

THE EMPTINESS OF THE LEXICON:
Critical Reflections on J. Pustejovsky's *THE GENERATIVE LEXICON*

Jerry Fodor and Ernie Lepore
Rutgers University Center for Cognitive Science

INTRODUCTION

A certain metaphysical thesis about meaning that we'll call Informational Role Semantics (IRS) is accepted practically universally in linguistics, philosophy and the cognitive sciences: the meaning (or content, or 'sense') of a linguistic expression¹ is constituted, at least in part, by at least some of its inferential relations. This idea is hard to state precisely, both because notions like metaphysical constitution are moot and, more importantly, because different versions of IRS take different views on whether there are constituents of meaning other than inferential role, and on which of the inferences an expression occurs in are meaning constitutive. Some of these issues will presently concern us; but for now it will do just to gesture towards such familiar claims as that: it's part and parcel of dog meaning dog² that the inference from x is a dog to x is an animal is valid; it's part and parcel of boil meaning boil that the inference from x boiled y to y boiled is valid; it's part and parcel of kill meaning kill that the inference from x killed y to y died is valid; and so on. (See Cruse, Ch. 1 and *passim*.)

IRS brings in its train a constellation of ancillary doctrines. Presumably, for example, if an inference is constitutive of the meaning of a word, then learning the word involves learning that the inference holds. If dog means dog because dog ---> animal is valid, then knowing that dog ---> animal is valid is part and parcel of knowing what the word dog means; and, similarly, learning that x boiled y ---> y boiled is part and parcel of learning what boil means, and so forth.

IRS constrains grammatical theories. The semantic lexicon of a language is supposed to make explicit whatever one has to know to understand the lexical expressions of the language, so IRS implies that meaning constitutive inferences are part of the semantic lexical entries for items that have them. Lexical entries are thus typically complex objects ('bundles of inferences') according to standard interpretations of IRS. It is this latter thesis that will primarily concern us in the present discussion. For reasons that we've set out elsewhere, we doubt that IRS can be sustained (Fodor and Lepore, 1992); a fortiori, we doubt the cogency of arguments that take IRS as a premise. The primary question in what follows will be whether there are any persuasive arguments for the complexity of lexical entries which do not presuppose IRS. Our main interest in Pustejovsky (hereafter JP) is that he purports to provide such an argument.

Here is how we will proceed: the natural alternative to the claim that lexical entries are typically complex is the claim that lexical entries are typically atomic (i.e., they lack internal structure). We propose to adopt a version of this claim as a sort of 'null hypothesis'; viz., that the only thing a lexical entry specifies is the denotation of the item it describes. Here again we scant the details for the moment; but, roughly: the lexical entry for dog says that it refers to dogs; the lexical entry for boil says that it refers to boiling...and so forth.³ We'll try to show that all the standard arguments for rejecting this null hypothesis, JP's included, either depend on assuming IRS or are independently unsound.

In Parts I and 2, we consider JP's account of the semantic lexicon; in particular, we will discuss and reject his argument that the complexity of lexical entries is required to account for lexical generativity. In Part 3, we'll consider whether our default theory, lexical atomism, mightn't actually be true. We'll argue that, strictly speaking, it's probably not; strictly speaking lexical entries are typically complex. But we claim that they are complex in a way that does not jeopardize either the thesis that lexical meaning is atomistic, or the identification of lexical meaning with denotation.

PART I: JP's THEORY

We take the theory of the lexicon outlined in TGL to be firmly within the IRS tradition. In particular, like other proponents of IRS, JP thinks that word meanings are constituted by inferences, hence that knowing what a word means involves knowing (some of) the inferences in which it participates. In fact, he apparently thinks (what IRS does not strictly require) that understanding a token (e.g., an utterance or inscription) of an expression involves actually drawing (some of) the inferences that the words that are tokened license: "the structuring of taxonomic information ... is not simply an exercise in domain modeling; it is necessary for driving the inferences that a language reasoning system must perform in order to understand a sentence" (18, our emphasis). As usual with IRS-motivated theories of meaning, it's taken for granted that the semantic lexicon should somehow formally specify the inferences by which meaning (/knowledge of meaning/sentence understanding) is constituted.

We stress this because JP occasionally writes as though it were not the specification of inferential roles but rather the specification of denotations with which lexical semantics is primarily concerned: "Lexical semantics is the study of how and what the words of a language denote" (1). But this is a little misleading. Roughly, JP thinks that lexical entries specify denotations via their senses; so that, for example, coextensive expressions may well be assigned distinct lexical entries.⁴ Thus, John bought the book from Mary and Mary sold the book to John are made true by the same event; but they differ, according to JP, in a property of 'headedness' or 'focus' which they inherit from the lexical entries of their respective verbs. Patently, however, 'headedness' isn't a property of denotations ("things in the world") but rather of denotations as represented. JP says that "Headedness is a property of all event sorts..." (72); but he can't really mean that. What he must really mean is that it's a property of all (semantically well formed) representations of event sorts (in English).

Similarly, one of JP's characterizations of an "event structure" specifies an event e_3 such that "...there is no other event that is part of e_3 " (69). But, clearly, all events have parts (down to whatever physics determines is the microstructure of the universe). What JP means is really that no other event is represented as (or implied as) part of e_3 by the form of words which expresses the event. So, here again, what the semantics is telling us about is not (or not just) what a word represents, but how it represents it; in effect, not just its denotation but its sense.

We don't mean to insist on what are, probably, not confusions but just JP being casual about the use/mention distinction. However, it's important to us to contrast JP's sort of project with a bona fide denotation semantics, according to which reference is the only semantic property of linguistic expressions, hence, a fortiori, the only semantic property of linguistic expressions that lexical entries specify.

In fact, JP enumerates a variety of constraints on semantic theories that a purely denotational lexicon clearly could not meet, hence which are supposed to motivate a richer notion of lexical semantic representation than a denotational lexicon could provide. With one exception, these are familiar from discussions in lexical semantics dating as far back as Katz and Fodor (1963). Consonant with our overall project, we propose to discuss them only briefly, putting aside those that are motivated either by IRS or what we regard as other tendentious assumptions. This will leave for Part 3 to consider what we take to be JP's main line of argument and the main contribution of his book: the claim that only if lexical entries are typically complex can the generativity of the lexicon be captured; a fortiori, that an atomistic lexicon would fail to explain "... how words can take on an infinite number of meanings in novel contexts" (42).

