

From impossible words to conceptual structure: the role of structure and processes in the linguistic lexicon

Kent Johnson,
 Department of Logic and Philosophy of Science, UCI
 johnsonk@uci.edu
<http://hypatia.ss.uci.edu/lps/home/fac-staff/faculty/johnson/>
 Draft: July, 2002

0 Introduction

An important finding of theoretical linguistics is that there are certain broad patterns of words that never appear in a language. For instance, although English contains the transitive verb *break*, as in *Mary broke the desk*, English does not contain a verb expressing the converse relation of *break*, as in **The desk blikked Mary*. The absence of a verb like *blik* does not appear to be an accident, in the way that it is an accident that there is no noun that picks out one's tallest friend's relatives. Instead, the absence of the verb *blik* appears to be due to the more general fact that whenever a transitive verb of English expresses a relation between the doer of an action and the thing that is acted upon – what linguists call the 'agent' and the 'theme' of the verb – the former is the subject of the verb and the latter is the object (e.g., *blacken*, *boil*, *kill*, *shrink*, etc.). Like all generalizations in linguistics, this one is subtle. Since I will use it to motivate the issue, let me say a few words about the generalization. First, the generalization does not apply to all verbs: if your coffee cup resembles mine, for instance, then mine may resemble yours. However, *resemble* is not a verb that clearly expresses that one thing acts upon or does something to another. I'm interested in the really clear cases, like those listed above (which is not to say there are no unclear cases; cf. Dowty 1989, 1991). Secondly, the passive form of verbs, as in *The desk was broken by Mary* are not relevant here, because the passive form is a morphosyntactic alternation of a root verb (for discussion cf. e.g. Baker, Johnson and Roberts 1989). The generalization applies only to root verbs, of which there are many like *break* and none like *blik*.

When a candidate word violates a generalization like the one just discussed, it is called an *impossible word*. Here, *impossible* just means impossible relative to the language in question, or perhaps relative to all natural languages. That is, impossible words are ungrammatical words.¹ The linguistics literature contains a variety of explanations why words like *blik* are impossible (e.g., Baker 1988, Grimshaw 1990). More generally, the phenomenon of impossible words has received much attention from widely differing standpoints (E.g., McCawley 1968, Carter 1976, Dowty 1979, Baker 1988, Grimshaw 1990, Hale and Keyser 1993, 1997, 1999, Fodor and Lepore, 1999, Fodor, 1998, Jackendoff 1990, Zubizarreta 1987).

The fact that there are impossible words is often taken to have important implications for psychology and the philosophy of mind. This is because it is common for linguists to explain why some words are impossible by appealing to complex structure within actual words. (All of the citations above give examples of this last claim, and I will illustrate how it works below.) Indeed, the existence of impossible words and the theories that explain them often provide some of the strongest linguistic support for the claim that many of our ordinary words contain a good deal of complex linguistic structure. The bridge from linguistics to psychology takes the form of an assumption that (*ceteris paribus*) the structure of our words is mirrored in the concepts they express. There are various ways of spelling out this word-concept bridging assumption. For definiteness, I will only deal with what I will call the ‘Isomorphism Assumption’, which says that the structure of a word is isomorphic to the structure of the concept it expresses (assuming that the word expresses a concept).² Many philosophers and linguists hold this view, or something like it. For instance, Jerry Fodor and Ernie Lepore write:

The idea that quotidian, middle-level concepts typically have internal structure – definitional, statistical, or whatever – plays a central role in practically every current approach to cognition. Correspondingly, the idea that words that express quotidian, middle-level concepts have complex representations ‘at the semantic level’

¹ Note that when a linguist claims that a word is impossible for a given language, she does not mean that it is metaphysically or nomologically impossible for the word to enter an arbitrary (artificial or natural) language. Rather, the claim is that the psychological laws that distinguish the highly restricted range of natural languages from all the logically possible languages also serve to prevent certain words from entering these languages (or from some relevant subset of them). That is, the linguist claims that it is ungrammatical for a similar reason that **Who did Mary kiss Bill* is. In the latter case, our linguistic abilities do not allow us to construct a certain structure (although we can override them by a nonlinguistic mental effort). In the case of impossible words, our linguistic abilities do not allow us to enter a certain structure into the lexicon. We can, of course, override the linguistic system, but the point is that we have evidence that any such structure would be a violation of the system’s rules. Some confusion would be avoided if linguists spoke of *ungrammatical words* instead of *impossible words*; e.g., FL 1999, pp. 450, 452.

² Although there are many interesting issues about the relation between words and concepts that might also be explored – e.g., what does it mean for word-structure to correspond to conceptual structure, what happens if we weaken the relation between words and concepts to a homomorphism, or strengthen it to identity – I will leave these issues aside, since they will not affect the main line of the paper.

is recurrent in linguistics; it is the defining thesis of what is often called ‘lexical semantics’ (Fodor and Lepore 1999, p. 445).

In a related spirit, the linguists Jaklin Kornfilt and Nelson Correa describe their theory of word-structure as ‘part of the larger system of conceptual structure which underlies human cognitive abilities’ (Kornfilt and Correa 1993, p. 79; cf. also Jackendoff 1990, esp. ch. 1). Similar views have been echoed in the philosophical and psychological literature (e.g., Peacocke 1992, p. 3, fn. 2, Fodor 1998, p. 2, Miller 1996, pp. 18-19). If such sentiments are correct, then we might be able to learn about at least some of our concepts – how many we have, which ones we have, how they are structured, etc. – by looking to the structure of language. Similarly, it might turn out that linguistics alone can justify the claims that some words have definitions and/or that some concepts have possession conditions involving other concepts (e.g., that *break* as in *Mary broke the desk* really does mean *cause to be broken*, and that to possess the relational concept BREAK, one must also possess concepts of causation). From here, it is only a short road to establishing the existence of analyticities, synonymies, and many other hotly contested issues from the history of philosophy. The Isomorphism Assumption also explains why people who care about concepts often care about the nature and structure of words. In particular, those who hold that ‘most lexical concepts [sc. the concepts expressed by our individual words] have no internal structure’ (Fodor 1998, pp. 2, 121; cf. Fodor and Lepore 1992, Laurence and Margolis 1999,) will care how the linguistic theory turns out. Such researchers will be strongly motivated to show that words have no internal structure. Thus, they will want to undermine the ability of theories of impossible words to supply evidence to theories about word-structure.³

Although the phenomenon of impossible words appears at first glance to be an obscure topic from a small corner of linguistics, we see that it is potentially quite important for the study of broader psychological issues. However, the empirical and conceptual issues involved in research into the nature of words are both delicate and complicated. In this paper, I will sort out a few of these issues. I will focus on a number of interpretations of linguistic practice which suggest that despite what linguists say, words contain no interesting structural elements. Some of these suggestions can be found in the literature, and others are novel, but have an intuitive

³ This way of putting things is a little loose. Everyone can agree that there is structure in a word that has no conceptual correlate; for instance, the phonological properties of words are probably not mirrored in our concepts. The issue is rather whether the concept that a word expresses contains all (and only, perhaps) the word’s semantic and semantically relevant syntactic structure. Since most of what I have to say will concern fairly straightforward examples like whether the transitive verb *break* contains a causal element, the admittedly rough characterizations given here will suffice.

appeal. If words have no interesting structure, then by the Isomorphism Assumption, it would follow that the concepts they express lack structure (or at least the kinds of structure that linguistics is sometimes thought to uncover). Such a result would be a major advance for proponents of conceptual atomism (e.g. Fodor 1998; cf. Fodor and Lepore 1992). However, I will argue that none of these approaches meets with any degree of success. Moreover, since these are the most plausible strategies I know of, it looks as though conceptual atomism does not have a correlate in natural language.

The paper is structured as follows. I begin by characterizing how a theory of impossible words seems to support theories about the internal structure of words. I then turn to three general strategies for avoiding the conclusion that words contain internal structure. The first strategy argues that there can be no linguistic explanation of impossible words, since any word, including impossible ones, can enter the lexicon as a primitive word. The second strategy attempts to dispose of one of the crucial mechanisms for creating structurally complex words. If we have no reason for positing that mechanism, then we have no reason for positing structure in words in the first place. The third strategy attempts to eliminate the need for positing structure in words by positing corresponding bits of structure in the syntax of language. I will argue that all of these strategies fail for multiple reasons. I conclude with some worries about the Isomorphism Assumption.

1 Impossible words and the arguments they figure into

In this section, I will provide a first-pass characterization of an argument purportedly showing that theories of impossible words evidentially support the existence of internal structure in many of our words. The rest of the paper will be based around attempts to undermine this argument. It will be easiest to illustrate how the argument works with an example, and I will use a simplistic theory of the causative construction in English. Before sketching this theory, let me briefly comment on the general role it will play. The objections to be considered are intended to hold against any theory of the lexicon whatsoever. Two of the objections come from Fodor and Lepore, who are adamant about this point: ‘We think there are principled objections to this form of argument’ (Fodor and Lepore 1999, p. 445), ‘impossible-word arguments are infirm in

principle' (p. 447; cf. also p. 446). So strictly speaking, the present theory of causative verbs will serve as a counterexample to the objections. However, I have deliberately chosen a well-known and traditional theory to illustrate in general how theories of this kind operate. So if your favorite view of the lexicon does not agree on every aspect of the theory, you should be able to adjust the offending aspects of the sample theory without altering the main points made here.⁴ Thus, I will not merely show that the objections in question are unsuccessful in principle, I will show that they are unsuccessful in practice and may be ignored.

