*Warps: Some Thoughts on Categorization**

Juan Uriagereka, Linguistics, University of Maryland

I. The Limits of a Derivational System

Uriagereka (forthcoming) tries to push some of Sam Epstein's ideas concerning derivations within the Minimalist system to a logical conclusion. The result of that paper suggests that we should go back to a traditional model, most thoroughly explored in the early seventies (a specific proposal exists in Jackendoff (1969)). The Minimalist version is slightly different, but the main point remains: LF and PF are accessed as the derivation proceeds. The question is what are these entities--LF and PF.

It is not clear that they exist as levels of representation, a notion borrowed from Concatenation Algebras. Since Chomsky (1955), Linguistic theory is taken to provide a (universal) array of these levels. Each is a system L based on a set of primes; the operation of concatenation (which forms strings of primes of arbitrary finite length) and an associated class of strings which is characteristic to L, or L-marker; and a mapping between various L-markers. For various reasons that I will not repeat now, these specific formal objects do not exist within Minimalism.

But how does this conclusion square with the first major premise of Minimalism, that a (virtually) conceptually necessary property of natural language sentences is their pairing intentional/conceptual structures with articulatory/perceptual ones? I have argued that one may adhere to this premise without being forced into having to admit levels of PF or LF. The key is to admit, instead, PF and LF components. This may seem like innocuous terminology; it is not. For if we have a level, we may dictate conditions on it; for example, standard Minimalism takes Full Interpretation to hold of PF and LF: derivations converge only if they meet certain specified requirements at the level. These conditions, then, are virtually defining the substantive character of linguistic objects--at the levels. If these go, there should not be any room left for derivations converging if they meet certain standards.

It is probably worth emphasizing this point. As things stand now, Minimalism offers two ways in which derivations may go wrong. One has just been described, and is called crashing. But apart from crashing, derivations may also be cancelled if they do not meet certain formal specifications. This situation appears more often than is perhaps realized. Imagine, for instance, a derivation that only uses part of a lexical array, but otherwise produces a perfect LF and PF.[1] What is wrong with that? How can we say that the derivation does not converge? Various other examples can be constructed. The point now is simple: what we may call the 'radical' Minimalist model that I have been exploring does not allow us to make the distinction between derivational crashing and cancellation; it can only admit the latter.

But then what are legitimate PF and LF objects? The conclusion is evident: if they are to be part of UG, PF and LF objects must be the inevitable results of derivations. While this is not necessarily hard to accept as a reasonable position to explore within the realm of PF, it looks a priori harder to accept as a claim about LF. Here is why:

---

[1] . For example, having the building blocks to say Socrates wondered whether the sophists were sane, but deciding to stop at the sophists were sane, a perfect (PF, LF) pair, albeit not one that corresponds to the relevant lexical array.

- LF would appear to be the place where such matters as reference, quantification, scopal and binding relations and perhaps others, are, if not established, at least partially determined.

- Indeed, LF would appear to be the place to check whether our predicates match our arguments, which presupposes a procedure to, if not establish, at least determine these notions enough for familiar thematic relations, and others, to allow for the relevant matchings.

- Possibly, LF might present the skeleton for a system of presupposition and entailment.

I am not going deny these important facts; I will however note that while the basic intuitions are moderately clear, that they should be expressed in terms of an LF level is less so, and may be where the facts turn into fiction.

I mean the last sentence more than stylistically. There is a very real sense in which language is about matching certain facts about the world (inside and outside ourselves) with certain fictions we create 'about' them. I do not know how well the fictions match, but we certainly use them; indeed are using them this very second. This is not controversial. What is controversial, though, is where the facts should be analyzed--indeed where more than how.

A majority of colleagues would still probably say that the basic facts are semantic in nature, syntax being reducible to semantics. This is clearly a coherent position to take, and if it is to be examined in light of naturalistic expectations, it would even be the one whose tenets are supported by the main stream of evolutionary biology, for which everything there is which is alive must have adapted.[2] Then one may say that semantics evolved for adaptive, communicative reasons, and syntax is at best a way of clearing up ambiguities, and at worst an effective way of theorizing in our path towards semantic understanding. In the limit, this is of course functionalism, and for reasons I have spoken about elsewhere (in particular, (1996)), I think wrong.

Another large chunk of the field would argue that the facts above are both syntactic and semantic. Suppose it is indeed true that radical Minimalism reduces to pure derivations. Then the syntax/semantic position must adhere to the claim that two separate competences exist, one semantic and one syntactic. Or if one believes that the semantic part is in simple correspondence with some articulate syntax, then two syntactic competences must exist: the derivational one, and a representational one that is responsible for irreducible formal conditions that are not expressible as transformations. It is in fact reasonable to say that Chomsky's practice in the Minimalist program, contrary to his rhetoric, is indeed of this latter type. Thus, he has no troubles in using licensing conditions for, in particular, an LF level, such as a thematic criterion, or the interpretability of specific features. These conditions differ from those without which a derivation may be argued to not be even one, such as restrictions on locality, last resortness, the uniformity and upward character of chains, and perhaps others.[3] Once

---

[2] . For various critiques of this position, see the popular works of Gould, Eldredge, Lewontin, and many others.

[3] . For a discussion on whether this is the correct interpretation of these conditions, see Uriagereka (1996:chapter 5).

again, this is the difference between converging as a derivation, and being one, to start with. If this intuitive distinction reflects something real, we plainly are speaking of two sorts of knowledge here--hence two sorts of competences. I am perfectly willing to embrace this position, and have until now. In this paper, however, I would like to explore a third, more radical possibility.

This third perspective asserts that the core facts above are a pure consequence of the derivational system. A version of this view seems to be close to Chomsky's rhetoric, if not his practice. In this view, derivational syntactic (henceforth, simply syntactic) knowledge interfaces directly with the articulatory/perceptual and intentional/conceptual mechanisms of performance. Simply put, if this view has any chances, it must have something to say about those properties above which, while not being plausibly semantic, play a role in the determination of possible LFs. I hasten to add that the view has little to say about a variety of phenomena that do not seem semantic at all--for instance, everything that does not involve command. This, then, leaves a 'no man's land' that one must worry about, and I have.[4] But I put this to the side now. The exercise I propose here takes this third intuition seriously and explores its consequences within a radically Minimalist model, with no levels.

Part of my previous work has attempted to show that some of the structural and relational properties of LF directly follow from what I have called a dynamically bifurcated model, with multiple Spell-out (see both Uriagereka (1996b) and (forthcoming)). For instance, command immediately appears this way, if Epstein is right in his initial premise that this notion follows from derivational Merger. Then no LF-principle needs to make reference to command, for this is all there is--or there could be. Until now, however, I have not had anything to say about the objects of LF (or PF, for that matter). Evidently, if we still need levels to define these objects--whether as labeled categories, role dependencies, chains, or whatever--we are still flirting with the standard notion. Consider, for instance, a @-criterion.

Such a criterion must be checked at LF--there is nowhere else to check it. It may seem that the matter can be taken care of transformationally, roughly in the same way that Case is checked; we could, for instance, assign a theta feature somewhere in the derivation, thereby meeting some requirement of the predicate or the argument. However, Chomsky has suggested that this is the wrong view, given the ungrammaticality of *John used t (which cannot mean that John used himself). If thematic relations were a matter of chains, why couldn't John simply move up to the subject of use, and there get its (second) role? Thus, we take thematic relations to be lexico-configurational, happening between words and their projections, and not between chains and their elements. Then why do we need a Theta criterion at all? The reason, Chomsky suggests, is empirical. The alternative, mentioned in Chomsky (1993), is the notion that @-dependencies are checked in performance (i.e., the semantics): the trivial reason use needs two arguments is because of its meaning. But this makes the prediction that two competing derivations, one violating, and one satisfying the Theta criterion, may in principle outrank one

---

[4] . For example, I am forced to assume that Weak Cross-over phenomena involve the Familiarity Condition (assuming hidden indefinites in all quantifiers that induce the effect). Likewise, I must rely heavily on Higginbotham's (1988) notion of 'context confinement', in order to speak of referential dependencies that do not obtain under command. And I basically have nothing to say about those instances of control that do not seem to occur under command, if these are not reducible to something else.

another, for they both converge--it is only in performance where they differ, in their intelligibility. However, Chomsky notes, this never happens.

The form of that argument is straightforward; however, it is hard to find actual examples that do not violate something else in the derivation. Consider, for instance, *John left vs. John left t, assuming the unaccusative hypothesis. We want the latter to be the correct structure, but we cannot reach this conclusion, since the alternative, in fact unintelligible derivation (for John is not receiving a role from left) outranks its counterpart, due to the fact that it involves no movement. Right? Well, except that the two derivations are not comparable to begin with, since at no point have they shared a partial numeration and did they have different alternatives at that point. The only examples that I know come even close to proving the point are pairs like *I believe t to be a new candidate about to knock on the door and I believe a new candidate to be t about to knock on the door. These two do share a partial numeration as the derivation proceeds, and one does make a more complex move--the good one (the movement of a new candidate). Therefore, if I in the first of these two is even allowed to move to subject position, the example should converge as gibberish, which is good enough to outrank the other. But even here, an alternative exists for the bad instance, as John Frampton points out: in the bad sentence, a new candidate does not get its Case features checked, and hence the derivation, in fact, crashes.[5] The point is, it may well be that the derivational apparatus takes care of the empirical problems that worry Chomsky, and we may then dispense with the @-criterion as a representational requirement on an LF component which, in any case, is really 'about chains', not configurations.

But that sort of reasoning would have to be applied elsewhere, to explain away potential representational conditions. The point above has been about the determination of saturation requirements, the very building blocks of syntactic structures and, ultimately, chains. These are the objects that, up to now, syntax constructs and LF filters. But note in this respect that we have considerably narrowed down the filtering effect. Chomsky (1993) still speaks of certain valid X'-theoretic structures which present linguistic stuff in a way which is recognizable to the computational system. That can only mean one thing. The good structures are recognizable by LF, the bad ones are not. However, an alternative exists: the derivation itself does not deal with the bad structures; to put the matter conspicuously, the derivation does not even construct the bad structures. This is precisely what happens with Chomsky's bare phrase-structure, whose output are the well-behaved objects of X'-theory, now deduced. And similar points arise with regards to the 'projection' of the targets of Merge and Movement, dumb moves yielding all sorts of impossible objects that cancel derivations (see Chomsky (1995:chapter four)).

The natural question then is whether everything else is like this. What else? We have started with projections, so why not ask the same question about categories themselves, functional categories like Comp and Tense, lexical categories like verb or noun, sub-categories among these, and so on. This is the stuff that syntax is made on, since Aristotle; the 'parts of speech'. An outsider to the field may well ask us how come central theoretical constructs have

---

[5] . Under the assumption, of course, that no other source of Case is available for the nominal. A possibility would be for the element to raise covertly to the object Case assignment position of the matrix verb, skipping the 'invisible' trace of the moved subject. For various technical reasons, this possibility may be ruled out as well.

lasted for so long. We may reply that the classics just had an approximate intuition of what nouns and verbs are, just as they had intuitions about atoms and all sorts of ur stuffs. However, we are not much better off than the classics, at least with regards to major lexical categories. We do have--these days anyway--literally dozens of intermediate categories. Perhaps too many, although we think they correspond to actual phenomena of the usual sort. Then again, our outsider might not be discouraged by our situation, and may even remind us that physics faced a very similar challenge in the first half of this century, with a notable inflation in the physics vocabulary for particles (at some point, numbers were added to the names). This did not lead to a crisis, though. It became apparent that researchers had to study families of particles, and then try a different level of analysis, within which each of these families cohered. Soon after, physicists postulated quarks.

The moral is simple; implementing it is not! At any rate, what follows can be seen precisely in this light. In a nutshell, I will be making a proposal whereby the computational system, the engine that powers all UG, literally creates everything syntactic. And I mean everything: interactions and basic objects. Of course, this will force me to go abstract in the level of analysis.

II. Hierarchies and Other Philosophical Matters

The attentive reader may have noted that I spoke, immediately above, of basic objects. I mean this in two different senses.

As Higginbotham has pointed out, Minimalism has an inductive character to it. Apart from some combinatorial apparatus, such a system crucially involves a basic clause of some sort, a step which is presupposed in any object generated by the syntax, no matter how complex. I will be interested in such basic clauses, which will correspond to certain basic objects. But then how could I even attempt to say that the system creates these basic objects? Aren't they what the definition is based on, and hence axiomatic in character?

