues have been more clearly articulated.

frame provides the grammatically relevant structure of the word.[2] Indeed, Levin and Rappaport Hovav see these structures as substantially determining a word's syntactic properties, and so, they constitute the interface point between syntax and lexical semantics.

It is the way these sorts of lexical entries interact with grammar, and indeed, how they determine syntactic properties of expressions (mostly verbs) that has been the main focus of attention within linguistics. Perhaps more importantly, it is considerations of grammar that inform the kinds of structural frames that are posited, and how they are composed of structural elements like CAUSE. Much of the linguistic work on the interactions between the lexicon and grammar has focused on so-called argument alternations: various patterns of projection of arguments of verbs, presented as pairs of sentences illustrating a change in the pattern of a verb's arguments. I shall illustrate this idea with a widely discussed example. Following Fillmore (1970), we may observe that verbs like *open* enter into the causative alternation, while verbs like *hit* do not:

(2) a. i. The boy opened the window.
ii. The window opened.
b. i. The boy hit the window.
ii. * The window hit.

We see here that what may have seemed to be similar sorts of verbs display different patterns of argument realization, and hence, different grammatical properties.

The leading idea behind the approach to the lexicon we are considering is that these sorts of grammatical properties are due to different structural frames. We have already seen the frame for a causative verb like *open*. According to Rappaport Hovav and Levin (1998), a verb like *hit* has a very different frame:

(3) a. *open*: [[x ACT] CAUSE [BECOME [y $\langle$*OPEN*$\rangle$]]]
b. *hit*: [x ACT$_{\langle HIT \rangle}$ $\underline{y}$]

The structural frame, according to this sort of theory, encodes the information needed to determine the patterns of argument realization for a verb, and

[2]The term 'root' is used somewhat differently in morphology (cf. Aronoff, 1994). Its usage here is attributed to Pesetsky (1995), and is taken up by Levin and Rappaport Hovav (2005).

so substantially determine the grammatical configurations in which the verb can appear. But the encoding in terms of event structure needs to be supplemented with some rules, called 'linking rules', which show how arguments may be realized on the basis of event structure. For instance, one version of a linking rule approach posits a rule that the immediate cause variable x in the frame of *open* is projected as the external argument, while the variable for the object undergoing a change of state y is projected as a direct internal argument (Levin and Rappaport Hovav, 1995). The intransitive member of the causative alternation is explained by a process in the lexicon which effectively binds the variable for the immediate cause. The result is an unaccusative structure, with no external argument.[3] In the *hit* case, we have only one event-structural position x. The additional variable y is selected by the root, but is not part of the event frame. Hence, this verb shows a much more limited range of intransitive occurrences.

The details of any such theory are quite involved, and will not be crucial to our discussion. What is important is that on this approach, grammatical properties, like argument alternations, are derived from semantic structure in the lexicon. We thus wind up with highly structured entries in the lexicon, which provide semantic content and grammatical properties together, via the combination of a structural frame and a root.[4]

A highly articulated even-structural frame in the lexicon is one way to explain such grammatical facts at those about argument realization. But it is not the only way. Indeed, many researchers have argued that facts about the projection of arguments of verbs should be accounted for within syntax. So, for instance, Harley (2007), building on important work of Hale and Keyser (1993, 2002), offers an analysis of transitive occurrences of *open* like:

[3]This is one of many options which might be explored with respect to these rich lexical structures. Another is that a null argument is projected in the intransitive case. Something along these lines is explored by Chierchia (2004), considering Romance cases where a reflexive morpheme appears to realize the external argument.

[4]We have focused on the event-structural differences between result-state verbs like *open* and manner verbs like *hit*. There may be other semantically motivated classifications that distinguish them. For instance, Levin (1993) classifies *hit* as a 'contact by impact' verb, while *open* is a change of state verb.

(4)

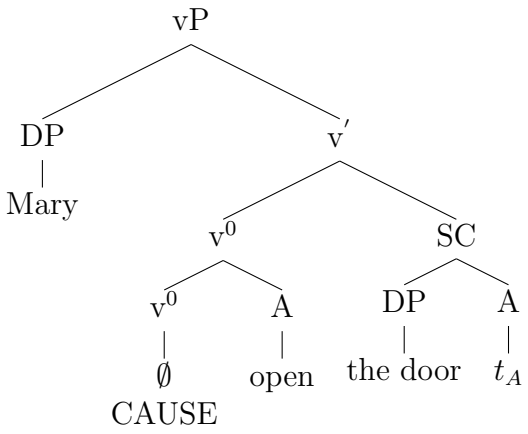

The external argument is selected by a 'causative' v head, and appears as its specifier. The intransitive variant is the result of an 'inchoative v':

(5)

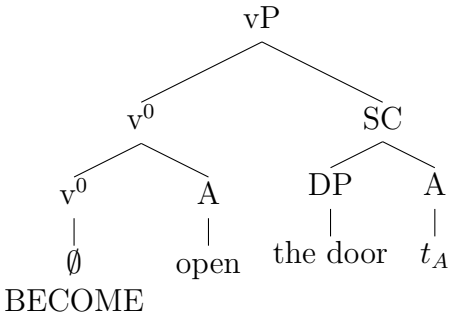

As with Levin and Rappaport Hovav's structured lexical entries, Harley's Hale and Keyser-inspired approach has many details which I shall not explore. For this discussion, it mainly offers us an example of how structure much like we saw a moment ago posited within lexical entries can be posited in syntax instead. We see similar elements, like CAUSE, and articulation of a structure in which they occur together with a root. In this example, the root is the adjective *open*, which combines with a structural frame provided by the syntax of v to both produce the intuitive meaning of the verb *open* and predict its syntactic behavior. Whereas before the structural frame was determined within the lexicon, here it is determined by syntax; but there is a structural frame and a root all the same.[5]

[5]The analysis I briefly sketch here is one instance of a very substantial body of work, notably by Hale and Keyser (1993, 2002), but also important work by Baker (1988), Chomsky (1995), Kratzer (1996), Larson (1988), Marantz (1997), and others. Hale and Keyser, in particular, posit a notion of l-syntax, working within the lexicon but governed

Whether argument projection is determined by a rich lexicon including event decompositions, or by syntactic processes, or something else for that matter, is a substantial empirical issue, on which it is not my goal to take sides. What I take from both the lexical and syntactic approaches to argument projection is the idea of what I shall call *packaging*. Both approaches posit a range of structural elements, like CAUSE, which combine with roots to form the intuitive meanings of verbs. On both approaches, there is a highly limited range of structural elements, determined by grammar. Likewise, on both approaches, the configurations in which the structural elements and roots combine is also determined by grammar (either via a list of structural frames in the lexicon, or by syntax). Thus, each approach exemplifies the idea that structural frames package roots. Important grammatical facts, like argument projection, are substantially explained by structural frames, i.e. by structural elements and their grammatically admissible configurations.

The content of a verb is provided by both its structural frame and its root. As the structural frame is determined substantially by grammar, and is marked by a limited range of possible configurations, it is the root which gives the verb its distinctive content. Such content is often labeled idiosyncratic (e.g. Grimshaw, 2005), as it accounts for the wide range of verbs we find in a language like English, and the variation of the verbal lexicon across languages. Structural frames package roots, and so form verbs with grammatical properties out of idiosyncratic root meanings.

