biolinguistics

volume 2, issue 4 fall 2008

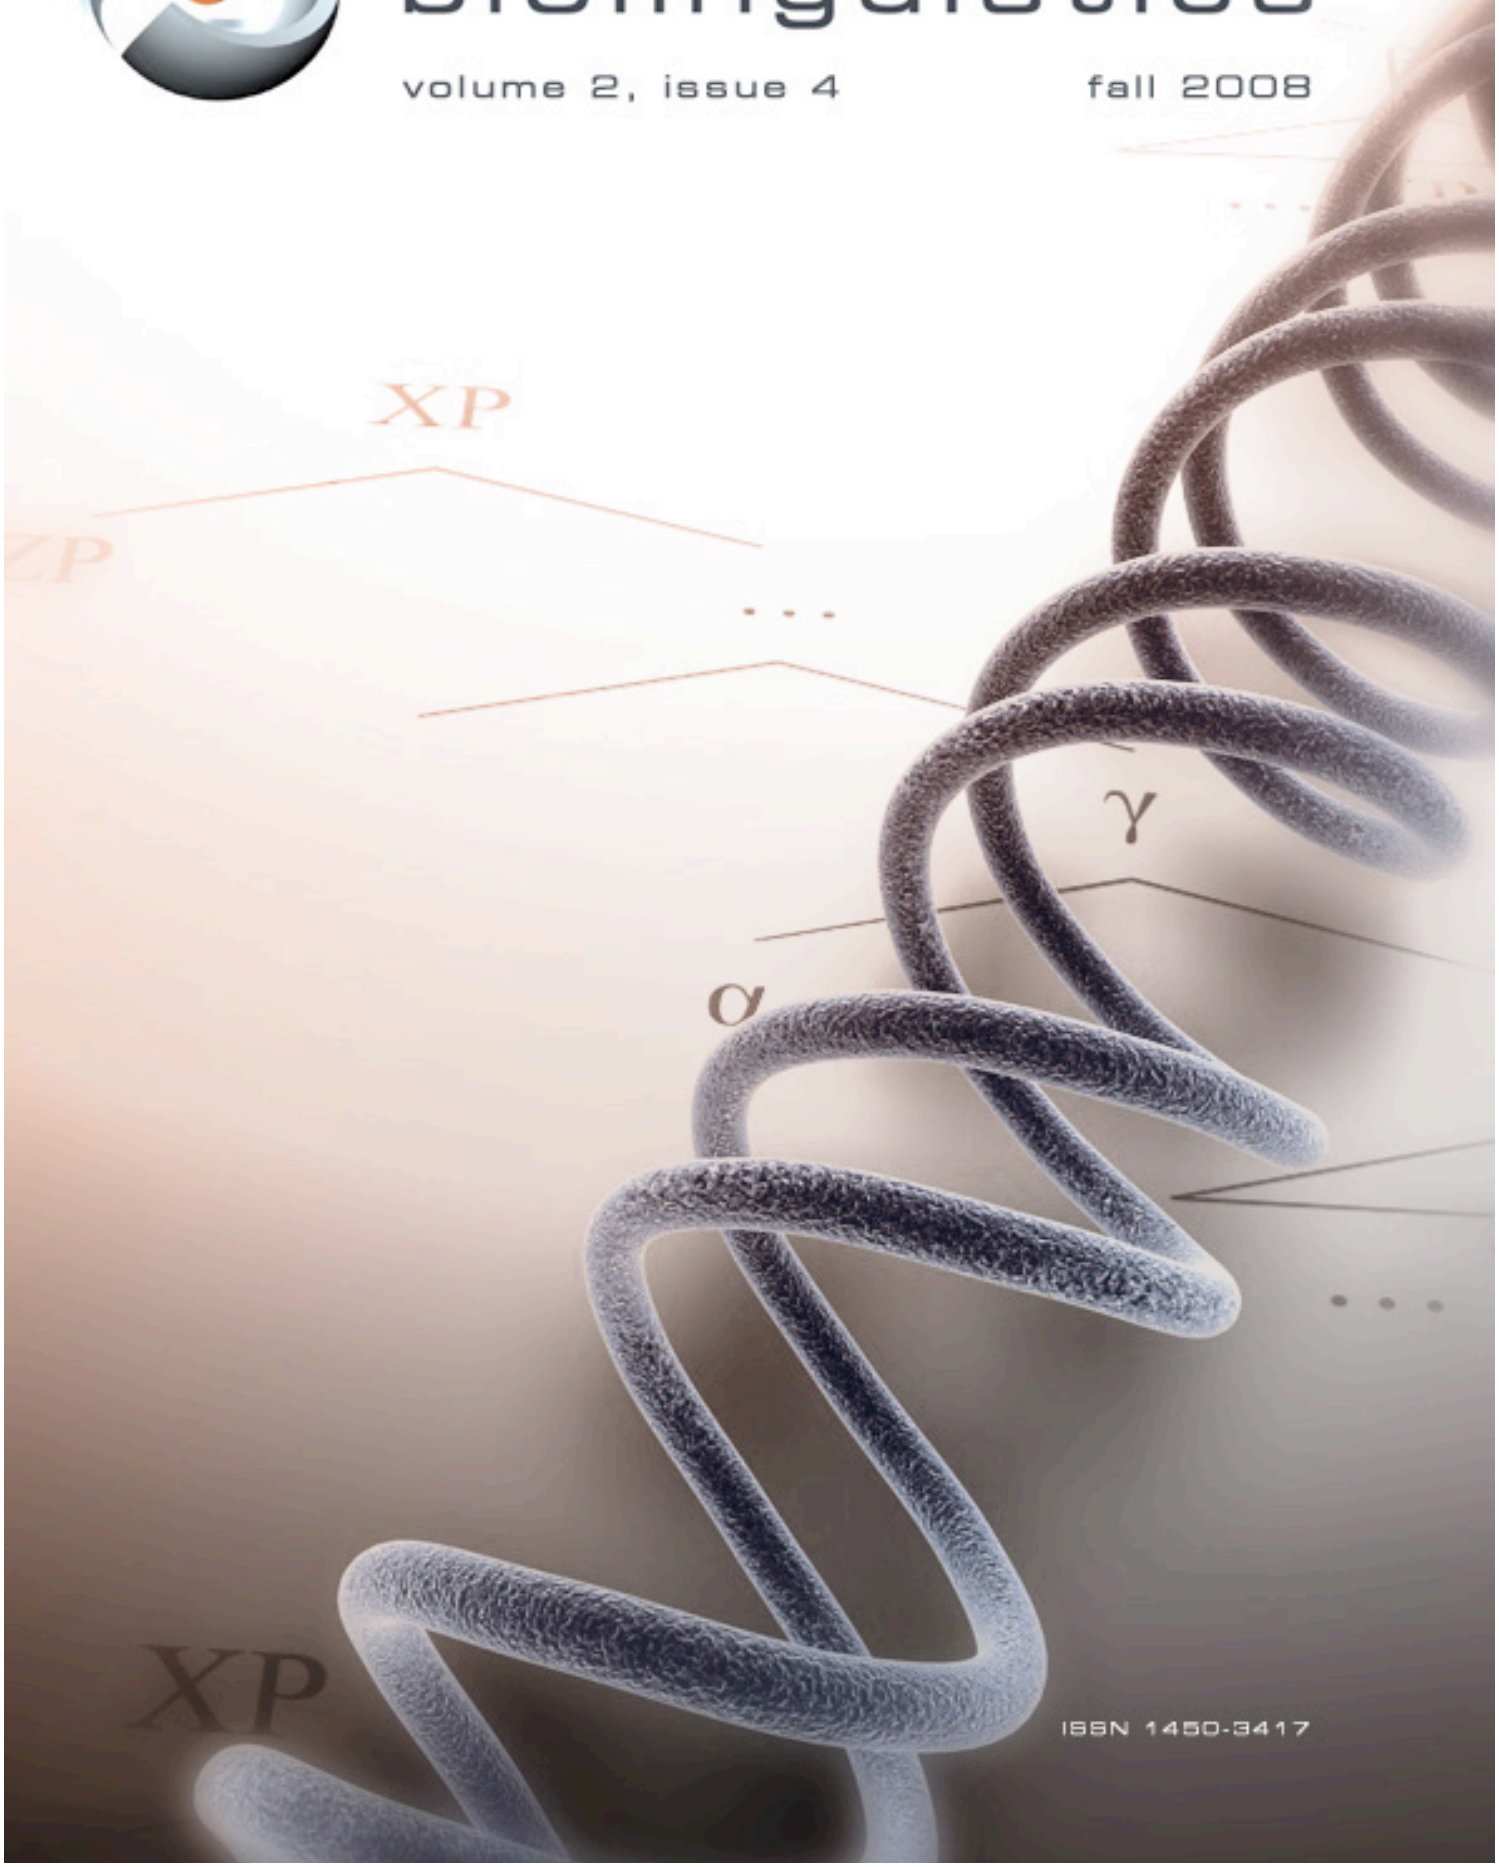

ISSN 1450-3417

The Asymmetry of Merge

Anna Maria Di Sciullo & Daniela Isac

This paper addresses the following question: What kind of properties must the structure-building operation Merge have such that, given a Numeration, the grammar will build the 'right' structure and avoid generating ill-formed configurations? The answer we will propose is that Merge should be seen as an asymmetric operation in the sense of relating two items whose sets of morpho-syntactic features are in a proper inclusion relation. In addition, we propose a partition of features into two stacks: categorial features and operator features. This distinction is independently motivated as it feeds into the definition of External Merge and Internal Merge (Chomsky's 2001). The proper inclusion condition will be assumed to hold for both of these operations, but the set of features under consideration for the evaluation of the proper inclusion relation differs: strictly categorial features for External Merge, and the whole set of features of lexical items for Internal Merge.

Keywords: asymmetrical relations, External Merge, feature valuing, Internal Merge, proper sub-set relation

1. Introduction

This paper is concerned with the properties of Merge, the operation that builds syntactic structures in the Minimalist program (Chomsky 1993, 1995, and related works). Our starting point is the observation that the existing definitions of Merge are relatively unconstrained, in the sense that many unattested structures may be built from a given Numeration.[1] Descriptively, what seems to matter in teasing apart the grammatical vs. ungrammatical structures that can be built from a given Numeration is *the particular order* in which the items are merged. For instance, if Merge first applies to the verb and the object, and the subject is merged later on in the derivation, a grammatical structure emerges. If on the other hand, Merge applies to the subject and little *v* before it incorporates the

This work is supported in part by funding from the Social Sciences and Humanities Research Council of Canada to the Asymmetry Project, grant number SSHRC 214-97-0016, as well as by a grant to the Dynamic Interfaces Project from the Fonds de recherche sur la société et la culture, grant number FQRSC 103690.

[1] This is also the case with Chomsky's (2001) definition of merge ramifying in External Merge (the merger of two syntactic objects that have not been previously merged) and Internal Merge (i.e. Move).

biolinguistics
© The Author (2008)

ISSN 1450–3417

Biolinguistics 2.4: 260–290, 2008
http://www.biolinguistics.eu

object into the *v*P, the resulting structure is ungrammatical. However, the *order* of application of Merge does not follow in any way from the current characterizations of Merge. In this paper we propose that such ordering constraints on the application of Merge follow from a view on *Merge as an asymmetric operation*. Our specific proposal is that Merge can apply only if the sets of features of the two merged elements are in a *proper inclusion relation*. Support for this view is offered not only by the fact that under this assumption we can block unwanted derivations, but also by a number of welcome consequences which follow from our view on Merge. One such consequence is that specifiers emerge as different from both complements and adjuncts. This is a desirable result, which converges with work by Nunes & Uriagereka (2000) and others. Another consequence of our view is that adjuncts also come out as having a special status with regard to the way they are integrated in the structure. Essentially, adjuncts are added to the structure that is being built by an operation that is not subject to the proper inclusion condition, and thus cannot be equated with Merge. This makes adjuncts different from both complements and specifiers. On the other hand, our proposal is also able to capture the fact that at some level of analysis adjuncts and specifiers pattern alike, and differ from complements. These asymmetries between complements/non-complements, specifiers/non-specifiers, adjuncts/non-adjuncts are thus related to a more basic asymmetry which is in-built as a property of the structure-building operation itself.

This paper unfolds as follows. In section 2, we describe the properties of Merge, as described by Chomsky (1995), and we discuss the predictions with respect to the order of application of Merge. In section 3, we discuss some possible solutions to these problems and show why these solutions fail to produce the desired effect. In section 4, we present our proposal, and in section 5, we lay out the consequences and the conclusions.

2. Merge

In the minimalist program (Chomsky 1995), syntactic structure is built bottom-up by the operation Merge, which has two crucial properties: (i) Merge is a *binary* operation, which combines two elements into one constituent, and (ii) Merge is *recursive*, where the output of Merge may subsequently be submitted to Merge with other elements yielding a further constituent. The items that are subjected to Merge are drawn from a list called the Numeration. A Numeration is defined as a set of pairs (*LI*, i), where LI is a lexical item and *i* is an index indicating the number of times that *LI* is selected. Every time a lexical item is selected from the numeration in order to enter the derivation, its index is reduced by one. The derivation terminates when all indices are reduced to zero. The first application of Merge selects two items from the numeration and all the other (subsequent) instances of Merge iteratively select items from the numeration, one by one, until the numeration is exhausted and a complex object is formed that contains all of the items that started out as individual elements. The iterative application of Merge is thus responsible for building up the structure from bottom to top; the initial input to the first application of Merge consists of terminal items, and the

final output of the last application of Merge is a hierarchical structure.

Whenever Merge applies to two items α, β, a new syntactic object is formed whose label is determined according to the following rule:

(1) *Merge* (Chomsky 1995)

Target two syntactic objects α and β, form a new object Γ {α,β}, the label LB of Γ(LB(Γ)) = LB(α) or LB(β).

Chomsky (1995: 243) distinguishes between two types of syntactic objects which can be subject to Merge: (a) lexical items and (b) objects of the type K = {γ, {α, β}}, where α, β are objects and γ is the label of K.

Initially, all objects in the numeration are of type (a), i.e. lexical items. Once Merge applies, it creates an object of type (b). The next stage of the derivation will thus include objects of type (a) and one object of type (b). Once an object of type (b) is created, Merge may assign a new element from the Numeration to this object (the workspace) and increase its complexity. Crucially, when an element from the numeration is merged to an already formed syntactic object, it can only be merged at the root of the structure existing in the workspace.

The implementation of Merge under this view presupposes the existence of a workspace: the space where the derivation unfolds and which will eventually contain the output of the recursive application of Merge.

To illustrate how the procedure works, consider the following Numeration:

(2) N = {(Mary,1), (*v*, 1), (loves, 1), (Peter, 1)}

If *loves* and *Peter* are selected from this list and submitted to the operation Merge, the result will be a binary structure whose label is determined by the projection rule in (1) above.

