

**Morphemes matter; the continuing case against lexical decomposition
(Or: Please don't play that again, Sam)***

**Jerry Fodor and Ernie Lepore
Center for Cognitive Science
Rutgers University**

Introduction

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" is recurrent in linguistics; it's the defining thesis of what is often called "lexical semantics," and it unites the generative and interpretive traditions of grammatical analysis. Recently, Hale and Keyser (1993) have provided a budget of sophisticated and persuasive arguments for the claim that `denominal' verbs are typically derived from phrases containing the corresponding nouns: `singvtr' is supposed to come from something like DO A SONG; `saddlevtr' is supposed to come from something like PUT A SADDLE ON; `shelvevtr' is supposed to come from something like PUT ON A SHELF, and so forth.¹ We think these are among the most persuasive arguments for lexical decomposition in the linguistics literature. Still, this paper is going to claim that they are finally unconvincing. In Part 1, we will show that there are quite serious arguments of a familiar kind against the decompositional analyses that Hale and Keyser (henceforth, HK) propose; in Part 2 we'll show that the arguments that HK offer in favor of their analyses are flawed.

A word on method. One way to argue against a putative lexical decomposition is to show that the decomposed word is not, *strictu dictu*, synonymous with the corresponding decomposing phrase; usually one does this by describing (actual or possible) situations to which the one but not the other would apply. If, for example, there are bona fide cases that would count as killing but not as causing to die, or vice versa, then that presumably shows that CAUSE TO DIE can't be the semantic level representation (hence the derivational source) of `kill'.

The status of this kind of argument is, however, moot. For better or worse, many lexical semanticists are unprepared to specify exactly what kind of meaning equivalence is supposed to hold between a lexical item and its phrasal source.² Moreover, the modal intuitions on which claims for synonymy often rest are notoriously insecure. We do not wish, in this paper, to rely on any general assumptions about what semantical properties of a phrase its lexicalization preserves. Accordingly, we won't, in what follows, ever use arguments of the form: verb V can't derive from phrase P because ...V... isn't synonymous with ...P...

Well, hardly ever.

Part 1.

In a 1970 paper that argued against the generative analysis of causative verbs, Fodor offered what he claimed is a general reason for denying that simple transitive causatives (like `kill' or `boil') are derived from structures that embed clausal complements (like CAUSE TO DIE or CAUSE TO BOIL). The basic idea was that the latter sorts of structures can exhibit ambiguities of the scope of modifiers (prepositional phrases and the like) which the former characteristically do not. The rule of thumb is that scope respects the surface lexicon, and this militates strongly for the view that the surface lexical inventory is intact at the semantic level of representation. So, (borrowing an example from Geoffrey Pullum) the "on purpose" in (1) can modify either `cause' or `sit'; but (2) has only the reading it was intentional of Smith to sit the students on the floor.

1. Smith caused the students to sit on the floor on purpose.
2. Smith sat the students on the floor on purpose.

This sort of asymmetry between surface causatives and their putative phrasal sources appears to hold quite generally. So it looks as though the derivational analysis of surface causatives should be rejected.

It's unclear to us exactly what HK think about this form of argument. It isn't discussed in (1993), and what they say about it in (1991) is enigmatic. Here is the relevant passage:

[Fodor's arguments]... had to do with the proposal that the simple verb 'kill' was derived from a "deep structure" syntactic representation underlying 'cause to die' -- and the arguments seem correct, for the position he was criticizing. The arguments do not carry over to the position we are entertaining here, however, since the verbs derived by incorporation in the lexicon are themselves input to d-structure. Thus, for example, the verbs 'shelve' and 'saddle', and the like, are lexical items in the true sense, and as such each necessarily involves a single [sic] "event position"...(p.12).

We find this text hard to interpret. For one thing, the main complaint in Fodor (1970) was not that the number of 'event positions' in lexical items generally doesn't match the number of event positions in their putative phrasal sources (though it's quite true that it generally doesn't). Fodor's main objection was that the derivational treatment of causatives incorrectly predicts that sentences like (2) exhibit scope ambiguities analogous to those in (1). Prima facie, HK's comments ignore this problem.