The precise nature of roots varies according to which theory is under consideration. In Levin and Rappaport Hovav's approach, roots are semantic elements within the lexicon, which are typically predicates of individuals or events. In Hale and Keyser's approach (and Harley's), roots are often syntactic elements, and often non-verbal. For instance, they emphasize the role of *adjectival* roots in creating causative verbs. The intuitive verb is created out of such an non-verbal root by a syntactic process, such as the one we illustrated above where an adjectival root incorporates with a higher 'little v'. In many frameworks, notably that of Marantz (1997), roots are fundamentally acategorial, and require some morphosyntactic process to form elements of the standard syntactic categories.

I shall not offer a full defense of the packaging approach to the lexicon, much less try to decide between its syntactic or rich-lexical variants. Indeed, I shall simply take it for granted that the idea of packaging captures an

by the same principles as other syntactic processes.

important trend in linguistics, and explore what consequences it has for the relation of word meaning to concepts. But, I shall pause to mention that the packaging approach has the virtue of offering an elegant explanation of one of the 'big facts' about lexical meaning: its distinctive combination of idiosyncrasy and structure. Lexical categories are in some ways highly idiosyncratic. Just what lexical items a given language has can vary widely, and seems to reflect the specific interests of individuals or groups. Perhaps more importantly, lexical categories are open, and can easily expand as we form new concepts or otherwise find new things we want to express. Just as we can be idiosyncratic in how we think and talk about the world, so to are our lexical categories idiosyncratic in what they contain.

At the same time, we find a great deal of structural uniformity in lexical categories. It is difficult to provide examples which do not involve some theory, but for instance, in spite of all the great variety of verbs, we still see a highly limited range of thematic roles—far more limited than our ability to conceive of types of participants in events.[6] Not only are thematic roles limited, there appears to be some organization to them. Many ideas about thematic hierarchies attempt to capture that organization.[7] Likewise, verbs fall into grammatically important classes, such as aspectual classes (aktionsart). They also fall into other sorts of classes, such as classes (mostly for verbs) like sound emission (*buzz*), contact by impact (*hit*), psych verbs (*frighten, fear*), etc. (cf. Levin, 1993).

Some of this structure might simply reflect structure in our concepts, or even in the world we describe. Perhaps such categories as sound emission might. (Though it is far from clear, at least insofar as not every salient grouping of concepts corresponds to a semantically robust class.) But a great deal of the structure we see seems to involve distinctively linguistic constraints. For instance, the limited number and distinctive organization of thematic roles does not seem to be the result of how we are able to think about events, but rather, an aspect of grammar.

We also see something about grammar at work in variation in lexical items across languages. Here are two examples, drawn from the literature. When we compare the English verb *blush* with its Italian counterpart *arrossire*, we find them to be nearly synonyms. But, they are not identical. They fall into different aspectual classes, and show different telicity properties.

[6]This point is widely discussed, but is the starting point of Hale and Keyser (1993).

[7]For an overview and references, see Levin and Rappaport Hovav (2005).

Blush an activity verb, and is atelic, while *Arrossire* is an achievement verb, and is telic. *Arrossire* means something more like 'become red in cheeks' (Levin and Rappaport Hovav, 1995; McClure, 1990). There is something in common across these two verbs, but they differ in grammatical properties. The packaging view accounts for this nicely, by giving these verbs the same root, but different structural frames; one appropriate for an activity verb, and one appropriate for an achievement verb.

Another example is drawn from Lhasa Tibetan, where the counterpart of *hit* takes a locative marker on the argument for the contacted object (Levin and Rappaport Hovav 2005, citing DeLancey MS):

(6) shing*(-la) sta=re-s gzhus-pa
tree-LOC axe-ERG hit
Hit the tree with an axe.

Again, there is presumably some core content in common between the English and Tibetan verbs, but they differ in grammatical structure, and perhaps more. Again, appealing to common root content with different structural frames offers a way to capture this.

The packaging approach, in any of its forms, helps to explain both the structure and the idiosyncrasy we see in lexical categories. Grammatically significant structure arises from the structural frames that package contents. Variation in linguistic structure, as between English and Italian or Tibetan, is explained by variation in which structural frames a given root appears in. Idiosyncrasy arises from the roots, which can vary idiosyncratically, as our thoughts about the world can. Presumably some classes of verbs arise because of natural groupings of roots, though we would expect few if any grammatical effects from similarities between roots alone.[8]

II Roots, Concepts, and Lexicalization

We have now seen a pair of approaches to the lexicon. Both embody, in different ways, the idea of packaging. Root elements are packaged by distinctively linguistic structure, either within lexical entries themselves, or by

[8]Grimshaw (2005) takes the position that roots, or idiosyncratic contents, are linguistically atomic, and so would have no role in grammar. In contrast, Rappaport Hovav and Levin (1998) posit a least a limited role for roots, as they can determine some arguments of transitive verbs.

syntax that builds up structures from roots. The packaging picture provides for a specific division of labor between structural and root elements, which in turn is reflected in the combination of idiosyncrasy and structure lexical categories display.

I shall assume the packaging picture, in some form or another, as we proceed. I shall use it to explore how the lexicon relates to our wider conceptual abilities. In this section, I shall show that it offers us a way to see word meanings as lexicalizing concepts. It does so via a model where roots mark interface points between language proper and concepts. In effect, roots are concepts, which get packaged by language into the meanings of lexical items. But I shall also show that the way roots are packaged places significant constraints on which concepts can be lexicalized to which sorts lexical items. The result, I shall argue, is that in many cases linguistic meaning is not built up from the most naturally associated concept for a word; but rather, by distinctively linguistic means, from a less obvious choice of concept to lexicalize. In subsequent sections, I shall further refine and defend my conclusion, and explore some of its consequences.

Before delving into these argument, we should pause to review some assumptions about concepts, and about the organization of cognition. Following the common practice, I take concepts to be metal representations. I shall avoid most commitments about what sort of structure, if any, concepts have.[9] But I shall assume they are part of our general cognitive repertoire, and typically not restricted to special-purpose cognitive domains like vision or language. I am here assuming there are such special-purpose cognitive domains. It will be useful, though not absolutely essential, to assume a Chomskian view, and assume there is a language faculty, among other modules of the mind. There is also a 'general intelligence' or non-special-purpose component to the mind.[10]

The assumption that concepts are not generally part of the language faculty seems to be well-motivated. For one thing, concepts do appear active in wider cognitive processes that are generally accepted not to be part of the

[9]For overviews of ideas about concepts, see Murphy (2002) or the papers in Margolis and Laurence (1999).

[10]Chomsky's views about the language faculty extend through much of his work. For a recent, philosophically oriented discussion, see Chomsky (2000). The classic on modularity more generally is Fodor (1983). As has often been remarked, there are differences between Fodor's and Chomsky's notions, though for the level of detail we will adopt here, they will not be significant. See Higginbotham (1987) for some discussion.

language faculty, such as metaphor comprehension. It is also widely accepted that some animals have concepts (though perhaps not just like ours), as do preverbal infants.[11]

Though concepts typically are not specific to the language faculty, words and concepts do relate. Our grasp of concepts and our grasp of word meanings tend to go together, and it is natural to assume that our words express our concepts. Thus, in many cases, concepts seem to be the best candidates for word meanings. At least, we would be surprised if concepts played no role in the semantic portion of the lexicon.