What I have to say here, in a sense, proposes to have the cake and eat it too. Perhaps an intuition will help understand what I have in mind. Evidently, inductive definitions are not restricted to any specific sorts of objects. Thus, it is common to define the set of natural numbers in terms of some basic clause (for some the empty set, for others a unit object), plus some natural operation (for some, a successor function, for others addition). In turn, once the set of natural numbers is so defined, one could use it as the basic clause for defining some other set; for instance, the set of integers, introducing subtraction. And the set so defined can be used as the base for defining yet another set; and so on--it's basic sets all the way up. I would like to suggest that the structure just described, which I like to call a 'warp' sequence,[6]

[6] . I use the term (which is defined in passing in fn. 23) as an homage to the most interesting concept of the Star Trek series: the 'warp drive'. A 'warp' is a shortcut out of the ordinary space-time fabric--in the same way that a 'worm whole' is--allowing access to a reality which is not defined somewhere, presumably (and crucially) by way of using some clever resource which is partially defined there. It is not accidental that the Star Trek writers also talk about different 'levels of warping'. The key is that once you warp out of ordinary space, you may then warp out of extraordinary space, and then again out of extra-extraordinary space, and so on. The reader who has trouble understanding this intuition may want to read the very amusing Krauss (1995).

is intuitively grasped by humans, and is furthermore the underlying structure of syntactic types, which thus have a certain 'dimensionality' to them.

These are major words. In a sense, though, if taken separately, the major words are old news. First, the topic of mathematical intuition is explored elsewhere in the philosophical literature. Parsons (1980) says:

The properties and relations of mathematical objects that play a role inmathematical reasoning are those determined by the basic relations of some system or structure to which all the objects involved belong, such as the natural numbers, Euclidean or some other space, a given group, field, or other such structure, or the universe of sets or some model thereof. It seems that the properties and relations of mathematical objects about which there is a "fact of the matter" are either in some way expressible in terms of the basic relations of this structure or else are "external relations" which are independent of the choice of a system of objects to realize the structure (149-50).

Indeed, there is a whole branch of mathematics, Category Theory, that is concerned with the structures that allow these sorts of relations.[7] It is important to emphasize here, again in Parsons words, that 'what is really essential to mathematical objects is the relations constituting the structure to which they belong.' I would call these relations basic. This is the second use of the word; it is directly related to the first, if as it seems these relations are 'dimensional' in the sense exemplified above--each dimension constituting the base for defining the next.

Where I differ in perspective with Parsons is in what he makes of the mathematical fact now being raised:

We are taking as a gross fact about arithmetic, that a considerable body of arithmetical truths is known to us in some more direct way than is the case for the knowledge we acquire by empirical reasoning. And this knowledge takes the form of truths about certain objects--the natural numbers. What is more natural than the hypothesis that we have direct knowledge of these truths because the objects they are about are given to us in some direct way? The model we offer of this givenness is the manner in which a physical body is given to us in perception (p.152).

Of course, a more direct way would be that our minds are built that way.[8] However, Parsons explicitly rejects that possibility:

---

[7] . For a very useful introduction, see Barr and Wells (1995).

[8] . Actually, the two positions may be made indistinguishable if one considers that, in order to perceive, one must have a mind built a certain way. Parsons leaves me puzzled, however, with his comment that we 'naturally think of perception as at least sometimes uncorrupted by thinking' (p. 157). Maybe this is just terminology, but how can one perceive without thinking? Probably, Parsons means conscious thinking, which he may also have in mind when talking of mathematical intuition. It should be clear, though, that the kind of thinking I am talking about, as a Chomskyan, is unconscious--the sort relevant for linguistic intuition. Certainly, whatever I have to say here about mathematical intuition is meant in this sense. And it is from this perspective that perceiving in a certain way is to have a mind of a certain kind.

We are not thinking of the capabilities of the human organism, and it may be extraneous to think of this "construction" as an act of the mind [his underscore]. . . I do not want my argument to rest on the notion of a priori knowledge (p. 158-159).

From a certain perspective, the latter statement is reasonable. If one were just interested in a certain psychological object, the position would be termed behavioristic. However, Parsons is ready to concede that in 'some cases, such as natural language, the concepts involved may be innate (p. 162),' and thus it would be unfair to call him behaviorist. Nevertheless, if one is interested in such questions as 'knowing truths about types by a certain kind of perception of tokens, which are then valid for any tokens of the types involved (p.160),' and more generally the validity of reasoning which rests itself on mathematical structures, we cannot eagerly blame the structure of mathematics--granted, as we intuit it--on the structure of our minds. If we do that without any qualifications, we may fall into the most blatant relativism.

Yet, there is a twist, and the main reason why I think this paper may have philosophical consequences--even if wrong. Uriagereka (1996) has tried to argue that Minimalism makes a lot of sense as a theory of a certain chunk of the natural world. For evolutionary reasons (of the exaptive sort), or developmental reasons (of the epigenetic sort, at very elementary levels), or yet unknown reasons, the structure of Human language resembles that of complex chemical systems and--I have tried to show--the formal structures of certain feature plans, like the Fibonacci patterns in the animal and vegetable kingdoms. Suppose this view is right, and the central discrete infinitude, underspecified plasticity, and dynamic economy of the linguistic system match similar properties in the physical word, in the broadest sense of the word (I am purposely trying to avoid the term 'biological world', subject to the fortunes of life and history). Then it is no wonder that the system might have the basic structure of mathematics, at least not more than the fact that systems of wave functions presuppose, all, natural numbers, integers, rational numbers, real numbers, and complex numbers. Surprise about the fact in the case of the mind, and not quantum physics, seems parochial.[9]

The reader may have noted that, in the last few paragraphs, I have used interchangeably the terms 'mind' and 'language'. I should have been more accurate. The philosophical point just raised, of course, only extends to the part of the mind which is responsible for language, abut which we have tons of evidence concerning the three major characteristics above. Then the philosopher could still, correctly, criticize my reasoning: How do we know that the structures of mathematics, which you are so eagerly placing inside the mind, do not fall into that part of the phenomenon (mind) which is not so clean and physical (say, the reptilian brain)? If so, relativism still looms threatening. And then of course I am forced into a position which I am happy to accept, and Chomsky has advocated: It is a linguistic mind that is capable of mathematical intuition (so much for the reptilian brain). Now the argument is complete: claiming that mathematical structures set up the dimensionality of the mind leads to no more relativism than claiming that they set up the dimensionality of the time-space fabric. Granted, we do not know what this means, but this is so in both instances.

---

[9] . The comment may remind the reader of certain speculations of Penrose's, concerning the quantum nature of thought. I have nothing to say about such a claim, which I do not understand. The point here is much more modest: from the present perspective, it need not surprise us that the 'warp' structures of use for quantum fields (translation: the need for expressions presupposing the familiar number system) would show up in language.

Coming back to how all of this might affect the linguist, here is the net result: if the basic structures of mathematical types are the basic structures of linguistic types, every time we talk about nouns, verbs, or about mass terms, states, and so on and so forth, and we mean this talk in a primitive sense, we really are using, to borrow Parsons terms '"external relations" which are independent of the choice of a system of objects to realize the [basic] structure.' In a very real sense, we are doing semantics, or as Higginbotham would put it, 'lexicography'--not syntax/mathematics.

I think lexicography is irrelevant to what concerns me here, which is the properties of the basic objects that enter a syntactic computation, and make it to interpretation. It should be as irrelevant as whether a verb starts with a consonant or has no secondary stress. Of course, nobody would deny this; however, many of us--these days anyway--are looking at the inner structure of such things as Aktionsart, noun systems, and a host of grammatical/functional systems (seventeen in the latest count). Every time we decide to place an Aspect Phrase below or above a Tense Phrase, and so on, we are making a decision which is either lexicographic--hence ungranted, for the syntax--or else should not have been made on those substantive terms, even if right.

What I want to propose here is a way of rationalizing those decisions (whose empirical validity, for the most part, I have no qualms with) in terms of the dimensionality of basic structures. To go on with the number intuition (nothing but a different instantiation of the same basic structure), I will be talking about relations among noun classes, functional categories, and the like, in much the same way that one would talk about relations among sets of numbers of different orders of complexity. If this exercise is correct, it should give us a rationale for why some familiar linguistic structures are hierarchical: verbal/nominal classifications, thematic expressions, auxiliary selection, and so forth. The proposal should also give us a basic structure for lexical entailments and robust presuppositions (as opposed to conversational implicatures), for the same reason that number theory allows us to run entailments concerning the layered structure of number sets.

III. Some Syntactic Background

In the early eighties, Anna Szabolcsi made an interesting proposal concerning the articulated nature of possessive NPs in Hungarian. Kayne (1994) recasts this proposal in current terms, and adapts it to other languages, showing its depth and scope. I am not going to review now either Szabolcsi's or Kayne's arguments for their possessive structure, but I do want to summarize the main point that concerns us here, regarding structures like (1):

(1) a. John has a sister

b. [SP be [$_{DP}$ SP D^0 [[$_{DPposs}$ John] AGR^0 a sister]]]

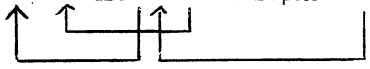

To obtain the surface order in (1a), the possessor John raises to the specifier (SP) position before be through SP D^0 and the D^0, which Kayne suggests is in some sense prepositional,

incorporates into be. The incorporated expression be+P surfaces as have, in the spirit of early work by Benveniste.

Hornstein, Rosen, and Uriagereka (1995) (henceforth HRU) suggest a small improvement for a structure like (1), for empirical reasons concerning part-whole expressions, as in my Saab has a Ford T engine. Given the structure in (1), this sentence would have to be analyzed as in (2a):

(2) a. [SP be [$_{DP}$ SP [D/P]^0 [[$_{DPposs}$ my Saab] AGR^0 a Ford T engine]]]
b. [My Saab$_i$ be+[D/P]^0_j [t$_i$ e$_j$ [t$_i$ AGR^0 a Ford T engine]]]

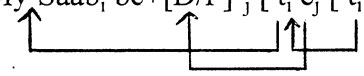

While this poses no particular problem, it prevents us from capturing an obvious paraphrase of (2): There is a Ford T engine in my Saab. HRU show that, indeed, this sentence should be derived from the same structural source as my Saab has a Ford T engine does, which poses an immediate question: Where do we fit the in? In order to respect the leading idea behind the Kayne/Szabolcsi analysis, the correct source structure should be something like (3), which incidentally redeems the intuition that the D^0 is somehow prepositional:

(3) [SP be [$_{DP}$ SP in [[$_{DPposs}$ my Saab] AGR^0 a Ford T engine]]]

But given (3), we must alter the details of the derivations. To obtain the surface order we must raise the predicate a Ford T engine to the SP of in and insert there in the matrix SP position. By the same token, (2) should be reanalyzed:

(4) a. [SP be [$_{DP}$ SP [D/P]^0 [[$_{DPposs}$ my Saab] AGR^0 a Ford T engine]]]
b. [my Saab$_i$ be+[D/P]^0_j [a Ford T engine$_k$ e$_j$ [t$_i$ AGR^0 t$_k$]

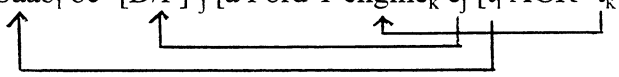

For mostly conceptual reasons, HRU also assume that the element in the SP of AGR is not base-generated there, but actually reaches this position in the course of the derivation, in order to fulfill some checking. Then, the complete, basic structure that the HRU proposal involves is (5). I use the notation X, instead of D/P, in order not to confuse the issue of what it is to have a category which can be either a determiner and a preposition.

(5)

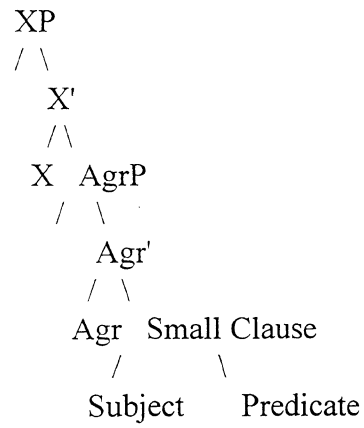

A powerful argument, I think, that the conceptual move in (5) is justified is provided by the paradigm in (6):

(6) a. There are very poor neighborhoods in the city.
b. The city has very poor neighborhoods.
c. The poor neighborhoods of the city.
d. The city's poor neighborhoods.
e. A city of poor neighborhoods.
f. *(A/the) Poor neighborhoods' city (cf. a poor-neighborhood-city).