It may be, however, that HK are indeed suggesting a reply: Namely, that the principles that semantically interpret scope relations apply after lexicalization in the course of grammatical derivations. In effect, that would be simply to concede that the semantic interpretation of scope respects the integrity of the surface lexicon. We have no objection to HK conceding this, of course, except that it appears to be ad hoc for them to do so barring some explanation of why scope -- unlike, presumably, the rest of semantics -- should be insensitive to the structures that prelexical syntax is supposed to generate. It matters that morphemes matter, and it needs to be explained. If 'seatvtr' is 'seatvtr' even at abstract levels of grammatical representation, that would make clear why the interpretation of scope never takes it to be CAUSE TO SIT. But if, on the other hand, the abstract levels of grammatical representation neutralize the distinction between 'seatvtr' and 'cause to sit', some explanation is surely required of why scope interpretation is unaware that they do so.

Or it may be that HK are suggesting something really quite radical: Viz., that all principles of semantic interpretation are ipso facto insensitive to prelexical syntactic representations. That proposal would have the virtue of making the analysis of scope not an exception to the rest of semantics, but we can't think of anything else to be said on its behalf. For one thing, if semantics is entirely blind to prelexical syntax, doesn't that rather, sort of, a little bit, suggest that maybe prelexical syntax isn't there? And, for another thing, consider the costs: If 'seatvtr' derives from CAUSE TO SIT, that would at least explain the intuitive semantic relatedness of the corresponding surface forms. But if semantics is insensitive to prelexical representation quite generally, then it is insensitive a fortiori to the fact that 'seatvtr' and 'cause to sit' are both derived from CAUSE TO SIT. The (putative) fact that 'seatvtr' means (or anyhow is close to the meaning of) 'cause to sit' must therefore count as an accident. Connoisseurs of heavy irony will find much here to please them since it is widely advertised as a (maybe the) main attraction of abstract treatments of the lexicon, generative or interpretive, that they allow the grammar to capture intuitions of word-to-phrase semantic relatedness. That we should buy into them because they do so has been a recurrent theme in the literature on lexical semantics from, e.g., Katz and Fodor (1963) to Pustejovsky (1995) inclusive.

So far as we can tell, then, HK offer no reason for doubting that the scope test is a reliable diagnostic of the clausal structure of abstract grammatical (including lexical) representations. For all that's been argued so far, if denominal verbs fail the scope test, that's a pretty good reason for rejecting the derivational analysis of denominal verbs. Our next point is that denominal verbs do indeed fail the scope test (most of the time; more about the exceptions presently).

Consider first the intransitive denominals. According to HK, something like (3) is the underlying source of (4). Notice that (3) exhibits an ambiguity of quantifier scope as between the readings (5) and (6).

3. JOHN DID A SONG TWICE.

4. John sang twice.

5. Some song x: (John sang x twice)

6. Twice (John sang)

In effect, either `twice' can have scope over `John sang', as in (6), or `some song x' can have scope over `twice', as in (5). However, just as one would expect, (4) permits only the reading (6). This is because, to get the reading (5), you need a direct object to quantify over; and (4) doesn't have one. Notice that this is so even though (6) entails that there are (possibly identical) songs, x and y, such that John sang x and John sang y. The fact that the scope of `twice' is unambiguous in (4) therefore suggests that these entailments are not to be explained by assuming that (4) has a complex semantic-level representation in which `John sang x' and `John sang y' are clauses.³

In short, precisely as with the causatives, the facts about scope suggest that surface lexical items (`singvtr') are in place at whatever the level is at which the principles of semantic interpretation apply.

In passing: We assume that HK would take `singvtr', as well as `singitr', to be derived from the corresponding noun; as far as we can tell, their derivational treatment of denominals is supposed to hold across the board. Notice, however, that (6.1) is ambiguous in just the same way as (3). This strongly suggests, contrary to what HK

6.1. John sang a song twice.

seem to say in the passage we quoted, that the difference between (3) and (4) has nothing to do with whether the semantic interpretation of scope is determined pre- or post-lexicalization. (6.1) works like (3) for the simple and obvious reason that both have direct objects for `twice' to quantify into. Correspondingly, (4) is univocal because it lacks a direct object and not because its verb has been lexicalized.