The theory I mentioned above focuses on explaining why we have verbs like *break* (as in *Mary broke the desk*) but not ones like *blik* (as in **the desk blikked Mary*). To account for this phenomenon, the theory hypothesizes that although the transitive verb *break* appears to be a simple lexical item, it is in fact a complex structure, derived from the intransitive verb *break* (as in *the desk broke*), and another structure that means roughly CAUSE⁵. Thus, there might be some processes in the lexicon that allow the intransitive verb *break* to combine with this causal prefix to form a complex structure that means roughly X CAUSE Y BE BROKEN. A crucial part of the linguistic construction of this verbal structure is a 'lexicalization' process, whereby CAUSE and BE BROKEN are fused into a single word, the transitive verb *break*.⁶ Thus, the process of lexicalization takes a structure of the form [X CAUSE Y BE BROKEN] as input, and yields a structure of the form [X BREAK Y] as output. Obviously, this process is in part phonological: *cause* is not a syllable of *break*. Below I will present empirical evidence that the process also has syntactic-cum-semantic effects as well. In addition to explaining how causative verbs like transitive *break* (and *boil*, *blacken*, *kill*, *shrink*, etc.) work, such a theory also appears to explain why there are no verbs like *blik*. The apparatus for generating transitive verbs from intransitive ones forces the agent (i.e., the subject of CAUSE) to appear as the subject of the resulting complex verb. Since *blik* places the agent in the object position, it is not generated by the apparatus, so the theory correctly predicts the word is not found in English.

⁴ In terms of a larger linguistic theory, I will be assuming a very mainstream form of Principles-and-Parameters (Chomsky 1981). However, the general strategy I offer should be amenable to a wider variety of linguistic theories, although some of them may need to fiddle with some the arguments a little. E.g., proponents of lexical-functional grammar will want to alter the argument in section 4.

⁵ Famously, the causal morpheme does not mean the same thing as the English word *cause* (e.g., Fodor 1970, Comrie 1986 Pietroski 1998, 2000, ch. 1). I will return to this point later.

⁶ Many researchers think that transitive *break* is not a single word, but is actually contains some phonologically unpronounced extra structure (e.g., Kratzer 1996, Travis 2000). Others disagree (e.g., Grimshaw 1990, ch. 2, Zubizarretta 1987). I chose the example because it is a very simple and intuitively plausible illustration of the kind of phenomena and explanations that go on in research into the lexicon. There are many additional examples present in the literature; e.g., the case of *shelve* below.

As a first pass, a theory of impossible words of the sort just given appears to license something like the following, which I hereby name the *Impossible Word Argument*:

- (1) According to theory T, there are processes in the lexicon for altering a word's structure, and for combining various structures into a single structure.
- (2) According to T, impossible words like *blik* are not the output of any processes in the lexicon. So T correctly predicts that the impossible words do not exist.
- (3) Thus, T *explains* why impossible words are impossible,
- (4) Thus, we are (*ceteris paribus*) justified in supposing that there are processes and structure in the lexicon.

Although this reconstruction does not capture every relevant aspect of the dialectic, it nonetheless allows us to pinpoint where the objections are focused. We will see that the first objection is directed at the inference to (3), and the next two concern the inference to (4). I turn now to the objections.

2 Objection one: Any word can be a primitive

The first objection to the Impossible Word Argument is due to Fodor and Lepore (hereafter FL). The objection concerns the strategy of explaining the impossibility of a word by appealing to a violation of a rule (or an illegal movement, etc.), where this violation occurs in the lexicon, before the word is inserted into the syntax. We saw such an example above with the case of *blik*. This strategy fails to explain why the language does not contain impossible words, FL argue, because such words might have entered the lexicon as primitive expressions. That is, if FL are right, then there is no reason that a speaker could not learn the word *blik* as a primitive transitive verb. ('Oh no, the sandcastle blikked Sue! That was not nice Sue! Sue ruined a beautiful sandcastle!') If a word can enter the lexicon fully-formed, needing no modification by processes, then the word itself will not violate any rules before, during or after it is inserted into the syntax.⁷ If this line of reasoning is right, then we have not yet explained why certain words

⁷ FL write: 'By definition, impossible-word arguments purport to explain intuitions of the form "There couldn't be a word *w* that means *E*" by showing that *E* couldn't be a derivational source for *w*. But, on the face of it, that doesn't show that there couldn't be a word that means *E*; the most it could show is that if there *is* a word that means *E*, then it must be primitive' (Fodor and Lepore 1999, p. 449).

are impossible. Why couldn't *blik* just enter the language as an unstructured primitive element, without undergoing any construction process within the lexicon or elsewhere?

As an argument about the methodology of linguistic and psycholinguistic theory, FL's objection contains a rather straightforward flaw. They assume that our linguistic abilities include an ability to add any new verb as a primitive expression in the lexicon. But no contemporary linguist or psycholinguist, as far as I know, has made (or would make) this assumption. To claim that any word can freely enter our language as a primitive expression is tantamount to leaving the existence of impossible words unexplained, because the assumption predicts that we will not find the broad patterns of non-occurring structures that are definitive of impossible words. That is, if any old word whatsoever could enter the lexicon as a primitive, we would expect there to be a significant population of words like *blik*, where the agent is in the object position of the verbs. The fact that we don't find any verbs like *blik* in our language would be utterly amazing if FL's assumption were correct.

In contrast to FL's proposal, linguists and psycholinguists commonly posit severe constraints on what can enter the lexicon. They often assume that the lexicon contains (i) processes for generating new words out of already existing elements, and (ii) processes for fitting a meaning onto a word, after the meaning has been extracted from the perceived environment (e.g., Grimshaw 1997, Gleitman 1990, Pinker 1989, Gropen et al. 1991, 1991a). Both these processes are commonly thought to observe strong constraints on lexical syntactic and semantic structure, in direct contradiction of FL's assumption. A theory of this sort, unlike FL's, can make strong and testable empirical predictions about impossible words. For example, the processes specified in (ii) entail that meanings are fit onto words in certain ways, and not in others. (E.g., theories allowing for agentive verbs to be formed in any of several ways still restrict the processes of word-formation to insure that the agent is always the verb's subject.) By insisting that a word-meaning be fit onto a word in only a limited number of ways, such theories can predict that words like *blik* do not occur. Moreover, these predictions take the form of explanations of the impossibility of the words in question. The prediction/explanation of such a theory has the form: *blik* is not a possible word because the agent is not in the subject position, and the theory entails that you can't form a verb whose agent is not in the subject position.

How does this response to FL's objection fit into the argument given in (1) – (4) above? Naturally, one could build in the assumption explicitly, by adding the premise:

- (1.5) If a word is not the output of a process in T, then T predicts that the word is impossible.

While I have no problems with (1.5), a couple points about it are worth noting. First, on our current assumptions, (1.5) is analytic (or virtually so). After all, in our list of processes, we are including those processes that fit a meaning onto a word, so in that sense every word, even the primitive expressions, are the output of one process or another. Second, claims like (1.5) are typically taken for granted in linguistics. This is because linguists frequently think of a linguistic theory as being total: i.e., the theory specifies every relevant structure (in the present case, every word-structure) there is in the language. (More precisely, linguists assume that the theory is total over a given relevant subset of constructions; nobody has a theory of how an entire language works.) Theories typically need to be evaluated not only by what structures they predict to be grammatical, but also by what structures they predict to be ungrammatical. Indeed, without attending to the latter form of prediction, linguistic theory would be trivial: one simply allows every structure there is into the language. Such a theory makes all the right predictions about what structures are in the language, but it is worthless because it makes no predictions about what structures are not in the language. Not only does such a theory give us no information about the general architecture of the language, it also renders irrelevant the nontrivial generalizations about what sentences are in the language, and these generalizations are the meat and potatoes of linguistics.

In sum: FL's suggestion that any word can enter the lexicon as a primitive is, as a matter of sociological fact, an assumption that linguists explicitly deny. Moreover, anyone interested in language has at least two good reasons to reject that assumption. First, the assumption hinders a theory's ability to make predictions about what it is not in the language, and such predictions are crucial to linguistic theory in general. Second, the assumption predicts that we should find a certain number of clear violations of some of the generalizations about the lexicon, where there is no explanation for these violations whatsoever, except that they just appeared on the linguistic scene as primitives. This assumption, however, appears to be empirically false.

The above argument strikes me as conclusive. FL are simply wrong to suggest that any word can enter the lexicon as a primitive expression. However, a number of philosophers have

expressed some residual unease about the claim that there are restrictions on the primitive expressions of language. Their concern goes like this:

Look, there is surely something right about the claim that any word (even the supposedly impossible ones) can enter the language as a primitive expression. If a mad dictator threatened a population of speakers with torture or death if they did not start using the verb *blik*, of course, they would begin to do so. Over time (a few generations, perhaps), the word might even become an accepted component of that language, just like *break* is an accepted part of English. (In more psychological terms, *blik* would be entered into the linguistic lexicon of typical speakers of the language.)

The general form of this argument is: for every (allegedly) impossible word, there is some possible circumstance in which the word comes to be a normal part of our language. The problem with the argument is that in the possible circumstances imagined, *blik* enters the language in a highly abnormal way, and there is no reason to suppose that the presence of the verb in these circumstances undermines the linguistic generalization that agents are always subjects. After all, *every* law about human languages and human linguistic behavior is highly *ceteris paribus*. For example, there are (ungrammatical) sentences that are not generated by human grammars. But the fact that your language will not allow you to utter (5a), whose meaning is given in (5b), does not mean that uttering (5a) is physically impossible for you.

- (5) a. *The child seems sleeping.
 b. The child seems to be sleeping.

Moreover, if a mad dictator demanded that you use (5a), it would surely enter our linguistic repertoires over time.⁸ But this fact does not mean that in accounting for human linguistic abilities, we shouldn't distinguish (5a) from (5b). We can and should distinguish the two simply because when the linguistic situation is normal (e.g., no mad dictators), English speakers do not utter (5a). They might, however, utter (5b) even if there are no mad dictators around. This is a well-known point that goes back to the foundations of linguistics. The situation is very analogous with respect to the lexicon, except that we focus on what words are (not) in the lexicon rather than looking at what sentences are (not) part of the language. Judgments about what words are actually impossible are notoriously subtle and difficult (although that of course

⁸ As XXX pointed out to me, we need not concede the general point about mad dictators' abilities to influence language. Just because we come to use a given expression X (or, less question-beggingly, a sequence of phonemes) along with certain linguistic elements, it does not follow that X is a part of the language. After all, if X observes none of the relevant laws of language, then we should ask what it is about X that makes it a part of language.

does not help these purportedly ‘in principle’ objections!) but some generalizations are pretty robust. The claim that agents are subjects of verbs (if the verb has a position for an agent) is one of them. In short, if the objection quoted above is any good, then the same considerations should lead us to dismiss (virtually) all of linguistics. But this, to put it mildly, would be extraordinarily bad science, since it would mean dismissing numerous striking and useful generalizations that appear when the *ceteris paribus* part of *ceteris paribus* generalizations are taken seriously.