It is very natural to think that concepts play a basic role in language acquisition as well. It is well established that very young, preverbal children have at least some concepts. At least, concepts of some sorts, thought it is also well established that in some cases concepts change as children develop (cf. Carey, 1985).[12] A natural picture of the acquisition of word meaning is that a child has to associate a word with a concept, and thereby be able to express that concept. As Clark (1983, p. 797) writes, "They have to work out which device—which word or expression—maps onto which concept before they can use particular devices to draw upon an instance of a particular concept for others." As psychologists have often noted, this task is strikingly difficult, as the range of concepts is large and the evidence the child has to work with appears rather small.[13]

This idea has been challenged in a number of ways. At the extreme end, we find neo-Whorfian views that hold that language structures and in some ways determines our concepts. Hence, they conclude, we cannot see concepts as there prior to language development.[14] But there are also, certainly, many

[11]See Gallistel (1990) on animal concepts and Carey (2009) on infant concepts, and the many sources they cite. For a survey of metaphor cognition, see Camp (2001). I have discussed issues of metaphor and its relation to the language faculty in my (2008).

[12]Just which concepts change, when, and how much, is a matter of intense debate. See the many references in Murphy (2002).

[13]Clark herself is most concerned with the wide range of concepts, and the limited ways our words express them. In anticipation of conclusions I shall argue for as we proceed, I should mention that she begins her paper by saying "Meanings—that is, word, phrase, and sentence meanings—and concepts are as different as apples and oranges. Although they are often discussed as if they were equivalent, the distinctions between them are crucial—not only in talking about development but also in considering the relation between language and thought" (Clark, 1983, p. 787). I shall argue for one version of this very point here. For further discussion of the role of concepts in language acquisition, see Bloom (2000), among many other sources.

[14]See the papers in Bowerman and Levinson (2001) and Gentner and Goldin-Meadow

ways in which our language and conceptual abilities develop together. Especially when we look at the functional elements in language, questions about whether they affect wider cognitive development are complex.[15] Even so, the idea that when children are learning lexical items in the 'naming explosion' of about 18 months, or when adults learn words, they are somehow mapping them to concepts, is an inviting one.

When we consider the relation between words and concepts with the packaging view of the lexicon in mind, a more nuanced picture than the simple association of lexical items with concepts emerges. The obvious conclusion from the packaging view is that *roots* map directly onto concepts. Whole lexical meanings (or syntactic constructions out of roots) need not, and indeed, we see that such full meanings or construction are composed of a number of language-specific elements, over and above their roots.

There are several reasons that can be cited in support of this claim. First, roots, under any of the views of the lexicon we reviewed in section I, have properties of concepts. Most importantly, roots do precisely the kind of categorizing job that concepts do. Each root functions as a predicate, typically monadic.

This claim can be refined a little. Roots come typed, and each type is a type of predicate. Different versions of the packaging view type roots somewhat differently (and sometimes in cross-cutting ways). We have already seen an instance of the typing of roots. Recall that Levin and Rappaport Hovav's lexical entries for *break* and *hit* involve very different structural frames. In fact, the frames themselves specify which types of roots can appear in them:

(7) a. $[[x \text{ ACT}] \text{ CAUSE } [\text{BECOME } [y \langle \mathit{STATE} \rangle]]]$.
b. $[x \text{ ACT}_{\langle \mathit{MANNER} \rangle}]$.

STATE roots are (stative) predicates of individuals, while MANNER roots are predicates of events. Levin and Rappaport Hovav also rely on types THING and PLACE. In the more syntactic approach of Hale and Keyser, the main roots types are syntactic. Many roots that Levin and Rappaport Hovav count as STATES are adjectives, which are also predicates of individuals. Similarly we see the verb type of root in many cases where Levin and Rappaport Hovav would have MANNER roots. (The two systems of root

(2003), and critical discussions by Carey (2009) and Pinker (1994b).

[15] See again the discussion of Carey (2009), who comes to the conclusion for the cases of quantifier and number systems and kind sortals that language acquisition has only weak effects on conceptual development.

typing do cross-cut each-other in some ways. Both nouns and verbs wind up sometimes behaving as predicates of events according to Hale and Keyser, and they gloss the root of *break* as nominal rather than adjectival.) Each root is an element of some type. Other ways of classifying verbs provide other schemes of typing root elements.[16]

This kind of superordinate categorization we see with roots is one of the pervasive features of concepts. It has often been observed that concepts are arranged hierarchically, and this structure is in many ways important to the way concepts behave in cognition.[17] Roots are typed at a relatively high superordinate level in the theories we are here examining, akin to the 'ontological' level discussed by Jackendoff (1983) and Keil (1979), but other classifications are possible, at different levels. Just like concepts, roots categorize, and fall into hierarchical arrangements of categories that provide typing.

The type of a root is very important to the linguistic properties of the lexical entry it creates. In many cases, the type of the root determines what structural frames it can appear in, which in turn determines the grammatical properties of the lexical entry. Just how this happens differs between the two frameworks we examined. For instance, in Levin and Rappaport Hovav's approach, frames specify types of roots, as in $[y\ \langle STATE \rangle]$. For Hale and Keyser, syntax imposes restrictions on the configurations in which roots appear, but categorial properties of root types (selectional properties) determine which syntactic configurations roots appear in.

So, we see that roots categorize and fall into hierarchical arrangements of categories, just as concepts do. Grammar, in some way or another, sees certain linguistically fixed superordinate properties of roots. This contributes to our theories of the lexicon, which help explain the structure in lexical categories. But also, allowing roots to indicate concepts helps to explain the variety and idiosyncrasy we see in the lexicon. Indeed, it offers two factors that can explain it. One is simply the variety and idiosyncrasy of our conceptual repertoire. We saw idiosyncrasy in the lexicon corresponding to idiosyncrasy in our ways of thinking about the world, which is readily explained if roots are concepts. Moreover, as has often been emphasized, we have more concepts than words (e.g. Clark, 1983). Thus, we can find

[16] In the distributed morphology framework of Marantz (1997), roots are fundamentally acategorial, but are in effect typed and built up to grammatical categories by further syntax.

[17] See e.g. Rosch (1978), and the many references in Murphy (2002).

variation in the lexicons of different languages simply from the variety of concepts to which we can map roots.

The further fact that in spite of this variety, we see limited and stable grammatical properties of lexical items is explained by the division of lexical entries into root and structural components. Roots, as we have observed, are more or less linguistically atomic, while limited structural frames, keyed mainly to the types of roots, explain most of the grammatical aspects of a lexical entry.

The model that captures all of this is one that treats roots as functioning like pointers to concepts, conceived of as mental representations residing outside the language faculty. The language faculty produces structural frames in the lexicon (or in syntax), and these contain root types. Filling in a root of a given root type involves, in effect, establishing a *pointer* to a concept, outside the language faculty. As it is outside the language faculty proper, we should expect a root to be linguistically atomic, except insofar as the lexicon already encodes aspects of its structure, like the type of the concept to which it points (and perhaps some features of its adicity). According to this view, roots thus mark points of interface between the language faculty and the wider cognitive makeup of a person.

This interface model captures nicely the idea that our concepts form the core of word meaning, while allowing for complexity in the word-to-concept relation. For instance, it nicely captures the intuition that our words do not fully express everything about a concept, by directing us to one particular type of concept within a family of concepts we might have. At the same time, it also captures why word meanings can show features of concepts, like typicality effects (see Murphy, 1991). Finally, it has an effect akin to 'syntactic bootstrapping' (e.g. Fisher *et al.*, 1994; Gleitman, 1990), of constraining the space of concepts that can be mapped to a given word. Once the word-leaner has identified a structural frame, it is constrained what type of concept can be mapped to the root.[18]

There is one further, and striking, fact about the concepts to which roots point. In most cases, they are *monadic*. We see this in two ways. For Levin and Rappaport Hovav, roots are typically monadic elements of types like STATE, THING, or MANNER. For syntactic approaches like Hale and Keyser's, they are typically syntactic elements like nouns or adjectives, which

[18]A number of other cognitive and developmental aspects of a (somewhat different) structural frame approach are discussed by Pinker (1989).

select one argument. Either way, roots are typically monadic.