(3)

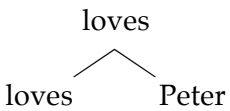

Next, one could select *v* from the Numeration and apply Merge again. The label of the resulting binary structure would again be determined according to rule (1).

(4)

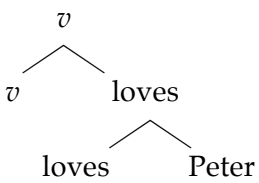

The last application of Merge will bring together *Mary* and the complex structure labeled as *v* in (4), and will produce the structure in (5).

(5)

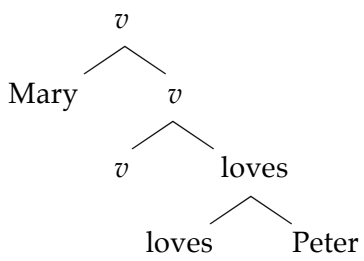

Since the Numeration now contains only items whose indices have been reduced to zero, the derivation will now stop.

The Merging procedure described above may indeed generate a well-formed structure, as in (5). However, this procedure may also build many unattested structures from the same set of terminals. To illustrate, let us start with the numeration in (6).[2]

(6) N = {Mary, *v*, winked}

If *Mary* and *v* are selected from the Numeration and if Merge applies to them, the result will be a binary structure labeled *v*, as in (7).

(7)

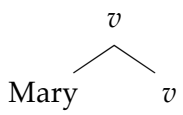

Next, Merge will assign the only remaining element in the Numeration, namely *winked*, and will build one of the structures in (8), which is ill-formed. (8a) yields the right word order but does not capture the predication relation between the subject and the verb, and (8b) fails both to capture the predication relation and to yield the right word order.

(8) a.

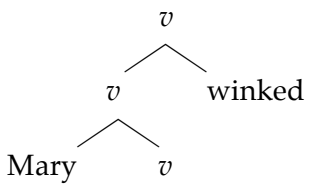

b.

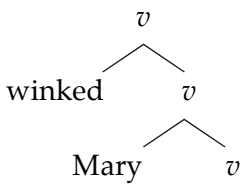

Similar considerations apply when deriving a transitive sentence. The 'right' derivation should first Merge the object with the verb and only then Merge the subject, but there is no property of Merge as defined above that would secure this order of application of Merge. Moreover, if the Numeration we start with is something like (9a), and the first instances of Merge produce a complex object like {*v*, {*v*, {is, {is, boring}}}}, there is no way to guarantee that what is Merged as a sister to this object is a phrase itself, and not just a head, as in (9b).

[2] In the rest of the paper, the numerations we will use for illustration purposes will ignore the numerical indices associated with the lexical items.

(9) a. N = {this, newspaper, *v*, is, boring}

b.

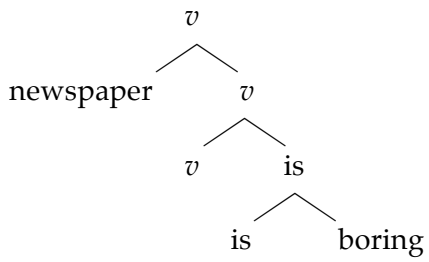

What seems to be needed is that the subject XP be built separately, through a successsion of Merge operations, and only then merged with {*v*, {*v*, {is, {is, boring}}}}. However, there is nothing in the description of the operation Merge above that would guarantee this.

This latter problem raises issues related not only to the order of application of Merge, but also to the notion of workspace. If at a certain stage in the derivation, there are two syntactic objects that were previously formed, as in (10), there is nothing in this procedure that flags {is, {is, boring}} as the workspace, rather than {this, {this, newspaper}}.

(10)

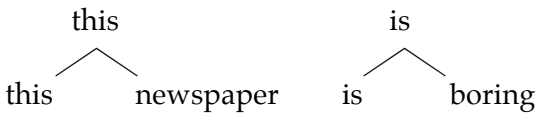

The choice is clearly important, since taking {is, {is, boring}} to be the workspace would allow a successful derivation (by selecting little *v* from the numeration and merging it to this syntactic object), but taking {this, {this, newspaper}} to be the workspace would lead to a crash, since the resulting structure will contain uninterpretable features. Such dead ends are undesirable from the point of view of economy. What seems to be needed is (i) a way to guarantee that objects are merged before subjects and that (ii) subject XPs are built separately, through a succession of Merge operations, and only then Merged to the workspace.

3. Possible Solutions

3.1. Renumeration

Johnson (2002) in an unpublished manuscript proposes that Merge is constrained by three factors: (i) the condition in (11), (ii) the Projection Rules, as stated in (12), and (iii) language-specific well-formedness conditions on the constituency of particular phrases.

(11) If an X^0 merges with a YP, then YP must be its argument.

(Johnson 2002: (12))

(12) *The Projection Rules* (Johnson 2002)

In $[_{\gamma}$ {α, β],

a. if just one of α and β is a phrase, then make γ a projection of the non-phrase;
b. if both α and β are phrases, then make γ a projection of the phrase that dominates the host.

These three factors will ensure in Johnson's view that certain phrases re-numerate after they are built in the workspace, and thus, that the right order of Merge follows. To illustrate, Johnson discusses the derivation of the constituent *flew after this talk*, and he considers two alternatives that can be chosen at the point where *after this talk* has been built. One possibility is for the derivation to proceed by selecting and merging the verb, as in (13) below. If this happens, a violation of the condition in (11) occurs.

(13) *Merge: (Johnson 2002: (19d))

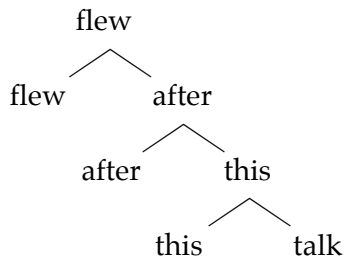

If, on the other hand, little *v* and the verb *flew* are selected and merged first, then renumerated, and then later merged with the PP *after this talk*, the derivation crashes because the application of the projection rules conflict with the well-formedness conditions on PPs in English (namely PPs cannot begin with a VP in English). In (14) below, it is the adjunct PP that projects according to the projection rules in (12), and this results in an ill-formed PP..

(14) *Merge: (Johnson 2002: (21g))

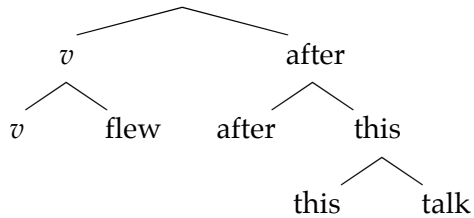

Instead of these two derivations, the successful one would first build *after this talk*, renumerate it, and then later merge it onto the independently constructed *v*P. Notice that the tree in (15) is identical to (14) except that the projected node is different in each case. This difference follows from the projection rules: In (14) above, the relevant projection rule is (12b) and it is the PP that is the host. In contrast, in (15), again the relevant projection rule is (12b), but this time the host is the little *v*.

(15) *Merge: (Johnson 2002: (22g))

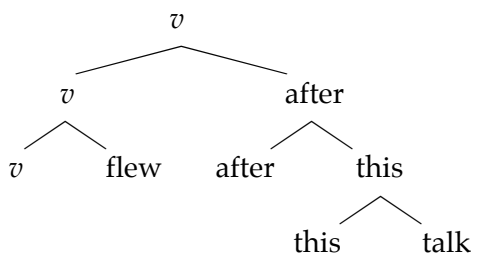

These constraints manage to account for Johnson's problem, namely why subjects and adjuncts are grouped together in disallowing extraction from them: Adjunct phrases and subject phrases, but not other phrases, will be forced to go through a stage in which they are renumerated. However, the question we are trying to answer in this paper is a different one: Given a numeration, is there any way to determine the order in which the items will be merged in order to form a complex syntactic object? This question does not receive an obvious answer in Johnson's terms. More specifically, once the derivation gets to the point where *after the talk* is built, two choices are in principle possible, according to Johnson: Either the derivation continues by further selecting items from the numeration and successively merging them to the already built syntactic object, or the derivation continues by first renumerating *after the talk* and then selecting other elements from the numeration. The choice depends, in Johnson's system, on two factors: (i) whether or not a violation of condition (11) occurs and (ii) whether or not well-formedness conditions of particular constituents are violated.

The first condition crucially makes reference to the argument status of a phrase: The verbal head can be merged to the PP only if the PP is interpreted as the verb's argument. The problem is that this decision must be taken at the point where the workspace contains the PP only and there is no intrinsic feature of the PP itself which determines whether the PP is an argument or not, and hence whether the PP is going to be renumerated or not. In order to decide this, the features of the verb need to be considered. Johnson's system does not exclude this possibility and is in fact compatible with it. However, his analysis does not explicitly spell out this direction. The reason is obviously related to the fact that his focus is different from ours; his aim is to provide an explanation of why adjuncts and subjects fall into a natural class, as evidenced by the fact that they are both islands for extraction.

The second condition, i.e. the observance of well-formedness conditions of particular lexical items, again depends on evaluating the features of the relevant head, in this case the preposition. Given the Projection rules in (12), not renumerating the PP would result in building an object whose label would be a P, since the PP would be the host for the next application of Merge. This in turn would violate the well-formedness condition on English PPs, namely the constraint that PPs cannot begin with a VP in English. The problem with extending this as a solution to our concern is that it is not clear where these well-formedness conditions come from. One could claim that the well-formedness conditions on English PPs are encoded as selectional restrictions on the head of the PP. If so, this indicates that such features must be taken into account.

3.2. Selectional Features

Another possible solution to the ordering question is to capitalize on the *c-selectional features* of lexical items. This notion goes back to Chomsky's (1965) 'strict subcategorization rules' which were meant to analyze a lexical category in terms of its local distributional context, or 'frames' where it can be inserted. In Chomsky's view, strict subcategorization rules are part of the set of Phrase Structure rules, but at the same time, they are also features that characterize some lexical subcategories, and as such, they are part of various items' lexical entries.

It is thus possible to determine whether an object in the Numeration is an argument or an adjunct by looking at the selectional features of the other items in the numeration. If an item has a selectional feature that selects another object in the Numeration, then the selected object is an argument. Hence, condition (11) above will apply to it. If, on the other hand, an item in the Numeration is not selected by any of the other items in the list, then it is an adjunct, and condition (11) will apply to it.

While this is clearly a possible solution to our problem, it comes with a price: Selectional features will have to be assigned a special status as compared to all the other features. The feature system will need to flag selectional features in a way that will guarantee that they will be involved in the operation Merge, to the exclusion of other features.

To discuss only a few examples of such views, in the *Aspects*-model, subcategorization features are contextual features that need to match corresponding Phrase Structure rules. In other approaches, such as Stabler's (1998), selectional features are again assumed to be different from other features. This difference is signaled by a special notation in Stabler's system: A selectional D feature on an item, for example, is encoded as '=d', and this is different from the way in which other features are encoded. Likewise, in Müller's (2007) view, subcategorization features are singular in that they are always at the top of an assumed hierarchy of features of lexical items. Moreover, if one lexical item has more than one subcategorization feature, the respective subcategorization features are ordered with respect to one another and are 'discharged' one after another, depending on their relative position in the hierarchy of features.

What these views share is that regardless of how the special status of selectional features is formalized, they are treated as different from other features. While the order of application of Merge can indeed be derived from the selectional features of the items in the Numeration, the 'cost' of this proposal consists in the burdening of the theory with a special type of features, which are different from the other types of features computed in the syntax. Ideally, one would need a system in which these features play a role but are not given a special status as compared to other features. There are several proposals in the literature to this effect. For example, in the approaches advocated by Svenonius (1994), Holmberg (2000), Julien (2000), Matushansky (2002), among others, c-selectional features are uninterpretable categorial features that must be checked against the categorial feature of the selected object. Insofar as selectional features are treated as members of a larger set of features (strong features, or uninterpretable features), these views are less stipulative, and therefore more desirable

than theories that consider selectional features as 'special' in any way. However, in these views it is not obvious how the order of Merge can be derived from selectional features.

One approach that seems to have both advantages, i.e. assume that selectional features are similar to other features and at the same time offer a solution to the ordering of Merge problem, is the one proposed in Adger (2003). In this approach, selectional features are uninterpretable and strong on a par with other (non-selectional) features. The notion of 'strength' of a feature is adopted from Chomsky (1993) and it essentially imposes a locality constraint on the feature that bears it: A strong uninterpretable feature must be checked under sisterhood.

In such a system, the ordering of Merge would be related to the locality requirement on the checking of strong uninterpretable features. Given that a verb like 'bring' for instance, has a strong uninterpretable [*u*N] feature, this feature will have to be checked under sisterhood, so Merge with an item bearing a matching [N] feature is the only possible choice.

However, in Adger's system, the ordering problem that we raise in this paper is still unresolved, we think. One point where the ordering problem becomes apparent is the case of items having more than one selectional feature, that is, more than one uninterpretable strong feature. One such case would be the verb 'show'. As discussed in Adger (2003), this verb has two selectional features, encoded as uninterpretable strong features: [*u*P] and [*u*N]. The fact that one of these features will be checked before the other one does not follow from the nature of uninterpretable strong features, since they are both of the same nature.

We thus conclude that views that would capitalize on c-selectional features in order to derive the ordering of Merge would not ultimately offer a solution to this problem.

4. Our Proposal

Our proposal is that the ordering of Merge follows from assuming that Merge is an asymmetric operation in the sense of (16).[3]

(16) *Asymmetry of Merge*

Merge is an operation that applies to a pair of elements in the Numeration whose sets of features are in a proper inclusion relation

Before we go on to illustrate how a derivation would proceed under this assumption, several remarks are in order.

First, we will assume that the numeration contains subarrays, i.e. sub-Numerations that define phases. The concept of Lexical Subarray was introduced by Chomsky (2000: 106ff.), who claims that Lexical Subarrays can be selected straightforwardly from the initial lexical array.

[3] See also Di Sciullo & Isac (in press) for the role of set inclusion in movement chains, and Di Sciullo (2005) for the role of set inclusion in morphological merger.

Second, when viewed as a procedure, Merge is an operation that involves sub-components, including select and search. Merge first selects an element from the Numeration, and then searches for an item whose features are in a proper inclusion relation with the initially selected one. Once search is complete, the two items are joined as sisters. The asymmetry itself is built into the sub-operation 'search'. Given that 'search' is part of Merge, we end up with a picture in which the proper inclusion relation is not a precondition on the application of Merge, but an organic part of Merge itself.