Similar points can be made about the family of denominals sometimes called location verbs. `John shelved a book twice' is ambiguous, just as one should expect if the ambiguity of (3) is a consequence of the fact that `singvtr' has a direct object. Consider, moreover, (7) and (8), the wh-question forms corresponding to (9) and (10) respectively:

7. Where did John put the book on the shelf?

8. Where did John shelve the book?

9. John put the book on the shelf.

10. John shelved the book.

11. where on the shelf did John put the book?

Among the various readings of (7) is (11); but notice that (8) has no corresponding interpretation. Here again the source of the asymmetry seems sufficiently obvious. (11) is contributed to (7) by the analysis on which `where' has scope over the prepositional phrase `on the shelf'. But there is no corresponding phrase in (8) for `where' to have scope over, so (8) can't mean (11). It appears that wh-scope too respects surface lexicalization, which is to say that `shelvevtr' must be present in the semantic level representations of (8) and (10). So the derivational treatment of location nominals should be rejected.

The same sort of point holds for almost all of HK's examples; we spare the reader further exposition. There is, however, one exception, and we turn to it now.

Consider (12) and (13), the wh-question forms of the locatum denominals (14) and (15). (Other members of the locatum family include *hobblevtr*, *'saddlevtr*, *'tattoovtr*, *'paintvtr*, etc.)⁴ Notice that (12) has the reading (16)

12. Where did Sam put the paint on the wall?
13. Where did Sam paint the wall?
14. Sam put the paint on the wall.
15. Sam painted the wall.
16. Where on the wall did Sam put the paint?

just as one would expect on the analogy of (7). However, (13) has this reading too; e.g., both (12) and (13) can get the answer 'on the upper left hand corner.' Cf., 'Where did Tom tattoo Bill?', for which where on Bill did Tom put the tattoo? is the preferred reading. Note also that, unlike (4), 'John dented his car twice' has the same sort of quantifier scope ambiguity as (3); it can mean either John put two dents on his car or John put a dent on his car on two occasions.

In short, locatum verbs, unlike the other denominals, pass the scope test for decomposition. Or, to put it the other way around, as far as the scope test is concerned, locatum denominal verbs could be derived from the corresponding nouns. We don't think, however, that it's very likely that they are so derived. Instead, we think the right moral is probably that for a verb to pass the scope test is necessary but not sufficient to warrant its decompositional analysis.

For one thing, it seems to us simply not credible that locatum denominals are the one exception to an otherwise enforceable prohibition against postulating syntactic structure underlying surface lexical items; surely, either there are lots of examples of lexicalization, or there aren't any. (We know that won't convince you, but we felt that we should say it anyway.) Second (more heavy irony here) it turns out that the one case where the scope test permits the derivational analysis of a denominal, there are quite considerable technical problems about how such a derivation might proceed. We want now to consider this.

HK don't offer detailed analyses for *'saddlevtr* and the like; they are explicitly aware that there may be problems. But they do suggest that an acceptable derivation might run somewhat as follows:

- 17.1 (PUT)_v (THE SADDLE)NP1 (ON (THE HORSE)NP2)PP
- 17.2 (PUT ON)_v (THE SADDLE)NP1 (THE HORSE)NP2 [raise preposition]
- 17.3 (PUT ON THE SADDLE)_v (THE HORSE)NP2 [raise NP1]
- 17.4 (*saddlevtr* THE HORSE)NP2)VP [lexicalize]

However, (17) strikes us as pretty unlikely. In the first place, we doubt that you can move the prepositional head from a PP leaving its NP object in place: Such movement would apparently license *'John saw with the girl his binoculars' <--- 'John saw the girl with his binoculars', etc. Likewise, it's unclear how to avoid *'he slept the couch' deriving from HE DID A SLEEP ON THE COUCH by raising and incorporation of the preposition. Finally, it looks like (17.3) ought to mean put the horse on the saddle, and it strikes us as not plausible that 'saddle the horse' should derive from something that means that. (We promised not to fuss about exactly what semantic properties lexical items are supposed to share with the phrases underlying them, but, surely, this is going too far.) The derivation one would prefer on semantic grounds would presumably be something like (18).