This is not an accidental feature of the views of the lexicon we have considered. Both structure complex lexical entries via structural frames, but the roots themselves are relatively simple and linguistically more or less atomic. They thus have no interesting argument-structure properties of their own, and merely perform the basic categorization role of concepts.

We should be cautious about overstating this fact. Both of our example approaches to the lexicon do have some polyadic roots. For Levin and Rappaport Hovav, some roots, particular manner roots like $\langle SWEEP \rangle$, can have additional argument positions. These do not correspond to positions in a structural frame, and thus are expected to show only a highly limited range of linguistic properties. Similarly, Hale and Keyser hold that prepositional roots are relational, and are in fact the only relational roots for them.

To illustrate the way monadic roots work in structural frames, let us look at one case more closely: that of causative verbs. We have already seen that these display interesting alternations that need to be accounted for by any theory of argument realization. Recall that for Levin and Rappaport Hovav, causative verbs are given an analysis in terms of result-state roots and a structural frame involving a CAUSE element (repeating 3a):

(8) a. *open*: [[x ACT] CAUSE [BECOME [y $\langle OPEN \rangle$]]]
b. *break*: [[x ACT] CAUSE [BECOME [y $\langle BROKEN \rangle$]]]

As we saw before, this sort of frame contrasts sharply with the frame for the non-alternating *hit* (repeating 3b):

(9) *hit*: [x $\text{ACT}_{\langle HIT \rangle}$ $\underline{y}$]

I shall not reproduce the corresponding syntactic analysis of causative verbs presented in (4), but it has the same elements: a CAUSE element occupying v, and a root, typically adjectival.[19] The root concepts for verbs like these are monadic *result states*, which is combined with structural elements like CAUSE to build up relational and causal meanings. Both the causal and agentive aspects of a verb's meaning (that an agent or something playing an agent role causes the result state) are contributed by the structural frame, not the root concept itself.

[19]Both Hale and Keyser (2002) and Harley (2007) emphasize the deadjectival nature of many causative verbs, but Hale and Keyser see *break* in particular as denominal. Regardless, the root is monadic.

What is surprising here is that it would be natural to assume that we have perfectly ordinary concepts OPEN and BREAK. And, as we discussed a moment ago, we might expect the verbs *open* and *break* simply to point to those concepts. But, at least if the views of the lexicon we are assuming are on the right track, that is fundamentally not what happens. Rather, it is monadic, result-state concepts that are packaged to form the meanings of these verbs; not the ordinary, polyadic concepts OPEN and BREAK. As we noted, the concepts lexicalized as roots are different from the ordinary ones in more ways than just adicity. Our ordinary concepts involve causality and agency, while the result states do not. The meanings of the verbs certainly also reflect causality and agency, but this is contributed by the structural frame, not the roots.[20]

This is enough to show that the appealing picture that words simply associate with concepts which are their meanings is misleading, if the packaging view of the lexicon is on the right track. Words do still map to concepts, via their roots. But roots do not always point to the concepts we might have expected, and can point to concepts with very different features than we intuitively expect, or see in the resulting packaged meaning. Lexicalization of a concept is, according to the packaging view of the lexicon, a much more complex process than merely mapping words to concepts. A type of root is fixed, and then a root of that type is selected. Often it cannot be the concept we would intuitively associate with the verb. The root is then packaged, and often this packaging reconstructs features of the intuitive concept, via a linguistically determined structural frame.[21]

III Concepts in and out of Language

We have reached my main point: at least if we assume the packaging picture of the lexicon, we find reason to think that concepts are lexicalized to form

[20] A similar point about monadic concepts and lexicalization is argued for by Pietroski (2010, forthcoming). Though we approach these matters somewhat differently and wind up with views that differ in some important respects, his position and mine have a great deal in common. This is no accident, as his work has been very influential on my thinking about the issues being discussed here.

[21] I thus take sides against Fodor (e.g. Fodor, 1998; Fodor and Lepore, 1999), though obviously Fodor would reject utterly the packaging picture of the lexicon I have assumed here. There have been a number of critical discussions of Fodor's views, including Collins (2011) and Johnson (2004).

word meanings, but not the concepts we would have expected. Rather than simply lexicalizing an ordinary concept by having a word associate with it, lexicalization is a complex process which selects root concepts of particular types and packages them. Often the root concepts do not themselves have important features of the ordinary concepts, and of the corresponding word meanings. These features are in effect added back by the structural frame in which a root concept appears.

This picture raises a number of questions. In this section, I shall consider a family of them bearing on the nature of the range of concepts available for lexicalization, and the structural elements that package them. In the following section, I shall consider some philosophical and conceptual issues the packaging view raises.

In this section, I shall discuss four issues. First, I shall ask how sure we are that we have ordinary polyadic concepts like BREAK. Second, I shall ask why we should think the structural elements are properly part of the language faculty. Third, I shall ask why language might wind up lexicalizing concepts in the surprising way I have suggested. Finally, I shall pause to show that one easy way out, which avoids the strong conclusion I have advanced, cannot be taken.

III.1 Ordinary Concepts and Polyadicity

The first issue I shall discuss is one about the nature of our (non-linguistic) conceptual repertoire. I assumed above that we have ordinary concepts like OPEN and BREAK, and that they are relational, imply causation, and imply there was an agent (in some suitably abstract sense) who did the causing. Most importantly, I have assumed we have such concepts independently of language. Is this really so?

To a certain extent, it may just seem obvious that we have such concepts. But, the conclusion of the previous discussion is surprising because we are inclined to assume that we have such concepts independently of our language faculty. Ordinary concepts like OPEN are presumed to be in our extra-linguistic cognitive repertoire, there for the lexicalizing; but the language faculty instead opts to lexicalize result-state roots for causative verbs. It is a substantial claim that we have such polyadic (causal, agentive) concepts independently of our language faculty. It still may seem obvious, but it would

be nice to be able to offer some clear evidence for the claim.[22]

Unfortunately, empirical support for this claim is not so easy to come by. The problem is that the most natural ways to probe for what concepts we have go through our linguistic abilities, and so, it is not trivial to determine whether we have concepts entirely independently of our linguistic resources. Even so, there is some empirical evidence which lends some credence to the claim. Specifically, there is evidence that non-verbal animals, pre-verbal infants, and children at the early stages of language development seem to possess some relational concepts. For instance, studies of reorientation in rats suggest they rely on geometric relations (Chen, 1986; Gallistel, 1990). Likewise for children at 1.5–2 years (Spelke, 2002). There is evidence that infants at 6 months represent some actions as goal-directed, which involves both relational representations and agency (e.g. Carey, 2009)

Another area where there is intriguing and suggestive evidence is in the study of the representation of cause. A long tradition, starting with work of Michotte (1963) on the perception of causality in adults, and later extended to studies of infants, indicates that infants at six or perhaps even four months have a concept of cause. Importantly for out concerns, this appears to be a genuinely relational concept, and even infants appear to be sensitive to its argument structure of agent and patient. Though this will raise a number of questions about the status of the element CAUSE in the lexicon, it does indicate that preverbal infants have relational concepts. Thus, as I think seemed intuitively plausible, we should not generally expect relational concepts only to arise as a result of language.[23]

None of these considerations offer evidence that we have concepts like BREAK or OPEN independently of language. But they do generally indicate that we have relational concepts independently of language. Thus, they at least add a little bit of hard data to the intuitive idea that such concepts are there to be lexicalized. If that is so, and I stress that it is conjecture, then it is indeed striking that the language faculty often selects result state concepts instead as it builds causative verb meanings.