The third observation concerns the feature sets of lexical items that are relevant for the condition in (16). Only morpho-syntactic features that are active in the syntactic computation are taken into account. In particular, we exclude phonological features, and semantic features. Moreover, in line with recent studies (Hale & Marantz 1993, Di Sciullo 2005, Bobaljik 2007), we also exclude Case features and phi-features, since we assume that such features are not computed in the syntax, but in a different space. Apart from these exclusions, we are making the following assumptions regarding the morpho-syntactic features of lexical items. To begin with, we assume that each lexical item has two types of features: interpretable and possibly uninterpretable ones. This is a fairly common view in the Minimalist literature. Moreover, we will assume that c-selectional features are not derivable from the semantic (theta) properties of lexical items and thus that they should be listed in the lexicon.[4] In line with other authors (see the discussion above, section 3.2), we will not assign any special status to selectional features: We will assume that they are just uninterpretable features, on a par with other uninterpretable features that lexical items might have. For example, the fact that little *v* selects a VP will be captured by positing an uninterpretable V feature, [*u*V], on little *v*. However, according to standard assumptions, little *v* also has another uninterpretable feature — [*u*T], meant to capture the morphosyntactic relationship between Tense and little *v*. This latter feature is clearly not a selectional feature, but in our system it is treated the same, i.e. simply as an uninterpretable feature.

Fourth, we will assume that both lexical and functional items have selectional features. This differs from other proposals (Adger 2003, among others), that tie in selectional features to theta marking properties of lexical items, and under which functional items do not have selectional features, given that they do not enter into theta marking. Thus we do not assume the existence of an independent hierarchy of functional projections and the hierarchical relation between little *v* and VP, or between T and little *v*, is treated the same as the relation between a V and its object, i.e. as a reflex of the proper set inclusion relation in our view.

Last, and perhaps most importantly, apart from the interpretable/uninterpretable distinction, we will assume an additional division between categorial features and operator features. The operator features that we will assume are features like [wh], [Topic] or [Focus] involved in A-bar movement and

[4] This assumption converges with views like Grimshaw (1991), but see Pesetsky (1982) for an acquisition based argument that the primitives of theta theory are epistemologically prior to the primitives of c-selection.

responsible for creating configurations that will be assigned an operator-variable interpretation at LF.[5] However, our proposal is that these features are involved not only in movement to an A-bar position, but in all movement. In our view, the distinction between categorial features and operator features underlines Chomsky's (2001: 7f.) distinction between External Merge and Internal Merge (i.e. Move): Categorial features are relevant for External Merge, whereas operator features are for Internal Merge. This distinction thus feeds into the definition of these two operations, which are independently assumed to exist. Characterizing External Merge and Internal Merge in terms of the types of features that get operated on makes it possible to characterize both the differences and the similarities between these two operations. The proper inclusion condition will be assumed to hold for both of these operations, but the set of features under consideration for the evaluation of the proper inclusion relation differs: In the case of External Merge the features under consideration are exclusively categorial features, whereas in the case of Internal Merge (i.e. Move) the whole set of features of the lexical items involved are considered. We can thus restate the proper inclusion condition in (16) in terms of this distinction, as follows:

(17) a. *Asymmetry of External Merge*

External Merge is an operation that applies to a pair of elements in the Numeration whose categorial features are in a proper inclusion relation.

b. *Asymmetry of Internal Merge*

Internal Merge is an operation that applies to a pair of elements in the workspace whose (total set of) features are in a proper inclusion relation.

What this means is that the relevant set of features that have to be evaluated for proper inclusion is different for the two operations. As it will be illustrated below, the point in the derivation where the shift from considering strictly the set of categorial features to considering the total set of features (including the operator ones) occurs is predictable. Internal Merge will be viewed as a last resort operation that applies only if external Merge cannot apply.[6]

A similar shift is discussed in Müller (2007), who proposes, together with Chomsky (2000, 2001, 2005), that some features are not obligatorily present on certain heads and that they can be added in the course of the derivation.[7] The relevant heads that can benefit from such 'expansions' of their features are edge heads in Müller's approach. Similarly, in our view, the expansion from the set of categorial features to the total set of features (including operator features) also occurs on certain heads only, namely those heads that have operator features to

---

5 See Rizzi (1997) on operators involved in topic configurations.

6 This is again not a new point, and it has been coined in the literature as the Merge over Move preference.

7 Chomsky (2000, 2001, 2005) proposes that this can happen *after* the phase is complete, but in Müller's (2007) view, the timing is different, i.e. a feature can be added on a phase head *before* the phase is complete.

begin with. In line with recent proposals, it is only edge heads that actually have operator features (see Butler 2004 and references cited therein) so in fact in our view too this expansion is only possible once the derivation reaches the point at which a phase head is merged. The 'last resort' flavor of this extension is not however limited to the availability of heads having operator features, but goes beyond that. In our view, this extension is warranted only if the numeration is exhausted and External Merge cannot apply. Moreover, as discussed by Müller, features can be added on edge heads only if this has an effect on the outcome of the derivation. What "having an effect on the outcome" means for Müller is being instrumental in implementing intermediary movement steps required by the Phase Impenetrability Condition. Similarly, what "having an effect on the outcome" means for us is triggering movement. The extension from the set of categorial features to the total set of features is licensed when the effect is to trigger movement. Alternatively, in Heck & Müller (2000, 2003), Fischer (2004), and Müller (2004), 'having an effect on the outcome' means 'balancing a phase'. In other words, under these approaches, features can be added on edge heads only if this is needed for balancing a phase. A phase counts as balanced if for every movement inducing feature on items in the numeration, there is an available matching feature. This idea is also paralleled in our proposal in that once the extension to all features is available for the edge head, it also becomes available to other heads in the c-command domain of the edge head. These latter heads will provide the matching feature for the operator feature in the edge head.

One question that arises is related to how 'costly' this distinction between categorial features and operator features is. In particular, even though we do not assign any special status to selectional features in our theory, we do introduce a distinction between categorial features and other types of features, and this seems to be just as 'costly' as a theory that grants singular properties to selectional features. It is obviously true that a system that would manage to predict the order of Merge without relying on this distinction would be simpler and more elegant. However, we think that this distinction is independently motivated by virtue of the fact that it underlines the distinction between the two structure-building operations: External Merge and Internal Merge. To the extent that these are two different operations, with distinct properties, and to the extent that their differences can be reduced to a difference in the type of features they compute over, the distinction between the relevant types of features is thus not that stipulative. Moreover, even if by some measures we traded a stipulation of selectional features with one on categorial features, our theory, unlike a theory based on selectional features, can account for the order of Merge in cases in which one item has more than one selectional features.

Below, we illustrate the feature specification that we assume for the most common lexical items that could enter a derivation:

(18)

nouns:	[N]	
Indefinite D:	[Num]	
	[*u*N]	
definite D:	[D]	
	[*u*Num]	
wh-D:	[D]	
	[*u*Num]	
	[wh]	
unergative V:	[V]	
transitive V:	[V]	
	[*u*D]	
unaccusative V:	[V]	
	[*u*D]	
v	[*v*]	
	[*u*V]	
	[*u*D]	
	[*u*Tense]	
unaccusative *v*:	[*v*]	
	[*u*V]	
	[*u*Tense]	
tense	[Tense]	
	[*uv*]	
	[*u*D] [EPP]	
	[*u*ClauseType:]	(operator feature)
C_1	[D]	
	[ClauseType]	(operator feature)
	[*u*Tense]	
C_2	[D]	
	[ClauseType]	(operator feature)
	[*u*Tense]	
	[wh]	(operator feature)

4.1. Sample Derivation

The purpose of this section is to provide an illustration of how asymmetric Merge can lead to the 'right' order of Merge. Taking as a starting point the lexical entries provided above, let us see how an analysis based on asymmetric Merge predicts the 'right' order of Merge. Let's assume that the Numeration consists of the following elements grouped into subarrays as indicated:

(19) N = {C, T, {D, Num, N, *v*, V, D, Num, N}}

The successive steps in building the structure are all in compliance with asymmetric Merge, as defined in (17). The derivational steps, given a Numeration like

(19), are given below:

Step 1. Select an item from Numeration that has interpretable features only
⇒ Select **N** {[N]}

Step 2. Select an item from Numeration that properly includes N
⇒ Select Num {[Num], [*u*N]}

Step 3. External-Merge N with Num.

(20)

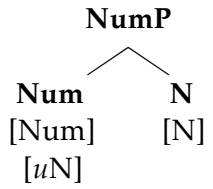

Given the Earliness Principle (Pesetsky & Torrego 2001), the uninterpretable feature of Num will get checked and erased as soon as possible[8].

(21)

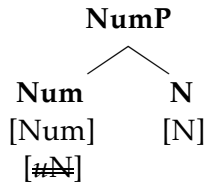

The newly created object is projected according to the rule in (1), i.e. out of the two items that merge, one will project its features. The newly created object thus has identical features with the item that projects. In our particular case, the projecting element is Num. The next step of the derivation will be to search for a new item in the Numeration whose features stand in a proper inclusion relation with the set of features of the object in the workspace.

Step 4. Select an item that properly includes Num
⇒ Select D {[D], [*u*Num]}

Notice that even though the set of features of D does not properly include the set of features of Num, as given in the Numeration, once Num is part of the derivation, its uninterpretable (selectional) feature gets checked and deleted, and thus its set of features comes to be properly included by the set of features of D.

Step 5. External-Merge D to the workspace and check uninterpretable features, as enforced by the Earliness Principle.

---

8 Notice that checking and deletion of uninterpretable features is subsequent to actual Merge. Even though after checking and deleting uninterpretable features, the proper inclusion relation does not hold, what is important is that it holds at the point where Merge applies, i.e. before checking and deletion of uninterpretable features of the newly merged item.

(22)

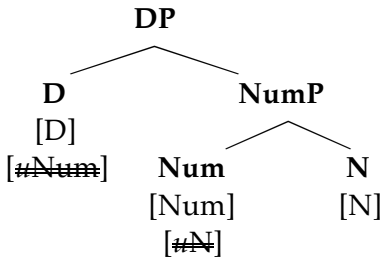

The workspace now contains a DP. If the Numeration contains more than one Noun, these 5 steps are taken for each Noun and all DPs are built in a parallel way. The result is that the workspace will potentially contain several DPs. The derivation will continue by Merging new elements selected from the Numeration to one of these DPs and the other DPs will be put on hold. What this means is that these DPs will be in the search domain for each of the next step of the derivation. The selection of the new element to be Merged to the workspace will again be made in compliance with the proper inclusion requirement.

Step 6. Select an item from the Numeration whose set of features properly includes the set of features of the DP in the workspace. Notice that at this point there are several choices, since not only the set of features of V properly includes the set of features of the DP, but also the set of features of little *v*, or T. We propose that the element that is selected must be the one whose features are the smallest superset of the set of features of the object in the workspace. In our particular case, the item that will be selected is V. This is a set theoretic version of 'locality'.

⇒ Select V {[V], [*u*D]}

Step 7. External-Merge V and DP and check uninterpretable features on V.

(23)

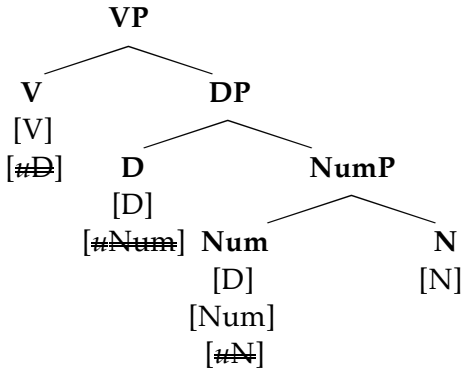

Step 8. Select an item from the Numeration whose set of features properly includes the set of features of the object in the workspace, i.e. VP.

⇒ Select little *v* {[*v*], [*u*V], [*u*D], [*u*Tense]}

Step 9. External-Merge little *v* to the VP which was in the workspace and check uninterpretable features on little *v*.

(24)

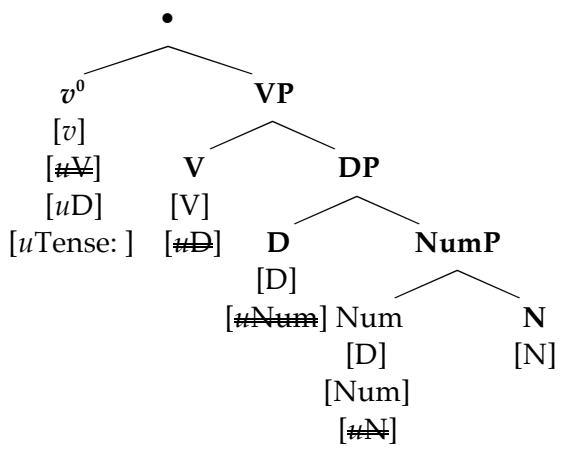

Step 10. The next step in the derivation should be to select an item from the Numeration whose set of features properly includes the set of features of the object in the workspace. There is one element in the numeration that fulfills this condition, i.e. Tense. However, this item is not part of the same subarray as all the items that have been already selected. The only syntactic object belonging to the same subarray is the DP that was put on hold in the workspace.[9] Even though this DP does not properly include the set of features of the syntactic object in (24), the proper inclusion condition can be met in the reverse direction, since the set of features of the object in (24) does properly include the set of features of the DP on hold.

⇒ Select DP {[D]} from the workspace/numeration

Step 11. External-Merge DP to the workspace and check uninterpretable features.

(25)

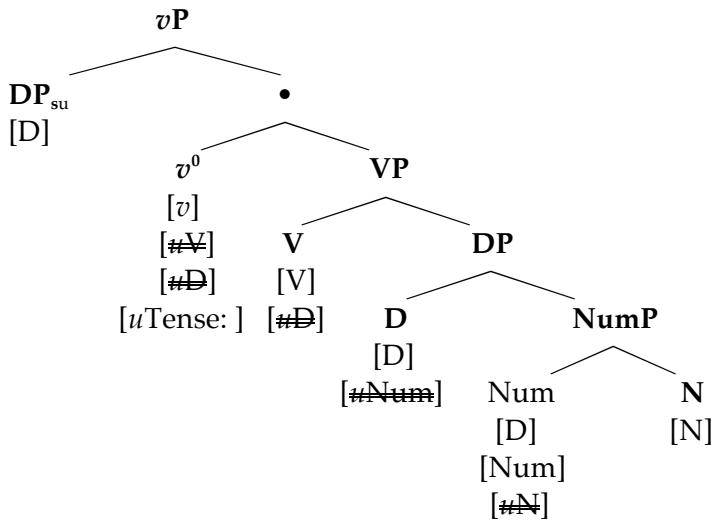

[9] Remember that several DPs can be built in a parallel way with the first five steps above, and that the derivation continues by Merging new items to one of these DPs, while the others are set aside for future derivational steps. Alternatively, the algorithm could have the operation 'renumerate'.