I imagine that at this point, the following objection might be made. ‘The analogy just made is inappropriate. Syntax is different from the lexicon in that syntax is *productive*. In a given language, there are a potential infinitude of sentences, but only finitely many words. We can extract the true regularities from syntax by looking to the infinite set of sentences.’ There are at least two replies to this worry. First of all, even though we know only finitely many basic words,⁹ there are nonetheless infinitely many possible words of English that could enter our vocabularies. Exploring how speakers react to various potential words is a source of evidence regarding the workings of the lexicon that is much like evidence from syntax. Indeed, empirical work on how children acquire certain words provides evidence that they do indeed obey numerous linguistic generalizations when fitting a meaning onto a word (Gropen et al. 1991, 1991a, Bloom 2000, Pinker 1999). Secondly, this objection doesn’t solve the original problem, but only makes it into one of evidence. To see this, note that the productivity of syntax is an empirical hypothesis about our linguistic abilities, and this hypothesis has been justified only by consideration of sentences of relatively short length. But for any string of less than 1,000 words, say, a mad dictator could force us to use it as a sentence so that it entered our linguistic repertoire, and perhaps he would also force us to refrain from using other allegedly grammatical sentences. In this circumstance, we would get a very different distribution of grammatical vs. ungrammatical sentences across the finite population of sentences that linguists actually study – linguists don’t study sentences 1,000,000 words long, after all. Moreover, in this circumstance, we would lose all of our empirical evidence about whether and how human languages are productive, and whether and why any given sentences are ungrammatical. In short, if considerations about what is barely possible are relevant to linguistic theorizing, then we are not

⁹ I leave aside the infinitude of words derived by productive derivational morphology, e.g., *great-grandmother*, *great-great-grandmother*, *great-great-great-grandmother*, etc. (cf. Pinker 1994, ch. 5).

entitled to make claims about the productivity of language, which means that a crucial assumption of the present is unavailable.

So far, we have been concerned with the issue of whether any words whatsoever can enter the lexicon, because we are interested in whether linguistic theories can support explanatory generalizations like ‘agents are always subjects’. I will end this section by commenting on the resiliency of such generalizations in the face of individual counterexamples. Suppose for the sake of argument that someone finds a couple English verbs where the agent appears as the grammatical object. Those two or three verbs will be enough to undermine the universal generalization that agents are always subjects. But they will not be enough to undermine the more guarded *ceteris paribus* claim that there is something about our abilities to fit words onto meanings that strongly encourages agents to appear as subjects. This more guarded claim can still supply the needed explanatory force for why *blik* is not found in English, in the same way that many other *ceteris paribus* claims of the empirical sciences are explanatory.¹⁰ Moreover, to undermine this more guarded *ceteris paribus* claim, one will probably need a *lot* of verbs that conflict with the generalization.

To illustrate this last claim, I used a Beth Levin’s (1993) lists of English verbs to collect some transitive verbs which clearly express agency, in the sense that something is clearly doing something to something. I stopped after I’d collected 464 such verbs. Although I was only looking for agentive transitive verbs, it turned out that all 464 of the verbs I collected clearly placed the agent of the action in the subject position and the theme in the object position. 464 is an extremely conservative estimate of the total number of agent-subject verbs; the typical speaker’s repertoire probably numbers in the thousands.¹¹ But in any case there are at least 464 transitive verbs expressing agency. I will call transitive verbs with the agent in the subject position ‘agent-subject’ verbs. I found no transitive verbs in English that have the agent as the grammatical object – which I will call the ‘agent-object’ verbs – nor do I know of any.¹² However, let us assume that there is at least one agent-object verb, and to be safe, let us increase

¹⁰ E.g., why was this dog born with four legs? Because *ceteris paribus*, dogs are born with four legs (although a few are not).

¹¹ To be safe, I did not include verbs that might plausibly be derived from some other linguistic source; e.g., *acidify* is probably a morphological compound of the noun *acid* and the derivational morpheme *-ify*. I did however include some verbs that have nominal counterparts. I assumed, for instance, that *perturb* is plausibly underived, and that the nominalization *perturbance* is probably derived from it. The claim that one verb is derived from some linguistic source and that another is underived is, of course, an empirical claim. This is as it should be.

¹² Cf. the earlier comments on verbs not expressing agency, and the passive voice of verbs.

this estimate by an order of magnitude, and assume there are 10 of them, so that we now have a total of 474 transitive verbs with agents and themes. Now suppose that there was nothing influencing the order in which agents and themes appeared in these 474 verbs. In other words, assume with FL that any word can enter the lexicon as a primitive. In that case, we would expect that somewhere around half (237) of them would be agent-subject verbs and the other half would be agent-object verbs. What then is the probability of only ten of these 474 verbs being agent-object verbs and the other 464 of them being agent-subject verbs?¹³ Since (we are assuming) any given verb is equally likely to be an agent-object verb as it is to be an agent-subject verb, a bit of elementary probability theory shows that any given subset of these 474 verbs is as likely to be the subset of agent-object verbs as any other subset. Thus, the probability that there would be 10 or fewer agent-object verbs can be computed as below:

$$(6) \quad \frac{\sum_{k=0}^{10} \binom{474}{k}}{2^{474}} = \frac{1.47 \times 10^{20}}{4.88 \times 10^{142}} \approx 3 \times 10^{-123}$$

(6) shows that the chance of there being no more than ten agent-object verbs is vastly less than one out of one googol.¹⁴ If there were 100 agent-object verbs, the probability of the null hypothesis would increase to only 2.65×10^{-57} . In order for there to be even a 1% chance that agent-object verbs are as likely to appear in the language as agent-subject verbs are, there would have to be at least 395 agent-object verbs. In short, unless the ratio of agent-subject verbs to agent-object verbs becomes vastly smaller than it appears to be, the only plausible option is to assume that there is indeed something about the language that encourages agents to appear in the object position. Another way to see this point is to hold fixed our set of verbs and to adjust the probability of a verb's being agent-subject. Let p be the probability that a given transitive verb

¹³ I restrict my attention to those verbs that express agents and themes on the assumption that that linguistic feature is determined by the semantic relation the verb expresses. Thus, when one coins a new verb to denote an action performed by one thing on another, one can create only an agent-subject verb or an agent-object verb.

¹⁴ The expression in the numerator is the sum of the binomial coefficients as k ranges from 0 (no agent-object verbs) to 10 (ten agent-object verbs). The binomial coefficient $\binom{n}{k}$ gives the number of size k subsets present in a set of size n . It is shorthand for $\frac{n!}{k!(n-k)!}$. The equation in (6) corresponds to a two-tailed test of a Bernoulli trial of size 474 with 10 successes or 10 failures where the probability of success is .5.

expressing agency is agent-subject. Notice that as p increases, so does the probability that at least 464 of our 474 verbs are agent-subject. We can then ask, how high would p have to be to ensure that there is a 1% chance that at least 464 out of our 474 verbs are agent-subject? This question can be answered by solving the equation (7), where p is the probability of a verb's being agent-subject, and $(1 - p)$ is the probability of its being agent-object.

$$(7) \quad \sum_{k=0}^{10} \binom{474}{k} p^{(474-k)} (1-p)^k = .01$$

(7) shows that p must be greater than 95.79%.¹⁵ Given this robust resistance toward agent-object verbs, it is a plausible empirical hypothesis that there really are constraints on how meanings are fit onto words. Although the hypothesis is empirical, it will not be easy to cook up an alternative story that not only accounts for the data as well as the present one but also integrates as well with other neighboring linguistic and psychological theories. Consider, for instance, the following alternative theory. Perhaps there are more agent-subject verbs because nowadays when we form new verbs, we have a tendency to structure them like other verbs that we already have, and perhaps we started doing this long ago with a few popular agent-subject verbs. Now there is a preponderance of agent-subject verbs, and the scarcity of the other kind is what makes them seem so awkward. The problem with this story is that it obviously relies on a notion of when two verbs are relevantly similar. The fact that the relevant similarity is expressed in terms of where the agent is located grammatically makes the story look like a causal explanation of the generalization, not an alternative to it. Moreover, if this were the end of the alternative story, it would leave open that other cultures should develop in which preference was given to agent-object verbs. Indeed, this possibility is so wide open that it would be rather surprising if we didn't find such cultures. But this particular linguistic generalization has been investigated in a wide variety of diverse languages. It simply does not appear to be the case that there are languages with agent-object verbs (cf. Baker 1988, 1996 for detailed discussion of many typologically distinct languages).¹⁶

¹⁵ For even odds that 464 out of 474 verbs would be agent-subject, we would need $p > .9775$, and for a 90% chance of this result, we would need $p > .9851$.

¹⁶ Two points: (i) The subject of a clause is not defined by its word-order. The notion of a subject is a relatively technical term of linguistics. Within the framework I adopt, the notion of a subject is defined purely configurationally, as e.g., the specifier of a

To sum up this section, FL claim that the linguist's argument in (1) – (4) is vitiated because although the finer-grained details of their theory don't produce the impossible words, the impossible words could still appear in the lexicon as primitive expressions. However, methodologically speaking, such an assumption would cripple a theory of the lexicon. Furthermore, the assumption makes false empirical predictions about the relative frequency of occurrence of certain kinds of words (e.g., agent-object verbs). Thus, the assumption in question should be rejected.