As far as we can make out, other than the considerations about generativity, JP offers three kinds of arguments for complex lexical representations:

1.1 Interlexical semantic relations.

Lexical semantics is required to specify "how words are related to one another..." (23) including, in particular, relations of "synonymy, antonymy, hyponymy and lexical inheritance, meronymy, entailment and presupposition".⁵

Clearly, a lexical entry that says only that dog refers to dogs will not thereby specify (e.g.) that dogs are animals (hyponymy); a lexicon that says only that bachelor refers to bachelors will not thereby specify that bachelor means the same as unmarried man (synonymy), etc. In fact, barring appeal to meaning postulates (of which JP disapproves, see, circa 54, 110), the only way of capturing such relations in the lexicon would seem to require complex lexical entries; e.g., one which includes ANIMAL in the entry for

dog and UNMARRIED and MAN in the entry for bachelor, etc.. In short, if the lexicon is constrained to capture "interlexical relations", then it is neither atomistic nor purely denotation.

The question, however, is how to motivate imposing this constraint on lexical entries; and here the problems are familiar and formidable. The basic issue is this: all the relations JP enumerates turn, in one way or another, on lexical entailment, a relation that is notoriously difficult to distinguish from other species of necessity. For example: Suppose it's true that dog ---> animal is valid, hence, that nothing could be a dog that wasn't an animal. Remember that semantic lexicons are supposed to be repositories of (and only of) meaning constitutive inferences; i.e., of those inferences that one has to be prepared to draw on pain of not understanding the word in question. Now, it seems clear, given this condition, that the necessity of an inference is not per se sufficient for its meaning constitutivity. two ---> prime is necessary by anybody standards, but it's implausible that you can't know what two means unless you know that two is prime. Similarly, it's presumably necessary that what's square isn't circular, but it's implausible that you can't know what square means unless you also know what circle means. Etc. Such examples are legion.⁶

What's worse for JP, is that many of the inferences that he'd like lexical entries to determine (e.g., from want a cigarette to want to smoke a cigarette; see below) are "defeasible" (45), hence not, presumably, strictly necessary. One might reasonably feel that if the necessity of an inference is insufficient to guarantee the semantic relatedness of its lexical constituents, no weaker modality is likely to suffice. JP's comments on these sorts of issues tend not to be illuminating. For example, "one fairly standard definition states that two expressions are synonymous if substituting one for the other in all contexts does not change the truth value of the sentence where the substitution is made. A somewhat weaker definition makes reference to substitution relative to a specific context. For example, in the context of carpentry, plank and board might be considered synonyms, but not necessarily in other domains" (23). But any two words can be substituted, salve veritate, in some context; e.g., elephants and asteroids substitute in the context are often bigger than a breadbox. JP could avoid this objection if he had an account of how to individuate semantic "domains"; but we bet he hasn't.

We suspect that JP hasn't really faced the problem that's implicit in making notions like 'defeasible inference' respectable as a basis for semantic theory; viz., the problem of distinguishing what the language tells one from what one knows about the world. "[My wife uses the subway every day]... is a near paraphrase of 'my wife travels on the subway every day,' an interpretation that is made possible by our knowledge of what the function of a subway is. (87)" But if the interpretation is indeed made possible by what we know about subways, then the conclusion should be that the two sentences are not "near paraphrases" in any sense that semantic theory is concerned with. Presumably what's in the semantic lexicon is information about words, and not information about the world. That's why it's perfectly possible both to understand Sarah likes to use the subway and to wonder what she likes to use it for.

If IRS is true, then there must be some answer to the questions 'what distinguishes linguistic knowledge from world knowledge?' and 'what distinguishes lexical entailment from mere necessity?': IRS guarantees that some interlexical relations are meaning constitutive, and should thus be specified by lexical entries, even if it's not quite clear which ones these are. In the present discussion, however, we are explicitly not taking IRS for granted. We're therefore impressed by the difficulty of saying with any clarity which interlexical relations are the semantically relevant ones, or how they should be individuated. Accordingly, in what follows, we'll assume that the failure of an atomistic, denotational lexicon to capture such relations as synonymy, hyponymy and the rest is not a decisive argument against it.

1.2 Semantic well-formedness.

"I will introduce a notion of semanticity analogous to the view of grammaticity [sic], but ranging over semantic expressions rather than syntactic structures. Semanticity refers to the semantic well-formedness of expressions in a grammar" (40).

No doubt, if there is such a property as semanticity, the lexicon should contribute to determining it. The trouble is that it's not clear what property it might be. JP provides no general characterization, and the few examples he offers are not transparent; in fact, we think they are a mixed bag.

Consider pleonasms like: ?Mary kicked me with her foot and ?Mary buttered the toast with butter. (40). On any account, what's wrong with these sentences is that the prepositional phrase adds nothing to what the rest of the sentence says, so a Gricean imperative ("be informative") is violated. We suppose JP must think that the locus of the redundancy is the meaning of the verbs; anyone who knows what buttering means knows that it's done with butter. This consideration does not, however, argue unequivocally that the lexical entry for butter is COVER WITH BUTTER or the like. Suppose that the dictionary says only that

butter_r refers to buttering. Still, anybody who knows this, and knows what buttering is, can tell that the sentences in question are redundant. It seems that this kind of "unsemanticity" doesn't, after all, argue for "a level of representation in the semantics, which operates according to its own set of constraints" (42).

Here are some of JP's other examples of semantic ill-formedness (41):

- (1) a. ?John began the dictionary.
- b. ??Mary began the rock.

According to JP, these are "...semantically odd because of what we normally associate with the semantic possibilities of a noun such as dictionary and rock" (41). We're inclined to think, on the contrary, that there's nothing wrong with (1a) at all; and that, if there's anything wrong with (1b), it's that one can't imagine what it was that Mary began to do with (/to) the rock. Given a clue (e.g., she began, in either sense, to paint it), the perplexity vanishes and so does the intuition that something is awry with the sentence.