Observe, given (6a-b), that the relation (city, neighborhood) is of the same sort we saw above for (car, engine); let us call it integral. Now witness what happens when considering the basic structure in (5) alone, without the matrix verb. Let us assume, following and adapting a proposal by Kayne, that the preposition or Case marker of is a realization of AGR, when some lexical material follows it. If so, in (6c) the city could not have left the small clause, while the poor neighborhoods must be at least as high as the SP of AGR:

(7)

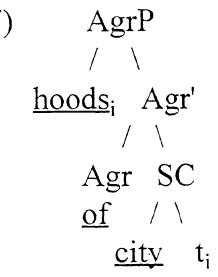

In turn, (6d) must involve movement of the city to the SP of X, taking 's to be the lexical realization of this element and abstracting away from definiteness:

(8)

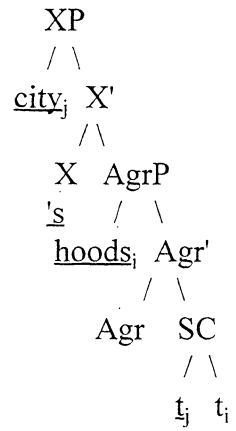

Surprisingly, but following the same reasoning, (6e) must be analyzed as in (9), with a city now having moved to AGR, and poor neighborhoods staying put:

(9)

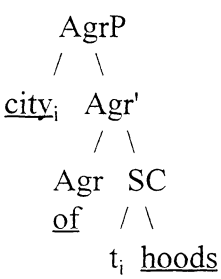

The conclusion is then very interesting: whatever has moved to the checking domain of AGR determines the reference of the whole expression.

There are two arguments for this. The first one concerns the fact that, regardless of the ultimate reference of the expressions in (6), they all express an integral relation (city, neighborhood). The present system expresses the fact in terms of a small clause, of which 'city' is in some pre-theoretical sense the subject, and 'hoods' is the predicate.

Traditionally, it was taken to be the case that parts of wholes and similar integrals are intrinsically relational. Obviously, there is no neighborhood without a city that this neighborhood is a part of. So suppose we represent the relevant notion as something like hood(x,y). Now what's the difference between Washington is several neighborhoods and Washington has several neighborhoods? Both are existential, but whereas the first quantifies over some city which is so-and-so (endowed with several neighborhoods, just as it is endowed with the Capitol and the White House), the second quantifies over several neighborhoods; it asserts, as it were, that there are several neighborhoods whose city Washington is.[10] The main traditional intuition is that these two possibilities are intrinsic to words like neighborhood, and all that be and have do is highlight one construction or the other.

But the first result of HRU is that this cannot be totally right, for the sentence there are several neighborhoods in Washington expresses an existential quantification over

[10] . Keenan (1987) has a very illuminating discussion of these topics, and an analysis of the sort I am about to mention.

neighborhoods as much as Washington has several neighborhoods does. That is, specifically, have is not at issue. What is appears to be more subtle, if the syntax above is on the right track.

The second result of HRU is to emphasize that the phenomenon is considerably more pervasive than it looks at first. Everything that has parts or is a part of something will appear in these constructions. Then it is hard to see what does not count... Even proper names allow us such statements as The Royal Family will always have Charles.[11] The alternative to making Charles relational is to push the relation into the syntax, as we have.

The second argument for the present view is syntactic. One thing is to say that two constructions manifest two quantifications concerning a relational noun; quite a different thing is to show how this is achieved. All of the expressions that concern us now make use, apart from different bes and haves, of elements like of or the genitive 's, which fits nicely in the Kayne/ Szabolcsi picture presented thus far. Furthermore, recall the ungrammatical (6f), which would have to be analyzed, in the present terms, as in (10):

(10)

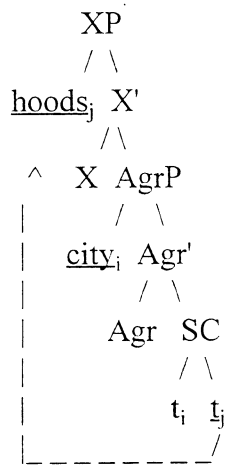

Here, the movement of 'hoods' all the way up, over 'city' (which by hypothesis must be in the checking domain of AGR, where it determines reference) clearly violates the Minimal Link Condition; it is 'too far away'. While we thus predict the ungrammaticality of (6f) with the analysis in (10), the standard analysis cannot account for the facts. Traditionally, it is the 'head' of the construction that provides its reference. From this point of view, the head of (6c/d) must be neighborhood, whereas the head of (6e) must be city, and the lexico-semantic relations between these elements must be expressed in their lexical make-up. But the only way in which the lexical structure could fall into the MLC is if the syntax is essentially as in (10).

The interesting consequence of our proposal is that the referentiality of a basic structure--when it comes down to interpretation--is determined in terms of something other than the intrinsic lexical properties of the elements in the structure. Just because city and

---

[11] . It may be thought that this is a different kind of have, but I am not convinced. In any case, in many Romance languages (e.g., Galician) it is perfectly appropriate to say such things as this:

(i) Na Familia Real hai Isabel, hai Carlos, hai Diana, hay moita xente.
'In the Royal Family, there's Elizabeth, there's Charles, there's Diana, there's many people.'

neighborhood are involved in a relation of a certain sort in these expressions, we cannot immediately determine whether we are talking about a neighborhood or we are talking about a city. The way to determine that, syntactically, is in terms of what has moved to AGR. We thus remove part of the mystique of reference from the lexicon, making this central notion a consequence of a syntactic process. Of course, we still have a substantive residue: to assume the Minimalist machinery, there must be a feature that is checked in AGR, which is presumably added at the point of lexical selection into the array that determines the syntactic derivation (a numeration); AGR attracts this feature, which must mean that, substantively, AGR is something like a referentiality site.

I must emphasize that what we have reached is an important conclusion, independently of the issues that concern me here. This is because the sort of structure that Szabolcsi and Kayne pursue faces an obvious difficulty in structures involving possessor raising, like the Spanish (11):

(11) Le corte la cabeza al pollo.
to-him/her cut-I the head to-the chicken
'I cut the chicken's head.'

The sort of structure we have considered takes a chicken and a head to stand in some sort of integral relation, where the head is a part of the chicken. However, while the syntactic expression of this relation yields all the results that HRU present, it is not less obvious that what the speaker takes himself to be doing in (11) is cutting a chicken's head, not a relation! In other words, the argument of the main verb cannot be a relation; it is a chicken, however we get that to happen. The proposal I have sketched provides a way.

One second thing to emphasize is that my conclusions would look very odd as corresponding to a semantic model like Montague grammar. There, all the notions of interest now are coded as semantic types, which are ultimately responsible for breaking down what in a sentence is referential and what is not, and how the different parts combine. It does not make a whole lot of sense to postulate a category (or corresponding feature) whose sole content is 'reference', for the Montague grammarian takes reference to be a property of a certain type. Indeed, the events that follow will emphasize this conclusion.

In contrast, the ideas presented so far are naturally interpreted in the Neo-Davidsonian program, if we think of the two elements that relate through the small clause as satisfying roles of some sort. Until we give content to these roles, let us just call them T1 and T2. Then, the nominal expressions in (6) will have the simple, rough semantics in (12):

(12) a. Thee [hood(e) & T1(city,e) & T2(hood,e)] ...
(The poor neighborhoods of the city, the city's poor neighborhoods)
b. Ee [city(e) & T1(city,e) & T2(hood,e)] ...
(A city of poor neighborhoods)

Note that whereas the reference is different in each of these instances (depending on what is predicated of the variable e), the thematic roles are the same, thus capturing the integral relation between city and neighborhood, whatever that relation turns out to be. The

corresponding syntax is simple. For the sake of mnemonics, we may think of AGR in (10) as a R(eferential) head, and may so label it; correspondingly, we may postulate an [+r] feature that is added in the numeration and is checked in the checking domain of R, via movement. Of course, this is the familiar practice from the literature on functional ghosts, but bear with me for a moment; we will soon go deeper.

We must also distinguish the poor neighborhoods of the city from the city's poor neighborhoods. One possible way of deciding on how these two differ is in terms of the substantive nature of the X element in (10), which Kayne felt had a status as either a P or a D. I suggest we treat this element as a two-place relation, which can be lexicalized either as a determiner (a two-place relation if the lexical basis for a generalized quantifier) or a preposition like in (also a two-place expression). This would have one extra advantage. In those instances in which X expresses a quantificational relation (henceforth represented as Q, even if may also be lexicalized as a preposition), we can invoke the fact, stressed in Higginbotham (1988), that this sort of element must introduce a context variable C, given that human quantification is contextually restricted. Then, it would be natural to suggest that whatever moves to the SP of Q has a contextual character. In particular, we may take speakers to confine the range of whatever quantification Q invokes in terms of the element C that moves to this SP:

(13) a.

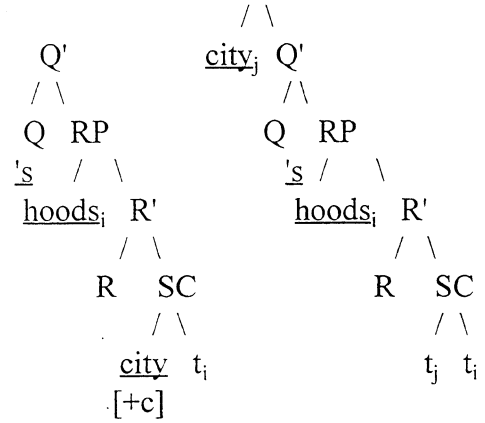

b.

c. [Thee: C(e) & hood(e) & T1(city,e) & T2(hood,e)] ...
Where the speaker confines the range of C to city

Note that the logical form in (13c) includes the context variable C, predicated of e. Whatever is placed in the syntax in a position to have to move to the SP of Q, for reasons that I return to, ends up being taken by the speaker as the element that anchors the context of the QP. This allows us to distinguish the city's neighborhoods from the neighborhoods of the city. In the latter, either no element has moved to the spec of Q (hence the contextual specification of the expression is open), or alternatively 'hoods' does. I do not know which of these options is best, but it does not matter for my purposes.

What counts is that we now have a complete syntax for the Kayne/Szabolcsi expressions, and furthermore we have a neo-Davidsonian semantics for them. In what follows, I plan to suggest that all syntactic structures reduce to these basic structures. In other words,

rather than having various categories with different labels, including multiple functional categories, I will suggest that what we encounter are recursions of the basic structures, as follows:

(14) [X^0 [R^0 [$_{SC}$[R^0 [$_{SC}$[R^0 [$_{SC}$...]]...]]...]]]

For this to make any sense, though, we must understand several things. First, what is the recursion in (14) based on? Second, what is the exact nature of the elements that enter into the small clause? Third, what is the substantive character of the R's in these basic structures? I would like to suggest that all of these questions are related.

IV. In and Out of Space

Let us suppose that what we usually refer to as the 'subject' of the small clause denotes a mental space with a dimensionality nD. Furthermore, let the 'predicate' of the small clause denote nothing other than what I have called a warp in that space at nD. If so, by the mere application of the warp to the space, we will force this space into an n+1 dimensionality. For example, in the instance discussed in (13), city would define a given nD space, and hood would define a warp to this space. By applying hood to city, we obtain a space at a new dimension n+1. Theoretically, we could then work within this space in terms of a warp for n+1D spaces, and so on.

There is every temptation to take the semantic value of the word neighborhood to be what relates to the semantic value of the word city. Whatever a 'neighborhooded' city is (some partition, say) is that because of properties of cities and neighborhoods; yet I am not saying anything of the sort. That would be like saying that when we divide two Spanish omelettes between four diners, it is a relation between omelettes and diners that is at stake. Surely there is one such relation, but it is irrelevant at the level that interests us. What we want to understand is what relation exists between two and four in terms of divisions, whether it is expressed through omelettes and diners, fortunes and inheritors, or whatever. To use Parson's term, these substantive relations are external, independent of the structure.

Of course nobody denies that, at least in the case of arithmetical operations; but I am actually extending the claim to lexical semantics, and indeed making a non-obvious move. Bluntly put, I am claiming that there is something crucial to the syntax of a part-whole relation, in making it a part-whole relation. You may not know what a zaptrack or a gropstench is, but if I talk of 'the zaptrack's gropstench', 'the gropstench in or of the zapstrack', 'a zapstrack of remarkable gropstenches', you will know that zaptracks and gropstenches stand in some sort of integral relation. The syntax tells you that. Of course, the syntax corresponds to a neo-Davidsonian semantics, where the T1 role above is now interpreted as a primitive space, and the T2 role as a primitive warp in that space. This is no more or less motivated than it is to call a role an 'agent', a 'theme', and so forth. In fact, I could have called T1 a 'whole' and T2 a 'part'. But that would be the mistake I am trying to avoid, making the relation substantive again. The real linguistic point of this article is the claim that a whole is to a part as a mass is to a measure, and so forth; just as the set of rational numbers is to division as the set of natural numbers is to subtraction: both operations warp you out of that domain, and cannot be fully defined there. The only difference between all of these is dimensional: the same process is happening at

different orders of complexity. If I am right, it is a syntactic tool that allows us to code these layered orders and their relations.