3 Objection two: Lexicalization is theoretically otiose

We have seen that there is good reason to hold that words do not enter the linguistic lexicon in just any way. Rather, it is a plausible empirical hypothesis that there are restrictions on the structure of the primitive words of the language. But acquiring a word as a primitive element of the lexicon is only one way to have a word in your repertoire. Another potential way is to possess some other words and to have a method for combining them into a structurally complex word. We have seen a potential example of this in the theory that says that *cause* and *be broken* combine to form the transitive verb *break*. But there are many other such examples. The adjective *black* combines with the morpheme *-en* to form the transitive verb *blacken*, as in *Ted blackened the fish*. An example I will use in what follows is the case of so-called 'denominal' verbs. Denominal verbs are verbs that are derived from a noun. For instance, the transitive verb *shelve* – as in *David shelved the books* – appears to be derived from the noun *shelf*. The linguists Kenneth Hale and Samuel Jay Keyser have suggested that the verb *shelve* is the result of the lexicalization¹⁷ of a complex syntactic structure, roughly the structure of the expression *put on a shelf* (e.g., Hale and Keyser 1993, 1997, 1999). Recall that lexicalization is the process that takes certain complex linguistic structures (like *put on a shelf*) as input and derives individual

VP or as the specifier of some kind of functional projection associated with the verb. (ii) A complete defense of this last claim would require a careful study of (inter alia) the ergative languages (cf. e.g., Bittner and Hale 1996). As I remarked at the beginning, the objections I am considering are supposed to hold in principle, so I am happy to allow my arguments to rest on some empirical assumptions.

¹⁷ Hale and Keyser refer to this process as 'conflation'; I will adopt FL's term 'lexicalization' throughout; any differences between the two processes should not matter here.

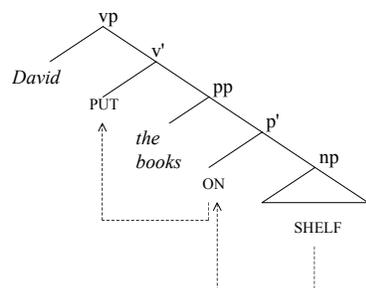
words (like the verb *shelve*) as output. (More precisely, the outputs of lexicalization act as individual words on the ‘surface’ syntax of the sentence.)

If Hale and Keyser’s theory is on the right track, then we should not find words like *shelve* that have a meaning and/or structure that would violate syntactic rules found in the overt syntax of natural language. If we do not find such words, then that is evidence for Hale and Keyser’s theory. On the other hand, if such words are common, then their theory will be vitiated. Leaving aside many linguistic details, Hale and Keyser argue that we do not find the sort of words that the theory predicts would have to be built out of ungrammatical structures. For instance, they note that we do not find words like **shelve*, as in (8)

(8) *David shelved the books on.

The verb **shelve* in (8) has roughly the meaning that David put the books somewhere with respect to a shelf, and the respect in question is given by the preposition *on* (as opposed to *near* or *under* the shelf). Hale and Keyser argue that the non-existence of words like the verb **shelve* in (8) is exactly what we would expect if their theory were correct. This is so because **shelve* has lexicalized a verb (roughly *put*) and a noun (*shelf* – the indefinite article *a* is not important here), but not the preposition *on* that modifies *shelf*. Hale and Keyser argue that the lexicalization of all but the preposition would require a structure in the lexicon that violates independently known laws of syntax.¹⁸ Thus, they claim, the non-existence of verbs like **shelve* supports the claim that there are structures and processes in the lexicon.

¹⁸For those who are interested in the details, HK’s proposal is that *David shelve the books* has the structure:



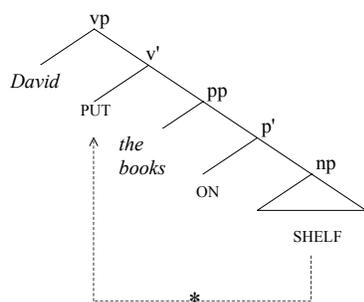
Here we see that there is head-movement from *shelf* to *on* to *put*. In contrast, the ungrammatical **David shelve the books on* has the structure:

FL offer a brief argument against this use of impossible words. They can be read as conceding for the sake of argument that the present theory really does provide a useful and important explanation of why some words are impossible (i.e., they can be understood as conceding that (1) – (3) are true). But for all that, FL argue, the theory still does not supply evidence that there are processes and structure in the lexicon. Their argument is as follows:

the assumption that lexicalization respects independent constraints on derivations is supposed to explain why intuitively impossible words *are* impossible; and the fact that *if lexicalization obeys the independent constraints on derivations, then it explains why impossible words are impossible* was supposed to be the evidence that there is such a process. [But] if [**shelve*] is ill formed ... because there is a mandatory constraint that [the lexical structure of **shelve*] fails to meet – then patently [**shelve*] is ill formed whether or not there is a lexicalization transformation.... But if [**shelve*] would be ill formed whether or not there is a grammatical process of lexicalization, how could the fact that it *is* ill formed be evidence for such a process? (FL 1999, pp. 451-2)

FL's point is this. Given that the input structure to the lexicalization process is ungrammatical, this fact alone should be enough to ensure that the corresponding word is ungrammatical. Thus, they claim, the lexicalization process itself is theoretically otiose. The theory explains the impossibility of **shelve* without really using lexicalization. So there is no need to posit a mechanism for turning structures into words.

But this is simply not so: without a process of lexicalization, there will be no way for any complex lexical structures to be formed into words in the first place. According to Hale and Keyser's theory, **shelve* is ungrammatical because it is underlyingly a complex ungrammatical structure, and there is a law of language that says that words, just like sentences, must not be ungrammatical structures. But if the process of lexicalization were absent from this picture, then there would be no linguistic connection between the verb **shelve* and the ungrammatical structure **put a shelf* from which it was derived. Since the only mechanism the present theory



Here, the head movement omits ON, violating the Head Movement Constraint, which has been independently motivated as a feature of grammar (Travis 1984, Baker 1988).

has for predicting the ungrammaticality of **shelve* is by associating it via lexicalization with an ungrammatical structure, the theory minus lexicalization would be utterly powerless to explain this impossible word. Thus, contrary to what FL say, lexicalization is not otiose, but rather is a crucial part of the linguistic theory.

Although FL's argument ends here, it may be possible to strengthen it with some augmentation. After all, the crucial role that lexicalization plays is to relate words to structures in the right way. It might be possible to posit some other mechanism that relates words to structures in a similar way, but which does not 'derive' words from structures as lexicalization does. For instance, whereas lexicalization squishes complex syntactic structure into a single word, we might suppose instead that all words are unstructured, but that there are nomic relations, underwritten by high-level psychological laws, that hold between various structures and potential words in the lexicon. For instance, there might be another law or collection of laws that allows a (linguistically primitive) meaning to be fit onto an unstructured word only if certain correlated structures are grammatical. Here the correlation between words and structures is given by the nomic relation. So the law in question might say, e.g., that *shelve* as in (s2) is a possible word only if the relevant correlated structure of *put on a shelf* is grammatical, and likewise, **shelve* as in (8) is possible only if the relevant correlated structure of **put a shelf* is grammatical. Since *put on a shelf* is grammatical and **put a shelf* is not, the theory produces the correct predictions about *shelve* and **shelve*.

At first glance, this sort of 'nomic connection' theory has three desirable features. First, the existence of nomic connections between words and structures allows for all sorts of explanatory theories to be generated and tested regarding the existence of impossible words. So such a theory would not be put in the uncomfortable position of predicting that any old words could appear in the lexicon. Second, the nomic connection theory appears to be linguistically equivalent to any particular rival linguistic theory. That is, it looks like the linguistic details of the nomic connection theory can be taken directly from a linguistic theory. This is because the nomic connection theory is at heart a kind of psychological interpretation of linguistic theorizing, rather than a contentful linguistic theory itself. A nomic connection view would support a generalization something roughly along the lines of (9)

- (9) If linguistic theory T predicts that X is a legitimate structure to be lexicalized into a word, then the lexicon can contain a primitive word Y that means what X means. Similarly, if T predicts that X is illegitimate, then the lexicon cannot contain a primitive word Y with X's meaning.

Notice that (9) doesn't say that Y *is* the structure X at some deep linguistic level. Rather, (9) only gives the bare lawlike statement that if structure X and word Y have the same meaning, then Y is a possible word if and only if X is a legitimate structure. No further relation (such as that X constitutes Y) is presumed by (9). The apparent linguistic equivalence of a nomic connection theory and any linguistic theory employing lexicalization makes the former theory evaluable without getting bogged down in the messy details of linguistics. We simply wait for linguists to give us the value of T, and then we have our theory. The third nice feature of the nomic connection view is that it allows individual words to be semantically and syntactically unstructured (although of course there might be constraints on how the meaning is fit onto the verb, e.g. agents being subjects, etc.). Thus, the nomic connection theory acts as a flexible theory that can make the correct predictions about impossible words and yet deny that there is any structure present in words. So it might look as if we could use the nomic connection theory as a plausible alternative to lexicalization, and thus deny the inference from (1) – (3) to (4) of the Impossible Word Argument. The denial of the inference would go as follows. Assume that (3) is correct: the linguistic theory T, which posits all sorts of structure and processes in the lexicon, provides an explanation of why impossible words are impossible. But the success of T is not thereby justified, because the nomic connection theory presents an alternative way of explaining the data at hand, and does not require positing structure and processes in the lexicon. Instead, the nomic connection theory only posits nomic relations between unstructured words and syntactic structures.

Unfortunately, the nomic connection theory may not be quite the panacea it appears to be at first glance. For one thing, not everyone who would like an alternative to the lexicalization process could accept the nomic connection theory. This is because many of the generalizations that the nomic relations will need to explain appear to be statable only with the help of the technical vocabulary of thematic roles and other sub-lexical semantic constituents (e.g., the notions of an agent or a theme, or of a specialized notion of causation or affectedness, etc.). But FL are skeptical about the very use of such notions (e.g., FL 1998, p. 275 – 278, Fodor 1998, pp.