Notice that little *v* has three uninterpretable features: [*u*V], [*u*D], and [*u*Tense:]. Two of these, namely [*u*V] and [*u*D], correspond to selectional features. However, the ordering of checking of these two features is not a problem in our system, since it is a reflex of how the derivation grows and not a function of the properties of the selecting head.

Step 12. Select an item from the Numeration whose set of features properly includes the set of features of little *v*. The item that meets this condition is T, with the following set of features: {[Tense:Pres], [*uv*], [*u*D] + [*u*EPP], [*u*ClType:]}. However, not all of these features are considered when the proper inclusion relation is evaluated, since the Clause type feature is an operator feature. In our proposal, it is only categorial features that are considered for External Merge, to the exclusion of operator features. Notice that in this particular case, the proper inclusion relation holds regardless of whether the whole set of features of T is considered, or only the subpart of categorial features. So, restricting the set of features to categorial features for External Merge seems to be irrelevant in this case. However, there are other instances, such as the merger of a DP bearing a *wh*-feature, in which this restriction will be shown to play a role.

⇒ Select T {[Tense:Pres], [*uv*], [*u*D] +[*u*EPP], [*u*ClType:]}

Step 13. External-Merge T to the workspace and check uninterpretable features.[10]

(26)

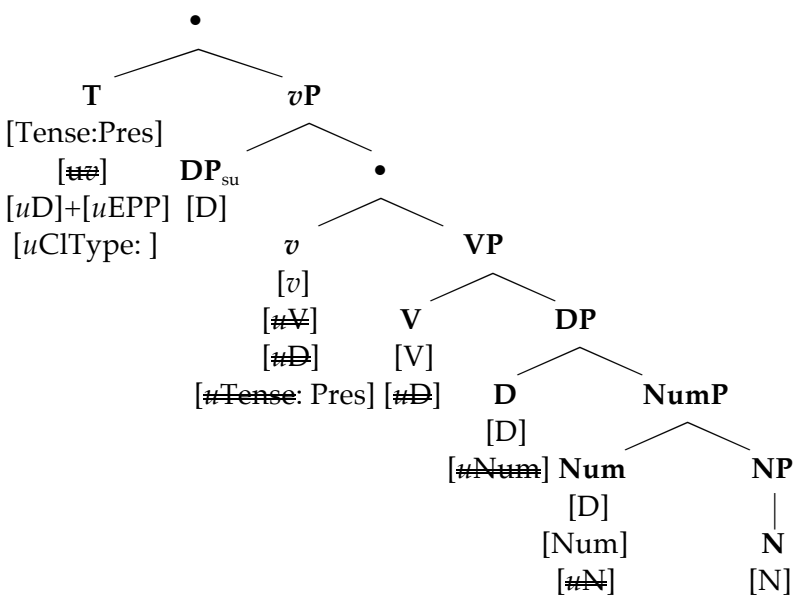

Step 14. The next step, according to our proposal, should be the selection of an item whose set of features properly includes the set of features of the object in

[10] Even though only the categorial features are considered for the purposes of External Merge, when External Merge applies, it applies to the lexical item, and so not only the categorial features are External-Merged, but all of the features that are part of the feature matrix of the respective item. The obvious analogy is pied-piping in the case of movement.

(26). However, at this point, there is only one item left in the numeration, namely the Complementizer, and clearly, C's set of features does not properly include the set of features of the object in (26). Moreover, there is no DP on hold anymore: The numeration only contained enough 'material' for the construction of two DPs, and both DPs have already been merged at this point.

Our proposal is that this is exactly the kind of situation when internal Merge becomes an option and the search domain for the next Merge is therefore switched from the Numeration to the list of terminal items in the already constructed object. In other words, rather than looking in the numeration in order to select the next item to be merged, the derivation can look within the tree in (26) for the selection of the next item to be merged.

The selected item will have to satisfy the same proper inclusion requirement, i.e. its set of features must be in a proper inclusion relation with the set of features of the highest node, i.e. T. Moreover, since External Merge is not possible at this point, and Internal Merge is the only option, in our view the sets of features that are evaluated will have to include operator features as well, as discussed above. So far, the only item in the numeration that had operator features was T itself. Since no element has a set of features that properly includes the set of features on T, the search will select items whose set of features is properly included by the set of features of T. The closest DP will thus be selected.

⇒ Select D(P) {[D]} in [Spec,*v*P]

Step 15. Internally Merge DP to the workspace and check uninterpretable features.

(27)

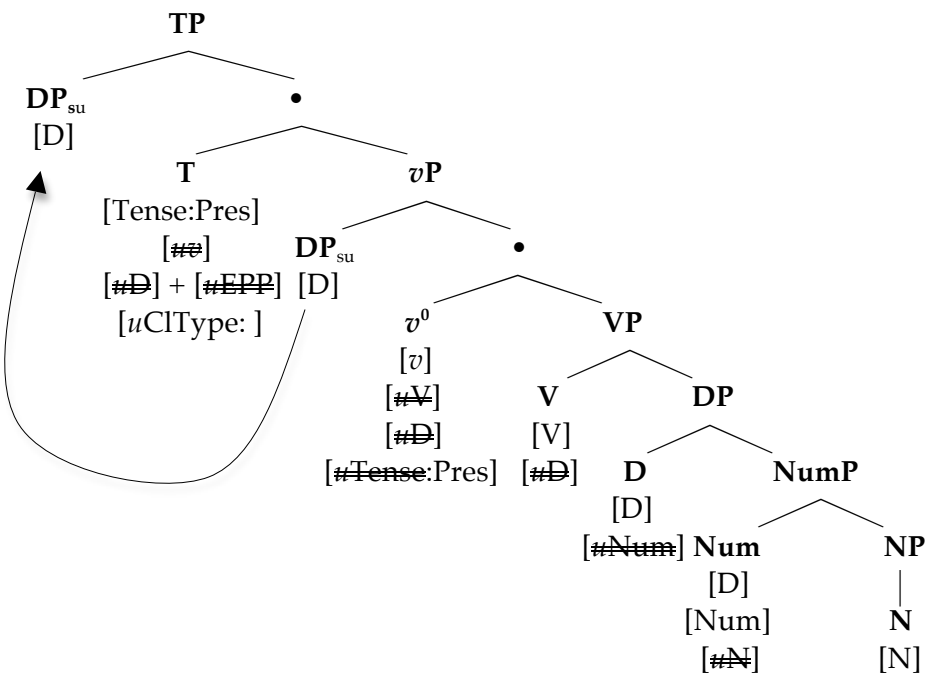

Step 16. The derivation continues by searching for items in the numeration whose set of features properly includes the set of features of T. Notice that the features of T that were taken into account in the previous step, i.e. Internal

Merge, included the operator feature on T, namely its Clause type feature. This operator feature will however be ignored in this next step, given that the search domain is the numeration and thus that the derivation is searching for an item to be External-Merged. The only item left in the numeration is the complementizer. Its total set of features includes {[D], [ClType], [*u*Tense]}. However, for purposes of External Merge, only the categorial features are considered, and thus the Clause type feature on the complementizer is ignored.

⇒ Select C {[C], [*u*ClType:], [*u*Tense]}

Step 17. External-Merge C to the workspace and check uninterpretable features.

(28)

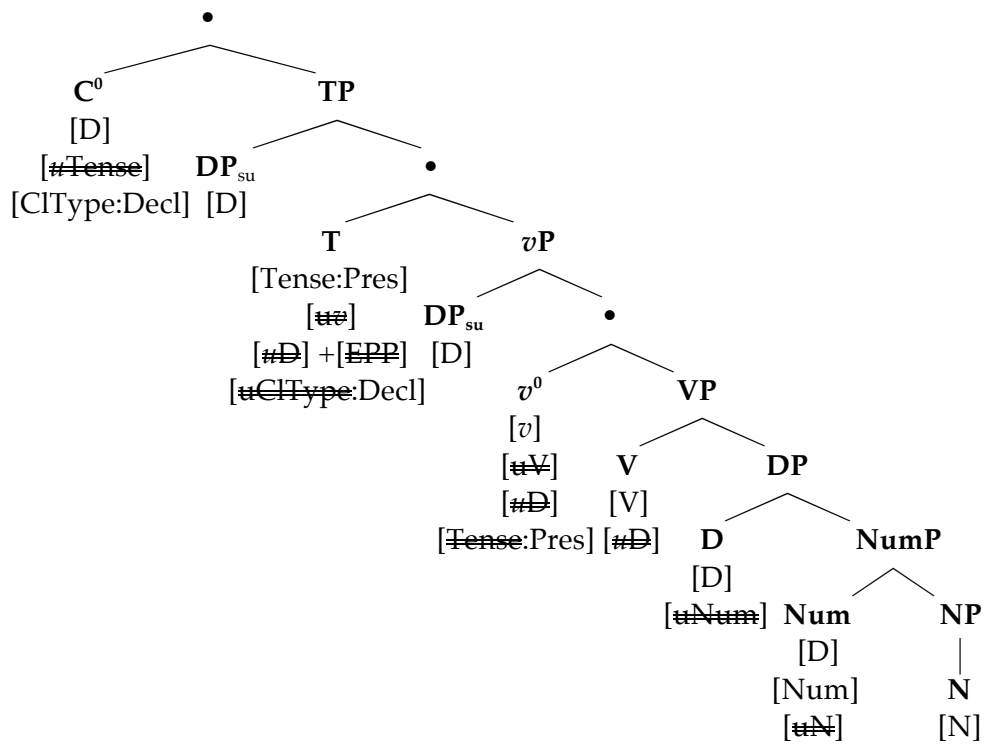

Once the C is External-Merged, T's [*u*ClType] feature will check against the matching interpretable [ClType] feature of the complementizer. The numeration is now exhausted, and all the uninterpretable features are checked. The derivation is now completed.

Let us now see how the derivation of a *wh*-interrogative would proceed. In this case, the Numeration will obviously contain different feature specifications for the complementizer and the *wh*-DP. The complementizer will have to include an uninterpretable [*u*wh] feature, which is an operator feature, and the *wh*-DP will be headed by a D with a [wh] feature, again an operator feature. Moreover, we propose that the feature specification of T will also be different in a *wh*-interrogative, in the sense that it will also include a *wh*-feature. Apart from C, T and the displaced *wh*-item, all the other items in the numeration will have the same feature specification as assumed above. The assumption that C includes a *wh*-feature in *wh*-interrogatives, or that the displaced constituent has one are not

new. The assumption that the T also has a *wh*-feature in *wh*-interrogatives can be supported by the overt manifestation of this feature in Bantu languages like Kinanda, as well as in other languages such as Berber, Celtic, Fiorentino and Trentino, Kikuyu, Palauan, Somali, Turkish, Ojibwe (see Schneider–Zioga 2002 and references cited therein). These languages show *wh*-agreement on the verb when the *wh*-DP is the object, and a preverbal *wh*- agreement affix when the moved *wh*-DP is the subject. What is crucial for us is that the overt *wh*-morphology *on* the verb is different depending on whether the *wh*-constituent is the object or the subject. This difference can be accounted for in our proposal in a natural way. As will become apparent below, we take *wh*-agreement with a subject to be a reflex of checking the [wh] feature on T against the *wh*- feature on the subject in a Spec-head configuration, whereas *wh*-agreement with the object is assumed to be the morphological manifestation of the checking relation between the same two features, but in a different syntactic configuration. More specifically, the object does not move and T's [wh] feature enters agreement (checking) with the [wh] feature of the object in situ.

In what follows we will not go through all the steps of a derivation of a *wh*-interrogative, but will focus only on the steps that differ from the derivation presented above. To begin with, we expect a first difference in the step that merges a *wh*-object and the verb. However, given that the *wh*-feature is an operator feature, and given that the Merge between the verb and the *wh*-object would be an instance of External Merge, the potential operator features of the verb or of the object are ignored. In other words, the features that are evaluated for proper set inclusion in this case are exactly the same as in the case of a non-*wh*-object, namely the categorial features, to the exclusion of operator *wh*-features. If the *wh*-item is not the object, but the subject, one expects a difference in the step that merges the *wh*-subject with a syntactic object like (24). Again, given that the *wh*-feature on the subject is an operator feature, and given that the merging of the subject in the specifier position of the little *v* is an instance of External Merge, the features under consideration will be the categorial ones. Thus, the proper inclusion relation between the little *v* and the subject DP will be unaffected by the presence or absence of a *wh*-feature on the subject.

The next steps in the derivation that are likely to be affected by the presence of *wh*-features are the steps involving T and C, since we assume that both T and C have *wh*-features. When T is merged to *v*P, the proper inclusion relation still holds, since T is supposed to be the superset and since the *wh*-feature is a feature posited on T. Moreover, since T is merged to *v*P as an instance of External Merge, the *wh*-feature on T will be irrelevant, as it is an operator feature. Now, in the next step of the derivation described above, the subject DP moves to [Spec,TP]. Since Move is Internal Merge, operator features become visible. If the *wh*-constituent is the object, T will attract the closest DP, namely the subject DP. The proper inclusion relation between the set of features of T and the set of features of the subject will clearly hold (the only difference from the equivalent step described above will be the presence of a *wh*-feature on T, hence an extra feature on what is supposed to be the superset anyway). If the *wh*-constituent is the subject, the proper inclusion relation will not change as compared to the equivalent step in the derivation above, since both T and the

attracted subject will have an additional *wh*-feature.

Let us now examine what happens when the derivation gets to the level of the *wh*-C. When a *wh*-C is Merged, the *wh*-feature is ignored, since this is an instance of External Merge. However, in contrast to the derivation described above, the derivation of a *wh*-interrogative will not end with the merging of the C. This is because the *wh*-feature on C is uninterpretable and moreover associated with an uninterpretable EPP feature. What is similar to the derivation above is the fact that at this point the numeration is exhausted in both cases. So the only way for the derivation to continue is by Internal Merge. The closest constituent with a matching *wh*-feature will be attracted — it could be either the subject or the object, depending on which of the two bears a *wh*-feature. Whichever is attracted, the set of features of C, which should now be assumed to include not only categorial features, but also the *wh*-operator features, properly includes the set of features of the *wh*-DP.

Before concluding, we would like to discuss an apparent problematic consequence of our proposal that was brought to our attention by one of the reviewers. More specifically, our system seems to have the undesirable consequence that skipping projections in the functional domain is impossible, since we need a selectional relationship between each Merge in the functional sequence. Even if selection is treated as an uninterpretable feature, the problem still seems to remain, since the occurrence of aspectual projections like PerfP or ProgrP will need to be encoded as uninterpretable features on the head that selects them, and their occurrence, on the other hand, is optional. The T head, for instance, will have to bear an uninterpretable [*u*Perf] feature, to *encode* the fact that PerfP is a complement of T, and the Perf head will bear a [*u*Progr] feature, in order to capture the fact that the ProgP is a complement of Perf. Given that these projections are optional, we will end up either with multiple T's, Perf's, and Prog's (each with a different uninterpretable feature encoding selection), or else with a number of optional selectional features on each of these heads. It therefore looks like our system needs to stipulate that a *v*P is obligatory in the structure, as is TP, while the remainder of the projections are optional. One way to deal with this problem would be to stipulate a hierarchy of functional projections that would take care of the ordering relation between these projections. However, once we do this, our proposal that the ordering of Merge can be seen as a reflex of the proper set inclusion relation will lose substantiation.