I will then write an expression like (13c) as in (15), where S stands for 'space' and W for 'warp':

(15) [Thee: C(e) & hood(e) & S(city,e) & W(hood,e)] ...
Where the speaker confines the range of C to city

(15) is homomorphic (and perhaps even isomorphic) with the corresponding syntactic structure, as discussed above.

Next, let me present the 'Aristotelian' use that Keiko Moromatsu has made of this system, which I believe will help us understand some of its more technical aspects. Muromatsu (1995) argues that nouns classes are naturally expressed in terms of the dimensional structures I have sketched. She substantively interprets the 1D dimension as expressing concepts (her qualia), the 2D dimension as expressing substance (her quanta), the 3D dimension as expressing form (her forma), and the 4D dimension as expressing change.

A few interesting properties that Muromatsu can predict include the following. Natural languages distinguish abstract and concrete, mass and count, animate and inanimate nouns, and others. Traditionally, this has been expressed in terms of binary parameters [+/-concrete], [+/-count], [+/-animate]. These parameters predict 2n possible noun phrases, for n the total number of parameters. However, less combinations exist. For example, one can imagine an animate expression which is a mass term (a lake, given traditional mythology) or an abstract expression that is countable (for example, a number); curiously though, human languages do not use specific morphemes to codify the properties of any of these. That is, a typical noun classifier tells us, essentially, that a given noun is of a certain countable sort, provided that it has a mass and a conceptual support. Thus, there are no noun classifiers which are specific to numbers, or alphabet symbols, or some such thing, even if these are obviously countable. Likewise, a typical measure element tells us that a given noun is of a certain measurable, mass sort, provided that it is seen as mere 'stuff'. That is, when I say 'here's a bucket of chicken for you', I cannot use that expression to mean that I am passing a living chicken on a bucket.[12]

The point is very simple: there are some obvious lexical entailments related to the nature of noun-classes, in such a way that if N is a grammatically animate noun, then N is also count, and so on. We can of course code these entailments through meaning postulates in the obvious way. However, if Muromatsu's classification is correct, the entailments directly follow from the way in which each dimension embeds into the next. Thus, for instance, all grammatically animate nouns have to be countable, all grammatically countable nouns have to have a substance, and so on.

---

[12] . The reader may think that this has to do with the nature of reality, but this is far from obvious. A large octopus, for instance, fits naturally into a bucket, and in fact possibly two separate buckets if the animal is very large. Nevertheless, passing two buckets containing each four legs and part of the body of a live, whole octopus does not accord well with the expression 'here's two buckets of octopus for you.' Typically, that expression would be used to refer to the octopus stuff, cut into pieces and deprived of life.

Second, given the hierarchy qualia > quanta > forma > ..., we should only start seeing systematic grammatical markings at the 2D quanta level. The reason for this has to do with the suggestion made on section II that nD warps to n+1D. This tells us that no warps are needed to stay within 1D. If so (and assuming Muromatsu's specific proposals), nouns used in a 'conceptual' fashion should not need to be accompanied by any grammatical mark of warping. This arguably relates to otherwise tricky notions, like the prototypical the lion or the generic a lion. These expressions pose two related problems.

One is grammatical: they lack significant markings, at best taking default gender or number. The other problem is semantic: How do they get to mean what they mean? The two problems are addressed as the same one if we take these sorts of expressions to be dimensionally trivial. What this means semantically is that the expressions do not denote things, or even entities whose presence is substantive; they are not reified or even concretized, which seems to accord with their meaning. The syntactic consequence is equally straightforward: they need not (hence cannot, given Minimalism) involve small clauses. Thus, as it were, in the lion or a lion what you see is what you get. If this does not get to denote a concrete lion or a set of lions (in relevant readings), this is because at 1D there are no things in the semantic model.

An important clarification is in order here. The reader is explicitly warned against any temptation to interpret the 1D lion as a set of lions, or--worse still--some such set in all possible counterfactuals. To insist, we have not built in the formal apparatus needed for things yet, hence could not have lions, which are at least things. In fact, what I am trying to say is that the whole question is ill-conceived in those terms, and thus it is pointless to argue whether it is best to talk about these or the other categorial constructs. At 1D, in the present model, the only categorial construct is raw space itself, with no strings attached even to a substance, let alone the sort of organized substance involved in a thing.

Third, Muromatsu's analysis has very interesting consequences for the syntax of classifiers, measure phrases, and the like. Specifically, Greenberg (1977) tells us that out of the six possible combinations of nouns, classifiers, and quantifiers, we only obtain four, and only two of those are possible in any given language:

(16)

Possibility A	Possibility B	Impossible
Q CL N	CL Q N	*CL N Q
N Q CL	N CL Q	*Q N CL

Observe that, as Greenberg notes, the relative order {Q, CL} never changes in the possible structures, within given languages. In turn the sequences <Q, CL> or <CL, Q> appear either before or after N, in each type of language.

A natural analysis of these facts in Muromatsu's terms is as follows:

(17) a.

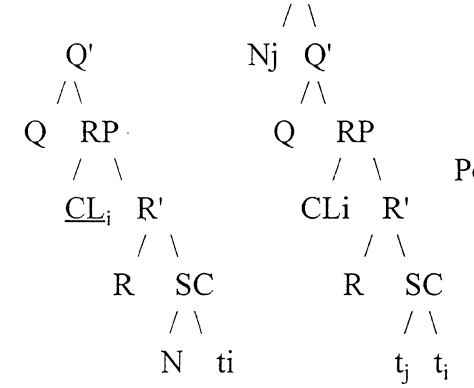

b.

Possibility A (cf. (13))

a.

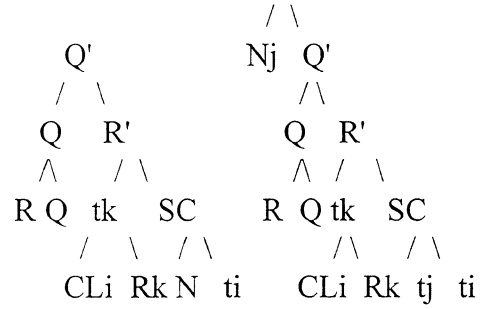

b.

Possibility B

Note that in all instances CL moves to the checking domain of R, and in that sense is what determines the reference of the classified expression. The classifier is, as it were, the reifier that turns the 2D substance into a countable expression which the quantifier can then take as one of its arguments. The main difference between Possibility A and Possibility B is simply whether the classifier is behaving as a maximal projection--in which case it moves to the SP of R--or as a head--in which case it adjoins to the head R. This is not surprising: each possibility is contemplated by universal grammar. In turn, observe that internal to the two types of languages that we have, N may or may not raise to Q, presumably for the same reasons that other categories may, as we saw above--to confine the contextual specifications.

The ungrammatical combinations can now be analyzed as in (18):

(18) a.

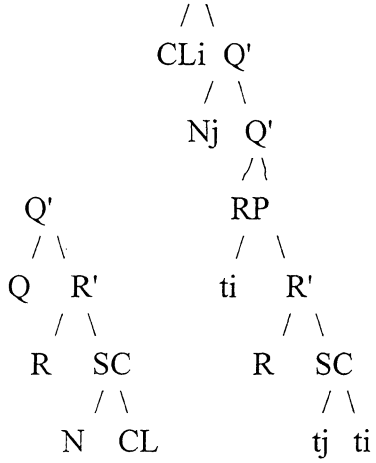

b.

Although Q N CL corresponds to the base structure in (18a), this particular structure is cancelled in the course of the derivation, since CL has not moved to the checking domain of R. This directly suggests that--at least in classifier languages--there is a strong feature in R attracting the interpretable [+r] feature that elements moving to R are given in the numeration.[13] In that way, a derivation that proceeds as in (18a) is simply cancelled, and does not go beyond the R projection.[14] In turn, The CL N Q ordering can only be obtained by moving CL and N to the checking domain of Q. However, there is no way of producing this movement in satisfaction of the Last Resort Condition of Chomsky (1995); doing it leads to another cancellation.

I think this last consequence is very important for the very same reasons that I gave in the previous section for the HRU analysis of integral relations. I know of no alternative explanation for Greenberg's generalizations.

Equally important, I believe, are the implication facts. The only alternative that I know of to the analysis above is in terms of meaning postulates. But these are merely observational devices; after all, why is there no meaning postulate stating that all ghosts are immaterial, which can be then coded by way of a classifier blah? The point has been: blah does not exist, although it could have if meaning postulates were for real. So something is missing, and the dimensional approach tells us what: you cannot have an individual without substance. This is not a fact about the world; we are told that muons are 'things' without mass, for instance. It is rather a fact about cognition; it makes as much sense to look for a classifier for a ghost, a number, a letter, or similar abstractions, as it does to look for the negative sign of 1 or the decimals of -2, or the i of 2/3. These are just category mistakes. In the case of numbers, this is so by definition. In the case of linguistic concepts, it seems to be so by nature.

---

[13] . Within the system that I am pursuing not checking a feature should lead to a cancellation, since there is no meaningful notion of crashing at LF. Then I cannot distinguish between checking prior to and after Spell-out, which must mean that all movement proceeds prior to this point. This is very much in the spirit of ideas in Brody (1995), although from a derivational perspective.

[14] . The sequence Q N CL could also be obtained in crazier ways, for instance moving CL to R first, and then subsequently moving N; this, though, will directly violate the derivational Last Resort Condition.

Perhaps this should be insisted on. It may be thought that Muromatsu's implications follow from the interpretation that she has given to 1D, 2D, and so forth. But this would be an error. Muromatsu may be right or wrong in thinking that 1D, in Human language, is used to denote certain abstractions, and 2D is used to denote masses, and so forth. Even if she were wrong about the correspondences, the implications would remain within the structures. To use an analogy, what she has done amounts to something like this: humans associate the structure of natural numbers to abstractions, the structure of integers to mass terms, the structure of rational numbers to things, and the structure of complex numbers to stages of things. This is no different from saying that the spin of a boson is expressed through an integer, and that of a fermion with a fraction, and so on. At any rate, suppose either of these empirical claims were wrong, and masses should be expressed through complex numbers or spins through imaginary numbers; it does not cease to be the case that rational numbers involve a system which presupposes integers, etc. The reason I think Muromatsu is right is not because of introspection (I do not have any insight into whether abstract terms should correspond to the structure of natural numbers). She may be right because the predictions she makes about entailments seem to accord with how human language works--for instance, in terms of licensing certain classifiers, and not others. In a similar way, we take the claim about spins to be right if it correctly predict that two fermions should not be in the same place at once, or whatever. There are no common sense intuitions to help us here, but simply the internal coherence of a system.

There is one other sense in which any simple-minded alternative to Muromatsu's view leads to some immediate embarrasments. Take the fact that one can say in Spanish and other Romance languages things like (19a), which is different from (19b):

(19) a. En donde yo naci hay mucho torero
'Where I was born there is much bullfighter' (sic).
b. En donde yo naci hay muchos toreros
'Where I was born there are many bullfighters.'

The latter has the standard meaning; the former is much more interesting. First, the term appears in the singular, and takes its own sorts of quantifiers; second, it is incapable of binding individual variables:

(20) a. Muchos toreros andan por ahi sueltos (desde que les pagan bien).
'Many bullfighters walk around (since they pay them well).'
b. Mucho torero anda por ahi suelto (*desde que le pagan bien).

In a word, mucho torero does not seem to be referring to a set of bullfighters, or a plurality of them. The temptation, of course, is to think that it denotes a mass term, like mucho vino 'much wine'. But what would that mean? Do speakers of Spanish have a kind of mind that conceives of individuals differently from the way in which English speakers do? This is not meant in jest; the literature is full of this relativism--the recent one.

Of course, there is a comparative grammar issue here; let's set it aside. From Muromatsu's perspective, it is not at all surprising that I should be able to apply a quantifier for mass terms to something which is typically used in a count fashion, simply because, for her, a count term is a 3D expression that has undergone two warps, each associated to a grammatical

marker. Who is to say, then, that when using the word torero I cannot stop at 2D? In a sense, what has to be explained is why English does not allow this more systematically, while most other languages do.

A particularly nasty example of Chomsky's illustrates that English speakers think like the rest of us: Expensive, hard-cover books that weigh more than five pounds are boring. Boring applies to an abstract entity; weigh, to a concrete entity; hard-cover, essentially to an individual entity; and expensive, contingently to an individual entity. Had we to classify book as [+concrete] (to take weigh) or [-concrete] (to take boring), we would fall into a paradox; the same paradox would arise if we had to classify book as [+stage-of-individual] (to take expensive) or as [-stage-of-individual] (to take hard-cover). However, if book is a 4D expression whose change potential is coded, and thus also involves a 3D coding of form, a 2D coding of mass, and a 1D coding as an abstract concept, then we can say that boring applies at 1D, weighs at 2D, hard-cover at 3D, and expensive at 4D. The argument is direct.