57 – 64). Thus, they could not allow these notions to be part of laws of linguistic psychology without recanting some of their other views.

In addition to being unavailable to the likes of FL, the nomic connection theory suffers from some more general problems. But to defend this last claim, it will be useful to recall the dialectic that took us to this point. We have been exploring whether we can resist the inference to (4) – the claim that we are justified in positing structure and processes in the lexicon. Many standard linguistic theories account for the phenomenon of impossible words by positing structure and processes in the lexicon. The present strategy has been to look for some method that is just as good at accounting for impossible words, but which does not force us to posit any structure or processes in the lexicon. (Although the nomic connection theory posits nomic structure within the lexicon, I will assume that since the nomic connections aren't linguistic structures, they aren't of the right sort to underwrite claims about the structure of words which would in turn support claims about the structure of concepts. If they are, then we needn't worry about defeating this option.) Notice, however, that it will not be enough to simply show that the nomic connection theory is *possibly* true. For it to be a viable alternative that we might adopt over the lexicalization theory, we need to make a case that nomic connections provide a *plausible* option.¹⁹ So to undermine the nomic connection theory, we would need to show that no considerations whatsoever could possibly make this theory possible. But given our relatively poor understanding of how psychological abilities like language are implemented in the brain, such an undertaking would not really be realistic at this point. What I will do instead is argue that from a linguistic perspective, the nomic connection view should not be adopted. Of course, even if my arguments are completely successful, some further considerations from outside of linguistics may nonetheless show that the view is overall a good one. My goal here is only to put the burden of proof on the proponent of the nomic connection view, and to show that adopting such a view will have to overcome some linguistic costs.

My first objection to the nomic connection theory is that it is an unmotivated addition to a linguistic theory. In particular, I suggest that contrary to appearances the nomic connection theory is not necessarily linguistically equivalent to a more typical linguistic theory. To argue

¹⁹ To take an extreme example consider a theory which says that there is only one word in the lexicon, but there's also a whole lot more syntax and phonology than we thought. Such a theory could be designed to show that it's all this extra syntax and phonology present that makes strings of this single word mean different things and sound different ways. We can and should grant that such a theory is perfectly possible. But we can and should (and, I presume, do) ignore such theories because they are simply too implausible to be taken seriously.

for this claim, I will compare the nomic connection theory to Hale and Keyser's view, discussed above. In showing that the former view is preferable, I will not of course have shown that the nomic connection view is inferior to every linguistic theory. However, the strategy I employ is generalizable to a large number of theories, and if the nomic connection view is supposed to be an all-purpose replacement for theories that posit structure and processes in the lexicon, it needs to be better than every linguistic theory.

The general reason for thinking that the nomic connection view is inferior to some linguistic theories is that it adds (at least) one new primitive element to the overall linguistic theory, namely a lawlike relation holding between structures and words. In contrast, Hale and Keyser's theory works by reusing mechanisms that have already been independently established within the linguistic theory. In particular, their process of lexicalization is really the process of 'incorporation', which is well-known from the syntax literature (e.g., Baker 1988, 1997). As it is originally introduced, incorporation is a syntactic process occurring when multiple words or morphemes combine to form larger, more complex words. Its syntactic effects are most clearly seen in languages other than English, such as the Bantu language Chichewa. Consider, for instance, the following two Chichewa sentences both of which express that the girl made the waterpot fall. (The examples are drawn from Baker 1988.)

- (10) a. Mtsikana a-na-chit-its-a kuti mtsuko u-gw-e
 girl do cause that waterpot fall
- b. Mtsikana a-na-gw-ets-a mtsuko
 girl fall cause waterpot

Notice that in (10a), there are two separate verbs, *its* ('cause') and *gw* ('fall'). In (10b), however, *gw* has been incorporated into the verbal structure for *its*. (The extra material surrounding *its* – *ana-* and *-a* – is inflectional morphology, similar in kind to the English past tense morpheme *-ed*, as in *kicked* or *kissed*. Also, *its* changes to *ets* for phonological reasons that are not pertinent here.) Crucially, when *gw* is incorporated into *its*, the result is a single complex verbal structure that quite literally contains both verbal elements as proper parts. There is, in particular, absolutely no reason to suppose that since *its* and *gw* can incorporate, there must also be a primitive verbal structure *anagwetsa*, related to the complex incorporation of *its* and *gw* by some

kind of nomic connection. So when incorporation applies, the constituents of the input are contained as proper parts of the output, contrary to the way that the nomic connection theory operates. Thus, when the process of lexicalization is identified with incorporation, as Hale and Keyser hypothesize, a single, independently motivated process is responsible for compressing various structures into single words.²⁰ Moreover, by identifying lexicalization with incorporation, Hale and Keyser's theory thereby predicts that the process of lexicalization will share all of the limitations inherent in the process of incorporation. This prediction turns out to have some rather pleasant consequences, because incorporation cannot take a structure like *put on a shelf* as input and yield **put-a-shelf on* as output. Thus, the identification of lexicalization and incorporation not only produces a theory of the lexicon whereby structures are squished into words by means of an independently motivated syntactic mechanism, but incorporation is also able to explain why certain relevant structures are ungrammatical. The fact that one independently motivated mechanism does these two further jobs in the lexicon is surely evidence in favor of the overall theory.

On the other hand, recall that the nomic connection theory works by taking the structures that a linguistic theory produces and associating them with individual, unstructured words. In the case of the present linguistic theory, the nomic connection theory has absolutely no work to do. The process of incorporation has already created a structure that will behave like an individual word, so the nomic connection is an otiose mechanism insofar as it associates the original (complex) word-structure with a second (simple) structure that also behaves like an individual word. Since incorporation gives us an item that will act like a word, why should we insist that the item in question is unacceptable and must be replaced with something else that acts like a word? In order to endorse the nomic connection view, one will need an answer to this question. Unfortunately, I know of no plausible and non-question-begging answer. Thus, it appears that the nomic connection theory can reproduce the results of linguistic theorizing only in the extensional sense that, with suitable tinkering, it can make exactly the same linguistic predictions as a theory that directly uses incorporation or some such similar mechanism. But it

²⁰ Two comments. First, it is worth mentioning in passing that this component of Hale and Keyser's theory has recently received further confirmation as part of the design structure of human languages. Lisa Travis (Travis 2000) has shown that this general framework is capable of accounting for a number of striking features of Western Polynesian languages like Malagasy and Tagalog. Secondly, for various reasons, Hale and Keyser hold that the verbs and prepositions used in their discussions are not the ordinary ones that are phonologically realized, but are similar lexical ones with no phonological properties at all. Their assumption is in line with much current research into other languages, where similar morphemes with phonological content are common (e.g., Travis 2000, Baker 1988, 1997).

does so by adopting all the machinery of a linguistic theory, and then positing nomic connections on top of that. By Occam's Razor, this last unnecessary step makes the nomic connection view less desirable than the more economical theory that lets incorporation do the work directly.

A second problem with the nomic connection theory concerns the motivation for the theory. We saw in (9) that the nomic connection view works by taking a structure that the linguistic theory provides, and allowing a word with the meaning of that structure to enter the lexicon. But this move is utterly ad hoc. Why should these nomic relations hold between exactly the right set of pairs of structures and words? Why should the first element of these two-place relations be restricted to grammatical structures only? Furthermore, why are these nomic relations restricted in their first position to only some grammatical structures? Why don't we have single words for other grammatical structures, like sentences or other clauses? Why, for instance, does English not contain a word like *plood*, which has the structure of *the person who said that* (cf. **Bill is plood Sarah is silly*)? Unless there is some deeper form of explanation of the way the nomic relations work, the theory's laws ensuring that (im)possible words will turn out on the theory to be (im)possible are simply stipulated solutions to a problem. (The incorporation view, by the way, has a very easy answer to this question: in the case of *plood*, the movement involved is not permitted by incorporation, and so it is predicted not to occur.²¹) There is another sense in which the nomic connection theory is ad hoc. It is hard to see why anyone would adopt it for reasons other than a desire to save some favored theory of word-structure. It is unclear what advantages accrue to the theory other than that it makes it possible to hold that words have no structure. Such a position is fine if one already has reason to hold that words have no structure. However, if you are looking to defend the claim that words have no structure, it is unclear how the nomic connection view will be of use.

A final problem with the nomic connection view concerns its focus on explaining impossible words. Although the nomic connection theory can (in an extensional sense) account for impossible words, it is unclear how it could also deal with other phenomena that motivate linguistics to posit structure in words. For instance, natural languages appear to be sensitive to whether the meaning of a transitive verb expresses that the subject somehow created or obtained the object of the verb (cf. e.g., Pinker 1989, Levin 1993). If you want to show some kindness for

²¹ This is *not* to say that there are not serious questions that can be raised for Hale and Keyser's theory. But most (if not all) challenges for a linguistic theory will also be challenges for a nomic connection theory that uses it.

your friend Marsha, you could either bake a cake for her, or you could perform some of her chores, such as brushing her horse. In reporting what you did, you might use the sentences in (11):

- (11) a. I baked a cake for Marsha.
b. I brushed the horses for Marsha.

However, only the first of these sentences has a variant in the ‘double-object construction’, where the preposition has been removed, and the two noun phrases permuted in order:

- (12) a. I baked Marsha a cake.
b. *I brushed Marsha the horses.