The trouble, in a nutshell, is this: if capturing semantic well-formedness is to be a constraint on representations of lexical meanings, hence on lexical entries, it has to turn out that what's wrong with a semantically ill-formed sentence is that what it means is defective; not just that it would be "normally" or "generally" hard to "readily" interpret without contextual support. It's not clear that any of JP's examples have this property. Indeed, it's not clear that there is such a property as defectiveness of meaning as distinct from ungrammaticality, necessary falsity, and any of a variety of kinds of pragmatic malfeasance. Barring more persuasive examples, or an independent argument for "a level of representation in the semantics" whose rules would determine the semantic well-formedness conditions for expressions, we are inclined to doubt that failure to mark (1a,b) and the like as defective militates very strongly against purely denotational lexical entries.

1.3 Distribution

It's a widely held view that the semantic properties of a lexical item "give rise to" (19) its syntactic properties: syntactic distribution is somehow determined or explained by, or is anyhow predictable from, meaning; so a lexicon which represents only denotation would miss linguistically salient generalizations. In this spirit, Jim Higginbotham (1994) remarks that "...the meanings of lexical items systematically infect grammar. For example,... it is a condition of object preposing that the object be in some sense 'affected' in the events over which the nominal ranges; that is why one has ... algebra's discovery (by the Arabs) [but] not *algebra's knowledge (by the Arabs)". Similarly, according to JP, "...the diversity of complement types that a verb or other category may take is in large part also determined by the semantics of the complements themselves" (10).

It is, however, notoriously difficult to assess the claimed correlations between lexical semantics and syntactic distribution, because one is almost never told what the semantic representations themselves mean. (What, exactly, is it for an object to be "in some sense 'affected' by an event"?). This imprecision tends to undermine the supporting examples. So, for example, JP says that the difference between "John is running (Therefore, John has run) [and] John is building a house (*Therefore, John has built a house)... is whether an action is homogenous in nature or has a culmination of some sort" (16). Well, one might have thought that the process of starting to run has a culmination of some sort; viz., running. But John is starting to run ---> John has started to run looks ok (Cf., John is starting to build a house ---> John has started to build a house, etc.) Our point isn't that this is a counterexample to JP's generalization; rather, it's that the generalization is formulated so imprecisely that one can't tell whether it's a counter-example. JP's emendation to 'having a culmination' (viz., "...culminates in a changed state, i.e., is an accomplishment" (16)) doesn't help a lot. Does it rule starting to run in or out? By what criterion, one wonders?

We're aware that this sounds like carping; but, in fact, we know of no cases in the lexical semantics literature where the semantic end of putative semantics/syntax correlations has been made out with sufficient clarity to permit the claims to be evaluated. In such a circumstance, what can one do but carp?

Here's another example from JP (10). Compare:

- (2) a. The woman ate her meal.
- b. The woman ate.
- c. The dog devoured the cookie.
- d. *The dog devoured.

The puzzle: Why can you drop the direct object of ate but not of devoured? The explanation, according to JP: "...while eat denotes an activity of unbounded duration (at least lexically), devour, one might argue, denotes a transition. [devour]... carries a completive implicate absent from eat" (11).

There is some uncertainty about how much JP is claiming here; if (as seems likely; cf. 9-12 passim) JP thinks that the fact that eat is lexically represented as denoting an unbounded activity explains why you can drop its object, then he must believe that there could not be a word which means what eat does but does not direct-object delete. This claim is remarkably strong and there is, to put it mildly, no reason to suppose that it is true. Notice that all of the following are bad, though unbounded activities are apparently involved in each: *John stroked; *John ground (/his teeth); *John pounded; *John smelled (/the salt air), etc.

But, anyhow, what is it to "denote an activity of unbounded duration (at least lexically)"? Nobody eats forever; and why doesn't the "change of state" from eating to not eating that occurs when one stops eating constitute "a transition," just as much as the cessation of devouring that occurs when one stops devouring? And, if devour carries a completive implication, why is she never finished devouring her meal ok?; and, even supposing that devour does carry a completive implication, so too, surely, does finish eating. But the woman finished eating her meal is perfectly ok though the *the woman finished devouring isn't.

The fact is, nobody knows what, if anything, it is for a word to denote a transition; so nobody knows whether devour does. Similarly, we believe, for the other claims about the semantic determinants of syntactic distributions with which the lexical semantics literature abounds. What's worse, it's possible to wonder, in such cases, whether there really are two different things to correlate. Perhaps the intuition that devour, but not eat, is semantically completive is just a hypostatic misconstrual of one's syntactic knowledge that the first but not the second verb takes a direct object mandatorily. You can't correlate a thing with itself; that, as the reader may recall, is why Pooh never caught a Woozle. (Cf. Milne, 1957. For further discussion, with further examples, see Fodor, (forthcoming)).

Finally, suppose that, despite the awfulness of the examples and the imprecision of the metalanguage that is used to describe them, it really is true, in God's eyes as it were, that (some or all) syntax is semantically driven. Would that show that there's something more to lexical representations than denotations? Not obviously. For, even if it's devouring's being a transition that explains why devour has to have an object, why should knowledge of this fact count as part of what the speaker/hearer knows about devouring rather than as part of what he knows about devouring? A priori, it seems as reasonable that syntax should be driven by what you know about what a word denotes as that it should be driven by what you know about what the word means. Barring some principle for choosing between these two interpretations, we don't see how the putative influences of semantics on syntax could count in favor of a semantics of lexical senses.

So much for the polemical background. Though one finds all these sorts of arguments in JP, they have a considerably older provenance, and JP's treatment of them is not more convincing than what earlier discussions offer. What makes JP's book interesting is the argument from the semantic generativity of the lexicon to the semantic complexity of lexical entries. We turn to that now.

PART 2: GENERATIVITY

A word about method: JP develops quite a lot of apparatus in setting out his story about lexical generativity, and neither his notation nor his exposition is strikingly perspicuous. We propose, wherever possible, to avoid the burden of reconstructing the technical details. Instead, we will proceed by considering several of JP's parade examples of the application of his theory to the analysis of linguistic data. We'll try to show that these analyses don't work, but our main purpose is not to impugn them; rather, it's to illustrate the character of the theory and of its defects. We claim that JP's failures to get the data right are principled and show that there is something wrong with the kind of account of the lexicon he's endorsing. We aren't, in short, just quibbling about cases.