In work in progress, Nobue Mori also pursues these sorts of matters, this time concerning the syntax/semantics of Aktionsart.[15] Some of her results adapt ideas from Pustejovsky (1993), which does not make use of dimensions or the particular syntax discussed here, but does use considerably similar notions. For reasons of space, and because the work is still in progress, I will not report on it now. It should be obvious that the sort of structure we are studying should not be privative of noun classes, just as it is not of numbers; in fact, this point is central to Category Theory. In any case, I take the noun classifications to illustrate what dimensions can do for us within lexical categories: we need not make reference to such substantive labels (or corresponding features) as abstract, mass, individual... terms. So this is a proposal about making (relevant parts of) lexical semantics trivially formal, with a given categorial status resulting from the sheer workings of the system, as derivations go by.

V. Learnability Considerations

No linguistic theory that is worth the name is such if it does not model language as (trivially) learnable by children. In this section I want to ponder the issues that dimensional warps pose for a language acquisition device. As I said before, I am taking the dimensions per se to be innate, in the physicalist sense that I raised. However, a Language Acquisition Device (LAD) must be able to associate a given dimensionality with a given (sound, meaning) pair, like sweet, water, man, and so forth.

This question is not new, and was first systematically raised by Quine (1960), and elsewhere. Interestingly, the sorts of mathematical constructs that I have discussed here might

---

[15] . Similarly, the theory should make predictions about the form and number of arguments of increasingly articulated eventualities. Thus, for instance, it is because the dimensionality of build is rather high that it can get into complex sentences like John built himself a house or John got a house built for him (in contrast, we do not expect this complexity for more basic states--*John loves himself Mary, *John got Mary loved for him--or achievements--*John climbed himself a mountain, *John got a mountain climbed for him). In the end, a full understanding of these matters should shed some light on the hierarchy of roles that arguments receive. If this picture is correct, so-called 'themes' should be the warps of a dimension which is considerably lower than the one 'agents' are the warps of, and so on.

meet with Quine's sympathies. However, I must emphasize, more so than in the case of Parsons, that we may have reached a similar mathematical conclusion (i) from a totally different--and diametrically opposing--starting conceptual point; and (ii) we are making totally different predictions concerning language acquisition. (ii) is not a necessary conclusion, but I believe it relates in Quine's terms to the reason why he probably went into the sorts of speculations I have tried: he is ultimately a behaviorist. For someone in this position, the mind should model reality, through some sort of pattern-recognition process. Inasmuch as a mind has no internal structure when it comes to being, it should start building its structure gradually, as the phenomenological world penetrates it. Early impressions may set up the more basic dimensions, and so on, until the work is complete. I think this theory is wrong, and I want to show both why nothing that I am saying forces us in that direction, and how what I am saying is actually in accord with the facts of linguistic acquisition.

We may present the facts in a narrative way--a famous Quinian example. A LAD observes a rabbit go by, when a native speaker points at the creature while uttering: 'Gavagai!'. How does the LAD know what the native meant? It could have been 'rabbit', or 'furry texture', or more exotic combinations; let us stick to those two to simplify. Let it be true that rabbit is an instance of a 3D expression, while furry is an instance of a 1D expression, therefore one which, in some definable sense, is informationally simpler. (There is, of course, an issue here about how the phenomenon 'rabbit' is taken to be 3D, while the phenomenon 'furry' is taken to be 1D; about such correspondences, irrelevant to my argument, I have nothing to say, and presuppose some psychological theory to do the trick.) What should the LAD conclude, in the absence of explicit training? Does gavagai mean 'rabbit' (or some such 3D term), or does it mean 'furry' (or some such 1D term)? We know what actual children do: they think gavagai means rabbit. Does this then prove the dimensional theory wrong?

It depends. If the dimensional theory is acquired behavioristically, I think the fact indeed disproves it. It would be impossible for the child to come out with the more complex dimensions, in the absence of the simpler ones. Of course, Quine never predicted that children would start their lives with dimensionaly complex structures; quite the opposite, he took the coherent position that children should start with the dimensionally trivial structures. It is not Quine's reasoning that was wrong. To salvage the theory from the facts, the Quinian would have to argue that, somehow, children are capable of acquiring the dimensions on the basis of non-linguistic early experiences. Very interesting psycho-linguistic experiments by Spelke et al. (1991), with three-month old children, strongly suggest that even such concepts as the material integrity of objects are understood by that early age. If the careful experiments are reliable, the Quinian is left with a window of about three months (assuming womb experiences are not helpful for these tasks) to feed into the LAD all the machinery that will then allow it to parse gavagai as 'rabbit', and not 'furry'. Certainly not an impossible task, but a very difficult one--and meaningless, in the absence of an explicit theory of dimensional acquisition. And once the machinery is in place, one wonders again why the child is not at the same quandary that the philosopher is, and why the child should resolve the dilemma by going with the dimensionally more complex (in Quine's terms, as well as mine) 'rabbit', instead of 'furry' meaning.[16]

---

[16] The quinian would have to assume something like the Subcase Principle in (21) below as a general fact about conservatism in learning, only after having already acquired the dimensions. A quinian learner

On the other hand, suppose the dimensional theory is hard-wired into the LAD. How does the present theory fare any better than Quine's; why is my LAD not facing a dilemma? The usual answer that psycholinguists of my orientation give for this is: 'Because the LAD is that way.' This, however, simply delays our understanding of what is going on. What in the LAD being that way is such that it forces it to go with 'rabbit' and not with 'furry'? For this question, I have not found a convincing answer. However, I believe that an answer already exists in the field if we take the dimensional theory seriously.

In fact, the answer is based on well-established psycholinguistic principles of language acquisition: the Subset Principle and the Mutual Exclusivity hypothesis. The former is probably the oldest principle within linguistics, and was proposed in some form already by Panini,[17] although its significance for acquiring language has been noted more recently, and made ample use of in various domains, ranging from phonology to syntax. Recently, Steven Crain has pushed the idea that the Principle is of crucial use in understanding certain aspects of the behavior of children with respect to semantic representations. To cut an important discussion short, I will go directly to a version of Crain's principle that may be of use for us here:[18]

(21) Subcase Principle

Assuming: a) a cognitive situation C, integrating sub-situations c1, c2, ... cn;

b) a concrete set W of lexical structures l1, l2, ... ln, each corresponding to a sub-situation c [19] ;

c) that there is a structure lt corresponding to a situation which is a subcase of all other sub-situations of C; and

d) that the LAD does not know which lexical structure lt is invoked when processing a given term T uttered in C; then: the LAD selects lt as a hypothesized target structure to correspond to T.

---

would never acquire the dimensions to start with, if acting conservatively--unless they get into the mind by osmosis, or some such process.

[17] . Specifically, Paninian rules are hierarchically ordered, in such a way that if Rule n teaches that such-and-such and Rule n+m teaches that so-and-so, then Rule n+m applies at the exclusion of Rule n. The ordering of the rules is, of course, a hypothesis about linguistic structure, and is based on the specificity of the rule, more specific (in a sense, weirder) rules being ordered last, hence applying first. I should say also that Panini was explicitly concerned with the hierarchies that concerned us here, and basically proposed the thematic hierarchy that we now assume, where agents are prominent.

[18] . See Crain et al (1994) for discussion and references.

[19] . In the examples discussed by Crain et al (1994), the issue are just propositions corresponding to sentences presented in circumstances that make them true for a speaker. Note, however, that I am going here into parts of propositions, hence only elements that may enter into the computation of truth, but are themselves more elementary. Hence, I use the term 'correspond' to indicate both the relation sentence(proposition)-situation, and whatever relation holds between terms like rabbit or fur and situations such that the speaker can successfully use those terms to refer in those situations, thus allowing the presentation of the containing proposition in those situations, to make the proposition true. In other words, the Subcase Principle is more general than the Semantic Subset Principle of Crain et al (1994).

Note that it is situations that enter into sub-case relations, and not the corresponding linguistic structures. This is one of the reasons I am using now the non-committal notion 'subcase', instead of the more traditional 'subset', since I do not know how exactly situations should be articulated set-theoretically. Nevertheless, we can make matters rather precise, as in (22):

(22) Given two cognitive situations c and c' obtaining at a super situation C, and for l and l' linguistic structures corresponding to c and c', respectively, and where d and d' are the dimensions where l and l' are expressed, we say that c' is a sub-case of c if and only if d'>d.

It is perhaps useful to consider these notions within concrete scenarios:

(23) Scenario 1: In fact, gavagai means 'furry', not 'rabbit'.
Scenario 2: In fact, gavagai means 'rabbit', not 'furry'.

Analysis (in terms of the Subcase Principle):
Assuming: a) a cognitive situation C [the perceived event], integrating sub-situations c_1 ['a 3D rabbit'], and c_2 ['1D furriness'];
b) a concrete set W of lexical structures l_1 and l_2 [the different possible interpretations of a word associated to the perceived event that universal grammar allows] each corresponding to a sub-situation c;
c) that there is a structure l_t [which involves three orders of the basic syntactic structure] corresponding to a situation c_t [concretely c_2, the '3D rabbit'] which [as per (22)] is a sub-case of all other sub-situations of C [concretely c_1, the '1D furriness']; and
d) that the LAD does not know which lexical structure l_t is invoked when processing a given term T [concretely, gavagai] uttered in C; then: the LAD selects l_t [which involves three orders of the basic syntactic structure] as a hypothesized target structure to correspond to T.

Before we analyze whether this result makes sense, notice two things. First, the '3D rabbit' situation is a sub-case of the '1D furriness' situation, if we accept the dimensional structure, and (22) as a measure of specificity. In general, the Sub-case Principle tells the child to go with the more specific analysis as the first hypothesis. Second, observe that the formal analysis in (23) is totally independent from the factual scenarios.

Let us see how well this conclusion fares with each of the scenarios posed. In scenario 1, the child is of course wrong, and a philosopher who might have decided to go with the 'Super-case Solution' would have been right. However, the child will not produce an erroneous expression when uttering gavagai in front of a rabbit (assuming all rabbits are furry). In turn, if the child assumes an Exclusivity Hypothesis concerning lexical meaning in the acquisition stages ('things only have a name'), the child can retreat from the mistake by either hearing the word gavagai used for any other furry object for which she already has a name, or by hearing another word used for 'rabbit'. So although the child is in fact wrong in the initial guess, correcting the wrong guess is trivial, and is in no need of explicit instruction.

In scenario 2, the child is right, and the 'Super-case' philosopher would be wrong. In fact, the philosopher would produce an erroneous expression when uttering gavagai in front of just any furry thing. It would not be easy for her to retreat from the mistake, for instance, by

hearing the word gavagai used in any other context, assuming all rabbits are furry.[20] Even if the philosopher heard another word being used in the context of a rabbit's presence, one for instance meaning 'furry', she would be at a loss: would that word mean 'rabbit'? Subsequently hearing the word meaning 'furry' in some other context might just mean that the word means 'object' or some such thing.

In sum, given the Sub-case Principle, coupled with the Mutual Exclusivity Hypothesis, the child always gets it right; it is a done deal, regardless of the learning situation. This state of affairs, if correct, not only addresses the Gavagai Puzzle, but it furthermore makes a prediction about the acquisition sequence by children. All other things being equal, human LADS should go with the structure corresponding to more specific situations, within alternative super-cases (perhaps within salience and informational limitations having to do with whatever cognitive restrictions a young child might have)[21] . Apparently this is true. Markman and Watchel (1988) experimentally demonstrated that children (i) assign new terms to objects (rather than to properties, substances, and other possibilities), as the first hypothesis about their meaning; and (ii) do not assign new labels to already labeled objects. This directly allows moves from a more general to a more specific form (for instance, deciding in scenario 1 above that gavagai denotes 'furriness', after having thought for a while that it refers to 'rabbit'), given the appropriate circumstance (e.g., univocally observing rabbits be referred to by another term). Needless to say, all of this makes sense only in situations where the learning alternatives stand in the appropriate sub-case relations, and only a few possibilities emerge, so that solutions are reached in a realistic time.[22]

In sum, the dimensional theory makes the right predictions concerning the acquisition sequence, while addressing Quine's most serious empirical worry, given familiar principles of language acquisition. But again, we succeed at the cost of proposing exactly the sort of solution that Quine rejects: an innately specified structure, and learning strategies that arise when a system with rich internal structure has to map this structure to some trivial outside stimulus--like the utterance of gavagai in a concrete situation. All that the Sub-case Principle is implementing is the conservativeness of LADs: they hypothesize only that which can be easily falsified in terms of positive data. This turns out to be not the simplest kind of structure, but the most complex (given some relevant circumstances and perhaps up to some limitations). Of course, postulating complex structure is trivially simple if the structure is already built in. The bottom line is that our LAD acquires language in the exact opposite sequence that Quine predicts--and, so it appears, according to fact.