Assuming this analysis is correct, a linguistic theory will need to appeal to the fact that some verbs express that their objects are created or obtained. Thus, verbs will need to carry information about some of their semantic properties. Although the nomic connection view might eliminate the need to posit structure and processes in the lexicon in order to explain impossible words, the strategy appears to be too narrow to generalize to cover the other evidence that words are not linguistically primitive (cf. Fodor 1998, ch. 3, esp. pp. 57-64 for some relevant discussion). On a similar note, by focusing on impossible words alone, the nomic connection view also leaves certain regularities unexplained. For example, it is widely known that the kind of causation relevant to causative verbs must be ‘direct’, at least in the sense that the causation does not go through the intentions of another agent. For example, suppose Jane warns Pete to be careful with the vase, and upon hearing this, Pete, who is secretly Jane’s sworn enemy, smashes the vase to bits. In this circumstance, Jane did cause the vase to break by her warning to Pete, but she did not break the vase. Importantly, direct causation does not seem to occur only in causative verbs like *break*. We also find it in e.g. ‘resultative’ constructions such as (13):

- (13) John whipped the cake batter smooth.

In such sentences, the predicate at the end of the sentence applies to the direct object of the verb as a result of the action of the main verb. (13) means that John whipped the cake batter, and the direct result of that whipping was that the cake batter was smooth. (13) would not be true if John

was whipping the cake batter, and the sound of his doing so attracted Sue, who took the batter from him and placed it in a blender and blended it until it was smooth, even though in this case John's whipping would have caused the cake batter to become smooth. The similarity between the types of causation in causative verbs and causative constructions is striking. If causative verbs like *break* are unstructured as the nomic connection view suggests, then why don't causatives and resultatives express different types of causation? There are many types of causation after all, such as Aristotle's four causes and the normal type of causation studied by metaphysicians. It seems very odd that both linguistic constructions (as well as others I have not discussed) employ the same sort of causation. In fact, it is odd that all causative verbs seem to employ the same sort of causation. On the nomic connection view, it would be unsurprising if some causatives employed one kind of causation and other causatives employed other kinds. Why this doesn't happen is a mystery. However, these problems disappear on a more typical linguistic theory. It is natural that causatives and resultatives all express direct causation, because they all contain (in the syntax or in the word) the concept of direct causation as a part of their meanings. None of these constructions express some other kind of causation because there isn't any other causal concept that is capable of being a deep part of the grammar in the way that direct causation is.²² Moreover, unlike the nomic connection view, this theory can explain why so many languages add a morpheme to their causative verbs that appears to express causation (cf (10)). The answer is of course that the morpheme really does express causation. It is utterly unclear what the nomic connection view would say about this phenomenon.

In sum, the process of lexicalization cannot just be dismissed, as FL suggest. It plays a crucial role in any theory that allows word-formation processes to occur. Attempting to avoid lexicalization by appealing to nomic connections is little more than an ad-hoc band-aid maneuver. The nomic connection view packs unwarranted complexity into a linguistic theory, and in return it fails to do much of the explanatory work that a more typical theory of the lexicon does with comparative ease. On balance, then, it appears that lexicalization is here to stay, contrary to FL's claims.

²² There are some notable differences between languages like English in which direct causation is not overtly realized as a morpheme and other languages where it is. Cf. Comrie 1986 for discussion.

4 Objection three: Put ‘everything in the syntax’ (EIS)

So far, we have seen two failed strategies for undermining the Impossible Word Argument. The first strategy failed because it was not sensitive to the details of linguistic theories, and thus made empirically false predictions, among other things. The strongest version of the second strategy was sensitive to the details of linguistic theories, but added some extraneous machinery that made the approach ad hoc and unnecessarily cumbersome. A natural thought, then, would be to see whether there could be a strategy that was both sensitive to linguistic theories and which did not add any new material to the theory. Given what we have seen so far, such a strategy might have the best chance of undermining the Impossible Word Argument without encountering new difficulties. The final type of objection I will consider is just such a strategy. We will begin by supposing that T is a good linguistic theory, and that T’s ability to predict the impossibility of certain words is genuinely explanatory. Moreover, unlike the previous objection, we also assume with T that lexicalization is needed in the linguistic theory. Despite all these concessions, the objection still contends that we needn’t posit processes and structure in the lexicon. Rather, it maintains that any processes and structures T posits in the lexicon might instead be posited in the overt syntax of natural language. According to this proposal – which I will call the ‘Everything in the syntax’ proposal, or ‘EIS’ for short – the lexicon is purely and simply the repository of the primitive elements of language. Thus, although the lexicon contains simple elements like *cause* and intransitive *break* (assuming these elements are indeed primitive), a word like transitive *break* is formed only after *cause* and intransitive *break* are entered into the syntax of the sentence. Once these elements are in a syntactic structure, they are then organized into transitive *break* in the syntax by the very processes that T holds occur in the lexicon.

Initially, it is very hard to find fault with EIS. Unlike the first two strategies, the proposal employs all and only the mechanisms of a successful linguistic theory. (Of course, EIS will not be available to the likes of Fodor and Lepore, who challenge the viability of the very mechanisms used in linguistic theorizing; e.g., FL 1999, Fodor 1998, pp. 58ff. However, even they can use EIS with any linguistic theory that does not contain the mechanisms they reject.) In fact, EIS seems to be merely a bit of relabeling of what counts as syntactic and what counts as lexical; what could be wrong with that? Furthermore, if EIS and the linguistic theory from

which it is drawn are otherwise equal, EIS should be favored. EIS is preferable because it is a simpler and more restrictive theory. The linguistic theory T, we may assume, posits structure and processes in the lexicon; but it will also surely posit structure and processes in the syntax (and semantics and morphology – I will leave these additional areas implicit). Thus, T allows for structure and processes in both the lexical and generative components of grammar, and so it is thereby a richer and more expressive theory. But, as a general principle of the philosophy of science, if two theories can account for the empirical phenomena equally well, then the weaker and more restrictive one should be favored.

Let us now turn a critical eye to EIS. The first thing to notice is that it seems almost *too* good. If EIS is right, then why haven't linguists already adopted it? Why do many linguists retain their more unwieldy theories according to which structure-forming processes are found in both the lexicon and the syntax? After all, linguists are the last people to be fooled into thinking that what appears to be a word is in fact a syntactically complex bit of structure. Moreover, it is a very common bit of linguistic methodology to look for simpler and more restrictive theories. Offhand, EIS looks like the sort of position a great many linguists would have naturally gravitated towards by now. I think that there are some good methodological reasons why linguists haven't adopted EIS, but my explanation will require some setting up and a few details.

A crucial feature of EIS is that by itself, it does not recover every explanatory advantage of a theory that divides processes between the lexicon and the syntax. We can see this by returning to Hale and Keyser's example of the verb *shelve*, as in *David shelved the books* (cf. Hale and Keyser 1993, 1999, Fodor and Lepore 1999). As I mentioned above, Hale and Keyser treat the words like *shelve* as having the same structure as larger syntactic expressions. In particular, they treat the individual word *shelve* as having all the internal structure of the verb phrase *put on a shelf*. Their theory also explains why there is no corresponding verb **shelve*, as in **David shelved the books on* (with the meaning 'David put the books on a shelf'). Details aside, a crucial part of their explanation is that the prepositional phrase *on a shelf* is lexicalized with the main verb within the lexicon, so that only a verb appears in the syntax. But according to EIS, we could just as well posit the prepositional phrase *on a shelf* in the syntax, instead of buried down in the lexicon. However, the EIS proposal cannot yet be complete. For it is well-known that in the syntax, there cannot be multiple independent fillers of the same thematic position (e.g., Chomsky 1981, 1986 calls this the Theta-criterion). For example, if Julie opened

the box using both a hammer and a saw, we would not express this by saying **Julie opened the box with a hammer with a saw*; rather, we would have to use something like *Julie opened the box with a hammer and (with) a saw*. In the present case, this generalization entails that there cannot be two separate prepositional phrases both of which specify where the book was put, as we see in (14):

- (14) a. David shelved the files on the bookcase;
b. *David put the files on a shelf on the bookcase.

(14b) is ungrammatical when both prepositional phrases modify the verb, yielding the meaning ‘David put the files on a shelf and on the bookcase’. (The reading of (14b) where *on the bookcase* modifies *shelf* is irrelevant here.) According to EIS so far, the grammatical sentence in (14a) should have the same structure as the ungrammatical sentence in (14b). So EIS will need to be augmented to avoid predicting that (14a) is ungrammatical, or that (14b) is grammatical. No doubt this can be done. We simply require that lexicalization operates so that *on a shelf* is absorbed into *put* in such a way that the location of the action of the verb remains unspecified. In other words, the process of lexicalization will solve the problem *as long as it works just as though it was a process of word-formation, taking place in the lexicon before the word is inserted into the syntax*. So far, so good. Lexicalization does its job in the syntax by operating before any other syntactic mechanisms do.

Were we to explore EIS further, we would probably handle other putatively lexical processes similarly. For instance, the mechanism that allows some verbs to have a middle form is often thought to occur in the lexicon. We have, e.g., transitive verbs like *sink*, as in *The captain sunk the fiberglass boats*; but we also have intransitive middle verbs like *sink* as in:

- (15) Fiberglass boats sink easily.

One of the reasons that the formation of middles is often thought to occur in the lexicon (at least for English) is that it seems to behave differently from other processes that are thought to occur in the syntax. For instance, notice that when *sink* is turned into a middle, the subject (= agent) position of the verb is no longer present. The absence of the agent of *sink* is robust, insofar as agency cannot be captured by a *by*- phrase, nor can *sink* supply a subject of a purpose clause:

- (16) a. *The fiberglass boats sink easily by crooks.
 b. *The fiberglass boats sink easily to collect the insurance.

(16b) is ungrammatical on the relevant reading ‘It is easy for one to sink the fiberglass boats in order that one may collect the insurance on them’. Moreover, speakers experience difficulty with (16b) at exactly the point when they begin to interpret *to collect the insurance* as a purpose clause in need of a subject somewhere in the sentence (cf. Tanenhaus et al. 1993). Contrast this behavior of middle-formation with the process of passivization of verbs, which we may assume is a syntactic process. When *sink* is passivized, the subject position of this verb appears to remain in the clause, at least in the sense that it can support *by-* phrases and purpose clauses:

- (17) a. The fiberglass boats were sunk by crooks.
 b. The fiberglass boats were sunk to collect the insurance.