- 18.1 (PUT) (THE SADDLE) (ON THE HORSE)
- 18.2 (PUT) (THE SADDLE ON) (THE HORSE)
- 18.3 (PUT THE SADDLE ON) (THE HORSE)

But though (18) seems to get the semantics right, it does so at the price of recognizing (THE SADDLE ON) as a constituent; which, surely, it's not. (If you think it is, what is its label?)

By the way, HK's suggested derivation for `saddlevtr' appears to undermine their explanation of why (19) but not (20) is allowed. (See circa p.89.)

- 19. We smeared mud on the wall.
- 20. *Mud smeared on the wall.

Notice that (21) is ok and must come from (22)

- 21. Mud smeared the wall.
- 22. MUD (SMEARED ON) THE WALL.

by raising and incorporation of the preposition. But then (22) should license (20). (But this argument is ad hominum, so we don't insist on it.)

Well, but if `paint the wall' doesn't come from PUT PAINT ON THE WALL, why is (16) a possible reading of (13)? Here's a speculation.

In Fodor and Lepore (1997) we suggested that lexical entries specify not just the meanings (in our view, the denotations) of morphemes, but also rules of composition which determine what effect a lexical item can have on the logical form of expressions that contain it. We took rules of `light verb' introduction to be paradigms. Thus, for example, the lexical entry for `want' specifies that surface configurations like `(wantvtr (a beer)NP)VP' correspond to logical forms like `(wantvtr (to have a beer)COMP)VP'. Such composition rules would be like `definitions in use'; i.e., they determine what logical form a lexical item imposes as a function of the surface syntactic configuration in which the item occurs.

This suggests a similar treatment for locatum denominals: The logical form corresponding to `((paint)vtr (the wall)NP1)VP' is something like `(PAINT NP1 ON NP2)VP', under the constraint that NP1=NP2.5 In effect, `paint the wall' means paint the wall on the wall, and the reason that (13) can have the reading (16) is that, at the syntactic level that determines semantic interpretation, the representation of (13) contains an explicit prepositional phrase.6

Three points about this proposal. First, we emphasize that it treats `paintvtr' (and, by extension, the other denominal verbs) as underived. `Paint' is a transitive verb (as well, of course, as a noun) at every level of representation that this kind of grammar acknowledges. That is incompatible with derivational treatments of the sort that HK proposed; or, indeed, with any form of lexical decomposition.

Second, we've been claiming that there are no well attested precedents for a derivational analysis of locatum denominals; for all the other denominals, the surface lexical inventory is the same as the lexical inventory of semantic level representations, at least if the scope test is to be believed. But there is arguably plenty of precedent for surface structures that are inexplicit about the polyadicity (the number of arguments) in the corresponding logical forms. For example, the LF representation of `John believed Mary' is plausibly something like `JOHN BELIEVED WHAT MARY SAID', and that of `John wanted a beer' is plausibly something like JOHN WANTED TO HAVE A BEER.

Finally, suppose that there is a rule of `light PP introduction', corresponding to the putative rule of light verb introduction, and that the reason that (13) can mean (16) is that wh- can have scope over a prepositional phrase about which LF but not SS is explicit. Then one might reasonably predict that transitive sentences with main verbs like `want' (which, by assumption, introduce light verbs at LF) should exhibit characteristic scope ambiguities. As, in fact, they do. Compare (23), which has both the readings (24) and (25), with (26), which univocally means (27).

23. John wanted a beer for a month.
24. John wanted (a beer for a month).
25. For a month (John wanted a beer).
26. John brewed a beer for a month.
27. For a month (John brewed a beer).

So much, then, for the first part of our discussion. We think there are pretty strong reasons for doubting that denominal verbs are decomposed at the semantic level. But, of course, the best that such considerations can buy us, all by themselves, is a stand-off; perhaps the evidence is equivocal and there are equally strong arguments that favor the derivational analysis. Clearly, HK think that there are. We turn now to assessing the ones that they offer.

Part 2.