While the problem of optional projections is obviously hard, we think that it is not impossible to offer a solution in terms of the proper set inclusion relation. What follows is a tentative solution to this problem, one that builds on a suggestion we found in Matushansky (2002) and Adger (2003). Based on a motivation that is independent from selection or from the need to provide an ordering relation between these projections,[11] Adger suggests that little *v*, Prog, Perf, etc. bear a more general [*u*Infl] feature, rather than more specific features like [*u*Tense] or [*u*Perf]. Similarly, Matushansky (2002), who builds on Julien

[11] Adger (2003) in fact assumes a hierarchy of functional projections and proposes a system in which the relation between little *v* and VP or between T and *v*P, for instance, is conceptually distinct from the selection relation that holds between a V and its object.

(2000), proposes that the lexical head determines the c-feature of all the functional heads in the extended projection of that lexical head. This proposal provides a means of formally expressing the notion of extended projection (Grimshaw 1991).

Building on this suggestion, we suggest the following feature specification for the relevant functional heads. For little *v*, we suggest that instead of the [*u*Tense] feature that we assumed above, a [*u*Infl:] feature could be considered. This [*u*Infl] feature could be valuated by any of the possible functional heads that could optionally appear on top of *v*P. If a Progressive head is present, its features will be {[Infl:Prog], [*u*Infl:], [*uv*]}. This set of features properly includes the set of features of little *v* (i.e. {[*v*], [*u*Infl:]}). Once the Prog head is merged, the [*u*Infl] feature on little *v* will be valuated as Progressive, and the uninterpretable [*uv*] feature on Prog will be checked.

(29)

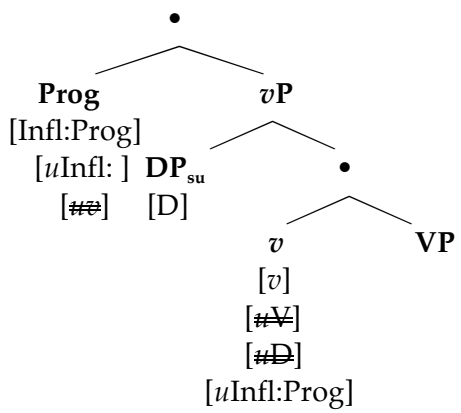

Now, if a Perfect aspectual head is also present in the numeration, its features can be assumed to include {[Infl:Perf], [*u*Infl:], [*u*Pol:]}. This set of features properly includes the features of the object in (29) above. Crucially, the [*uv*] feature on the Progressive head has been checked and erased by the time the Perf head is merged. Notice also that positing a [*uv*] feature on both the Prog and the Perf head does not commit us to having a *v*P as an obligatory complement to any of these heads. Having the *v*P as a complement is possible for these heads, but not obligatory. This is because selectional features are treated as uninterpretable features in our system, and uninterpretable features are under no restriction to be checked locally: They could be checked locally, under sisterhood, but they could also be checked at a distance, under c-command.

Once the Perf head is merged, the [*u*Infl:] feature on the Prog head is valued as Perf, and the [*uv*] feature on Perf is checked.

(30)

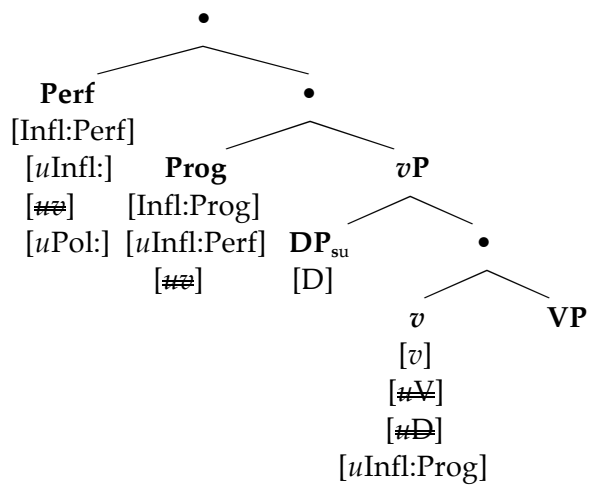

It is easy to see how the Prog head could be missing in this configuration without disturbing the proper set inclusion relation between the Perf head and the little *v*.

Apart from these two aspectual heads, and possibly a Voice head, which we have not considered here, but which would work along the same lines, there is another projection that for simplification purposes we have not taken into account in the sample derivation presented above, but which we nevertheless take to be an obligatory projection. This projection (that we will call a SigmaP, following Laka 1990 and others) hosts a polarity feature that can bear either the Negative or the Affirmative value. The total set of features of the Sigma head that we are assuming is: {[Pol:neg/aff], [*u*Infl:], [*u*Infl], [*uv*]}. This set properly includes the set of features of the object in (30) above, that is, {[Infl:Perf], [*u*Infl:], [*u*Pol:]}.

When the Sigma head is merged, there are two features on the Sigma head that get checked: its uninterpretable [*u*Infl] feature that captures the selectional property of Sigma, as well as its [*uv*] feature. Likewise, when the Sigma head is merged, the [Pol:] feature on the Perf head is also checked.

(31)

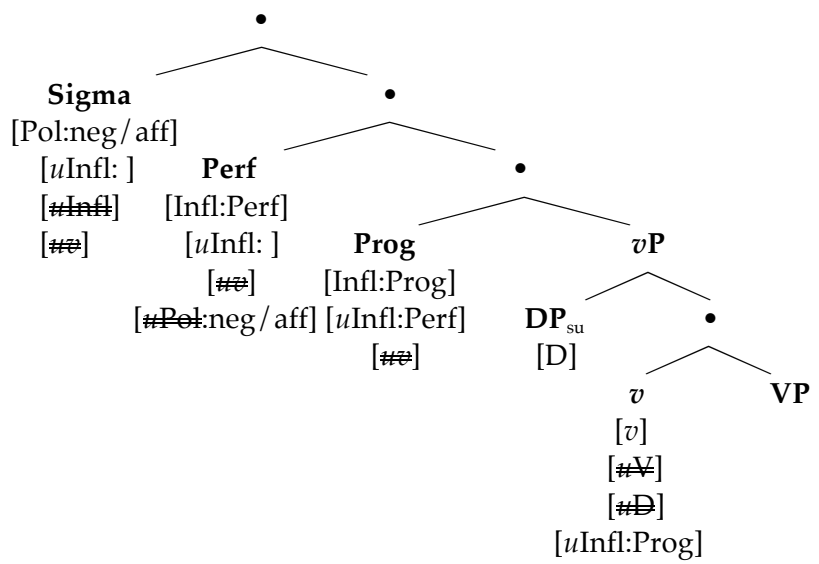

If the PerfP is not projected, the proper inclusion relation between the features of Sigma and the features of Prog still obtains, and if both Perf and Prog are absent, the proper inclusion relation between the features of Sigma and the features of little *v* is likewise preserved.

Finally, the T head in the numeration will have the following features: {[Infl:Tense], [*u*Pol], [*uv*], ~~[*uv*]~~, [*u*Infl], [*u*ClauseType:], [*u*D], [EPP]}. T's set of features properly includes the set of features of the object in (31) above, even under the restrictive assumption that only categorial features should be taken into account when evaluating the proper inclusion condition. When T is merged, its ~~[*uv*]~~ [*u*Infl] feature is checked against the interpretable [*v*] feature of Perf, and its [*u*Pol] feature is also checked against the interpretable [Pol] feature of the Sigma head. Moreover, the [*u*Infl:] feature of the Perf head and of the Sigma head can now be valuated as Tense.

(32)

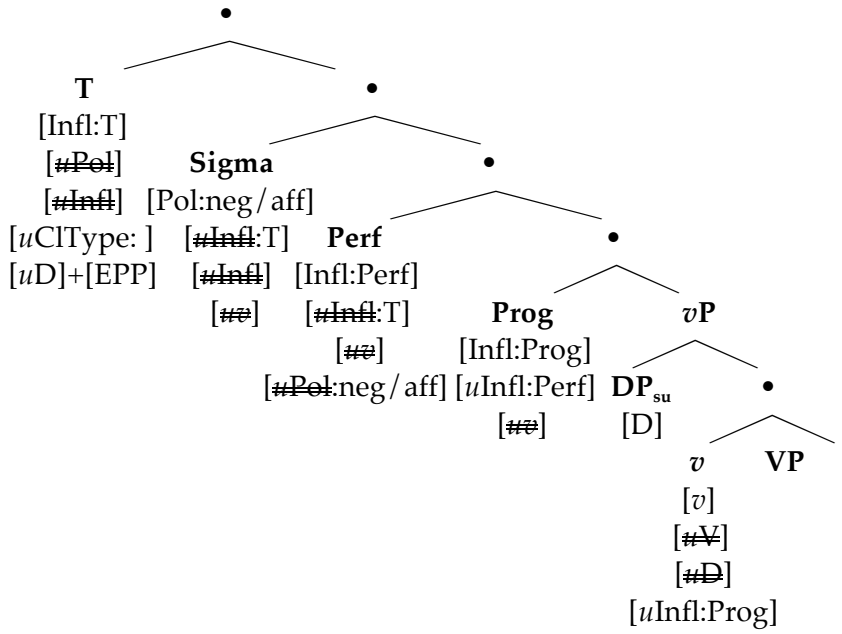

Notice that at the point in the derivation where ProgP has been built, and a new item has to be selected from the numeration, the choice between Perf and Sigma could be switched. The proper set inclusion condition would still be met, but PerfP would be higher than SigmaP.

(33)

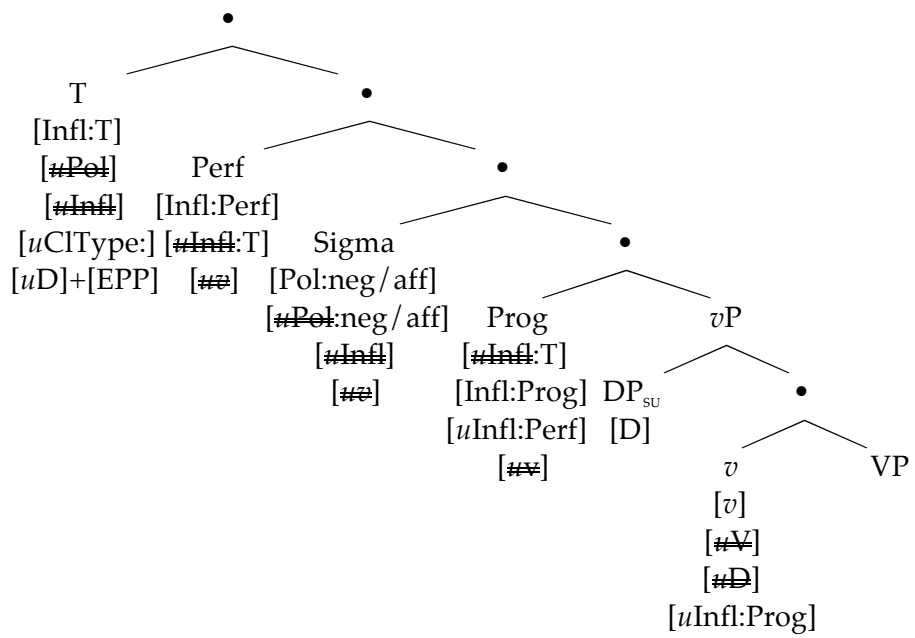

If Sigma is merged first, its [*uInfl*] feature is checked against the [Infl:Prog] feature of the Prog head, and its [*uv*] feature is checked against the [*v*] feature of the little *v* head. The remaining set of features — {[Pol:neg/aff], [*u*Infl:]} — is properly included in the set of features of the Perf head which is merged next. Once Perf is merged, its [*u*Infl] feature is checked against the [Infl:Prog] feature of the Prog head, its [*uv*] feature is checked against the [*v*] feature of the little *v* head, and its [*u*Pol:] feature is checked and valued by the [Pol:neg/aff] feature of he Sigma head. The remaining features of the Perf head — {[Infl:Perf], [*u*Infl:], [~~uPol~~:neg/aff]} — is properly included in the set of features of the T head, which is merged next.

This is in fact a desirable result, in spite of the fact that it seems to introduce optionality in our theory. A Sigma phrase carrying polarity features can indeed occur on either side of the Perf head.

(34) a. He claims to have not understood the instructions.
b. He claims to not have understood the instructions.

This is predicted to be possible under a theory such as Butler's (2004) who argues that phases should be defined as any domain that has a predicative core, a layer of functional structure and a quantificational layer (including modality and polarity phrases). In this view, not only *v*P and CP are phases, but also aspectual phrases, such as PerfP and ProgP. Given that the periphery layer can be reiterated on top of each of these phases, we end up with a picture in which the Sigma Phrase could show up both higher and lower than the Perf head.[12] This is

[12] Sigma Phrase can also be projected on top of TP, according to Butler, as part of the CP layer on top of TP. The higher Sigma head will have different features than the Sigma in (32) and (33). In particular, its set of features will include {[*u*Pol:], ([EPP]), [*u*Infl], [*u*ClType:]} and will thus properly include the set of features of T. Crucially, we follow Zanuttini (1994) and

captured in (32) and (33) above by the fact that the Perf head is assumed to have a [*u*Pol:] feature which can be checked and valued either by a lower Sigma, or by a higher one.

As already mentioned, the discussion above is tentative. Questions arise for example as to why can the [*u*Infl] feature on Perf or Sigma for instance not be valued from below, i.e. from the Prog head or from the Perf head respectively, and why do they have to be valued from a c-commanding head. We leave this for future research.

On the other hand, the exact feature content of each item in the numeration plays a crucial role in our account. It could be claimed that even though we have eliminated stipulations related to selectional features, as well as to the hierarchy of functional projections, we have moved the stipulations into the lexicon, in the sense that to a certain extent we have stipulated the feature content of the items in the numeration. At least some of the features we have assumed are features that have been used elsewhere in the literature, and are independently motivated. For the rest, ideally, it should also be possible to show that they are motivated independently from our proposal on Merge, something we have not done in this paper. Rather, we have treated the feature specification of lexical items as a hypothesis and then explored the consequences of this hypothesis. To the extent that the order of Merge can be derived from this hypothesis, this can be taken as motivating the latter.