---

[20] . Needless to say, I am assuming that explicit corrections are pointless, at least in the case of children.

[21] . I mean this as an answer to an obvious question: why doesn't the child go with the most complex conceptual structure that is possible for a concept when hearing 'Gavagai!'. Suppose that is the sort of structure involved in something like 'a would-be fake rabbit'; that is, something which, with a bit of magic here and there, could turn in an counterfactual situation into something which is not a real rabbit. I would have to say that this sort of concept is not immediately considered in the pool of candidate hypotheses, perhaps because it is just too hard to understand (even for an adult).

[22] . Which is not meant as an off-hand remark. Poorly understood matters are relevant here, starting with how situations are analyzed, or what are the paradigms that are relevant for lexical learning. These questions are familiar, and I have nothing to add to what I said in Uriagereka (1996:chapter 6)).

VI. Some Questions of Symmetry and Harmony

In the last two sections I want to raise the stakes. This is because of a simple fact: I have said nothing so far about functional structure. What I want to say about it presupposes understanding what I have called a 'warp'.

Models M and M' to describe how, for example, the natural numbers and the integers are generated, respectively, (technically, semigroups with specific generators) relate to each other in very much the same way that something like number 2 relates to number 1. In particular, the warp operation that carries you from nD to n+1D is, at this level of extraordinary space, an ordinary generating function, just as something like addition is at the lower level, for a generator like the empty set or one. I use the word 'level' as in Figure I (and see fn. 23). The relation between any of the models generating various sets of numbers, and the super-model for them is one of levels of complexity. The super-model is of course a model of models, and as such lives on its own meta-dimension. And if it warps to super-models of a higher order, this will imply yet a higher level of complexity, and thus a meta-meta-dimensions.[23]

Figure I:

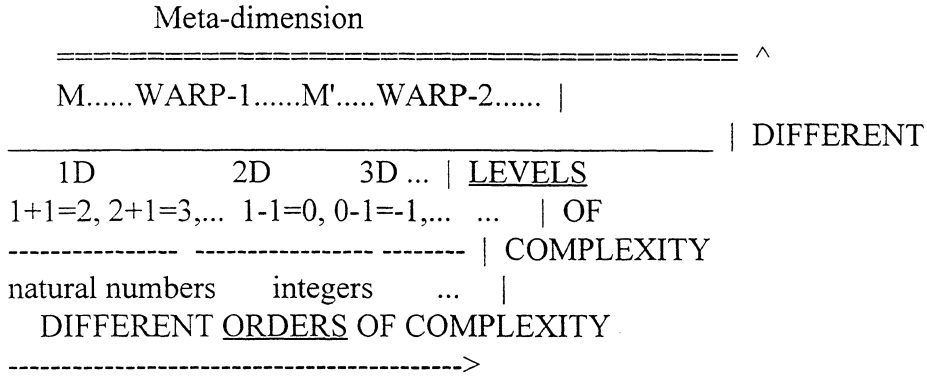

I think this state of affairs corresponds nicely to the relation between functional and lexical categories. Just as lexical categories were shown to relate to the basic 1D, 2D, ... dimensions, I would like to suggest that functional categories are one level removed from all this, in their own meta-dimension. Before I provide some mechanics of how this might work, I want to remind the reader of a fact that several people have noticed, but perhaps no one has emphasized as much as Derek Bickerton: When functional categories are missing, everything

---

[23] . For those interested in these matters, natural numbers are a first order consequence of a certain generating function f, while the negative integers are a second order consequence of f, by using the inverse of the generating function in the previous order. This inverse--say, subtraction for natural numbers--produces results which are outside of the natural numbers (the negative numbers), which is precisely what carries us to the next set, by using the set of natural numbers as a generator. It is in this sense that the order of the sets of integers and the set of natural numbers may be said to differ. In general, we may say this: for a given function F (like addition, subtraction, etc.) that naturally generates some sort of n-th order space S (like the set of natural numbers, etc.), the warp of F (or for S) is its inverse. (I do not know whether this is more than a conjecture.) In the text, I crucially distinguish orders of complexity from levels of complexity. The model M which is responsible for describing the generation of natural numbers is a sort of mathematical structure, and as such we can ask how it relates to model M' which is responsible for describing the generation of integers. The warping function can be seen as a generating function whose arguments are models, and which yields a semi-group of models. We may then say that the structure of a model is at a different level than the structure of the object that the model describes. This is the intuition that Figure I tries to capture. For reasons of space, I cannot go into these matters any further.

else goes. If we think of it, this is truly remarkable. Take the forbidden experiment and ask the obvious question: how is it possible for a child raised outside of a linguistic community not to acquire trivial little words like for or to or the, or perfectly simple morphemes that any two year old masters, like the plural forming -s, so that she may utter I want elephant eat peanut,[24] instead of the obvious alternative. In what sense is the more complex (or whatever) than elephant, peanut, eat, or want?

I would like to suggest it is all a matter of levels. An item like the is to elephant, in my view, just as the entire number system is to 2; it is just an accident that we happen to use the same linguistic articulators to pronounce the as we do to pronounce elephant; actually, in some languages: in others you pronounce elephant, but not the. So in a sense, with regards to the there may be nothing to acquire, really; it is a dimension, or perhaps more to the point, a model of dimensions. It is thus that it is fundamental. Acquisition at this level may perhaps be seen as growing, in the sense that royal jelly makes a bee grow into a queen, affecting certain fundamental anatomical processes in the developing larva. So too, input linguistic data (we do not know which) trigger the growth of mental dimensions for language, the models within which more standard lexical dimensions, of the sort above, are expressible. Miss those triggering stimuli and you will miss more than the words you heard, just as a bee-larva will miss more than a tasty dinner if she misses royal jelly.

Of course, even if that story is coherent, two other steps are necessary to make it scientific. First, one must argue the case, beyond the plausibility, philosophical comment above. Second, one must address the issue of why functional categories should appear where they do, interleaved with lexical categories as the derivation proceeds. The second question is too ambitious to even sketch an answer for at this stage, and is in any case the topic of my current inquiry with Ian Roberts--so I must set it to the side. The last roundabout I invite the reader to take is intended to address the first matter.

There is a stubborn problem of syntactic theories, since the fifties, that I would like to comment on. Early on, the problem was appropriately turned into an axiom: Derivations start with the symbol S, which is necessarily rewritten as NP (the subject) and VP (the predicate). When rewrite rules were given up, one remained: S --> NP VP. Work in the eighties always ended up adding an extra stipulation concerning the fact that all sentences need a subject, even if it is not the bearer of a semantic role. In the nineties, the problem is turned into a feature. Tense attracts, prior to Spell-out, a D feature from somewhere down in the tree; by another name, this is a subject.

Now consider the matter from a different perspective, capitalizing on the fact that two major schools of thought within semantics tell us either that sentences are crucial semantic units, or rather propositions are. The latter is the traditional view, and is concerned with studying the internal make-up of propositions. The neo-Davidsonian program partially reacts against this view when centering semantic discussions over the claim that sentences are descriptions of events. Everything I have said here can be seen within this project, for all I have done is propose (the structure of) a lexical semantics, in terms of dimensional warping. As the

---

[24] . This sentence is a literal one uttered by Geenie, the famous child who has tragically raised without access to language.

semantics in (15) show, the descriptive apparatus of the neo-Davidsonian project remains untouched. So far, all the present theory does is restrict the class of possible roles in a sentence (they are all warps to increasingly more articulated spaces), and provide a syntax to do the job. But now comes the point of asking whether there is anything more than structured events (or whether the traditional intuition that propositions are made-up of something like subjects and predicates is also right).

This view was expressed in Raposo and Uriagereka (1995), where it was argued that while the neo-Davidsonian structure of Brutus stabbed Caesar is identical to that of Caesar, Brutus stabbed, these are different expressions, and perhaps even different propositions. Consider (24):

(24) a. A band of kids from Liverpool recorded this great song.
b. This great song, a band of kids from Liverpool recorded (it).
c. And then, there recorded this great song a band of kids from Liverpool.

In a language like English, some of these are less than perfect, but let us try and put that to the side.[25] Intuitively, (24c) is about something that happened (the recording of a great song), while (24b) is about a song, and (24a) is most saliently about a band of kids, although it can also be about something that happened. Can this change the truth conditions?[26]

Imagine we are in Hamburg, the day the young Beatles had their first recording. In the real world, Stu Sutcliffe did not make it to this recording, but suppose he had. Sutcliffe's specialty were ballads, including Elvis's Love me tender. So suppose the Beatles record Love me tender in this their first studio appearance, as a mere try out song that never makes it beyond the studio session. Now, it is certainly felicitous to say (24a) either to report an event or to tell us something about the Beatles early years (assuming we refer to the group as a band of kids from Liverpool, and to the song as this great song). The event reading is highlighted in (24c), marginal as this sentence is in English. In contrast, consider (24b). I do not think it is felicitous to say that about Love me tender. Whether it is in fact false depends on whether the rest of the expression is meant exhaustively. If it is, the proposition is directly contradicted by

---

[25] . For instance, in German we have:

(i) a. Ein paar Burschen aus Liverpool haben dieses tolle Lied aufnehmen
b. Dieses Lied haben ein paar Burschen aus Liverpool aufgenommen
c. Es haben ein paar Burschen aus Liverpool dieses Lied aufgenommen

(Thanks to Elena Herburger for the data and useful discussion.) For some reason, topicalizations as in (24b) are not very felicitous for some English speakers. The same is true about transitive existentials, as in (24c), although examples like (ii) are marginally acceptable:

(ii) And then,
a. there entered the room a man from England.
b. there hit the stands a brand new journal.
c. there visited us a group of students from India.

[26] . When considering this question, the reader must abstract away from focus considerations, which also alter truth-conditions, but in a way that is irrelevant for my purposes here.

Elvis also having recorded the song. If, on the other hand, the comment about Love me tender is not meant exhaustively, the proposition is non-cooperative. Either way, the expressions in (24) are clearly different, and the neo-Davidsonian structure does not capture this fact.[27]

Once we assume that something more than standard neo-Davidsonian structure is necessary to account for the subtleties of propositional interpretation, the question is where to add this something. Raposo and Uriagereka suggest that this be done through a 'higher level' predication, much in the spirit of early work on these matters that distinguished categorical judgments from thetic judgments. We can take the former to be about some event-participant, while the latter are about an event itself. A judgement, in turn, can be thought of as a proposition presented from the contextual perspective of some judge, usually the speaker.

Semantically, it appears clear that we have to assume that judged propositions are about something. But how do they get to be that way? Of course, in my terms it would make sense to think of the matter thus:

(25) A proposition is a warped eventuality.

What this is saying is that propositions relate to events just as individually classified nouns relate to mass terms, or for that matter integers relate to natural numbers. I should emphasize, however, that if (25) is true, it must be taking place at a level different from that were noun classes relate, not at a mere higher order (and see Figure I). That is to say, if we have reached order n with a given lexical structure, applying functional structure to it will not warp us to level n+1. For the purposes of the functional warping, we may as well be in a dimension 1, 2 or whatever, at a different level.

One advantage of (25) would be that we could take its syntactic expression to correspond to the sort of syntax I have motivated above, as in (26). Here the element labeled 'VP' for ease of reference is possibly very complex, internally. This 'VP' determines an eventive space s_e, which is warped by T, in a small clause relation, as before. The reason I dot the space above 'VP' is that material within 'VP' is not of the same level as material outside. Other than this (central) fact, the architecture of what we are now seeing is identical to that of what se saw before:

---

[27] . Actually, Elena Herburger's dissertation is written, in large part, to demonstrate that the Neo-davidsonian project can in fact capture facts of the relevant sort without being enriched in the sorts of ways that Raposo and Uriagereka propose. I put this possibility to the side, since the work is still in progress.

(26) a. Q'

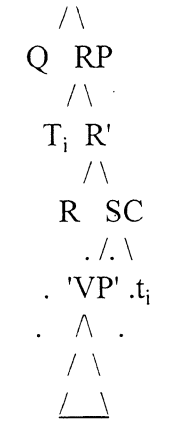

Brutus kill Caesar

b. Q'

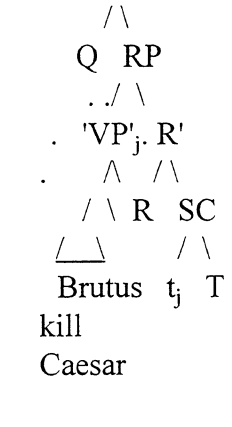

This is in fact the key to what I have to say in this section, and the reason why I think it may be useful to explore these matters, since we have not enriched our descriptive apparatus when going into (26). True, we have used a new label like T, but this is merely expository. What T is, in this system, is a warp for 'VP', at a level which is higher from anything we have seen thus far. Of course, the relation between T and 'VP' is not neo-Davidsonian, in that it is not expressing any traditional role of the sort explored up to now; it is instead introducing a T anchorage, by somehow delimiting $\underline{s}_e$ through an indexical time which (in some non-trivial sense) integrally coincides with it. Finally, if what we saw for the attracting potential of R is true at this new level as well, then either T or 'VP' must move to R, to 'check' reference. While in both instances the relation between 'VP' and T is integral (one is not understood without the other), two different objects are produced.