[22]The claim is complicated, as has often been noted, by the fact that it is not always transparent which concepts are polyadic. As Pietroski (forthcoming) notes, we might view the concept TRIANGLE as monadic (x is a triangle) or triadic (lines x, y, z form a triangle).

[23]For overview and discussion of the large literature on these issues, see Carey (2009), Saxe and Carey (2006), and Scholl and Tremoulet (2000).

III.2 The Status of Structural Elements

In discussing linguistic packaging, I have emphasized the important linguistic role of a small collection of structural elements, which package roots. According to the packaging view, these elements, unlike roots themselves, are not linguistically atomic, as they are fundamental to explaining the grammatical properties of lexical items. But the full force of my conclusions of the last section come with the assumption that they are part of the language faculty proper, and not, say, part of a rich extra-linguistic conceptual system.

The main argument for this is that their primary role is grammatical: grammar determines their function and organization. But some reflections on the meanings of structural elements adds further support this conclusion. Let us again focus on the element CAUSE. It has frequently been argued that it does not have exactly the same content as our word *cause*, and perhaps, not our intuitive concept of causation either.

First, there is a much-discussed observation that CAUSE is more restricted than our concept of causation (e.g. Dowty, 1979; Parsons, 1990; Pietroski, 2005). Consider the following pair:

(10) a. I caused the window to become broken, by hiring a kid to throw a brick through it.

b. # I broke the window, by hiring a kid to throw a brick through it.

The second is unacceptable (as it is obviously false). This shows that CAUSE seems to be closer to some notion of 'direct' causation than our ordinary talk about cause.[24]

In other ways, the content of CAUSE seems wider than that of causation. Consider:

(11) a. Dissension toppled the government.

b. Time abated the damage.

In both these cases, we have a relation that is more abstract and inclusive than canonical instances of causation. At least, most philosophical analyses of causation would suggest that neither dissension nor time enters into causal relations. This is not conclusive, as it is not clear just how such philosophical

[24] Of course, Fodor (1998) sees this as undermining the use of elements like CAUSE in lexical entries, whereas I take it as reason to think they have a different role and source in our cognition than an ordinary concepts.

considerations might relate to our intuitive concept of cause. Our use of overt expressions like *cause* does not clearly settle the matter, but it does sit awkwardly with these sorts of cases:

(12) a. ? Dissension caused the government to become toppled.
b. ?? Time caused the damage to become abated.

The second of these strikes me as clearly anomalous, because time cannot cause. The first one is only somewhat marginal to my ear. Even so, I think we should cautiously suppose that CAUSE is more inclusive than our concept of causation. We also see a very inclusive notion of agency at work in these cases, as time and dissension are not strictly speaking agents at all.[25]

The element CAUSE helps to explain commonalities in grammar and in content across a very wide range of causative verbs, including *abate, decrease, expand, grow, soak, topple*, as well as *drop, break, bend*, etc. (cf. Levin, 1993). It plays a substantial role in the projection of arguments in syntax. At the same time, it appears to differ in some respects from our notion of causation. It is more restricted in some ways, and more permissive in others. Though its content is abstract and hard to state simply (it is something akin to bringing about a result state in the right way), this content appears to be common across a wide range of verbs. As CAUSE plays a fundamental role in grammar and has a content distinct from, and more abstract than, our ordinary concept of cause, we get some explanatory mileage from concluding that it is part of language proper, not identical to any extra-linguistic concept of causation we might have.

The case of CAUSE suggests we should see structural elements as having distinctive content, and we have frequently observed that they have specific linguistic effects. It thus makes sense to suppose they are part of the language faculty proper, and not simply drawn from our wider cognitive repertoire, the way roots are. We thus have the surprising picture that the meaning of a causative verb is not an ordinary relational concept from our general cognitive repertoire, but one built up from a monadic result-state concept by distinctively linguistic means.

[25] We should ask if some of these involve metaphorical or otherwise figurative uses of causal expressions. I do not see any clear evidence they are metaphorical; at least, they do not show the open-ended quality of many living metaphors. Moreover, the very wide distribution of sentences like this with many different causative verbs would make the claim of metaphor very strong.

This does not mean that elements like CAUSE can have no connection to our wider cognitive faculties; rather, it only means that whatever mechanism language acquisition employs does not simply import a fully formed concept of causation the way it imports roots. As I mentioned above, there is strong evidence that some concept of causality is formed very early in childhood, and though it is highly controversial, it is even suggested that it may be modularized. There are thus plenty of causal resources for the language faculty to draw upon in creating CAUSE. It is possible that the distinctive features of CAUSE are features of an early concept of causality, or of a modularized system that is not identical to our mature concept. Or it is possible that language acquisition modifies a causal concept to produce a more general, structural element which can figure into a wide range of verb meanings.

III.3 Why Monadicity?

Next, we should pause to consider why language might consistently lexicalize monadic roots rather than polyadic concepts they relate to. I do not have anything like a full answer to this question, and indeed, I am not sure there always is an answer. (Sometimes language just does surprising things!) But there are some facts about language acquisition that lead to a natural speculation.

The facts involve the difficult task children face in learning verbs. As has been much-discussed in the literature, there are many possible verb meanings children could associate with events they observe. To add to the difficulty, there is often a delay between events and use of verbs to describe them, so children's linguistic and observational inputs are not easily correlated (cf. Fisher *et al.*, 1994; Gleitman, 1990; Pinker, 1989, 1994a). Though it is controversial, it is frequently suggested that there is a bias towards nouns over verbs in early vocabulary.[26]

These facts invite the speculation that simplifying the kinds of roots that can be lexicalized might help the child in their difficult task. I already observed in section II that the typing of roots in a structural frame might simplify the language-leaner's task, by a kind of bootstrapping. But there is a further possibility, that narrowing the space of concepts to be lexicalized to monadic ones could have a much more significant simplifying effect, as

[26]See Poulin-Dubois and Graham (2007) for an overview of some of these issues.

discussed by Pietroski (forthcoming) and Hurford (2007). It is well-know from logic that monadic systems are usually significantly less complex than polyadic ones, and it is tempting to think the language faculty might exploit some similar simplifications.

I offer these as speculation. But they do remind us that there could be good reasons from language acquisition for the distinctive combination of monadic roots and packaging we see with causative verbs.

III.4 An Easy Way Out?

Finally, we should pause to rule out one easy way to get monadic roots, which might mitigate the surprise of finding monadic roots for polyadic concepts. It would be easy to get monadic roots from relational concepts, if all that happened was that in lexicalization, one of the arguments of a relational concept was existentially bound. This would allow, for instance, that in the meaning of *break* it is the intuitive concept BREAK that gets lexicalized, rather than the result state concept, but in the process, it gets turned into $\exists x\text{BREAK}(x, y) = \langle BROKEN \rangle$.

There are several reasons this does not seem to be what happens. One is simply that the root of a causative verb is typically a result state, while $\exists x\text{BREAK}(x, y)$ is not stative. Moreover, it predicts the wrong meaning for *break*. It would give it the meaning of causing there to be someone who in turn causes something to be broken. But this is not what *break* means. To make this vivid, consider someone who is contributing to the delinquency of a minor who goes on to break a window. According to the existential analysis, it would be true to say he breaks a window himself, but he does not.