5. Conclusion

In this paper we proposed a procedure for deciding the order in which items in the Numeration must be Merged. The criterion that we proposed is the proper inclusion relation: The set of features of the merged item must stand in a proper inclusion relation with the set of features of the object derived in the workspace. In addition, we assumed that the proper inclusion relation is 'local' in the sense that the element that is selected must be the one whose features make up the smallest superset of the set of features of the object in the workspace.

Our proposal does not assign any special status to selectional features. The latter are assumed to be uninterpretable features, on a par with other uninterpretable features, which cannot be seen as selectional. In addition, the order of Merge in the view presented in this paper is linked not to selectional features but to a relation between sets of features of lexical items, namely a proper inclusion relation.

We have proposed that the sets of features that are evaluated for the proper inclusion condition is different in the case of External Merge vs. Internal Merge. For External Merge, we have assumed that the relevant set of features is reduced to categorial features, whereas for Internal Merge, we have proposed that there is

---

Culicover (1996) in assuming that the concept of sentential negation is syntactically distributed across two functional projections: one under T and one higher than T. For language like Italian, where the negative marker occurs higher than T, the negative marker is assumed to move from the lower Sigma position to the higher one.

a last resort shift from the set of categorial features to the total set of features of lexical items, including both categorial features and operator features. Characterizing External and Internal Merge in terms of relevant sets of features involved allowed us to capture both the similarities and the differences between the two structure-building operations.

The content we provide to the notion of asymmetric Merge is different from what other authors have proposed. In particular, it is different from a notion of asymmetry related to which of the two elements that enter Merge will project, as expressed by Chomsky (1995: 246), or Langendoen (2003: 310).

> The operation Merge (α, β) is asymmetric, projecting either α or β, the head of the object that projects becoming the label of the complex formed.
> (Chomsky 1995: 246)

Our notion of asymmetric Merge is also different from the derivational view advocated in Jaspers (1998), Zwart (2006), and Johnson (2002). Under this latter view, Merge is asymmetric in the sense that it does not simply join two elements A and B, but it actually transfers one element at a time from the numeration to a workspace (derivation). Element A is thus merged to a workspace B, instead of A and B merging together. The result of Merge can then be described as an ordered pair <A,B>, where B is the current stage of the derivation and A the newly added element. In Zwart's terms, merge turns the current derivation into a dependent. Dependency is a semantic relation which must be syntactically realized. The core dependency relations are (a) head-complement: The complement is the dependent of the head, and (b) subject-predicate: The predicate is the dependent of the subject. The latter relation might be surprising, but the argument against considering the subject as being the dependent, rather than the predicate, is that the subject is not directly related to the verb. Arguments are related to the verb, but a subject can be any type of argument and even a non-argument. Therefore, the asymmetric relation of dependency is not between head and dependent but between dependent and non-dependent.

Our proposal on Merge as involving a proper inclusion relation between the two elements that undergo Merge does not exclude this view, but is different from it. Our claim is that the asymmetry of Merge is reflected in the morpho-syntactic properties of the members of <A,B>, in the sense that the set of morpho-syntactic features of A properly includes the set of morpho-syntactic features of B. This view is thus compatible with a derivational approach of the theory of grammar, in that we are not assuming that derivations are driven by some internal global syntactic architecture, but we assume instead that derivations proceed on a strictly local basis, caring only about the syntactic relations between members of sister pairs.

Crucially, the 'derivational' view on Merge cannot account for the ordering problem: given a Numeration, how can we predict the order in which the elements of the Numeration will be Merged to the workspace?

Under the assumptions that we are making in this paper about the nature of selectional features and about the feature specification of lexical items in the

numeration, the ordering of Merge operations follows. In addition, our proposal also has the desired effect of building subject DPs separately and of merging them as phrases, rather than as lexical items. This result is obtained by our assumption that all DPs are built in a parallel way and simultaneously in the workspace. In our view, after these DPs are built, the derivation continues by selecting one of them as the host for the next Merge and leaving the others on hold. Moreover, Specifiers emerge as having a special status not only in the sense that they need to be built separately, but also from the point of view of the directionality of the proper inclusion relation. In the step-by-step derivation described above the proper inclusion condition always holds between the workspace and the merged item, but the directionality of proper inclusion is not consistent. Given that the object in the workspace is extended by the Merge of a new element, the expectation is that the direction of proper inclusion should be from the Merged element to the object in the workspace. In other words, the set of features of the merged element should properly include the set of features of the object in the workspace. However, even though in most of the steps shown in the derivation above this is the directionality of proper inclusion, sometimes this directionality is reversed and the set of features of the object in the workspace properly includes the set of features of the merged item. This change of direction of the proper inclusion relation is apparent only when items are merged in a 'specifier' position.

An additional result of our analysis is that adjuncts are integrated into the derivation by an operation that has different properties than Merge. More specifically, the operation that puts together an adjunct and the XP that is its adjunction site is not subject to the proper inclusion condition. Clearly, there is no proper inclusion relation between the set of features of, say, a PP adjunct, and the set of features of *v*P.

Lastly, the question arises as to why there should be a proper subset relation built in the definition of Merge. One possible answer to this question could be that this is required by conceptual necessity, in order to ensure legibility of the interfaces by the external systems. Interestingly, there is a morphism (structure preserving mapping) between asymmetric Merge and the semantic operation of functional application. In functional application, generally, the function must be of a higher type than its argument. In asymmetric Merge, generally, the grammatical features of the selector must properly include the ones of the selectee. We see this morphism as a conceptual motivation as to why the proper subset should hold for the merger of predicates with arguments. The morphism ensures the mapping of the expressions derived by the core operations of the grammar and the ones derived by the operations of the conceptual-intentional system.

References

Abney, Steven. 1987. The English noun phrase in its sentential aspect.

Cambridge, MA: MIT dissertation.
Adger, David. 2003. *Core Syntax*. Oxford: Oxford University Press.
Adger, David & Gillian Ramchand. 2003. Merge and move: Wh-dependencies revisited. In Lisa L.-S. Cheng & Norbert Corver (eds.), *Wh-Movement Moving On* (Current Studies in Linguistics 42). Cambridge, MA: MIT Press.
Bennis, Hans, Norbert Corver & Marcel den Dikken. 1998. Predication in nominal phrases. *The Journal of Comparative Germanic Linguistics* 1, 85–117.
Bobaljik, Jonathan David. 2007. Where's phi? Agreement as a post-syntactic operation. In David Adger, Susana Béjar & Daniel Harbour (eds.), *Phi-Theory: Phi-Features Across Interfaces and Modules* (Oxford Studies in Theoretical Linguistics 16). Oxford: Oxford University Press.
Butler, Jonny. 2004. Phase structure, phrase structure, and quantification. York: University of York dissertation.
Carstens, Vicki. 1991. The morphology and syntax of determiner phrases in Kiswahili. Los Angeles, CA: University of California dissertation.
Carstens, Vicki. 2000. Concord in minimalist theory. *Linguistic Inquiry* 31, 319–355.
Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge MA: MIT Press.
Chomsky, Noam. 1993. A minimalist program for linguistic theory. In Ken Hale & Samuel Jay Keyser (eds.), *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger* (Current Studies in Linguistics 24), 1–52. Cambridge, MA: MIT Press.
Chomsky, Noam. 1995. *The Minimalist Program* (Current Studies in Linguistics 28). Cambridge, MA: MIT Press.
Chomsky, Noam. 2000. Minimalist inquiries: The framework. In Roger Martin, David Michaels & Juan Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, 89-155. Cambridge, MA: MIT Press.
Chomsky, Noam. 2001 Derivation by phase. In Michael Kenstowicz (ed.), *Ken Hale: A Life in Language* (Current Studies in Linguistics 36), 1-52. Cambridge, MA: MIT Press.
Chomsky, Noam. 2005. Three factors in language design. *Linguistic Inquiry* 36, 1–22.
Culicover, Peter. 1991. Polarity, inversion, and focus in English. *ESCOL* 8, 46–68.
Di Sciullo, Anna Maria. 2005. *Asymmetry in Morphology* (Linguistic Inquiry Monographs 46). Cambridge, MA: MIT Press.
Di Sciullo, Anna Maria & Daniela Isac. In press. Movement chains at the interfaces. *Canadian Journal of Linguistics — Special Issue on Interfaces*.
Fischer, Silke. 2004. *Towards an Optimal Theory of Reflexivization*. Tübingen: Eberhard-Karls-Universität Tübingen dissertation.
Grimshaw, Jane. 1991. Extended projection. Ms., Brandeis University.
Hale, Morris & Alec Marantz. 1993. Distributed morphology and the pieces of inflection. In Ken Hale & Samuel Jay Keyser (eds.), *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger* (Current Studies in Linguistics 24), 111–176. Cambridge, MA: MIT Press.
Hauser, Marc D., Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298, 1569–1579.

Heck, Fabian & Gereon Müller. 2000. Successive cyclicity, long-distance superiority, and local optimization. *WCCFL* 19, 218–231.

Heck, Fabian & Gereon Müller. 2003. Derivational optimization of Wh-movement. *Linguistic Analysis* 33, 97–148.

Holmberg, Anders. 2000. OV order in Finnish. In Peter Svenonius (ed.), *The Derivation of VO and OV* (Linguistik Aktuell / Linguistics Today 31), 123–152. Amsterdam: John Benjamins.

Jaspers, Dany. 1998. Categories and recursion. *Interface* 12, 81–112.

Johnson, Kyle. 2002. Towards an etiology of adjunct islands. Ms., University of Massachusetts, Amherst.

Julien, Marit. 2000. Syntactic heads and word formation: A study of verbal inflection. Tromsø: University of Tromsø dissertation.

Kayne, Richard S. 1993. Towards a modular theory of auxiliary selection. *Studia Linguistica* 47, 3–31.

Kennedy, Christopher & Jason Merchant. 2000. Attributive comparative deletion. *Natural Language and Linguistic Theory* 20, 553–621.

Laka, Itziar. 1990. Negation in syntax: On the nature of functional categories and projections. Cambridge, MA: MIT dissertation.

Langendoen, Terence. 2003. Merge. In Andrew Carnie, Heidi Harley & Mary Ann Willie (eds.), *Formal Approaches to Function in Grammar: In Honor of Eloise Jelinek* (Linguistik Aktuell / Linguistics Today 62), 307–318. Amsterdam: John Benjamins.

Matushansky, Ora. 2002. Movement of degree/degree of movement. Cambridge, MA: MIT dissertation.

Matushansky, Ora. 2006. Head movement in linguistic theory. *Linguistic Inquiry* 37, 69–109.

Müller, Gereon. 2004. Phrase impenetrability and wh-intervention. In Arthur Stepanov, Gisbert Fanselow & Ralf Vogel (eds.), *Minimality Effects in Syntax*, 289–325. Berlin: Mouton de Gruyter.

Müller, Gereon. 2007. On deriving CED effects from PIC. Paper presented at the *30th Annual Colloquium of Generative Linguistics in the Old World (GLOW 30)*, Tromsø, Norway. [12–14 April, 2007]

Nunes, Jairo & Juan Uriagereka. 2000. Cyclicity and extraction domains. *Syntax* 3, 20–43.

Pesetsky, David. 1982. Paths and categories. Cambridge, MA: MIT dissertation.

Pesetsky, David. 1987. Wh-in-situ: Movement and unselective binding. In Eric Reuland & Alice ter Meulen (eds.), *The Representation of (In)definiteness* (Current Studies in Linguistics 14), 98–129. Cambridge, MA: MIT Press.

Pesetsky, David & Esther Torrego. 2001. T–to–C: Causes and consequences. In Michael Kenstowicz (ed.), *Ken Hale: A Life in Language* (Current Studies in Linguistics 36), 355–426. Cambridge, MA: MIT Press.

Picallo, Carmen. 1991. Nominals and nominalizations in Catalan. *Probus* 3, 279–316.

Ritter, Elizabeth. 1991. Two functional categories in noun phrases: Evidence from Modern Hebrew. In Susan D. Rothstein (ed.) *Perspectives on Phrase Structure: Heads and Licensing* (Syntax and Semantics 25), 37–62. San Diego, CA: Academic Press.

Rizzi, Luigi. 1997. The fine structure of the left periphery. In Liliane Haegeman (ed.), *Elements of Grammar: Handbook of Generative Syntax* (Kluwer International Handbooks of Linguistics 1), 281-337. Dordrecht: Kluwer.

Schneider–Zioga, Patricia. 2002. The case of anti-agreement. *MIT Working Papers in Linguistics, Proceedings of Austronesian Formal Linguistics Association (AFLA 8).* Cambridge, MA: MIT Press.

Siloni, Tal. 1997. *Noun Phrases and Nominalizations: The Syntax of DPs* (Studies in Natural Language and Linguistic Theory 40). Dordrecht: Kluwer.

Stabler, Edward. 1998. Acquiring languages without movement. *Syntax* 1, 72–97.

Svenonius, Peter. 1994. C-selection as feature checking. *Studia Linguistica* 48, 133–155.

Szabolcsi, Anna. 1989. Noun phrases and clauses: Is DP analogous to IP or CP? In John Payne (ed.), *Proceedings of the Colloquium on Noun Phrase Structure*. Manchester: University of Manchester, Department of Linguistics.

Zanuttini, Raffaella 1994. Re-examining negative clauses. In Guglielmo Cinque, Jan Koster, Jean-Yves Pollock, Luigi Rizzi & Raffaella Zanuttini (eds.), *Paths Towards Universal Grammar: Studies in Honor of Richard S. Kayne,* 427–451. Washington, DC: Georgetown University Press.

Zwart, C. Jan-Wouter. 2006. Local agreement. In Cedric Boeckx (ed.), *Agreement Systems* (Linguistik Aktuell / Linguistics Today 92), 317–339. Amsterdam: John Benjamins.

Anna Maria Di Sciullo
University of Quebec at Montreal
Department of Linguistics

P. O. Box 8888, Downtown Station
Montreal, Quebec H3C 3P8
Canada
di_sciullo.anne-marie@uqam.ca

Daniela Isac
Concordia University
Department of Classics, Modern Languages and Linguistics
H-663 1455, de Maisonneuve West
Montreal, Quebec H3G 1M8
Canada
disac@alcor.concordia.ca

Symmetry in Visual and Linguistic Perception

Brent Henderson

Like linguistic perception, visual perception is an active process in which the mind makes use of innate structural principles in the computational process. It is therefore useful to ask whether the visual and linguistic computational systems make use of the same or similar principles. This article describes the role played by principles of symmetry in visual perception as suggested by researchers in that field, and suggests that a subset of those principles play a strong role in the perception of linguistic structure. It is claimed that a distinction should be made between the construction of linguistic structure and its perception in the computational system, and that principles of symmetry apply in subtly different ways in each. It is argued that movement's inherent locality, successive-cyclicity, has a bipartite nature, being sensitive only to certain barrier nodes in the construction of structure while adjoining to every intermediate projection in the perception of structure.