Note that while (26a) introduces reference to tenses, (26b) introduces reference to events. Possibly, this is the main difference between Brutus kill-ED Caesar and Brutus's kill-ING of Caesar; the latter denotes a given event, while the former existentially quantifies over a time which is univocally associated to a concrete event. *Adverbs of quantification* vs. referential determiners separate these readings:

(27) a. *Brutus never killed Caesar.*
b. *Brutus's never killing of Caesar.
c. *A/the Brutus killed Caesar
d. A/the killing of Caesar by Brutus

At this juncture, we could point out the similarity between (26a) and (13), repeated now as (28), with the more accurate notation in terms of S/W:

(28) a.

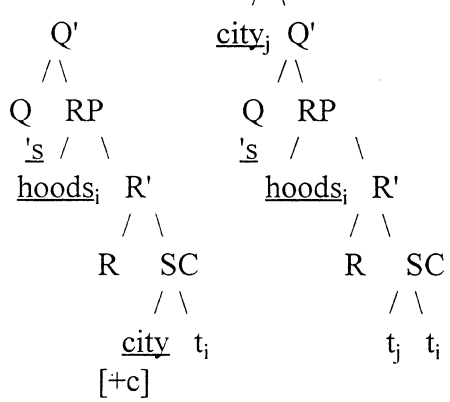

b.

c. [Thee: C(e) & hood(e) & S(city,e) & W(hood,e)] ...
Where the speaker confines the range of C to city

Suppose we were to invoke the type of movement in (28b)--which we suggested has the contextual effect noted in (28c)--in (26a). This would carry 'VP' all the way up to Q SP, just as city moves in (28b). Perhaps this can explain the sort of situation we find in a language like Basque:

(29) Jonek Miren maite du.
Jon-subject Miren-object loved present-agreement
'John loves/has loved Mary.'

Observe that Tense is post-verbal, directly suggesting (in the spirit of Kayne (1994)) that the entire 'VP' has moved up to a site that commands Tense.

Let us for now set aside the interpretive consequence of moving 'VP' to the Q SP in (26a), and assume some reason R for why the movement is legitimate. An interesting descriptive question is whether, in situations of the sort in (26a), just the entire 'VP' must move to the SP of Q, or rather a sub-constituent of 'VP' may also move there. This may be the situation in languages involving overt topic marking, like Japanese, where we could take the main sentential topic to be a reflex of T. Thus:[28]

(30) a. Osamu-wa Keiko-o mi-ta.
Osamu-topic Keiko-object see-past
'Osamu, he saw Keiko.'

---

[28] . The data that follow are courtesy of Keiko Muromatsu. The possibility in (i) is not included in the paradigm in (33) because the judgement in this instance is not categorical, but thetic:

(i) Osamu-ga Keiko-o ni-ta.
Osamu-subject Keiko-object saw-past
'The event took place that Osamu saw Keiko.'

b. Keiko-wa Osamu-ga mi-ta.
Keiko-topic Osamu-subject see-past
'Keiko, Osamu saw her.'

c. Osamu-ga Keiko-wa mi-ta.
Osamu-subject Keiko-contrast see-past
'It was Keiko that Osamu saw.'

Observe that either the thematic subject (30a) or the thematic object (30b) can be topic-marked. When this is the case, though, the constituent marked with a topic that the proposition is about must move to the left edge, the Q SP in (29a) by hypothesis. When this does not happen, as in (33c), the reading is different, involving matters concerning focusing that I cannot go into.

So far we have seen a parallelism between 'verbal' (26) and 'nominal' (28) structures. The parallelism extends to 'clausal' structures as well, as (31) shows. Just as 'VP' relates to T, so too the resulting expression, which we may think of as 'IP', relates to a truth declaring item (as suggested in work in progress that I am conducting with Esther Torrego)):

(31) a.

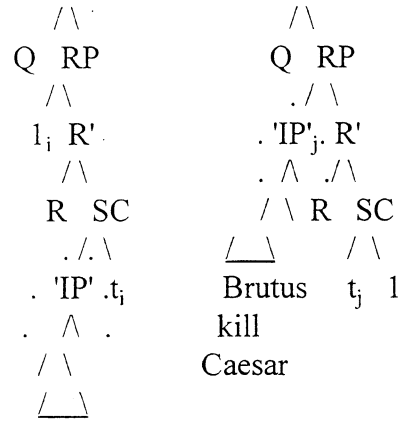

b.

'IPs' are quantified tense-event spaces $\underline{s}_e$ which can be warped through 1. As before, 1 may move to the referential projection, or 'IP' may. In the first case, we will have a standardly judged proposition like Brutus killed Caesar when uttered by a speaker as a judgement. Technically, the sentence will be some sort of speech act quantifying over 'the truth', which is integrally associated to a given quantified tense-event space. In the second case, we will have, instead, reference to a proposition. This is probably the sort of expression involved in Oh, to kill Caesar!, and might also be relevant for propositions with imperative force like kill Caesar!.

Continuing with the parallelism, either the entire 'IP' should move to the SP of the next level of Q, leaving 1 behind, or a sub-constituent of 'IP' should. The first possibility is arguably what happens in Japanese when empathy particles are invoked, particularly in conversation:[29]

[29] . Data again courtesy of Keiko Muromatsu. Thanks also to Nobue Mori for helpful discussion on these topics. On these issues, see Kuno (1973).

(32) Osamu-wa Keiko-o mi-ta-yo.
Osamu-topic Keiko-object saw-past-empathy
'Osamu, he saw Keiko, I declare.'

The second possibility may be illustrated through Quechua:

(33) a. Paqarin-ta-n hamu-nqa.
tomorrow-object-empathy come-agr-future
'He will come tomorrow, I declare.'

b. Mana-n hamu-n-chu.
not-empathy come-agr-negative
'We/she does not come, I declare.'

Muyskens (1995), where these data are taken from, observes that empathy markers are always constituent external, and typically attached to the first constituent,[30] and in any case never in constituents to the right of V, as expected given the picture that (31a) provides, with 1 hosting the empathy marker and the 'edge' constituent in the Q SP. Importantly, imperative and exhortative clauses cannot take these particles, and must occur in the domain of Tense, which suggests that the import given to (31b) is probably correct, with this sort of structure lacking the import of a judgement.

To make these parallelisms complete, we must also ask whether, just as T and 1 have a special status in bounding the 'verbal' and the 'clausal' structures, there is a similar bounding element within the 'nominal' structure. In fact, there is. No matter how complex the internal make up of a nominal expression is (how many dimensional layers it involves which license given lexical entailments), it ultimately invokes reference to something in the last layer of structure. This suggests the presence of an indexical equivalent to T or 1, some element P which may in fact translate as ostension, pointing, or even locative elements, as in the expressions below:

(34) a. That [speaker points] guy is my friend.
b. This here guy is my friend.

Now we can take an 'NP' structure to determine an eventive space s_e, which is bounded by P (hoping again that this substantive interpretation will be deduced). If so, the structure in (28) has to be sharpened to (35):

[30] . Or to focused constituents, which pose questions I am not dealing with here. Muyskens also observes that empathy-marked constituents can be preceded by topics; this would seem to contradict what I have said above about topics in Japanese, which are supposed to move at a lower level. However, if the empathy-marked constituents are focused on that instance, then the situation is no different from that arising in (33c). Alternatively, the topics may also be, in some instances, conversational or 'hanging' topics, with an 'as for' import.

(35) a.

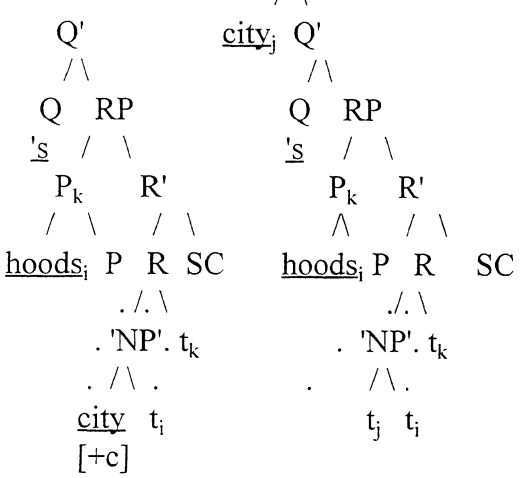

b.

Observe that an element like hoods does not move directly to R; rather, it first moves to P, and then the complex element hoods-P moves to R. The movement of a nominal like hoods to P, however, does not seem to be general, as (34b) arguably shows if P is lexicalized as here. Moreover, it is reasonable to suppose that whereas a warp like hoods may associate to another (more abstract) warp like P, the same is not true about the spatial city, which may need to move to R--*when it does--all by itself (in a parallel fashion to what we saw in (31b) for 'IP')*. Ultimately, though, these questions have to be addressed separately and more carefully, and some of them are familiar: What motivates head movement and how does it vary across languages? I have little to add now.

Summing up, we have the following picture so far, at the higher level:

(36) a. There are three domains of structural regularity that behave in symmetric ways: a 'nominal', a 'verbal', and a 'clausal' zone.

b. These zones are successively embedded, and may relate in terms of different orders of abstract spaces and their warps.

c. Functional elements like T, 1, or P may be warps to the 'verbal', 'clausal', and 'nominal' zones, respectively.

d. Zonal warps attract (for reason R) given constituents to the edge (Q SP) of the structure, where they themselves are attracted (R).

e. Constituents in a structural edge appear to confine context ranges.

(36a) is useful as a generalization. (36b) is a comment about symmetry; it appears that the domains in (36a) are, in some intuitive sense, alike. (36c) commits to the interpretation that functional structures are warps of some sort. (36d) and (36e) merit some separate discussion.

Suppose the 'attracting' Reason R is a result of something deeper than the system thus far sketched would lead us to expect. Note that in order not to fall into generative semantics,

we have to sharply distinguish domains of structural conservation, where we expect structural qualities to be preserved in the course of the derivation. It is very important to identify these domains, in the quest for understanding inner properties of the system which are often quite abstract; through the observation of these pockets of regularity we may be able to hypothesize corresponding structures. We may think of the distinctions between words and syntactic structures precisely in this light: each of these determines different domains of conservation. Syntactic structures are productive, transparent, and systematic, in the sense of Chomsky (1970), while words are none of these; conversely, words align themselves in paradigms subject to such learning procedures as the Sub-case Principle or the Exclusivity Hypothesis discussed above. These are just different systems. Word-level units do not talk to syntactic units, and are instead submitted to the Word-Interpretation component. Assuming all of this, it is then natural to ask how this rhetoric corresponds to the formal apparatus discussed so far. Looking at a structure like (35), how do we know what is syntax and what is not?

VII. On Alignment

I propose (37) as an answer to the question above:

(37) Alignment Principle
For X warped, X referring syntactic expression s, match s with a PF boundary anchored in X, and treat the result as a WI unit.

In a nutshell, if (37) is right, it is the same impulse that drives the system to create WI units (complex words, idioms, and perhaps others) that drives constituents to the edge of 'nominal', 'verbal', and 'clausal' zones. This intuition, although coming from a different perspective and involving different mechanisms, is very much in the spirit of Steven Anderson's old dictum that 'clitics are the morphology of phrases'. In fact, clitics (including verbs) typically appear placed in the sorts of edges that such phrases as topics, subjects, or genitives mark. What is the purpose of these edges from the point of view of the syntax? If (37) is right, it is what determines the cut between syntax and morphology, in some abstract sense.[31]

Note also that the Alignment Principle seeks the harmony of edges. Intuitively, the principle matches semantic edges with phonological edges. I have little to say here about what are phonological edges. Obviously, though, the left periphery of a phrasal unit should be an edge. In turn, I have had something to say about semantic edges; these are the warps. The point can be made more concrete on mathematical grounds, using notions of Topology; I find this unnecessary. It is intuitively obvious that what a warp does is carry us 'out of the edge' of a given space. In doing so, the warp also delimits an inner space. The intuition behind (37) is then that the space delimiter in this categorial sense must align itself with a perceptual edge, in the PF component.