Moreover, if the idea is that existential binding of arguments is available in lexicalization, it threatens to over-generate. If we can existentially quantify out either argument of a relational concept in lexicalization of the concept, we will find verbs with roots like $\exists y\text{BREAK}(x, y) = \langle BREAKER \rangle$. We would then be able to form a causative construction like:

(13) [[x ACT] CAUSE [BECOME [y $\langle BREAKER \rangle$]]]

This seems unlikely as a verb, as it requires an agentive object (and not an experiencer object).

Thus, the easy way out of lexicalizing a relational concept via a process of existentially binding an argument away does not seem to be available.

Rather, it appears that it is the result state concept itself which is lexicalized. It appears we really do lexicalize a monadic result state rather than an intuitive relational causative concept in cases like *break*, and that lexicalization packages such root concepts with distinctively linguistic elements to reconstruct the meaning of a word like *break*.

IV Meaning and Concepts

I have now articulated and defended my main thesis. To conclude, in this section, I shall consider some consequences of this thesis for our understanding of meaning and the relation of mind to language. First, I shall suggest that our grasp of meaning is in fact grasp of fairly complex linguistic structure, and consider what this tells us about the role of grammar in our understanding of meaning. Second, I shall return to the question of how the meanings of our words relate to our concepts, and to questions of linguistic determinism.

IV.1 Grasp of Meaning

Much of the motivation for the packaging view of the lexicon comes from considerations of grammar. But we should pause to ask how it relates to a core philosophical idea about meaning. Meaning, according to a long tradition, is closely tied to what we understand when we understand a word or sentence.[27]

Returning to the case of causative verbs, it is clear that what we grasp cannot simply be the root. We clearly recognize in our grasp of meaning the causal and agentive aspects of causative verbs. *Open* and *break* do not have meanings simply providing result states, but indicate that an act of opening or breaking was carried out.

If the packaging picture is right, then the place in the lexicon we find these aspects of verb meaning is the structural frame, and especially, elements like CAUSE. As we explored above, these elements do have contents, but often abstract and distinctively linguistic contents. Our grasp of the meaning of a verb is grasp of the content provided by the whole lexical entry, including structural and root elements, and how they are configured.

[27]This idea may well be attributed to Frege (1892), and is certainly important to Dummett (e.g. Dummett, 1991), and to Higginbotham (e.g. Higginbotham, 1989).

We see something similar in other cases. Consider an intransitive (and apparently unaccusative) occurrence of a verb like *open*:

(14) The door opened.

Our grasp of the meaning of the verb here includes a sense of change of state, witnessed by our willingness to infer that the door was previously not open. This again is not captured by the root concept $\langle OPEN \rangle$, which is purely the state of being open. Again, it is provided by an element of the structural frame, in this case, something like BECOME.

It follows that if the packaging picture is right, grasp of the meaning of a verb is grasp of a highly structured and distinctively linguistic object. It might be grasp of a highly structured lexical entry. According to more syntactic approaches, it is grasp of a syntactic structure and its terminal nodes. In this respect, grasp of many words turns out to be essentially like grasping words with overt morphological complexity. Grasp of intransitive *open* becomes like our grasp of *-en* verbs like *redden*:

(15)

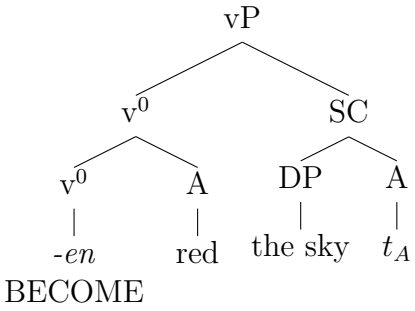

It is not remarkable that we grasp such structures, as we clearly grasp the meanings of phrases and sentences. But according to the packaging picture, virtually all of our grasp of word meaning is like this.

With this in mind, let us return to a long-standing issue in philosophy of language: the status of Frege's context principle. Recall that Frege instructs us that "... it is only in the context of a sentence that words have any meaning ..."(Frege, 1884, p. 73). Many philosophers have found this not only plausible, but important. On the other hand, many have rejected it. Especially, philosophers of mind thinking about the role of concepts in our grasp of words are inclined to reject it, as are many philosophers motivated by considerations of compositionality. How could we understand a sentence compositionally if we have no independent grasp of the basic building blocks of sentences? And surely our words provide just such building blocks.

Interestingly, the packaging view we have explored here provides a somewhat modified version of Frege's principle, that I believe is immune to these objections. In place of Frege's principle, let us propose:

(16) **The modified context principle:** Only in the context of a grammar (language faculty, etc.) does a word have a meaning.

The principle reminds us that according to the packaging view, grasp of a word's meaning is in fact grasp of something grammatically complex (but not like a sentence). The complex lexical entries or pieces of syntax that build meanings only exist in virtue of the grammar of natural language, that determines which structural elements are available and how they combine. One cannot get the word meanings we grasp without a grammar.

The modified context principle is immune to some standard objections to Frege's original version. Phrase and sentence meanings are allowed to be built up compositionally from independent elements, and indeed, we can now see the compositional construction of meaning in the lexicon as well. Speakers can grasp individual lexical items independently of any particular sentences or phrases they might appear in (though not independently of the grammar which builds up such sentences and phrases). So, many of the happy consequences of compositionality are preserved, even though there remains a strong sense in which meaning is determined by wider grammar. Finally, there remains a role for our concepts in informing our word meanings. It may not be the one originally supposed by opponents of the context principle, but concepts independently of grammar still provide idiosyncratic content via roots. These form the cores of word meaning. But the modified context principle holds that even so, we do not have a genuine word meaning outside of a grammar, which determines how roots may be packaged.

Before concluding this subsection, I shall pause to mention one difficult issue for the packaging view I have advocated, pertaining to compositionality; though, I shall mention it only to put it aside. I assume that compositional semantics is truth conditional semantics, and so, lexical entries must provide appropriate contributions to truth conditions. This means roots will have to contribute appropriate contributions to truth conditions as well, and in many cases, will have to contribute extensions (or properties, intensions, etc). The difficult issue is that a great deal of thinking in psychology and philosophy of mind makes the question of whether concepts have extensions complex (at best). For instance, it is not a simple matter to map a prototype to an extension, if it can be mapped at all. Perhaps a Putnam-esque externalism (e.g.

Putnam, 1970) could guarantee extension are available, in some cases? Perhaps in addition to selecting monadic roots, lexicalization might select roots which can provide extensions? At the same time, we observed in section II that words can show typicality effects and other non-extension-based aspects of concepts. So, whatever the mechanism by which concepts are lexicalized as roots, it will have to provide extensions, but not ignore other aspects of concepts. Just how this can be achieved is difficult question, but one the view of lexicalization I am defending makes important.

IV.2 Language and Concepts Revisited

So far, I have assumed, and tried to offer at least a little evidence for the claim, that we have ordinary concepts like BREAK and OPEN. Against this background, I argued that they do not provide the roots for the words *break* and *open*, and so the simple relation we might expect between words and corresponding concepts is not found.

Instead, we find a kind of language-specific reconstruction of causal concepts like BREAK in the lexicon. It is language-specific, as it packages a result state root by distinctively linguistic means, with distinctively linguistic elements like CAUSE in language-specified configurations. Thus, we observed, it is only in the context of a grammar that a word has a meaning.