JP's discussion of generativity starts by invoking the distinction between 'polysemy' and mere accidental homonymy.⁷ Intuitively, the paradigm of polysemy is the kind of quasi-ambiguity exhibited by the believe in John believes Mary, on the one hand, and John believes that 2 is less than 3, on the other. However, according to JP, polysemy can also be exhibited by expressions which, unlike believe in this example, are syntactically homogeneous; e.g., the window that JP says means aperture in he crawled through the window and the window that means physical object in the window is rotting. Similarly, JP thinks good, though it has only one lexical entry, means something different in good knife and good secretary.

Many linguists have shared JP's sympathy with Weinreich's (1959) complaint that standard treatments of the semantic lexicon fail to reconstruct the difference between polysemy and mere homonymy. In effect, by postulating distinct lexical entries for each meaning of a polysemous term, 'sense enumeration

lexicons' (SELS) miss "the fact that ... the two senses [of a polysemous expression] are logically related..." (37). However, on JP's view, the fact that believe or window (or use, see above) seem to take on different, but related, senses in their various contexts, is itself just a special case of the perfectly general fact that "...words can assume a potentially infinite number of senses in context" (105). Correspondingly, the theorist's goal is to capture the context sensitivity of lexical meaning "while limiting the number of senses actually stored in the lexicon" (105), i.e., without proliferating lexical entries beyond necessity.

To see how this is to be accomplished, let's start with JP's treatment of the polysemy of bake.

Prima facie, bake is polysemous between "a creative activity" as in bake a cake and a "change of state predicate" as in bake a potato. "Intuitively we'd like to capture the fact that the former objects are prototypically brought about by the activity they are in composition with" (123), i.e., that "certain objects come into being by virtue of an activity which can otherwise simply change the internal state of an object" (98).

The kind of "sense enumeration" analysis that JP wishes to reject might distinguish two lexical entries for bake; e.g., MADE vs. HEATED UP, or something of the sort. JP's account, by contrast, has a single lexical entry for bake and distinguishes the semantics of bake a cake from the semantics of bake a potato on the basis of a difference in the lexical representations of the object nouns cake and potato: "...we can derive both word senses of verbs like 'bake' by putting some of the semantic weight on the NP. This view suggests that, in such cases, the verb itself is not polysemous. Rather, the creation sense of 'bake' is contributed in part by the meaning of 'a cake', by virtue of it being an artifact" (124).

This treatment is entirely characteristic of JP's general approach; indeed, quite a lot of the book consists of applying it to a wide variety of sorts of examples. Notice, in particular, the idea that the sense of "governing" expressions is inherited partly from their lexical entries, but partly also from the semantics of the expressions that they govern in a context. According to JP, this is true not just of the relation between verbs and their complements but also, for example, of the relation between adjectives and the nouns they modify (thus, good inherits FOR CUTTING from knife in good knife, etc. (129)). Notice also that pursuing this strategy places constraints on the lexical contents of the governed expressions. Use can inherit TO RIDE ON from subway in use the subway only if the "telos" (function) of a subway is part of the entry for subway. Likewise, the creative sense of bake a cake can emerge from the occurrence of ARTIFACT in the entry for cake only if that entry specifies the information that cakes are artifacts; etc. In effect, the more context sensitive the meanings of governing expressions are supposed to be, the richer must be the lexical entries for the expressions that they govern. So, JP's standard tactic for arguing that there is a certain piece of semantic information in the lexical entry for a noun is to show that there is a context in which the information is inherited by a verb to which the noun functions as a complement; or by an adjective that modifies the noun, etc.

In all these respects, to repeat, JP's treatment of bake is absolutely paradigmatic. Correspondingly, the failure of JP's analysis of bake is characteristic of the general inadequacy of the theory. We turn to considering this analysis at some length.

To begin with, suppose that JP is right and the reason that bake in bake a cake means create and bake in bake a potato means warm up is that cakes are artifacts and potatoes are "natural kinds". JP thinks that that would explain why the two bakes aren't merely homonyms. But would it? How, exactly, does the explanation go? "[artifacts] such as cookies, cakes and bread are typically baked. The process of baking, modulo such objects, is a creative activity, while relative to objects such as potatoes, carrots and other natural kinds, it is simply a change of state predicate... to classify the verb bake as having both senses lexically specified is to miss the semantic generalization..." (98). (See, also, 47.) Apparently, on this account, what makes the two bakes not merely homonyms is that both express the same "process of baking" which, in one case, constitutes a creative activity and, in the other case, merely effects a change of state.

But now: By what criterion do both kinds of baking count as the same process? What decides that the bake in bake a cake (hence creating one), denotes the same activity as the bake in bake a potato, (hence heating one?) (Whereas, presumably, bank is homonyms because the bank in bank a check counts as a different process from the bank in bank a plane.) This is just the polysemy problem all over again; all that's happened is that it's been kicked upstairs from the semantics to the ontology: Whereas we used to worry about how to count senses, we are now invited to worry about how to count processes. The total yardage gained would appear to be negligible or nil.

But, for the sake of the argument, let's put ontological qualms aside and go back to the linguistics.

If the creative sense of bake is determined by something that it inherits from its direct object ---and if bake and cake are themselves univocal--- then bake a cake must have only the "creative" reading. But, in fact, bake a cake is ambiguous. To be sure, you can make a cake by baking it; but also you can do to a (pre-existent) cake just what you do to a (pre-existent) potato; viz., put it in the oven and (noncreatively) bake it. Since bake a cake is ambiguous and cake is univocal, it must be that bake is lexically ambiguous (specifically, polysemous) after all, contrary to JP's analysis.

The ambiguity of bake a cake shows that JP is wrong about the polysemy of bake. But it doesn't show that he's wrong about cake being lexically marked as an artifact. Indeed, it looks, first blush, as if only distinguishing artifacts from "natural kinds" could account for the differences between bake a cake and bake a potato, since the latter appears to have only the noncreative reading. Maybe some of JP's program can be saved after all.