HK's arguments for the derivational analysis of denominal verbs are all variations on the same general theme: The derivational account predicts/explains the intuitive impossibility of certain verbs by showing that the syntactic processes that would be required to derive them are prohibited by independently well-confirmed grammatical constraints. (Typically, these are constraints on movement transformations.) We have no principled objection to this form of argument. And, rather than squabble, we're prepared to take HK's word as to what constraints on derivations are well-confirmed; if they say that a derivation is independently blocked, we'll assume that indeed it is. There remains, however, a caveat that we want to enter.

You might expect that the predictions/explanations that HK's arguments license would be of the form 'there couldn't be a verb that means V'. In fact, they are all of the following form (see, pp.60-64): We're given a sentence that contains the neologistic verb, together with a sentence that expresses a paraphrase; and what is explained is why the former couldn't be derived from (i.e., from the representation that directly underlies) the latter. HK put it that "...English simply does not have verbs... [that]... have meanings corresponding more or less to the... paraphrases given here" (pp.59-60). So, for example, according to HK, there can't be a verb 'screen' such that there is a sentence *'they screened clear' corresponding to the 'they cleared the screen', and the reason there couldn't is that there's an independently motivated block on the derivation THEY CLEARED THE SCREEN --> 'they screened clear'.

Here, then, is our caveat: This sort of argument goes through only if (i) the sentence with the neologistic verb really would be ill-formed; (ii) the sentence with the neologistic verb would not be ill-formed on grounds independent of the blocked derivation; and (iii) there is no independent reason why, if the neologistic verb existed, the sentence in question couldn't have the proposed paraphrase. As it turns out, however, there are reasons to deny that all three of these conditions are met in any of the examples that HK offer.

Thus, in the present case, it's plausible that 'they screened clear' would be ill-formed even if 'screenitr' were derivable; viz., for the boring reason that 'clear' is a predicate adjective, not an adverb. (Compare 'they saw *clear/clearly'.) Intuitions about words that don't exist but could are, no doubt, labile. But, as far as we can tell, there's no reason why there shouldn't be a verb 'to screen' which means to clear the screen(s), such that, e.g., 'they all screened right away when the fire alarm went off' is well-formed. (*'They all screened clearly' would be out for whatever reason *'they cleared their screens clearly' is.)

In all the other cases, the problem is not that the blocked sentences would be ill-formed on grounds independent of whether denominals are derived, but that there are independent grounds why they couldn't have the paraphrases that HK propose. Consider (28) on the reading (29). HK imagine a

28. *It cowed a calf.
29. A cow had a calf.

derivation of (28) in which the subject 'a cow' is lowered and incorporated with the main verb 'had', leaving an expletive 'it' in the surface subject position. They claim that this derivation is illegal since "it is well known that a subject... that originates as an external argument... cannot incorporate into the verb that heads its predicate" (p.60). But there's a plausible alternative account of the '*' on (18). What's wrong with reading (28) as (29) is that, so interpreted, there isn't any way to make it compositional.

Patently, one can't suppose that it's the 'it' in (28) that contributes a cow to (29). That would violate the principle that mandatory constituents ('syntactic constants') do not express semantic content. But if, on the other hand, you think of the 'cowed' in (28) as what corresponds to the 'cow had' in (29), you thereby permit a lexical item to contribute a nonconstituent to the semantic representation of a sentence that contains it. Presumably this is not allowed. This restriction isn't gratuitous, of course. As far as anyone knows, it's required to insure that the semantic representations of well-formed sentences are themselves well-formed, as (29) would not be on the parsing a (cow had) a calf.⁷

In short, what's wrong with (28) on the reading (29) is that there is no way to assign that reading consonant with the demands of compositional interpretation; there is no acceptable way to pair the elements of the sentence with the elements of its presumed canonical representation. A sentence that has the syntax of (28) and means (29) is thus ruled out whether or not denominal verbs are derived.⁸

Compositionality problems of a similar sort infect HK's treatment of the examples (31), in which the 'on' in (32) is

- 31. *He shelved the books on.
- 32. He put the books on a shelf.

unincorporated. HK remark that there is plenty of precedent for (31) "...such as 'take (the business) over, take (a stray cat) in....[etc.]'" (p.60). But, in fact, when sentences end in unattached prepositions, the latter always alternate with the corresponding verb+particle constructions (('take+over (the business)' 'take+in (a stray cat)', etc.); and the choice of the preposition is conditioned by the identity of the verb (*'take through the business', *'take with the cat', etc.). So, the question that HK are raising must be why there can't be a sentence (33), of which (32) is the canonical representation.