(17b) means something like ‘The fiberglass boats were sunk by X in order that X could collect the insurance’. Thus, although the subject position of *sink* is not overtly at hand, it still retains some residual presence (cf. Baker, Johnson and Roberts 1989). We can see that the subject position of *sink* supplies the subject of the purpose clause by noting that (17b) cannot mean ‘The fiberglass boats were sunk by X in order that Y could collect the insurance on them’, where $X \neq Y$. In the case of middle formation, EIS will export this process from the lexicon and place it in the syntax. Despite this exportation, note that the process of middle formation behaves unlike other syntactic processes: when transitive *sink* is turned into a passive, the thematic position for the subject remains, whereas when *sink* is turned into a middle, that position disappears, just as if it were never present in the syntax in the first place.

(The preceding examples suggest that it may not be quite correct to say that EIS captures absolutely every generalization from linguistics. Now that middle formation takes place in the generative syntactic component of language, there may be a level of syntax at which verbs like *sink* and *cut* are represented as transitive verbs, and then at some later development in the syntax, their subject position is eliminated as per (16). But if such a scenario is realized, then such principles of language as the Projection Principle will be false. The Projection Principle says that the selectional requirements of a lexical item are present at every level of syntactic

representation (e.g., Chomsky 1981, pp. 29 – 38, Baker 1988, pp 49 – 51). But if a transitive verb is inserted into the syntax before the process of middle formation has occurred, then there will be an initial level of syntax where something that later becomes an intransitive verb is represented as a transitive verb, requiring two arguments. So at an early stage of the syntax, the verb needs both a subject and an object, and at a later stage, it needs only an object in the subject position. Of course, EIS can simply reformulate the Projection Principle so that it applies everywhere except for that initial level of syntax that looks like a level of word-formation. However, such a move would not be welcomed by someone who thought that it was not the job of syntax is to organize the various elements of language into some form suitable for phonological and semantic interpretation, and not to alter the fundamental structure of clauses determined by the various lexical items.)

Beyond the two cases we have just surveyed, there will be other processes that we will remove from the lexicon and put in the syntax. Call such processes ‘L-processes’. L-processes are distinctive in that they have a property I will call ‘Locality’: the total contribution of the output of the process seems to be highly localized around the ‘visible’ (i.e., phonologically pronounced) output of the process, leaving little evidence of any deleted material. For example, if *shelve* has the structure PUT ON A SHELF, there is no evidence in the syntax of this additional structure (and so the locational thematic role can be specified by another prepositional phrase) and, unlike the passive, when the middle voice of *sink* is formed, the agent position of the relation *X sink Y* is unavailable for certain operations elsewhere in the sentence. L-processes can be contrasted with many others in syntax which do not display Locality. For example, consider how an English question word like *who* or *what* appears at the front of a sentence, as in *who did Mary kick?* (cf. *Mary kicked John*). Even in such a simple sentence, we see that when *who* switches from the object position to the front of the clause, it also produces a change in the location of the past tense morpheme, leaving it at the front, attached to the dummy verb *do* (cf. **Who Mary kicked?*). Furthermore, when a question word is ‘moved’ out of an embedded sentence, there must be an unoccupied position at the front of each clause between the *wh*-word and the position it is associated with.²³ If that position is already occupied by another question word, ungrammaticality results.

²³ Strictly speaking, this is not quite right, as can be seen by the counterexample of *Who do you think that the man who knows where the treasure is will befriend?* However, above crude characterization will suffice for now.

- (18) a. Who did Peter know Mary gave some candy to?
 b. *Who did Peter know what Mary gave to?

Many more examples of Local and non-Local processes could be given. A survey of these processes shows that according to EIS there is a bifurcation of processes: some processes have the feature of Locality, and others do not. Further research has shown that (what I am calling) L-processes also tend to share other features, such as occurring before and independently of any non-L-processes. This is of course exactly what would be expected of a process that operates just on words, before they are entered into the syntax. However, merely noting that EIS posits two distinct kinds of processes in the syntax is not inherently problematic. There is nothing intrinsically wrong with stipulating that some syntactic processes share a collection of features which the other syntactic processes do not have.

Even though EIS's bifurcation of syntactic processes is innocent in principle, we can begin to see an important methodological discrepancy between EIS and much actual linguistic practice. For many linguists, the fact that a process is an L-process is itself (*ceteris paribus*) a *sufficient* reason to posit the process in the lexicon instead of in the syntax. Thus, it is *because* a process has features like Locality that linguists posit them in the lexicon in the first place. In contrast, EIS must hold that being an L-process is not sufficient for a process to be lexical, since it places L-processes in the syntax. In fact, the main difference between EIS and linguistics appears to be no more than how the lexicon is defined. On the one hand, linguists define the lexicon functionally: it is where L-processes occur, which is why adopting EIS would not produce a more economical theory. ON the other hand, EIS defines the lexicon to be the storehouse of linguistic primitives. So perhaps the difference between EIS and more typical linguistic theories is merely terminological, and so one should be able to freely decide whether to adopt EIS or a more typical linguistic theory. Perhaps by adopting EIS, one can clearly avoid the conclusion of the Impossible Word Argument, which says that we are justified in supposing that there is structure and processes in the lexicon. However, EIS offers only a hollow victory. EIS avoids the Impossible Word Argument by defining the lexicon so that all the problematic structures are to be found in the syntax, but the problematic structures were the primary issue in the first place! We began with the question of whether items like transitive *break* are structurally complex. Linguistic theory and EIS agree that it is, but EIS insists that *break* is actually a

syntactically complex structure composed of several primitive expressions, which individually mean something like *cause*, and *be broken*.²⁴ EIS's maneuver of putting all the structure into the syntax does nothing to reduce the number of semantic primitives contained in transitive *break*. In fact, EIS agrees with the Impossible Word Argument in spirit, since it also maintains that items like *break*, which at least appear to be individual words, are actually complex constructions. So although EIS might win a small local battle, it ends up conceding the war.

As another way to view this problem, consider whether EIS can offer any new help for the atomist about concepts. It appears that it can't. On the one hand, if by the term *word*, we mean an element of the lexicon, EIS does allow us to conclude that all words are unstructured, and so by the Isomorphism Assumption (which says that the structure of a word is isomorphic to the structure of the concept it expresses) we can conclude that all concepts expressed by words are unstructured, too. But in this case, we have left out all of the interesting cases, such as *break*. Simply redefining the hard cases like *break* so that the Isomorphism Assumption doesn't apply to them is unlikely to answer any interesting questions. If, on the other hand, by the term *word*, we mean something that will include the controversial items like *break*, then regardless of whether (4) is true, the scope of the Isomorphism Assumption encompasses more than just the elements of the lexicon, and this widened scope contains precisely the expressions that matter.

In sum, EIS does appear to provide a viable form of linguistic theory. However, it does so by rearranging things so that the area of language that we are interested in is now spread across both the lexicon and parts of the generative component of grammar. This rearrangement does not address any of the questions about the complexity of the items of interest, although it does relabel the items as syntactic structures instead of words. Similarly, EIS fails to provide any help for understanding the structure of various concepts by studying the words (or structures) that express them.

5 The Isomorphism Assumption

²⁴ In the particular case of causatives, EIS is actually in line with much contemporary thinking (e.g., Hale and Keyser 1987, Kratzer 1996, Travis 2000). As I have discussed, though, we are using a simplified theory of causatives to make a methodological point that applies to many further cases which would be too complex to easily present.

Throughout this paper, I have not really questioned the Isomorphism Assumption, which links word-structure to conceptual structure. I will close with a few words about it. On the one hand, something about the Isomorphism Assumption is surely correct: no doubt there is some useful sense in which our words express concepts, and that at least some structurally complex words tell us something important about the concepts they express. But at the same time, there is much about the Isomorphism Assumption that we can criticize. The first thing to notice is that it presupposes some prior individuation of words. That is, you can't state what the Isomorphism Assumption even says until you say what counts as a word. But as we have seen, saying what counts as a word is a highly non-trivial matter. Are words just the elements in the lexicon? If so, then by what principle should we decide which elements are in the lexicon, and which are not? If not, then what non-lexical elements should count as words, and what lexical elements should count as non-words? (For further discussion see also e.g., Di Scullio and Williams 1987, Pinker 1999). And once we've figured out what a word is, we will have to face the harder project of saying how to individuate concepts. And once we've done that, we'll need to give some reason for supposing that words and concepts (however we've individuated them) are so similar that structure in one predicts a corresponding bit of structure in the other. Since I don't have much of a sense about what should count as a word or a concept in this domain – or what should count as a bit of structure in a word or concept, for that matter – I confess that I do not know what the Isomorphism Assumption says, much less whether it is true.

Let me leave aside some of these worries about the Isomorphism Assumption, and assume that some natural characterization of it will be found. Even so, why should we suppose that this characterization will be true? A glance at linguistic methodology suggests that it is an open question whether there is structure in a concept that is not present in the corresponding word, or vice versa. On the one hand, we could easily have structure in a concept that was not present in the corresponding word. Suppose for instance, that concepts turn out to be the kind of highly structured entities that prototype theories postulate (e.g., Medin 1989, Smith 1995). Much of the structure of the prototype will not be linguistically relevant, and indeed may simply not be part of the structure of language. I assume that if a bit of conceptual structure has no syntactic, morphological or phonological effects, and no compositional effects on the way that the concept organizes with other linguistically realized concepts to form the meanings of sentences, then that is good evidence that the structure is not part of language. Whether or not this assumption is true

in the long run, it is at least a viable possibility, which is all that we need to put the Isomorphism Assumption in question.