But, in fact, even assuming that the lexicon distinguishes artifacts from natural kinds isn't enough to do the trick. Notice that although knives and trolley cars are artifacts, bake a knife and bake a trolley car resist a creative reading quite as much as bake a potato does. But if bake a cake is heard as creative because cake is marked as denoting an artifact, then bake a trolley car should be heard as creative too. Clearly, something has gone wrong.

In fact, there's a striking difference between JP's informal account of bake and what he actually puts in the examples of lexical entries that he offers. According to the informal account "the creation sense of 'bake' is contributed, in part, by the meaning of 'a cake,' by virtue of it being an artifact" (124). But the lexical entry for 'cake' specifies not that cakes are artifacts, but, in effect, that cakes are made by baking (123). (Cf., JP's discussion of 'co-compositionality'). Presumably, this is required in order to distinguish cake from potato, on the one hand, and from knife, on the other: the entry for potato says they are a "natural kind", and knife is presumably lexically specified for some means of production other than (and exclusive of) baking.⁸ In short, in order to explain why knife doesn't select the creative sense of bake, JP is required to claim that you don't know what any term for an artifact means unless you know how that kind of artifact is made. If, like the present authors, you don't know how they make pencils, you don't know what pencil means.

The right story is surely this: as far as the language is concerned, bake is polysemous and bake a potato and bake a knife are both ambiguous. What makes bake a potato and bake a knife sound funny is a thing about the world, not a thing about the words: everybody knows that you can't make either a potato or a knife by baking them. If you didn't know this, you would hear the ambiguity, as indeed you do in John is baking something or what is John baking?⁹ Contrary to JP, bake is lexically ambiguous and the semantics of bake NP offers no argument that the lexicon contains ontological information.

But is the bake case representative? We'll have a quick look at a scattering of JP's other examples. We claim that JP has no convincing cases where the meaning of a governing expression is modulated by the lexical content of the expressions that it governs. In fact, it's our view that this never happens.

Consider JP's treatment of begin. The idea is that begin (/finish) picks up the "telic role" of its direct object (115-117). This is to account for such inferences as:

- (3) a. ...begin a book ---> ...begin to read a book
- b. ...finish a cigarette ---> ...finish smoking a cigarette
- c. ...begin a beer ----> ...begin drinking a beer etc.

It's this sort of example that provides JP's primary evidence that information about function is part of the lexical entry for book, beer, and the like.

On second thought, however, the analysis of begin doesn't work; it is not, in general, the case that if what NP denotes has a function, then begin NP means begin to use NP to perform its function. ...begin a car doesn't mean begin to drive a car; `...begin a thermometer doesn't mean begin to measure the temperature... etc. Enjoy, which is another of JP's favorite examples of a verb that incorporates the telic role of its direct object (88), fails for the same sort of reasons. JP notes ...enjoyed the meal --> ... enjoyed eating the meal, etc. But his account incorrectly predicts the well-formedness of ...enjoyed the doorknob; ...enjoyed the Federal Government¹⁰ enjoyed the carpet tack, all of which are bad. The last example is especially embarrassing for JP since enjoyed using the carpet tack is fine.

In fact, use may seem more suitable for JP's purpose, since if X uses Y, where Y is something that has a (conventional/typical) use, then the invited inference is that X uses Y for what Y is used for. (Even so, there's something odd about ...used a glass of beer, `...used a meal, etc.) In fact, however, use raises a nasty problem for JP since it's perfectly possible to use things that don't have uses (...a rock to break a window;

...snow to make a snowball, etc.). Enjoy works that way too since one can perfectly well enjoy things that have no function (the clement weather; dancing the eightsome reel, etc.) These examples raise a serious question about which JP hasn't much to say: what happens if a verb makes a demand on an argument that the lexical entry of the argument doesn't satisfy?

JP's view may be that, if the governing expression demands X in the lexical entry for the governed expression and that entry isn't marked for X, the resulting phrase is interpreted as existentially generalized in respect of X. We infer this from JP's treatment of Mary believes John (120-122). According to JP, believes wants a proposition in the lexical entry for its direct object, which, however, the lexical entry for John presumably doesn't contain. Though we don't claim to understand the mechanism by which it's achieved, JP's solution is that the interpretation of Mary believes John contains an existential quantifier over the required propositional argument; something like MARY BELIEVES WHAT JOHN COMMUNICATED.¹¹ Assuming that believe John constitutes a precedent for use a rock, then John used a rock should come out meaning something like John used a rock to perform some function.

_____ If, however, that is the treatment that JP has in mind, it's certainly untenable. The problem is that, given interpretive mechanisms that strong, it becomes unclear how a combination of a verb with an NP argument could ever fail to be interpretable. For example, JP's own analysis of begin a rock is undermined; it should be (but isn't) heard as meaning begin to use a rock for something. In like wise, why doesn't John asserted a rock mean for some rock-involving proposition P, John asserted P? Conversely, why doesn't Bill used that John is tall mean for some function F, Bill used the proposition that John is tall to perform F (compare Bill used John's being tall to illustrate the effect of diet on growth). What JP really needs in order to argue for telic roles in the lexical entry for NP arguments, is a verb which picks the telic role when there is one and which yields unsemanticity when there isn't one. But, as far as we know, there is no such verb.¹²

We pause to summarize the argument so far. JP's general strategy is to explain the apparent polysemy of governing expressions by appeal to the semantic heterogeneity of the objects they govern: if X is a governing expression that is prima facie A/B polysemous, assume that the expressions that it governs are lexically cross-classified as being either A or B.

We know of no cases where this strategy works. For one thing, it predicts that an A/B polysemous expression should be univocal in the A direction when it governs a univocal A-expression; whereas, in all the cases we can think of, A/B polysemous expressions are A/B ambiguous in both A and B contexts (cf. bake a cake). For another thing, the claim that governing expressions inherit semantic content from and only from the lexical entries of the expressions they govern fails in both directions: Begin inherits a "telos" from a cigarette but not from a car, and the propositional construal of believes John can't derive from the lexical entry for John. And so on.