- 33. *He shelved on the books.

Possible answer: because the reflex of 'shelve+on' in the canonical paraphrase of (33) would be the nonconstituent put...a shelf, and, to repeat, lexical items can't contribute nonconstituents to semantic level representations.⁹ Compare (34) and (35): the

- 34. John put+out the cat.
- 35. John put out the cat.

surface lexical item 'put+out' in (34) introduces put out into the semantic representation (35), of which put out is indeed a constituent. So all is well.

The last of HK's examples we'll discuss are sentences like (36), of which the canonical representation is supposed to be (37); and (38), of which the canonical representation is supposed to be (39).

- 36. *She churched her money.
- 37. She gave her money to a church.
- 38. *He bushed a trim.
- 39. He gave a trim to a bush.

HK take sentences like (36) and (39) to violate movement constraints on PPs, and they have a rather delicate story to tell about how to block these sentences while allowing derivations like `put the books on the shelf' ---> `shelve the books'.¹⁰ But we're pretty sure that no solution of that kind can be right. What's wrong with deriving (36) and (38) from (37) and (39) has nothing to do with constraints on extraction of (or from) PPs. Rather, it's that these derivations would require the incorporation, into the underlying verb, of its indirect object (e.g., the incorporation of `to the church' in `give' in (36)). And this, as far as we can tell, is never allowed; at least, we can't think of any examples to the contrary. Compare `give a trim to the bush' ---> *`bush a trim' with `give water to the bush' ---> `water the bush', where a denominal verb alternates with a DO, and with `put the money in the bank' ---> `bank the money' where a denominal verb alternates with a PP.¹¹ We're claiming that, unlike constraints on incorporating PPs, or Direct Objects, both of which are pretty clearly lexically governed (so, `water the bush' <--- `put water on the bush' but not *`mother the hospital' <--- `put mother in the hospital'; and `banked the money' <--- `put money in the bank' but not *`trash the journal' <--- `put trash in the journal'), the constraint that rules out (36) and (28) is purely structural. In which case, the treatment of (36) and (38) should precisely not be assimilated to the treatment of constraints on incorporating DOs and PPs.

It's not, in fact, hard to see why the prohibition against incorporating IOs is absolute. We remarked above that what's wrong with (28) is that a word in a sentence can never correspond to a nonconstituent in the canonical paraphrase of the sentence.¹² Exactly this same principle applies to block (36) and (38) (and `bank the money' on the reading give the money to the bank.) To get any of these derivations, you would have to lexicalize an expression consisting of a verb together with its indirect object, leaving its direct object in place. But (V+IO) is not a possible constituent (presumably in any language). Ergo, no word can contribute the semantic content of an underlying verb and its indirect object to a canonical paraphrase. Notice that `what did John give Bill?' isn't a counter example to the claim that (V+IO) isn't a possible constituent. Not, at least, on the (standard) assumption that the VP of this sentence contains a phonetically unrealized trace which constitutes the direct object of `give'; i.e., the constituent structure of the VP is (give (trace)DO (Bill)IO)VP rather than (give (Bill)IO)VP.

In short, the principle that governs (36) and (38) is that interpretation works only on well-formed expressions. This treatment unifies our analysis of (36) and (38) with our analysis of (28) (and, indeed, with our analysis of (18), where we argued that `saddle the horse' can't come from `put the saddle on the horse' if only because that would require lexicalizing the nonconstituent `the horse on.'). HK are themselves aware that examples like (36) and (38) are equivocal. They remark that "... the verbs here may be impossible for a variety of reasons" (p. 61). But we think that the principle that excludes them is, in fact, perfectly general, and that the HK treatment simply misses the generalization.