Keywords: linguistic structure, successive-cyclic movement, symmetry, visual perception

[A]ny particular representation makes certain information explicit at the expense of information that is pushed into the background and may be quite hard to recover.
(Marr 1982: 21)

1. Introduction

A central insight of linguistic inquiry over the past half a century is that language interpretation is not a passive task, but an active one. Language, as it exists in the form of everyday speech, is understood to be a partially-specified phenomenon upon which the mind imposes additional requirements (hierarchical structure, for example). Much of linguistic inquiry is an effort to discover and better understand these requirements. Similar statements can be made about the cognitive understanding of visual perception. Visual perception is widely recognized as an active process wherein visual systems often compute fully-specified properties (such as motion or shape) given only partially-specified

This paper develops a presentation given at the Biolinguistics Investigations conference in Santo Domingo, Dominican Republic in February 2007. I thank the attendees for comments, especially Alberto de Almeida, Kamil Ud Deen, Charles Reiss, Karin Stromswold, and Ken Wexler. Thanks also to Cedric Boeckx, Michael Leyton, and three anonymous reviewers for very helpful discussion.

biolinguistics
© The Author (2008)
ISSN 1450–3417
Biolinguistics 2.4: 291–316, 2008
http://www.biolinguistics.eu

stimuli. Just as in linguistics, there are therefore strong poverty of stimulus arguments for environment-independent, innate principles of visual perception.

Given these parallels, we might fruitfully ask the question whether some of these environment-independent principles in linguistic and visual perception might be the same. From both a biological and evolutionary perspective, to arrive at such a conclusion would not be surprising. After all, both language and vision are cognitive computational systems and are, with little doubt, the two that most dominate the human cognitive landscape: Humans are very visually and linguistically guided creatures. Furthermore, it is very likely that the human visual system predates by far our linguistic system in our evolutionary history. We suspect this because, physically speaking, other animals including apes and cats, have visual systems similar to our own (though with important differences), while of course no animal has a corresponding linguistic system akin to that of humans. This sequence of evolutionary development, therefore, makes the visual computation system a very likely source for the exaptation of computational principles that could be employed for linguistic computation as well.

In the present work, I suggest that this is the case: that a certain subset of computational principles employed by the human visual system in perception is also employed in the perception of linguistic structure. If this conclusion turns out to be correct, we will have uncovered some "third factor" principles in the sense of Chomsky (2005); that is, principles employed in linguistic computation that are not unique to language.

The principles I will consider here are principles of symmetry involving simple mathematical concepts from geometry and basic group theory. To find such principles playing a strong role in our cognitive processes is surprising;[1] nevertheless, that they play a role in visual perception is well-known, even if controversial and poorly understood. In this paper, I will suggest that they also play a role in the how we perceive and generate linguistic structure. I will specifically apply these principles to one linguistic phenomenon, successive-cyclic movement, that historically has been difficult to naturally implement in derivational theories of syntax. Briefly, successive-cyclic (SC) movement is the notion that overly long distance movement in syntax does not take place 'all-at-once,' but in a series of shorter, successive movements through intermediate positions. While there is a variety of empirical evidence that suggests S-C movement exists, syntactic theory has struggled for decades toward a natural implementation, often resorting to unmotivated ad hoc features to encode it. At the center of the difficulty has been disagreement about exactly which intermediate positions S-C movement targets: Some argue all available intermediate positions between the extraction site and final landing site must be targeted; others argue that it is only certain positions (barriers or phase edges, for example) that S-C movement is sensitive to. I will flesh out this history below and suggest that the role for principles of symmetry that I describe here is what is behind the lack of consensus. In doing so, I present an account that understands S-C

[1] Surprising even though, as a reviewer points out, the push toward greater simplicity and elegance are driving forces of the minimalist program. But as Chomsky as pointed out, if some strong version of minimalism turns out to be true, it is indeed surprising since one does not expect this of biological systems (Chomsky 2000: 9).

movement to have both a derivational and perceptual nature, each with its own properties. To provide a preview, my claim will be that, perceptually, S-C movement targets each available intermediate position while, derivationally, S-C movement targets only certain barrier-like projections (namely, [Spec,CP]). To the extent that the perceptual side of the account is true, I argue, most S-C movement needn't be explicitly encoded into the derivation of syntactic structures, thus eliminating the need for ad hoc features in the grammar.

This paper is organized as follows. In section 2 I outline the basic mathematical principles of symmetry that concern us. In section 3, I discuss a small subset of findings from the literature on vision that suggests these principles play an important role in visual perception. Then in section 4, I suggest that the same principles come into play in the perception of linguistic structure in similar ways. Indeed, I hope to show that this is a central finding of the linguistic enterprise over the past 50 years, though things are seldom discussed in these terms. Finally, in the second part of this section, I suggest that the principles discussed here offer an understanding of successive cyclicity that does not rely on unmotivated movement or current theories of phases. In tackling these issues, I hope to encourage others to think of linguistic theory in the terms employed here and to encourage the search for and refinement of general organizational principles that might be common to various human cognitive faculties.

2. Principles of Symmetry

The terms symmetry and asymmetry are used in various ways in various contexts, so it is important to define exactly what I mean by their use. When I speak of principles of symmetry here, I have in mind the meaning of the term as used by mathematicians. Unlike the common use of the term, the mathematic property of (a)symmetry is never inherent. That is, objects themselves cannot properly be said to be symmetric or asymmetric. Rather, the symmetry of an object can only be defined with regard to a *transformation* that relates the object to another image of itself. Transformations are simple, single step operations (such as rotation or reflection) that can relate two images. If the two images related by a transformation are identical, then that object is said to be symmetric under that particular transformation. For example, a simple square related to its image by a 90 degree rotation results in two squares that are exactly alike in every way. Thus, a square can be said to be *symmetric* under 90 degree rotation, or, to put it another way, a 90 degree rotation can be said to be a symmetry of a square.

The set of transformations under which an object is symmetric is said to be that object's *symmetry group.* This group will always be a closed set of transformations, since the combination of any two member transformations will always yield another member of the set. For instance, 90, 180, 270 and 360 degree rotations belong to the symmetry group for a square since the square is unchanged under any of those transformations. A vertical reflection transformation, which relates the square to its mirror image along a vertical axis, will also be a member. Combining this reflection transformation with, say, the reflection transformation with a 90 degree rotation, however, will yield nothing new since

the resulting image will be exactly equivalent to a 180 degree rotation.

Given that a square is symmetric under so many transformations, we speak of it as having a high level of symmetry (though not as high as some other shapes, such as a circle). Now consider the relationship between the two images below. The image on the left is our highly symmetry square. The image on the right looks like a square that has had two triangular pieces cut from it. Unlike the square, the image on the right is not symmetric under a 90 or 270 degree rotation. Nor, unlike the square, is it symmetric under diagonal reflections along axes that bisect its right angles. Note, however, that the two images do still share quite a few symmetries: Both are symmetric under 180 degree rotations as well as vertical and horizontal reflections along axes that bisect their sides.

Figure 1: Symmetry breaking

Figure 1 thus illustrates the important concept of *symmetry breaking*. Some symmetries of the square on the left have been broken in the image on the right. Nevertheless, even when symmetry is broken, it is usually the case that a large amount of symmetry remains in common between related images. In fact, whenever two images are related via transformations, it will be the case that the symmetry group of the image with fewer symmetries will be a subset of the symmetry group of the image with more symmetries. It is this relationship between their respective symmetry groups that allows the two images to be related.

This basic understanding of symmetry preservation and symmetry breaking will be essential to the following discussion. Below, I will suggest that the mind makes use of the concepts of transformations, symmetry groups, and symmetry breaking when perceiving structural shape, and that these principles also form the basis for the construction and perception of linguistic structure.

3. Symmetry in Visual Perception

In any area of inquiry, scientists often find it helpful, in fact crucial, to take measures to tease apart the governing principles of a system from peculiarities of its performance in any particular application. Often this means removing a system from its typical functioning environment and seeing how it functions 'in a vacuum.' Linguists do this, for example, when they ask for grammaticality judgments on sentences without providing a discourse context. By doing so they sometimes find speakers have sharply negative judgments about a construction that they might otherwise accept in natural conversation. If one were to only examine natural corpora, one might never know about such judgments. The same kind of isolation and abstraction approach was taken regarding research in visual

computation by the Gestalt school of psychology beginning over 100 years ago.[2] In undertaking this approach, researchers found that subjects had judgments about their visual experience that did not always line up with reality. To preview one case I will discuss below, one well-known effect the Gestalt school discovered was that objects can be perceived to undergo motion even though no actual motion has taken place, so-called 'apparent motion' effects.

In searching for possible explanations for these judgment phenomena, both vision researchers and linguists have adopted the very reasonable view that psychological experience in the absence of proper stimuli must be a reflection of the deep organizing principles of the mind, principles often obscured in usual cognitive performance. That is, vision and language perception are taken to be active rather than passive processes. In both cases, fully-specified experience is computed from only partially-specified stimuli, providing poverty of the stimulus arguments for innate principles of vision and language.

One of the most interesting insights into the nature of these basic organizing principles of vision is the central role that the above notion of symmetry seems to play. In this paper, I would like to focus on three general conclusions from the literature that illustrate this role. These are listed below:[3]

(1) a. Visual computation makes use of symmetry principles.

b. Where transformations are imposed by the cognitive system, it is the simplest possible transformations that are employed.

c. Symmetry principles are employed by the mind to infer past states of an object from the present one.

In the following subsections, I illustrate each of these conclusions in turn.

3.1. Computing Similarity

One subfield of visual computation is concerned with how judgments of similarity between objects are formed. The dominant view for the past twenty-

---

[2] A reviewer points out that the approaches taken to be analogous here are in practice somewhat different. Typically linguists study an object abstractly in order to determine its properties, only later and secondarily considering how the object might be integrated into its natural context. In vision, the typical approach is to study an object abstractly to determine its properties so that perception of the wider context in which the object is identified might be better understood. While the difference is significant, the basic methodology is similar enough to make my point: Studying natural objects in isolation reveals facts and generalizations not possible from looking at them in a natural context.

[3] In the present discussion of visual perception, I am necessarily simplifying what are intense and complex discussions in a field I am only peripherally familiar with. While I believe that most researchers in the field agree that principles of symmetry and simplicity can be fruitfully applied to the kinds of phenomena I mention here, the precise role these principles play and what their origins might be are hotly debated topics. My purpose here is not to make a statement about what the conclusions of these debates should be. Rather, I wish only to draw attention to evidence that suggests that at some levels linguists and psychophysicists seem to be coming to similar conclusions about the workings of the mind, hopefully encouraging more discourse between these two important fields.

five years is the contrastive model of Tversky (1977) which takes similarity to be computed as a function of common and distinctive features of the objects that are being compared. In general, the more features two objects have in common, the more similar they are understood to be. But as Hahn & Chater (1997) have pointed out, the representations of natural objects cannot be fully specified by a list of features alone, but rather must also contain information about how such features are related to one another. It is therefore reasonable that object comparison would involve not just compiling lists of features to compare, but also consideration of relationships between corresponding features in the objects being compared

The view that such relationships are essential to computing similarity is taken up in Hahn *et al*. (2003) who develop a 'representational distortion' approach to similarity, arguing that similarity, rather than involving a comparison of relevant feature lists, is a function of transformational distance. Put simply, the simpler the transformational operations are that it takes to turn one object into another, the more similar the two objects are judged to be. I illustrate below with stimuli of the sort used by Hahn *et al*. in one of their experiments (based upon an experiment in Imai 1977) involving similarity judgments between sequences of black and white blobs that are transformationally related in various simplex and complex ways. In the table below, the pairs of representations are related either via one simple transformational operation, or a combination of two or three transformations.

No. of Trans.	Type	Stimuli: Item One	Stimuli: Item Two
1	Reversal	●○○○○○●●○●●●	○●●●●●○○●○○○
		●●●○●○●○●○●●	○○○●○●○●○●○○
1	Mirror	○●●●○○○●○●●○	○●●○●○○○●●●○
		○○○●○○○●●○●●	●●○●●○○○●○○○
1	Phasic	●●○○●●○●○●○●	●○○●●○●○●○●○
		●●●●○○○○●●●●	○●●●●○○○○●●●
2	Reversal + Mirror	●○○○●○●○●○○●	○●●○●○●○●●●○
		●○○○●●○○●○●●	○○●○●●○○●●●○
2	Reversal + Phasic	●●○●●○●●○●●○	○●○○●○○●○○●○
		○○○●○○○●○○○●	●○●●●○●●●○●●
3	Reversal + Phasic + Mirror	●○●○○●●●●○●●	○●○○○○●●○●○○
		○●●○○○●○●●○●	●○○●○●●●○○●○

Table 1: Sample stimuli used in Experiment 1 of Hahn *et al.* (2003)

Hahn *et al*. tested such pairs of stimuli (which also included deletion and insertion as basic transformational operations in addition to reversal, mirror, and phasic operations) with psychology students and found an almost linear relationship between the number of transformational required to relate two items and how similar they were judged to be. On a scale of one to seven, subjects assigned pairs of items related by a single transformation an approximate score of 5 while those related by two transformations were assigned an approximate

score of 4. Those related by three transformations received a score of approximately 3.5 while controls (which could not be related by three or fewer transformations) were around 2.5. Thus, the major finding of Hahn *et al*. is that a strong correlation exists between similarity judgments and the number of transformations relating a pair of stimuli: The more transformations involved in relating one object to its counterpart, the less similar the stimuli are judged to be.

Hahn *et al*.'s experiments provide strong evidence that principles of symmetry are central to the cognitive computation of similarity.[4] But how does Hahn *et al*.'s approach to these results compare with a feature-based account? The question is not so easy to answer since one of the well-known problems (or strengths, depending upon one's perspective) of feature-based approaches is that exactly what constitutes a relevant feature is only defined for particular contexts. In the present case, however, Hahn *et al*. show that a feature-based approach can only fare as well as their own if, trivially, the features that are assumed to be relevant are the shape properties of the items that remain unchanged under the applied transformations; in other words, only if the symmetries of the objects are counted as features. This strongly suggests that how similar two shapes are judged to be is a function of the number of symmetries preserved by the transformational relations that relate them. The conclusion is that symmetry is a central organizing principle in similarity judgments.[5]

3.2. Symmetry and Apparent Motion

The second domain of inquiry in visual computation to be discussed here has to do with what is known in the literature as apparent motion, a phenomenon first described by Wertheimer, Kroffka and other Gestalt psychologists (e.g. Koffka 1935, Wertheimer 1912). In modern work, the name most associated with the psychophysics of apparent motion is Roger Shepard, whose work on the topic is both highly admired and highly controversial. In apparent motion experiments, two images are flashed on a screen in different positions and in close temporal sequence. Typically, subjects report experiencing the object moving across the screen from one position to the other, even though no motion actually took place. This alone points to the active nature of the perceptual process. As Shepard notes, "Quite apart from questions about the particular type of movement experienced, the fact that any connecting movement is experienced is presumably the manifestation of an internalized principle of object conservation" (Shepard 2001: 582). That is, Shepard views apparent motion as resulting from the implicit assumption that two objects viewed in close temporal and spatial proximity are assumed to be the same object, even if they appear in different locations and with slightly different shapes. Apparent motion is a solution to the spatial disparity between the two objects while transformational operations are employed to relate any shape or orientation disparities. In other words, Shepard is arguing that

---

4 See also Kemp *et al*. (2005) for a 'generative' theory of similarity that encompasses and expands the transformational approach. Briefly, Kemp *et al*. argue that two objects are judged to be similar to the extent that they are judged to be formed by the same process.