We conclude:

i. Apparent polysemy is generally real; the reason bake seems to be lexically ambiguous is that it is. This is compatible with a denotational semantics on the assumption that baking (creative) and baking (warming up) are different processes. JP provides no grounds for doubting this assumption.

ii. There is no evidence that the meaning of governing expressions is ever modulated by the semantics of the expressions that they govern. For all the arguments show so far, bake behaves the same way in bake a cake and in bake a potato (i.e., it's ambiguous in both); enjoy means the same thing in enjoy a cigarette as in enjoy the sunset (i.e., it means enjoy), good means the same thing in good knife and in good car (i.e., it means good), etc. If bank means something different in bank a check and bank a plane, then bank is not polysemous but homonymous, and the context effects are not modulation but selection.

PART 3: COMPOSITIONALITY AND LOGICAL FORM.

We start this section by distinguishing two kinds of issues that JP's treatment generally runs together: on the one hand, questions about how complex lexical entries are; and, on the other hand, questions about the generativity of the lexicon. As we've seen, JP typically argues for lexical complexity by claiming that it's needed to account for generativity. (Cake must contain ARTIFACT in order to account for the polysemy of bake; telephone must contain a "telos" in order to account for the polysemy of use, etc.) However, in principle, the issues are dissociable. Someone who agrees with us that JP's arguments for lexical complexity are unconvincing and who shares our suspicion that content is just denotation, could wonder, nonetheless, whether a denotational lexicon mightn't be generative. So let's turn to this.

Suppose, for the sake of the argument, that the lexical entry for cake contains MADE BY BAKING. What happens next? Actually, JP tells two different stories (sometimes on the same page), apparently not noticing that they are different. First story: the lexicon is generative; cake contributes MADE

BY BAKING to the meaning-in-context of bake, thereby bestowing a "creative" sense on the verb. "Thus, we can derive both senses of verbs like 'bake' by putting some of the semantic weight on the NP. This view suggests that, in such cases, the verb itself is not polysemous [i.e., it has only one lexical entry]. Rather, the creation sense of 'bake' is contributed in part by the meaning of 'a cake'.... The verb appears polysemous because certain complements... add to the basic meaning by co-composition" (124).

That's one of JP's stories; here's the other: "The semantics for the VP [our emphasis] 'bake a cake' results from several operations.... The operation of co-composition results in a qualia structure for the VP that respects aspects of both constituents..." (124). "The result of co-composition is a semantic representation at the VP level [our emphasis] that is identical in structure to the lexical form for a creation verb such as 'build'" (125). According to this second account, the lexicon isn't generative after all; cake never contributes anything, at any level, to the representation of bake. Rather, cake contributes its lexical content (atomistic or otherwise) to the interpretation of the VP phrase bake a cake (as does, mutatis mutandis, bake).

The polemical reason for insisting on the difference between these two stories is that it's news if the lexicons of natural languages are generative; everybody always thought they are just lists. But it's no news that there's an infinity of VPs; only a connectionist could doubt it.¹³

However, we have a second reason for stressing the distinction: while we're very dubious that the meaning of a word is ever a function of its context, we have no doubt at all that (barring idioms) the meaning of a phrase is always a function of the meanings of its lexical constituents. In short, everybody, whatever he may think about the generativity of the lexicon, has to face the problem of compositionality; the problem, that is, of saying how lexical semantics contributes to determining the semantic interpretation of phrases. And, though we take the question to be very largely moot, we're inclined to think that the exigencies of the compositionality problem really do require that lexical entries can't just specify denotations even assuming denotation is all that there is to content.

Consider the phrase want a beer; and assume ---what, to be sure, is tendentious--- that the right interpretation of this phrase is want to have a beer¹⁴ whereas, by contrast, the right interpretation of drink a beer is just drink a beer. Question: how does the theory that derives the meanings of phrases from the meanings of their lexical constituents insure this difference of interpretation?

One way would be to agree with SEL that there are two wants, differing somehow in meaning, one of which takes an infinitival complement, and one of which takes an NP. But here we agree with JP; this treatment is too unrevealing to be plausible. In particular, it misses the equivalence of wants a beer and wants to have a beer; and it fails to explain why, if you want to drink a beer and I want a beer, then we both want something and we both want to have something,¹⁵ and it fails to explain why the two wants are in complementary distribution. So, what's the right story about want?¹⁶ Roughly, we think, something like this: want denotes the relation that holds between x and y when and only when x is a creature & y is a state of affairs (or whatever kind of thing infinitive expressions like to have a beer denote), and x wants y.¹⁷ The way to avoid the lexical polysemy is to maintain that want denotes this relation both in wants a beer and in wants to have a beer; i.e., it denotes this relation whether its complement is an infinitive or an NP.

There is, of course, an immediate problem for this proposal: a beer isn't a "state of affairs" i.e., it isn't the sort of thing that infinitival complements denote. So, how could want a beer and want to have a beer denote the same relation? This is, of course, just the polysemy problem over again: how can the semantics insure that expressions of the form want NP get the same semantic value (viz., the same denotation) as corresponding expressions of the form want INF without assuming two lexical entries for want? Or, to put the same question in slightly different terms, how does a compositional semantics operate to assign the interpretation wants to have NP to expressions of the form wants NP?

Here's the heart of our proposal: lexical entries are allowed to be complex. When it is, the lexical entry for an item specifies (i) its meaning (viz., content, viz., denotation) and (ii) a rule of composition which contributes to determining the logical form of the phrases of which the item is a constituent.

A (very) schematic derivation of the interpretation of expressions of the form wants X (where X is a lexical item) should serve to introduce the general idea. We assume that the semantic interpretation proceeds, node by node, from the bottom to the top of a (surface) syntactic tree. In the present case, the crucial steps are:

Domain:

.VP

.V .NP
wants X

Stage 1 input: <wants_V, X_{NP}>

operation: Assigns lexically specified semantic interpretations.

output: Interpretations of (viz., assignments of denotations to) the lexical nodes: Assigns to the V-node the set of ordered pairs <y,x> such that y wants x, and assigns to the NP node the lexically specified denotation of X.

Stage 2: input: the domain tree with the lexical nodes interpreted as per stage 1.

operation: Interprets the node VP

output: VP is assigned {y: y wants to have F(X)} where F(X) designates the interpretation that X receives in Stage 1.¹⁸

note: We assume that the operation in stage 2 is driven by a composition rule that is part of the lexical entry for want. Viz: if the constituents of VPi are <wants_V, X_{NP}>, then the interpretation of VPi is want to have F(X)

Several comments, in no particular order.