Conclusion

No doubt, the last word on denominals has yet to be uttered. For one thing, we are far from certain that our story about light constituent introduction is right; more about that in later publications. Also, as we've just been suggesting, it strikes us that the relation between denominal verbs and the corresponding nouns are really often quite specific to particular pairs of items. For example, there are noun-verb pairs where it is intuitively much more plausible to treat the noun as a variant of the verb than to go the other way around. We have trouble getting it down that `laughv' derives from `a laugh', `sing' from `a song' and so forth.

We're inclined towards the following picture. Denominals are `derived' only in the sense that, in the lexicon, some nouns have the feature possible verb (`saddle'), and some verbs have the feature possible noun (`laugh'). If there really are powerful generalizations about which kinds of lexical items have which kinds of features, they should be written as lexical redundancy rules. On this old fashioned and familiar view, there are no lexical decompositions and morphological derivation is entirely an intralexical affair.

We don't, of course, at all claim to have shown that the old fashioned and familiar view is right. Suffice it that if the denominal verbs constitute the argument to the contrary, then the argument to the contrary is weak. Morphemically simple denominal transitives don't behave, in respect of scoped elements, as though they have two-clause semantic representations; no more than morphemically simple causatives do. And the evidence for HK's claim that the derivation of such denominals respects the general constraints on syntactic movement looks to be pretty comprehensively confounded. We think the moral is that God wasn't just fooling around when He made morphemes. The commandment that Moses forgot: THOU SHALT NOT TAKE THE SURFACE LEXICON IN VAIN. Occam's Razor commends this, and so, it appears, do the data.

NOTES

*We're grateful to Ken Hale and S.J. Keyser for a running conversation about the issues discussed in this paper. We didn't convince them or vice versa.

1. We use capital letters for canonical names of semantic representations and italics for canonical names of meanings. (So, according to the decompositional view of causatives, 'kill' derives from CAUSE TO DIE and means cause to die.) However, for convenience, we will sometimes not distinguish between a semantic representation and the corresponding surface phrase; e.g., we'll say that 'kill' is derived from 'cause to die' to abbreviate the view that 'kill' and 'cause to die' are both derived from CAUSE TO DIE.

2. For worse, we think. Philosophers care about the issue of decomposition largely because it connects with questions about analyticity, linguistic necessity, a prioricity and the like. If, however, to claim that 'kill' derives from CAUSE TO DIE does not entail that 'kill' and 'cause to die' are synonyms (or, at a minimum, necessarily equivalent), then it would seem that all these connections are moot: Perhaps, for example, it's compatible with the derivational view that there is no meaning relation at all between the 'boil' in 'boil the water' and the 'boil' in 'the water boils'.

But linguists too should care about the semantics of lexicalization. If there are semantical properties that lexicalization is required to preserve, evaluating candidate lexical derivations depends on knowing which ones they are. And if there are no semantical properties that lexicalization is required to preserve, then why should a level at which 'kill' and 'cause to die' fall together be supposed to be concerned with the representation of meaning? Indeed, if you really, really don't care about the semantics of lexicalization, why derive 'shelve the books' from 'put the books on the shelf,' rather than, say, from 'put the books near the shelf,' or from 'put the books a mile from the shelf'.

3. This is just the sort of situation that meaning postulates were invented for, and perhaps one should invoke them. On the other hand, it mustn't be taken simply for granted that such entailments are semantic at all, rather than, e.g., metaphysical. The issues here are familiar but very deep.

4. For purposes of intuition building: The difference between location denominals and locatum denominals is that, whereas the former incorporate underlying prepositional phrases ('shelve the book' comes from 'put the book on the shelf'), the latter incorporate the heads of underlying prepositional phrases, together with underlying direct objects ('paint the wall') comes from 'put the paint on the wall.' We'll see later how to draw the location/locatum distinction if a nonderivational account of the lexicon is assumed.

5. Presumably the requirement that NP1=NP2 is lexically governed. It's plausible, for example, that 'donate' can introduce a 'light' indirect object (so that 'John donated the bookNP1' has the logical form 'John donated the bookNP1 to NP2') under the constraint that NP1 and NP2 not be identical.

6. We can now cash fn. 4: The relevant difference between location verbs and locatum verbs might be that the latter but not the former introduce light elements.