But as least as we have set things up so far, this means that in many cases we have two distinct concepts where we would have expected one. For instance, we have both:

(17) a. Ordinary concepts like: BREAK.
b. Complex lexical entries like: [[x ACT] CAUSE [BECOME [y $\langle$*BROKEN*$\rangle$]]].

(Of course, we could have a more syntactic version instead.) It is not immediately obvious that we should call the lexical entry a concept, but I believe we should. It certainly seems to be a representation we have available in cognition, that functions like a concept. We can categorize and reason in words, presumably on the basis of their meanings. So, I shall continue to call a word's meaning a separate concept. The two sorts of concepts run in parallel. One is the linguistic meaning of a word, and the other an associated non-linguistic mental representation.

Is the conclusion that there are two separate parallel tracks of concepts reasonable? Very tentatively, I suggest that it is. To see why, let us return to the observation from Clark (1983) that the language learner has a very wide

range of concept with which they might associate a word. If we assume, as I have, that we have (at the appropriate stage of cognitive development) a rich repertoire of ordinary concepts, then the range will have many reasonable contenders to be the meaning of any word. But the range comes not merely from families of related concepts, like BREAK, DAMAGE, HURT, etc. We also have at our conceptual disposal options that would describe an event as one where an object is caused to become broken, versus options that would describe an event as one where an act is taking place which has the manner of breaking (say, a motion towards an object with sufficient force, hitting the object, etc.). We do not attend regularly to such variation in concepts, but I take it to be clear enough that we are able to think in these various ways.[28]

If that is right, then some fairly subtle selecting of concepts must be done in lexicalization. Not only are such concepts as manner and result state concepts available, but language makes active use of just these sorts of distinctions (or correlate distinctions in other frameworks). We see how language does this in the different grammar of *hitting*, which is lexicalized as a manner verb, and *breaking*, which is lexicalized as a causative. But as such distinctions are subtle and hard to recognize in ordinary thinking, it would not be surprising that the language faculty itself should do the lion's share of the work in selecting between these sorts of concepts, rather than leaving it up to the broader cognitive abilities of the agent. If so, then we would precisely expect much of the structure of kinds of events to be represented in the language faculty proper, leaving only pointers to the appropriate concepts to function as roots. This is just what the packaging picture provides. In this light, the presence of 'linguistically reconstructed' concepts in parallel with ordinary ones does not seem particularly surprising. It rather seems to be a side-effect of the complexities we find in lexicalization.

According to this picture, we have two parallel concepts only insofar as we have a linguistically specific reconstruction of one of many concepts we have at our disposal. If anything, this amounts to exploiting the resources of language to draw distinctions that we are often not consciously aware of. Our lack of awareness is witnessed by the natural assumption that both *hit* and *break* express the same kinds of concepts. Attention to language proves that is not so, but we often miss such distinctions in introspection of our concepts. This reminds us that differentiation and selection of concepts is not always

[28]Again, it would be nice to be able to cite some solid empirical evidence here, but I shall leave the claim as one of what seems plausible.

something that is transparent to us upon easy introspection. If it seems odd that we have a linguistically specific reconstruction of e.g. BREAK, it may be because we are not attending to the many variants of breaking we can think about, and the specificity of the meaning of the word *break*.

It is easy to speculate that a substantial amount of our most sophisticated thinking makes direct use of the meanings of lexical items, rather than the associated non-linguistic concepts. At least, such thinking seems aptly expressed in language. Attending to certain structural features of concepts might wind up being about the meanings of words that express them. If you ask if a breaking event involves one subevent or two, you might well be attending to features of the meaning of *break*. Even so, one might still rely on the non-linguistic concept, say, in actions to avoid breaking a plate while putting it in the dishwasher.

I thus suggest that implying parallel concepts should not be seen as a drawback of the packaging picture. That being said, there are other options for how lexical meaning and our broader conceptual faculties could relate, and some of them raise some interesting questions about linguistic determinism. To conclude, I shall mention some of them.

One option is that what we see in a lexical entry is not really language-specific structure, but rather, an articulation of the 'ordinary' or intuitive concept.[29] I argued against this option in section III, where I gave some reasons for taking elements like CAUSE to be distinctively linguistic. Even so, the arguments I presented there are more suggestive than conclusive, so this option remains a possibility.

There are two more dramatic possibilities. One is to reject the supposition that our ordinary concepts like BREAK lack distinctively linguistic structure. The idea would be that we never really have such a concept prior to language development. In cases of variation across languages in how roots are packaged, of the sort we saw in section I, this would imply that speaker of different languages have different concepts. It thus amounts to a form of linguistic determinism. It is a more limited and more nuanced one than is associated with Whorf and Sapir (cf. Carroll, 1956), but it is a form of linguistic determinism none the less.

A third option is one of refinement. Perhaps we start with some concept

[29]I believe this idea is reasonably attributed to Jackendoff (e.g. Jackendoff, 1983, 1990), though Jackendoff is certainly committed to the existence of a language faculty. It is also common in various forms of cognitive grammar, which would reject many more of the assumptions I have made throughout this paper.

of, e.g. BREAK, perhaps not clearly distinguishing causal breaking events from alternatives like breaking manners of action. Input from the language faculty allows us to refine and replace such concepts with more specific ones, relying on linguistic structure to specify types of events and other aspects of concepts. This is a somewhat Vygotskian picture (Vygotsky, 1962). It still implies a modest amount of linguistic determinism, but not in our basic concepts.

I have argued that the parallel concepts position is plausible, and I have offered some arguments against the view that semantic structure is identical to conceptual structure. So, I prefer the parallel tracks view to any of these options. But, I believe a more thorough exploration of all the options, including those which imply some modest degree of linguistic determinism, is in order. Deciding among them is a complex matter, involving some substantial empirical issues. I hope to pursue this in other work. I shall rest here with my claim that the packaging of roots in the lexicon allows word meanings to involve mappings to concepts, but opens up a surprising amount of space for complex relations between the meanings of words and concepts.

References

Aronoff, M., 1994. *Morphology by Itself*. Cambridge: MIT Press.

Baker, M. C., 1988. *Incorporation*. Chicago: University of Chicago Press.

Bierwisch, M. and R. Schreuder, 1992. From concepts to lexical items. *Cognition* **42**:23–60.

Bloom, P., 2000. *How Children Learn the Meanings of Words*. Cambridge: MIT Press.

Bowerman, M. and S. C. Levinson (eds.), 2001. *Language Acquisition and Conceptual Development*. Cambridge: Cambridge University Press.

Camp, E., 2001. Metaphor in the mind: The cognition of metaphor. *Philosophy Compass* **1**:154–170.

Carey, S., 1985. *Conceptual Change in Childhood*. Cambridge: MIT Press.

———, 2009. *The Origin of Concepts*. Oxford: Oxford University Press.

Carroll, J. B. (ed.), 1956. *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*. Cambridge: MIT Press.

Carruthers, P., 2002. The cognitive function of language. *Behavioral and Brain Science* **25**:657–726.

Chen, K., 1986. A purely geometric module in the rat's spatial representation. *Cognition* **23**:149–178.

Chierchia, G., 2004. A semantics for unaccusatives and its syntactic consequences. In A. Alexiadou, E. Anagnostopoulou, and M. Everaert (eds.), *The Unaccusativity Puzzle*, pp. 22–59. Oxford: Oxford University Press.

Chomsky, N., 1986. *Knowledge of Language*. New York: Praeger.

———, 1995. *The Minimalist Program*. Cambridge: MIT Press.

———, 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Clark, E. V., 1983. Meaning and concepts. In P. H. Mussen (ed.), *Handbook of Child Psychology*, vol. III: Cognitive Development, pp. 787–840. New York: Wiley, 4th edn.