On the other hand, there could also be linguistic structure in a word that is not present in the corresponding concept. In one sense this is almost trivially true, since words have phonological structure, and concepts presumably don't. But the linguistic lexicon may even contain semantically relevant information that is restricted to the language processing systems, and which serves no purpose in a psychological theory of concepts. For instance, we've seen that a verb's thematic structure is an important grammatical property that is determined by the semantic properties of the verb's meaning. But this thematic structure may not be present in the structure of the concepts these verbs express. For all we know, thematic structure may be restricted to the language faculty. For instance, we saw above that the formation of middles out of a transitive verb (like *sink*) hides the agent of the verb from the syntax. In order for the process of middle formation to work, the process must be able to identify the agent in a verb like transitive *sink*, or perhaps the agent in a concept like X SINK Y. (Middle formation is much harder for verbs in which no agentivity is present; e.g., **Peter resembles easily*.) Further research has suggested many reasons for supposing that agentivity is somehow part of the linguistic structure of a verb like *sink*. But a linguist who posits this structural property in the word *sink* is not thereby committed to claiming that the concept of sinking contains a similar bit of structure. The concept of sinking may well be an unstructured unit of meaning, even though our linguistic abilities recognizes some linguistically relevant properties in it. The thematic structure of the verb may have been encoded in the lexical entry for *sink* when this verb was acquired, at which time the concept of sinking was linked to the verb. In fact, this may be a quite general phenomenon: as part of how we fit meanings onto words, our linguistic ability may screen the meaning (i.e., the concept) for various properties. If the concept has some of these properties, this information may be recorded as part of the structure of the word, and this structure may serve to partially define the word's grammatical role in the language. (The total interaction between syntax and semantics at the acquisition stage is bound to be substantially more complex than this, however, cf. e.g., Gleitman 1990.)

The above speculations suggest that the Isomorphism Assumption could be false. Without something like the Isomorphism Assumption, it is unclear what connection there is between word-structure and conceptual structure. At the same time, there is undoubtedly some

relevant sense in which our words express concepts; the tricky part is to develop a clear and well-justified theory of what that connection is. Even some of the people I have criticized at length are in agreement about this. Fodor, for instance, admits that ‘Getting clear on the word-concept relation is no small matter’ (Fodor 1995, p. 34). More importantly, recent work by Robert Matthews casts doubt on the general question of whether our propositional attitudes have a structure that is mirrored in any interesting way by our natural language attitude reports (Matthews 1994, 2002). Much of Matthews’ positive view, a measurement-theoretic account of propositional attitudes, could be co-opted to form a general architecture for thinking about concepts and the words that express them.

Whatever the final status of the Isomorphism Assumption, it may well be useful as a simplifying assumption in the various developing areas of inquiry, as a ladder to eventually be kicked away. Then again, it may not. But by getting clearer on the nature of language and of concepts, we can hope to get clearer on just where (if ever) such simplifying assumptions are useful, and more importantly, what more accurate theories to replace them might look like.

6 Conclusion

In the course of this paper, we have seen that there is surprisingly little wiggle room for avoiding the linguist’s conclusion that many of our words are structured, and that there are processes in the lexicon that produce these structures or otherwise operate on them. One can’t maintain, for instance, that such theories will never explain impossible words on the grounds that any word can enter the language as a primitive. Such a theory results in methodological impotence and empirical falsity. Similarly, one can’t do away with the process of lexicalization. Such a process is a crucial component of linguistic theorizing, and the attempt to replace it with a nomic connection view shows no promise on various fronts. Likewise, one can’t maintain that words are unstructured simply by redefining the lexicon so that everything happens in the syntax. Such a strategy forces us to use new terms instead of the old ones like *word* and *lexicon*, but it does not solve any of the original problems. Finally, we saw that a fundamental assumption linking theories of words to theories of concepts is badly in need of clarification and justification. This presents us with an interesting challenge, since our words surely express concepts.

The general morals of this paper are perhaps pretty familiar to many, at least in outline. Human languages are complicated for all sorts of different yet closely interrelated reasons, and this fact extends to the realm of individual words. Many of our ordinary words are structured and there are processes that operate on them. The complexity of theories of words (and other parts of language) is so great that regardless of the true nature of concepts, the precise relationship between linguistic theories and theories of cognition is bound to be extremely complicated. By examining a few of the linguistic details that should be relevant to such a theory, I hope to have made a small contribution to our getting clear on this relation. Once the word-concept relation is made clearer, we may be able to address more rigorously some of the bold claims made at the beginning of this paper. Until then, we are not in a position to make justified claims about what word-structure, and linguistic structure more generally, teaches us about concepts.

References

- Baker, M. 1988: *Incorporation*. Chicago: University of Chicago Press.
- Baker, M., Johnson, K., and Roberts, I. 1989: Passive arguments raised. *Linguistic Inquiry*, 20, 219-52.
- Bittner, M and Hale, K. 1996: Ergativity: toward a theory of a heterogeneous class, *Linguistic Inquiry* 27, 531 – 604.
- Carter, R. 1976: Some constraints on possible words. *Semantikos*, 1, 27-66.
- Chomsky, N. 1981: *Lectures on Government and Binding*. Dordrecht: Foris.
- Comrie, B. 1985: Causative verb formation and other verb-deriving morphology. In T. Shopen (ed.), *Language Typology and Syntactic Description*. Vol. III. Cambridge: CUP.
- Di Scullio A. and Williams, E. 1987: *On the Definition of Word*. Cambridge, MA: MIT Press.
- Dowty, D. 1979: *Word Meaning and Montague Grammar*. Dordrecht: D. Reidel.
- Dowty, D. 1989: On the semantic content of notion of ‘thematic role’. In G. Chierchia, B. Partee and Turner, R. *Properties, Types and Meaning*. Volume II. Dordrecht: Kluwer.
- Dowty, D. 1991: Thematic proto-roles and argument selection. *Language*, 67, 547-619.
- Fodor, J. 1995: Review of Christopher Peacocke’s *A Study of Concepts*. *The London Review of Books*. April 20. Reprinted in J. Fodor 1998, *In Critical Condition*. Cambridge, MA: MIT Press.
- Fodor, J. 1998: *Concepts*. Oxford: OUP.
- Fodor, J. and Lepore, E. 1999: Impossible words? *Linguistic Inquiry*, 30, 445-53.
- Fodor, J. and Lepore, E. 1998: The emptiness of the lexicon: reflections on James Pustejovsky’s *The Generative Lexicon*. *Linguistic Inquiry*, 29, 269—288.

- Fodor, J. and Lepore, E. 1992: *Holism: A shopper's guide*. Oxford: Blackwell.
- Gleitman, L. 1990: The structural sources of verb meanings. *Language Acquisition*, 1, 3-55.
- Grice, P. 1989: *Studies in the Way of Words*, Cambridge, MA: HUP.
- Grimshaw, J. 1990: *Argument Structure*, Cambridge: MIT Press.
- Grimshaw, J. 1997: Locality and extended projection. In P. Coopmans, M. Everaert and J. Grimshaw (eds.), *Lexical Specification and Insertion*, Hillsdale, NJ: Lawrence Erlbaum.
- Gropen, J., Pinker, S., Hollander, M. and Goldberg, R. 1991: Syntax and semantics in the acquisition of locative verbs. *Journal of Child Language*, 18, 115-151.
- Gropen, J., Pinker, S., Hollander, M. and Goldberg, R. 1991a: Affectedness and direct objects: the role of lexical semantics in the acquisition of verb argument structure. *Cognition*, 41, 153-195.
- Hale, K. and Keyser, S. 1999: A response to Fodor and Lepore, 'Impossible words?'. *Linguistic Inquiry*, 30, 453-66.
- Hale, K. and Keyser, S. 1997: On the complex nature of simple predicators. In A. Alsina, J. Bresnan, and P. Sells (eds.), *Complex Predicates*. Stanford: CSLI Publications.
- Hale, K. and Keyser, S. 1993: On argument structure and the lexical expression of syntactic relations. In K. Hale and S. Keyser (eds.), *The View from Building 20*, Cambridge, MA: MIT Press.
- Hale, K. and Keyser, S. 1987: A view from the middle. *Lexicon Project Working Papers*, 10, Cambridge, MA: Center for Cognitive Science, MIT.
- Jackendoff, R. 1990: *Semantic Structures*. Cambridge, MA: MIT Press.
- Margolis, E. and Laurence, S. 1999: Concepts and cognitive science. In Laurence and Margolis (eds.), *Concepts: Core Readings*. Cambridge, MA: MIT Press.
- Matthews, R. 1994: The measure of mind. *Mind*, 103, 131-146.
- Matthews, R. 2002: Logical form and the relational conception of belief. In G. Preyer and G. Peter (eds.) *Logical Form and Language*. Oxford: Clarendon Press.
- McCawley, J. 1968: Lexical insertion in a transformational grammar without deep structure, in B. Darden, C. Bailey and A. Davison (eds.) *Fourth regional meeting of the Chicago Linguistics Society*. Chicago: University of Chicago Press.
- Medin, D. 1989: Concepts and conceptual structure. *American Psychologist*, 44, 1469 – 1481.
- Miller, G. 1996: *The Science of Words*. New York: Scientific American Library.
- Pietroski, P. 1998: Actions, adjuncts, and agency. *Mind*, 107, 73-111.
- Pietroski, P. 2000: *Causing actions*. Oxford: OUP.
- Pinker, S. 1989: *Learnability and Cognition*. Cambridge, MA: MIT Press.
- Pinker, S. 1999: *Words and Rules*. New York: Basic Books.
- Smith, E. 1995: Concepts and categorization. In E. Smith and D. Osherson (eds.) *Thinking: An Invitation to Cognitive Science*. Cambridge, MA: MIT Press.
- Tanenhaus, M., Boland, J., Mauener, G. and Carlson, G. 1993: More on combinatory lexical information: thematic structure in parsing and interpretation. In G. Altmann and R. Shillcock (eds.) *Cognitive Models of Speech Processing*. Hillsdale, NJ: Lawrence Erlbaum.
- Travis, L. 1984: *Parameters and Effects of Word Order Variation*. Ph.D. diss. MIT: Cambridge.
- Travis, L. 2000: Event structure in syntax. In C. Tenny and J. Pustejovsky (eds.) *Events as Grammatical Objects*. Stanford: CSLI Publications.