1. The derivation of wants to drink a beer assumes that the surface syntax of that sentence is identical to its logical form. The composition rule assigns to the VP node the set {y: y wants to drink a beer}. In this case, the composition rule is presumably not lexically governed; i.e., it's the unmarked treatment for phrases consisting of a verb with an infinitive complement.

2. Notice that, according to this proposal, want never means (denotes) anything except a relation between a creature and a state of affairs; not even in wants a beer; i.e., not even in an expression where its surface complement fails to denote a state of affairs. So, want isn't polysemous; it's content is absolutely context invariant. We're as far from a generative lexicon as it is possible to get.

The cost of this univocality is complex lexical entries, which determine not only the content of an item but the logical syntax of the phrases to which they contribute their content.

3. The proposed mechanism does the same sort of job that JP's notion of "type coercion" is designed to do. There are, however, differences other than the fundamental one which distinguishes a semantics of complex lexical senses from an atomistic lexical semantics of denotations. For example, it's part of what we take to be the context insensitivity of want that it always introduces the same "light verb" into the VP it governs. This repeats the remark we made in fn 14: We assume that wants a pretzel means wants to have a pretzel not wants to eat a pretzel. The meaning of 'want' doesn't decide what one can want a pretzel for.

Indeed, we think, exactly contrary to JP, that it is this consideration that distinguishes verbs like want, believe, keep,¹⁹ which introduce light verbs, from verbs like use, enjoy and begin, which don't. According to us, the semantic interpretation of use NP is the set {y: y uses NP}. Part of the evidence that this is so is that there is no relation (other, of course, than using) that y has to have to the NPs that y uses. Contrast the necessary truth that every y that wants NP ipso facto wants to have NP.

4. There is a variety of lexically governed effects on logical form other than light verb introduction. For example we think it's plausible that good introduces a quantifier into the interpretation of good NP; roughly, a good NP is one that is good for whatever it is that NPs are supposed to be good for (cf., Ziff 1960). Notice that this treatment makes good context insensitive; good quantifies over the function of the NP it modifies, and the way it does so is independent of which NP it is. Since the meaning of good is context independent, the lexical entry for the NP needn't specify a "telos"; so, JP to the contrary notwithstanding, the semantics of good knife provides no argument that knife has a definition that includes its function. (This is just as well since, as JP himself remarks (43) good children and good weather are perfectly OK, though neither children nor the weather have functions. See above for the corresponding point about enjoy and use.)

5. Our discussion hasn't assumed that there is a level of logical form at which, for example, want a beer is represented as WANT TO HAVE A BEER. On our account, all that happens is that wanting to have a beer (material mode) is assigned as the denotation of the expression wants a beer. Our treatment is,

however, compatible with positing an explicit level of logical syntax should there prove to be any reason to do so; we have no views on the matter.

We close with two pending questions: the first strikes us as not very urgent but the second is vital.

First: What about polysemy? We don't have a theory of polysemy beyond the suggestion, implicit in the preceding, that where it is sensitive to the syntactic structure of the context, polysemy belongs not to the theory of content but to the theory of logical form. That leaves lots of residual cases like lamb (meat v. animal), window (the opening v. what fills the opening), newspaper (the thing you read v the organization that publishes it) and so forth.

We suspect, actually, that there is nothing interesting to say about such cases; the meanings of words can partially overlap in all sorts of ways, so there are all sorts of ways in which polysemous terms can differ from mere homonyms. Nothing in the literature convinces us that there are powerful generalizations to state.

Surprisingly, JP apparently shares this view so far as the polysemy of nouns is concerned (see, circa 90-95). In cases like lamb, window, etc. JP does exactly what SELs do: he has branching lexical entries which allow him to say, for example, that lamb always means physical object, but doesn't always mean animal or always mean food. This is really unavoidable, given JP's architecture: Governing expressions get their ambiguity from what they govern, so the governed expressions have to get theirs from the lexicon. Of course, as JP himself points out when discussing SELs, appealing to branching entries to distinguish polysemy from homonymy is a merely notational solution. It... "accounts for the data, but in a post hoc fashion, without making any predictions as to whether a particular datum should be possible or not" (42). Bank river vs bank building is presumably homonymy rather than polysemy, but both banks, like both lambs, mean physical object.²⁰ Nor is it clear why, if nouns can be really (viz., lexically) polysemous, verbs can't be too.

Second: Prima facie our notion of coercion is more constrained than JP's. For example, it can't turn out that the content of one expression ever depends on the content of another; it's not allowed that 'want N' sometimes means want to eat N and sometimes means want to drink N.

Nevertheless, we do allow that the logical role of an expression can be determined by the lexical entries of the forms that govern it; beer is the logical object of drink in want to drink a beer, but it's the logical object of have in want a beer. This departs from the most rigorous notion of compositionality, according to which each constituent contributes its content and only its content to its hosts, and the effect of a constituent on its hosts is absolutely context independent.

The entirely rigorous notion of compositionality seems to us almost certainly not attainable. For example, want and the like to one side, it is a truism that logical vocabulary is typically defined 'in use', so that the lexical entry for (as it might be) the determines the logical role of the NP in the NP and does so in a way that is specific to 'the'; presumably the logical role of the NP in an NP is quite different. Indeed, it's arguable that the lexical content of the is exhausted by what it says about the logical form of the NP.²¹

But if pure compositionality can't be had, it also can't be that just anything goes. What we've said is tantamount to: There can be context effects on what a lexical item contributes to logical form but not on what it contributes to content. But since we don't know how to be rigorous about the form/content distinction, we're not persuaded that the notion of coercion that we've sketched actually succeeds in constraining the ways that a constituent can affect its hosts. Which is to say that, though we're sure that language is compositional, we don't know what the claim that it is amounts to. That, and not polysemy and the like, seems to us to be the main problem of lexical semantics.

FOOTNOTES

1. The corresponding doctrine is generally supposed to be true of the meaning (content/'sense') of mental entities like concepts and beliefs. Our discussion is intended to apply, *mutatis mutandis*, to either version of IRS.