Kiparsky (ms) has suggested a semantic test for this distinction; locatum verbs are such that "putting x in y is a canonical use of x"; location verbs are such that "putting x in y is a canonical use of y" (p.10). But we doubt that this will do much work if the goal is to predict which verbs are possible. Notice *`Granny mantleshelved the little china dogs' even though putting little china dogs on mantleshelves is a canonical use for both. Also: *`page the book'; *`fireplace the logs'; *`goldfish the bowl'; *`clothe the clothes closet'; *`vault the valuables'; etc.

In any case, a purely semantic test for the difference between locatum verbs and location verbs would be no comfort for theorists like HK, who want to argue that such distinctions are bona fide linguistic on the grounds that they correspond to an independently defined syntactical distinction between location and locatum verbs.

7. On an atomistic treatment of the lexicon (see Fodor and Lepore, 1997), barring only logical and grammatical constants and maybe items that introduce light elements (see above), each surface morpheme contributes a correspondingly unstructured item to the interpretation of the sentence; in effect, each contributes itself. So the `cow' in (29) comes from `cow', the `had' comes from `had', and so forth. To that extent, the requirement that surface lexical items must correspond to constituents of semantic representations is satisfied trivially and automatically. This strikes us as a not inconsiderable advantage of atomism.

8. Much the same considerations explain why you can't have a sentence (30) that means (29) and derives from A COW HAD A CALF by raising and incorporating `HAVE' and leaving the empty verb `do'.

30. *A cow did a calf.

As far as we can see, the HK treatment of (28) does not, in and of itself, predict that (30) is ill-formed.

9. Our complaint isn't, of course, that put...a shelf is a discontinuous constituent in (32). It's that there is no precedent for a constituent, continuous or otherwise, that consists of (just) a verb together with the nominal object of one of its prepositional phrase complements. The `(the phone up)' in `John hung the phone up' is not a counter-instance since there is no constituent `(the phone up)' in `John hung the phone up'.

10. We are pussyfooting because we don't really understand HK's proposal. If we are reading them right, they think (36) and (38) would involve extracting the specifier of a PP (e.g., the `money' in `money to a church') and that this would violate the ECP principle. But we don't see how that story could square with the well-formedness of, e.g., `bank the money', which HK would presumably derive from the subject of `in' in `put the money in the bank'; nor is it clear to us why moving the subject of a PP violates ECP if moving the object (as in `saddle the horse' <--- `put the saddle on the horse') does not.

11. The latter example is especially striking; notice that `bank the money' can't mean give the money to the bank. I.e., although the denominal form `bankvtr' exists with the semantic force of an incorporated PP, it does not exist with the semantic force of an incorporated IO.

12. If you are thinking of generating lexical items from syntactically complex expressions, then the relevant constraint is that you can't lexicalize a nonconstituent. But, of course, we don't believe in lexicalization; so what we're really saying is that no word can alternate with a non-constituent of its canonical paraphrase. When we talk informally of constraints on lexicalization, we always really mean constraints on such alternations.

REFERENCES

- Fodor, J. (1970) "Three reasons for not deriving 'kill' from 'cause to die,' *Linguistic Inquiry* 1; 429-438.
- Fodor, J. and Lepore, E (forthcoming) "The emptiness of the lexicon: Critical Reflections on J. Pustejovsky's *THE GENERATIVE LEXICON*." *Linguistic Inquiry*.
- Hale, K and Keyser, S. (1991) "On the syntax of argument structure," *Lexicon Project Working Papers*, No. 34, MIT Working Paper in Linguistics.
- Hale, K. and Keyser, J. (1993) "On argument structure and the lexical expression of syntactic relations," in Hale, J and Keyser, S. (eds.) *The View from Building 20*, MIT Press, Cambridge Ma.
- Katz, J. and Fodor, J. (1962), "The structure of a semantic theory," *Language*, 39, 170-210.
- Kiparsky, P., "Remarks on denominal verbs," Alex Alsina, Joan Bresnan, Peter Sells, Center for the Study of Language and Information, Stanford, 1997.
- Pustejovsky, J. (1995) *The Generative Lexicon*, MIT Press, Cambridge, Mass.