Collins, J., 2011. Impossible words again: Or why beds break but not make. *Mind and Language* **26**:234–260.

Dennett, D. C., 1996. *Kinds of Minds*. New York: BasicBooks.

Dowty, D. R., 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dummett, M., 1989. Language and communication. In A. George (ed.), *Reflections on Chomsky*, pp. 192–212. Oxford: Blackwell.

———, 1991. *The Logical Basis of Metaphysics*. Cambridge: Harvard University Press.

Fillmore, C. J., 1970. The grammar of *Hitting* and *Breaking*. In R. A. Jacobs and P. S. Rosenbaum (eds.), *Readings in English Transformational Grammar*, pp. 120–133. Waltham: Ginn.

Fisher, C., D. G. Hall, S. Rakowitz, and L. Gleitman, 1994. When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua* **92**:333–375.

Fodor, J. A., 1975. *The Language of Thought*. New York: Crowell.

———, 1983. *The Modularity of Mind*. Cambridge: MIT Press.

———, 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Fodor, J. A. and E. Lepore, 1999. Impossible words? *Linguistic Inquiry* **30**:445–453.

Frege, G., 1884. *Die Grundlagen der Arithmetik*. Breslau: Koebner. Citations to the translation as *Foundations of Arithmetic* by J. L. Austin, second edition, Northwestern University Press, Evanston, 1953.

———, 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* **100**:25–50. References are to the translation as "On Sense and Meaning" by M. Black, reprinted in Frege (1984).

———, 1984. *Collected Papers on Mathematics, Logic, and Philosophy*. Oxford: Basil Blackwell. Edited by B. McGuinness.

Gallistel, C. R., 1990. *The Organization of Learning*. Cambridge: MIT Press.

Gentner, D. and S. Goldin-Meadow (eds.), 2003. *Language in Mind*. Cambridge: MIT Press.

Glanzberg, M., 2008. Metaphor and lexical semantics. *The Baltic International Yearbook of Cognition, Logic and Communication* **3**:1–47. Http://thebalticyearbook.org/journals/baltic/article/view/ 14/13.

Gleitman, L., 1990. The structural sources of verb meanings. *Language Acquisition* **1**:3–55.

Grimshaw, J., 2005. Semantic structure and semantic content in lexical representation. In *Words and Structure*, pp. 75–89. Stanford: CSLI Publications.

Hale, K. and S. J. Keyser, 1993. On argument structure and the lexical expression of syntactic relations. In K. Hale and S. J. Keyser (eds.), *The View from Building 20*, pp. 53–109. Cambridge: MIT Press.

———, 2002. *Prolegomenon to a Theory of Argument Structure*. Cambridge: MIT Press.

Harley, H., 2007. The bipartite structure of verbs cross-linguistically. Ms.

Higginbotham, J., 1987. The autonomy of syntax and semantics. In J. L. Garfield (ed.), *Modularity in Knowledge Representation and Natural-Language Understanding*, pp. 119–131. Cambridge: MIT Press.

———, 1989. Knowledge of reference. In A. George (ed.), *Reflections on Chomsky*, pp. 153–174. Oxford: Blackwell.

Hurford, J. R., 2007. *The Origins of Meaning*. Oxford: Oxford University Press.

Jackendoff, R., 1983. *Semantics and Cognition*. Cambridge: MIT Press.

———, 1990. *Semantic Structures*. Cambridge: MIT Press.

Johnson, K., 2004. From impossible words to conceptual structure: The role of structure and process in the lexicon. *Mind and Language* **19**:334–358.

Keil, F. C., 1979. *Semantic and Conceptual Development*. Cambridge: Harvard University Press.

Kratzer, A., 1996. Severing the external argument from its verb. In J. Rooryck and L. Zaring (eds.), *Phrase Structure and the Lexicon*, pp. 109–137. Dordrecht: Kluwer.

Larson, R. K., 1988. On the double object construction. *Linguistic Inquiry* **19**:335–391.

Levin, B., 1993. *English Verb Classes and Alternations*. Chicago: University of Chicago Press.

Levin, B. and M. Rappaport Hovav, 1995. *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Cambridge: MIT Press.

———, 2005. *Argument Realization*. Cambridge: Cambridge University Press.

Marantz, A., 1997. No escape from syntax: Don't try morphological analysis in the privacy of your own lexicon. *University of Pennsylvania Working Papers in Linguistics* **4**:201–225.

Margolis, E. and S. Laurence, 1999. *Concepts: Core Readings*. Cambridge: MIT Press.

McClure, W., 1990. A lexical semantic explanation for unaccusative mismatches. In K. Dziwirek, P. Farrell, and E. Mejías-Bikandi (eds.), *Grammatical Relations*, pp. 305–318. Stanford: CSLI Publications.

Michotte, A., 1963. *The Perception of Causality*. London: Methuen. Translation by T. R. Miles and E. Miles from the French edition of 1954.

Murphy, G. L., 1991. Meaning and concepts. In P. J. Schwanenflugel (ed.), *The Psychology of Word Meaning*, pp. 11–35. Hillsdale: Erlbaum.

———, 2002. *The Big Book of Concepts*. Cambridge: MIT Press.

Parsons, T., 1990. *Events in the Semantics of English*. Cambridge: MIT Press.

Pesetsky, D., 1995. *Zero Syntax*. Cambridge: MIT Press.

Pietroski, P. M., 2005. *Events and Semantic Architecture*. Oxford: Oxford University Press.

———, 2010. Concepts, meanings and truth: First nature, second nature and hard work. *Mind and Language* **25**:247–278.

———, forthcoming. Semantic monadicity with conceptual polyadicity. To appear in the *Oxford Handbook of Compositionality*.

Pinker, S., 1989. *Learnability and Cognition*. Cambridge: MIT Press.

———, 1994a. How could a child use verb syntax to learn verb semantics? *Lingua* **92**:377–410.

———, 1994b. *The Language Instinct*. New York: HarperCollins.

Poulin-Dubois, D. and S. A. Graham, 2007. Cognitive processes in early word learning. In E. Hoff and M. Shatz (eds.), *Blackwell Handbook of Language Development*, pp. 191–211. Oxford: Blackwell.

Putnam, H., 1970. Is semantics possible? *Metaphilosophy* **1**:187–201.

Rappaport Hovav, M. and B. Levin, 1998. Building verb meanings. In M. Butt and W. Geuder (eds.), *The Projection of Arguments*, pp. 97–134. Stanford: CSLI Publications.

Rosch, E., 1978. Principles of categorization. In E. Rosch and B. B. Lloyd (eds.), *Cognition and Categorization*, pp. 27–48. Hillsdale: Erlbaum.

Saxe, R. and S. Carey, 2006. The perception of causality in infancy. *Acta Psychologia* **123**:144–165.

Scholl, B. J. and P. D. Tremoulet, 2000. Perceptual causality and animacy. *Trends in Cognitive Science* **4**:299–309.

Spelke, E. S., 2002. Developing knowledge of space: Core systems and new combinations. In A. M. Galaburda, S. M. Kosslyn, and Y. Christen (eds.), *The Languages of the Brain*, pp. 239–258. Cambridge: Harvard University Press.

Vygotsky, L. S., 1962. *Thought and Language*. Cambridge: MIT Press. Edited and translated by E. Hanfmann and G. Vaker from the Russian edition of 1934.

Wunderlich, D., 1997. Cause and the structure of verbs. *Linguistic Inquiry* **28**:27–68.