Technically, the Alignment Principle first finds a major semantic X warp in a referring position R; second, it seeks a major syntactic constituent to move to the phonological edge

[31] . Of course, we are dealing with objects that are definitely larger than some words are, and might be thought of as lexical units at LF. See Uriagereka (1996:chapter 6, section 4) on this.

associated to R. Syntactic conditions restrict this movement in the obvious way, although here a great deal of variation must be expected.[32] Likewise, the syntax relates R to the Q element that takes it as its internal argument, either overtly or covertly. This places R in the domain of Q, which may be seen as the justification for why the phonological edge of R is actually the SP of Q, not the SP of R.

Finally, consider (36e) and what it implies. We could motivate movement to the SP of Q in purely syntactic terms, through some kind of contextual feature. However, that would be against the spirit of the Alignment Principle, and I believe also wrong. Consider, in this respect, (38):[33]

(38) a. To have idle time! -

b. [John], I like. +-

c. He just arrived, [John (did)]. -+

d. [John], she loves, [Mary (does)]. +-+

e. (*)[A man] +

f. *[Mary], [John], she loves. ++-

g. *She loves him, [Mary (does)], [John]. -++

h. *[John], I like, Mary hates. +--

i. *He likes Mary, he loves Sue, [John (does)]. --+

j. *[John], she likes, she hates, [Mary (does)]. +--+

On the right hand side, I have interpreted elements in the edge as + warps, and elements in the center as - spaces, for now without committing to any deeper interpretation. Formally, it is interesting that only the first four combinations are possible, and (38a) may or may not be depending on what one takes vocatives to be. The combinatorics strongly resemble syllables, a fact that is highlighted if we take these elements as abstractly as Perlmutter (1993) does when speaking of ASL hand movements and positions.

---

[32] . For instance, if Chomsky's recent treatment of Bantu languages is on the right track, these allow movement of objects the 'IP' edge, a sort of derivation argued for, also, by Raposo and Uriagereka (forthcoming) for some instances in Romance. Then the familiar question is why this is generally impossible for English. Chomsky has spent a considerable amount of time dealing with this matter, and I have nothing to add to his conjectures.

[33] . Many of the facts below are inspired by the discussion in Lasnik and Saito (1992). Claims concerning the impossibility of two edges directly relate to Muysken's (1995) observations: in Quechua it is completely impossible to have two empathy marks.

This may be a coincidence, or it may be telling us something about the role of edges in syntax, and if true would allow syntacticians to have some common ground both with phonologists and pragmaticists. Researchers like Lisa Selkirk have often talked about syntactic edges matching phonological edges in instances involving focus, a sort of discussion that is now common in the Optimality literature. In turn, pragmaticists have often emphasized the role of discourse in configurationality. There may be something right to that sort of claim, although I still would argue that it is backwards. If the Alignment Principle is right, it is for formal reasons that semantics aligns with phonology. It is a harmony seeking system that tells the semantic edge to go with the phonological edge. This entails the need to create a phonological edge. Simply put, if nothing moves to the domain of the Q-R complex, the phonological edge will be random, assuming that the Q and R elements are clitic-like in themselves; we need to displace some lexical structure to Q-R to create a perceptually robust phonological edge. The question is what.

The pragmaticist would perhaps want to argue that it is the, in some sense, 'salient' element that gets so displaced. But this would be giving teleological power to derivations. Rather, it seems more plausible to say that something, whatever is near-by and meets conditions for movement, is displaced to form an edge--and then... it anchors a context. It is because the element has been appropriately displaced to a domain of prominence that the context encoding device, a performative matter, can be triggered. I admit that in some instances the grammar may have featural elements, or morphemes, to designate the element to be displaced. But even in that instance, I would argue that these features are assigned randomly, with the sole purpose of guaranteeing a displacement to the edge, thereby entering into specific interpretations.

A related question is how the syntax knows what is a phonological edge and what is a semantic edge, to proceed to match them. If what I have said is right, the phonological edge is literally created in the course of the derivation: the moved, 'lexically heavy' constituent produces the edge. But how about the semantic edge? Apparently, the syntax recognizes an interpreted edged--that is, a warp--only if this element is in the checking domain of R. Taking R-dependency to be what allows movement for reason R purposes is more than a pun: it would justify the important contrasts below:

SEMANTIC Warp IN R	SEMANTIC SPACE IN R
(39) a. Brutus killed Caesar.	b. A/the killing of Caesar.
(40) a. Xwan-mi hamu-nqua. Juan-empathy come-future-agr 'Juan will come.'	b. Mijhu-y(*-mi). eat-imperative-empathy 'Eat!'
(41) a. The/*a poor neighborhood of the city	b. The/a city of poor neighborhoods

(39) involves a 'verbal' zone; the Quechua (40), a 'clausal' zone; (41) involves a 'nominal' zone. The question in (39) is why (39b) does not force the subject Brutus; in other words, why gerundials and similar elements do not have to obey the Extended Projection Principle. The

question in (40) is why (40b) does not take a empathy particle (similarly, perhaps, one could have asked why imperatives do not trigger standard 'verb second' effects). The question in (41) is why (41b) does not require a definite 'head', whereas (41a) does. I suspect these questions are related. In the (b) examples, the element in R is not the semantic warp (T, 1, or P), but some semantic space ('VP', 'IP', or 'NP'). If what I am trying to establish about the syntax identifying warps is right, the derivation will not take the (b) examples as involving a semantic variable, and hence will not force phonological alignment, since the Alignment Principle will not apply. These should be dependent expressions.

In sum, what we typically call an argument is a space which is bounded by P; what we call a proposition is a kind of space which is bounded by T; what we call a judgement is a kind of space which is bounded by 1. Arguments are warps to basic spaces which are, in some definable sense, expressed at a more elementary than these very abstract spaces. The former are lexical, the latter are functional.

We need the Alignment Principle because phrases are not words. In the latter, the pairing of the PF and LF formats is listed in the lexicon, after being acquired in ways already mentioned. But there is no such pairing for phrasal-level units. Of course, in the old days, the pairing was given through an axiom of the form 'S-->...' In the sort of dynamically split model that I have been exploring, something else must ensure that PF be locally compatible with LF, as multiple Spell-out proceeds. The Alignment Principle has that effect. And for those phrases that do not meet it, we make a straightforward commitment: they cannot be root phrases, and must constitute larger units which do satisfy the alignment.

One virtue of these speculations should be pointed out. The Minimalist program simply does not expect semantic effects that are not syntactically driven. But a variety of 'discourse-dependent' phenomena pose a serious challenge for this view. Everything beyond neo-Davidsonian structure--conversational and sentential topics, thematization, focusing--seems to involve phonological information in crucial ways, either through overt displacements whose character looks quite different from that of standard movement, or from a variety of prosodic cues (intonation, stress, pauses, etc.). There has been, to my knowledge, no natural way of fitting these phenomena within the current Minimalist picture. One can deny the LF significance of the processes, but there is every reason to believe that it exists. Then the puzzle is immediate. My answer to this puzzle has invoked the Alignment Principle, which instantiates some specific assumptions about harmony across levels of representation. The Alignment Principle can be thought of as part of Spell-out, just as Kayne's LCA. The grammar must match up semantic and phonological boundaries. From this perspective, it is not surprising that phonological and semantic ends meet, at the point where they should--at Spell-Out.

The point, though, is that if the analysis above is right, we have justified treating functional categories with the syntax that we had motivated for lexical categories. This is only a step, though. Now we must understand what they are doing there, and also what happens with the other fourteen categories--or however many there is. These are the specific matters that Roberts and I are trying to clarify.

VIII. Concluding Remarks

It is not unreasonable to say that we have:

- Chartered a partial semantic representation for some unfamiliar lexical notions, such as abstract and mass terms, or classified nouns.

- Provided a framework to capture lexical entailments among these elements.

- Provided a solution for Chomsky's puzzle concerning multiple predications involving the same lexical item with different senses.

- Given an acquisition account of Quine's gavagai puzzle.

- Given an explicit Minimalist syntax to carry these lexical semantics.

- Proposed a type of solution for the puzzle that 'discourse configurationality' poses for Minimalism.

- Proposed a way of restricting 'functional structure' while at the same time leaving its descriptive potential open.

- Suggested a way in which the level of lexical structure relates to the level of functional structure.

- Suggested a way to interpret the notion 'local domain for interpretation' at WI.

More importantly, we have seen through the HRU and Raposo and Uriagereka proposals that these theoretical results are grounded on standard linguistic data. This is, to my mind, clearer still in the on-going work of Muromatsu and Mori, as is in the incipient work of several others which, lacking space, I could not report now.

But while empirical gains and consequences are central to any natural science, we should not loose sight of the forest. I have started this paper by asking what a radically Minimalist, derivational proposal may have to say about syntactic objects. We have now seen how the dynamics of the derivation may produce a variety of objects, all of them of a very similar 'character', in the sense of involving some space and a warp to this space. Although we have no criterion of admissibility for these objects at LF or PF--lacking these levels altogether--the very way in which the objects are formed is their passport to existence. If I am correct, standard linguistic categories are forced to exist, given derivations. The creation of impossible objects and interactions leads to an immediate derivational cancellation. No derivations crash;

the notion of a divergent derivation does not make sense in this view.[34] If so, the right question could not be what sorts of semantics our syntax will meet.

References

Barr, M. and Ch. Wells (1995) Category Theory, Prentice Hall.

Brody, M. 1995. Lexico-Logical Form, MIT Press.

Crain, S., Weijia Ni, and Laura Conway (1994) 'Learning, Parsing, and Modularity,' in Ch. Clifton, L. Frazier, and K. Rayner (eds.) Perspectiveson Sentence Processing, Lawrence and Erlbaum Associates.

Chomsky, N. (1970) 'Remarks on Nominalization,' in R. Jacobs and P. Rosenbaum (eds.) Readings in English Transformational Grammar, Ginn and Co.

Greenberg, (1977) 'Numeral Classifier and Substantial Numbers,' in A. Makkai, V. Makkai, and L. Heilmann (eds.) Linguistics at the Crossroads, LivianaEditrice.

Higginbotham, J. (1988) 'Contexts, Models, and Meanings,' in Mental Representations. The Interface between Language and Reality, R. Kempson (ed.), Cambridge.

Hornstein, N., S. Rosen, and J. Uriagereka (1995) Integral Existentials, Proceedings of the West Coast Conference on Formal Linguistics 1994.

Jackendoff, R. 1969. Some Rules of Semantic Interpretation for English. Doctoral dissertation, MIT.

Kayne, R. (1994) The Antysymmetry of Syntax, MIT Press.

Keenan, E. (1987) 'A Semantic Definition of Indefinite NP,' in The Representation of (In)definiteness, MIT Press.

Kuno, S. (1973) The Structure of the Japanese Language, MIT Press.

Kraus, L. (1995) The Physics of Star Trek, Basic Books.

Lasnik, H. and M. Saito (1992) Move @, MIT Press.

Markman, E. and G. Watchel (1988) 'Children's Use of Mutual Exclusivity to Constrain the Meanings of Words,' Cognitive Psychology 20.

Muromatsu, K. (1995) 'The Classifier as a Primitive,' to appear in The University of Maryland Working Papers in Linguistics, 3.

Muyskens, P. (1995) 'Focus in Quechua,' in K. Kiss (ed.), Discourse Configurational Languages, Oxford.

Parsons, Ch. (1980) 'Mathematical Intuition,' Proceedings of the Aristotelian Society, March 17, 1980.

Perlmutter, D. (1993) 'Sonority and Syllable Structure in American Sign Language,' Linguistic Inquiry 23, 3.

Pustejovsky, J. (1993), 'Type Coercion and Lexical Selection,' in J. Pustejovsky (ed.) Semantics and The Lexicon, Kluwer.

Quine, W. (1960) Word and Object, MIT Press.

Raposo, E. and J. Uriagereka (1995) 'Two Types of Small Clauses,' in A. Cardinaletti and M.T. Guasti, Small Clauses, Academic Press.

Raposo, E. and J. Uriagereka (forthcoming) 'Indefinite se,' NLLT.

Spelke, E., K. Breinlinger, J. Macomber, and K. Jacobson (1991) 'Origins of Knowledge,' Psychological Review, 99, 4.

Uriagereka, J. 1995. 'Formal and Substantive Elegance in the Minimalist Program,' in M. Bierwisch, H.M. Gaertner, and C. Wilder (eds.) The Role of Economy Principles in Linguistic Theory. Berlin: Akademie Verlag.

Uriagereka, J (1996) Rhyme and Reason: A Minimalist Dialogue on Human Language MIT Press.

Uriagereka, J. Forthcoming. 'Multiple Spell-out' to appear in M. Browning (ed.) Working Minimalism, MIT Press.

[34] . Which is not to say that an ungrammatical, interpretable object, of the sort Higginbotham (1985) first seriously discussed, could not be produced in performance, and furthermore given some interpretation.