2. Cited forms are in italics (single underline in ms). Expressions in small caps (double underline in ms) stand for semantic values; e.g., meanings or denotations. Semantic representations are in full caps. However, citation conventions in quoted passages are as per the quoted text.

3. If it's assumed that it's necessary that dogs are animals, then, of course, whatever denotes a dog denotes an animal. But, according to the present view, the lexical entry for dog does not provide this information; it says only that dog denotes dogs. Correspondingly, according to the present view, knowing that dogs are animals is not necessary for knowing what dog means.

4. Whether JP thinks that the semantic representations that the lexicon specifies actually determine denotations is unclear from the text, though there are passages that suggest he does. For example, discussing the "qualia structures" which constitute the lexical entries for nouns, JP says that they "...contribute to (or, in fact, determine) our ability to name an object with a certain predication" (85).

5. JP says nothing about how to decide which lexical relations are on the list of interlexical relations to which grammars must be responsive, or whether the enumeration he provides is supposed to be complete.

6. Transitivity severely compounds the problem. Since it's necessary that what's circular isn't triangular, it follows that you can't know what either square or circle means unless you know what triangle means. And so on ad infinitum.

7. Terminological stipulation: We take ambiguity to be the generic property of which polysemy and homonymy are species.

8 How one gets such exclusion (/inclusion) relations to hold among semantic representations without, in effect, resorting to meaning postulates is a traditional, and unsolved, problem in the literature on lexical semantics.

9. Assuming that gapping is semantic, JP's theory requires that John baked the potatoes and Mary the cookies means that Mary heated up the cookies, not that she made them. We don't think that prediction is true either.

We suspect that gapping is semantic, and that it distinguishes polysemy from true homonymy; hence *he banked the plane and the check.

10. Readers who doubt that the Federal Government has a function may vary the example to taste.

11. In particular, according to JP, the semantic representation of Mary believes John is the existential generalization of the p_2 in 'BELIEVE (\wedge_{p_2} JOHN) (MARY)' (122). The problem is: Where does the propositional constant p_2 in this formula come from? Surely not from the lexical entry for John since John doesn't mean anything propositional? As JP himself says "unlike the case of ... the book... the [proposition]... type required by selection is [not part of] the [lexical] interpretation of the complement" (121). As far as we can tell, JP never does suggest an answer to this question. (About raises a comparable problem, since it seems to govern an event in a book about Vietnam [i.e. about The Vietnam War] but not in a book about Nixon. Where does this event come from? JP tells us only that about "covers" objects of two different ontological types in the two cases; which, we suppose, is just to say that it's polysemous. Ah well.)

12. On 137ff, JP provides an (unconvincing) explanation of John asked me (/wondered) the temperature. But he doesn't face the question why, given that John wondered what the temperature is is well formed, doesn't John wondered the temperature have a reading on which the temperature is "coerced" to an interrogative proposition?

13. We're not, by the way, denying that JP's two stories are compatible. Perhaps he thinks that cake adds to the interpretation of bake a cake by modulating the meaning of bake. Our point is that evidence that cake contributes to the meaning of bake a cake is neutral as to whether the lexicon is generative.

14. JP, of course, does not think this is the right interpretation; it's part of his story that want is generative and that want a beer means want to drink a beer. We'll return to this presently.

15. You should be prepared to admit this even if you doubt there is something (viz., to have a beer) that we both want or both want to have.

16. We're concentrating on whether want NP is polysemous and ignoring the questions that are raised by its intentionality (viz., by the opacity of the NP position to existential generalization and substitution of coreferentials.) If you like, assume that wants expresses an 'intentional relation', and plug in your favorite account of them.

17. Actually, the talk of creatures, states of affairs and the like is entirely heuristic. That's because, on the present view, the lexicon isn't required to specify a sense for want but only required to say what 'want' denotes. Perhaps, if x wants y, then, necessarily, x is a creature and y is a state of affairs. But a purely denotational lexicon isn't required to say that this is so, any more than it is required to say that dogs are animals or that 2 is prime. A denotational theory thus avoids the commitments to unanalyzed semantic entities and properties ("agents", "patients" "events", "transitions", "being in some sense affected," etc.) by which, as we've seen, standard views of lexical semantics are plagued.

18. If the sentence under interpretation is of the form NP1 wants X, a later cycle will identify y with the denotation assigned to NP1.

19. We think that 'keep' always expresses a relation between an individual and a state of affairs, just like 'want'. Thus, 'John kept the crowd happy' is true of John iff he maintained (i.e., kept) a state of affairs which consisted in the crowds being happy, and 'John kept the money' is true of John iff he maintained a state of affairs which consisted in John's having the money. For extensive discussion of 'keep', see Jackendoff (1992); Fodor (forthcoming).

20. We're not at all sure what JP's dot operator does (see 92ff), but we suspect that it has the effect of allowing components of lexical entries to combine freely, in the way that features do. If that's the right reading, then JP's lexical entries are even less constrained than those of classical SELs, like Katz and Fodor (1963).

21. If this is right, then there are perhaps three kinds of lexical entries: ones which specify only a composition rule (the, and), ones specify only of a denotation (Tom, eat), and ones which specify both (want, believe, good).

REFERENCES

- Cruse, D.A. (1986) Lexical Semantics Cambridge University Press, Cambridge.
- Fodor, J. (forthcoming) Concepts, Oxford University Press.
- Fodor, J. and Lepore, E. (1992) Holism, A Shopper's Guide, Basil Blackwell, Oxford.
- Higgenbotham, J. (1994) "Priorities of Thought," Aristotelian Proceedings Supplement.
- Jackendoff, R. (1992), Languages of The Mind, MIT Press, Cambridge, Mass.
- Katz, J. and Fodor, J. (1963) "The Structure of A Semantic Theory," Language, 39, 170-210.
- Milne, A.A., (1957) The World of Pooh, E.P Dutton, N.Y.
- Pustejovsky, J. (1995) The Generative Lexicon, MIT Press, Cambridge Mass.
- Weinreich, U. (1959) "Travels through Semantic Space," Word, 14, 346-366.
- Ziff, P. (1960) Semantic Analysis, Cornell University Press, Ithaca.