5 For a strong view of the role transformations may play in perceptual processes, see Foster's (2001) commentary on Shepard (1994/2001).

internal principles of symmetry and transformation play an important role in visual perception.

Perhaps even more interesting than the fact that apparent motion exists and employs principles of symmetry, however, are the details of exactly what sorts of transformations the mind prefers to impose in apparent motion experiences. In simple cases where the only difference between the two objects is its position in the visual field, the facts are not very interesting: Subjects experience direct, rectilinear translational motion between the objects' positions (that is, movement in a straight line). More interesting cases involve differences not just in the position of the object, but also orientation. Take, for instance, the objects in Figure 3. In order to see these objects as being the same object in an apparent motion experiments, subjects must not only relate them spatially, but must also perceive a clockwise rotation of 90 degrees (illustration taken from Todorovič 2001).

Figure 2: In apparent motion experiments, subjects experience movement between one position/orientation and another, even though movement has not really taken place

Obviously the two transformations that relate the object on the left to the one on the right are translation (movement) and rotation (orientation), but one can imagine a number of different combinations of these transformations that subjects might experience. One possibility is a sequence of translations followed by a single 90 degree rotation. This is option (a) in Figure 4 below. Another is a sequence of combinations of translation/rotation transformations in which the object rotates a little with each movement along a straight path from A to B, shown in option (b). However, typically subjects experience neither of these. Rather, they report experiencing the motion in option (c).

(a)

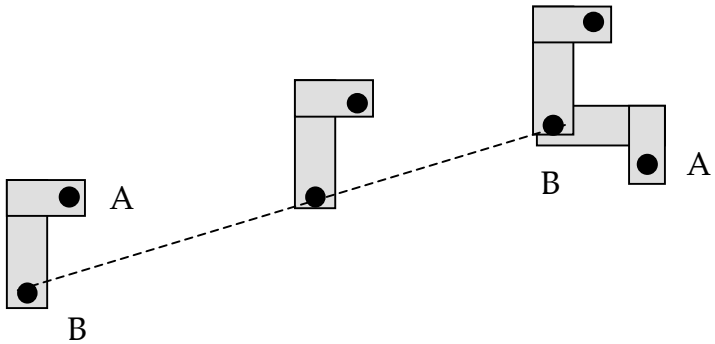

(b)

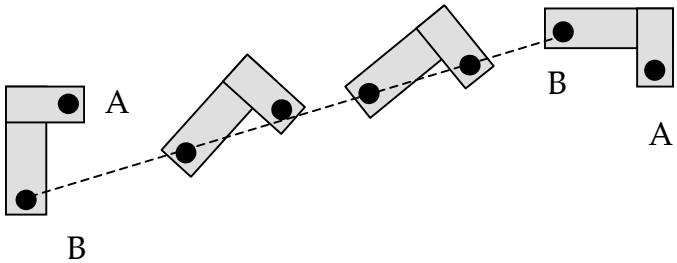

(c)

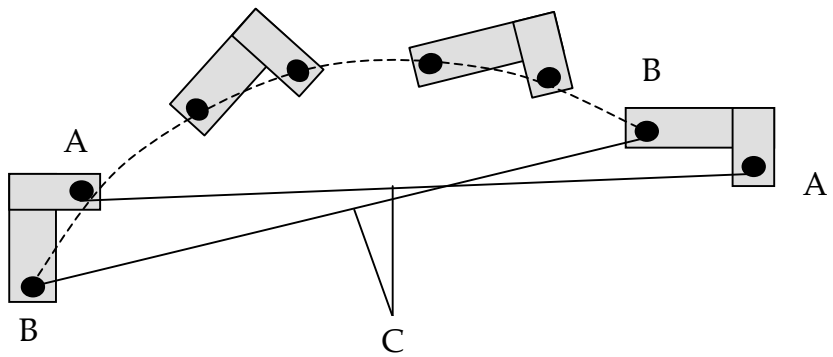

Figure 3: Possible transformation sequences relating object in first to object in second position (in an apparent motion experiment, (c) is most commonly experienced motion)

The motion in (4c) can be described in two ways: One is a sequence of translation/rotation combinations (like the option in (4b)) that just happens to follow an arced path. Another, simpler way is as a transformation of *pure rotation* that involves no translation at all. This rotation takes place around a point C which can be found by the intersection of perpendicular bisectors of the distance between A and B.[6] Discussing this finding, Shepard (1984) argues that the preference is based on economy concerns; he argues that in visual perception the cognitive faculties have a strong preference for employing a unique, simplex transformation (those that involve a single transformational operation) over complex transformations (those that involve a combination of such operations). This conclusion is particularly striking in the present case since the preference for a simplex transformation makes the overall computation more complicated (from one perspective, at least) since it requires computing point C.

I hasten to note that Shepard's conclusions are not uncontroversial and are in fact hotly debated. [7] However, his conclusion from apparent motion studies builds on the conclusion from Hahn *et al.*'s similarity experiments: Not only are principles of symmetry relevant for visual computation, but the computational system shows a preference for the simplest possible transformational sequences.[8] In the area of visual perception at least, this evidence suggests that cognitive judgments about similarity and object conservation are based on the simplest application of symmetry transformations.

3.3. Symmetry and Shape History

The third conclusion I would like to discuss here is concerned with how judgments are made about the past states of an object, or its shape history. That

[6] Note that, if one described this motion in the first manner mentioned, one would still need to compute point C, but there would be no explanation for why the translation/rotation is experienced along an arc rather than in a straight line. For this reason, pure rotation must be the preferred interpretation of the facts.

[7] See volume 24 of *Behavioral and Brain Sciences* for a reprinting of Shepard's important 1994 paper (= Shepard 2001), and a variety of papers and comments reacting to it as well as the interesting research it has encouraged.

[8] See Todorovič (2001) for arguments that determining what is meant by 'simplicity' in these cases is not in itself a simple matter.

is, given the present shape of an object, what do we infer about its previous states and how do we infer it? In a way, this topic is closely related to the topic of similarity judgments except in the present case one is extrapolating from a presently observed object what its similar, past state must have been. If we are right in following Hahn *et al*. (2003) in supposing that principles of symmetry play a central role in similarity judgments, we should not be surprised to find them operating in the domain of object shape history as well. Indeed, Leyton (1992) asserts that this is the case, arguing that when we observe an object with a low level of symmetry, we automatically infer that in the past the object must have had a higher level of symmetry. Leyton suggests that one common experience illustrating these principles is the observation of a dented can. Observing the shape asymmetries of the can, we infer that at some point in the past the can did not have a dent and that its shape was symmetric in the relevant ways. Moreover, we commonly assume that the can obtained its dent from a single causal event that introduced a symmetry-breaking transformation – a fall from a shelf, for instance.

Leyton (1992) reports on experiments he conducted in the 1980s which suggest the psychological reality of this symmetry-inferring process. Subjects in one experiment were provided only with a rotated parallelogram. Asked to construct a previous state for parallelogram, subjects typically constructed a non-rotated parallelogram. When subjects were then asked to construction a previous state for this shape, a rectangle was typically produced. Finally, asked to construct a previous state for the rectangle, subjects constructed a square. The entire sequence is seen in (6). Note that what subjects have done here is apply simplex transformations, one at a time, in such a way as to gradually increase the symmetry group of the object and restore maximal symmetry to the shape, taking it from a rotated parallelogram, which is symmetric under very few transformations, to a square which is symmetric under many more transformations in a step-by-step fashion.[9]

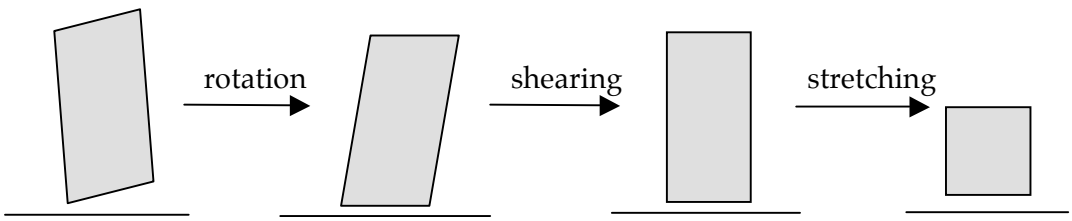

Figure 4: Leyton's subjects inferred the rotated parallelogram's shape history by applying simplex transformations one at a time to create progressively more symmetric shapes

The general conclusion Leyton draws from these results is that, psychologically, present asymmetries in the shape of an object are understood as having resulted from past symmetries.[10] Part of visual computation involves computing a

[9] As a reviewer points out, the rotated and non-rotated parallelograms in Figure 4 only have different symmetry groups with respect to the horizontal line below them. This line was included in the original study.

[10] Leyton applies his conclusions to domains outside the realm of visual computation alone, even showing how these principles apply to Transformational Grammar.

sequence of past states for an object, each with a symmetry group larger than the one it precedes. Furthermore, as in the case of apparent motion, these states are related by simplex transformations where possible: Rather than seeing a rotated parallelogram as resulting from a single past state (the square) that underwent a single complex transformation composed of stretching, shearing and rotation in Figure 6, subjects instead infer three previous states for the object, each related to its predecessor by a single simplex transformation. This is further confirmation of the conclusions drawn in the previous two sections, that human cognition makes use of principles of symmetry with special preference given to simplex transformations over complex ones.

Finally, another important (though on the surface, trivial) conclusion here is that while present asymmetries are taken to result from past symmetries, present symmetries are assumed to always have existed and no differing past state is taken to have existed for them. Again, the dented can serves as an example: If someone dents the can and then flawlessly repairs it, an observer will infer that the can has always existed in this state. Previous states in which an object has a smaller symmetry group (is more asymmetric) than its present state are never inferred.

3.4. Conclusions

In the previous three sections I have introduced three ways in which principles of symmetry under transformation have been found to be important for visual psychology and computation. I summarize them here. First, these principles explain some judgments of similarity which are difficult to account for using feature-based models. Second, when there is a choice in relating images, single-step transformations are preferred over multi-step transformations, even if this might result in a more complicated computational load as in the case of apparent motion. Third, part of computing the present state of an object involves computing its past states which are always computed as being more symmetric than the present ones: Past symmetry is inferred from present asymmetry. Present symmetry, on the other hand, is assumed to be present in all past states of the object.

4. The Perception of Linguistic Structure

Given that the three conclusions described above are relevant for the computation and perception of the structure of visual objects, we might ask whether similar principles and conclusions are also relevant for the psychological perception of linguistic structure. In this section, I argue that this is the case and that, in fact, linguists already implicitly recognize the centrality of these principles. As I will show, many of the basic ideas of our theories about syntactic structure have been based on principles of symmetry, though things have seldom been discussed in the present terms. Finally, I will argue that being more explicit about the role being played by symmetry principles suggests a solution to an enduring problem for syntactic theories, namely successive-cyclic movement.

4.1. Symmetry in Linguistic Theory

To apply the principles we have been discussing, we must start with some basic questions and answers. Unlike in visual perception where the variety of forms that must be processed is wide-ranging and nebulous, linguistic structure is well-defined and variation is highly constrained. Therefore, we do not expect the full range of principles active in visual perception to be active or required in linguistic perception. Rather, we might at most expect to find a subset of those principles whose functions are compatible with the requirements of linguistic structure.

To begin, we must ask what the basic shape of linguistic structure is and what the relevant transformations that preserve/break its symmetry are. Though not typically stated in these terms, mainstream generative syntactic theory can be thought of as a central state whose symmetry is broken under two general simplex transformations. I take this central symmetric state to be the basic system of a predicate and its arguments commonly refer to as argument structure.[11] The two transformations that, when applied, break the symmetry of this state, are projection and movement, which I take to be linguistic terms for the simplex transformations dilation and translation (more on this below). The application of these transformations to argument structure introduces asymmetries that obscure the core symmetries of the latter. [12] Projection, for instance, introduces functional structure that may be irrelevant for the interpretation of argument relations while movement often makes selection relationships obscure for the listener. Given a full-fledged linguistic structure to which projection and movement have applied, then, the task for linguistic perception is to reconstruct the past symmetric state of argument structure from the present less symmetric state created by projection

---

[11] By 'argument structure' I refer to a verb and its internal arguments, taking subjects to be an argument of a functional level of structure projected above the lexical VP level. I further assume that if a verb has more than one internal argument, they are each related to the verb via symmetric complement relations. Given standard tree representations, this requires that ternary branching be permitted at least at the lexical level. I put aside the wider implications of this assumption since they are not central to the theme of this paper.

[12] Another relevant debate in linguistic theory can also be stated in these terms, namely the debate about whether headedness is a property of syntax or not; in present terms, whether binary branching linguistic structure is symmetric under the reflection transformation. Much work beginning with Kayne (1994) has argued that it is not, imposing uniform right-branching structures on all languages. Others have maintained that it is symmetric under reflection and that whether a language has right or left branching structure is a matter of parametric variation. A synthesis of the two approaches is presented by Moro (2000) who argues that anti-symmetry is imposed by the interface with the phonological component. In the same spirit of the present work, Moro refers to reflectively symmetry binary structures as 'points of symmetry' that must be broken, made anti-symmetric, in order to be linearized (and thus pronounced). He argues that movement occurs as a function of spell-out to break these points of symmetry. Moro's view conflicts with the present account in which movement is taken to be motivated by morphological (feature checking) considerations. Whether the two could be compatible is an interesting question that would seem to hinge on whether or not Moro's approach to movement can derive a uniform paths approach to successive-cyclicity, the chief fact that I argue the present account enlightens.

and movement.[13]

Though projection and movement are symmetry breaking operations for argument structure, the resulting structure still maintains a large amount of symmetry, in particular its binary-branching and hierarchical structure. Since these two structural properties are not disrupted by projection or movement (only one level of functional structure at a time can be projected at any given step in a derivation, and movement does not result in structure being destroyed), they are symmetries of the system. Furthermore, it is the maintenance of these symmetries that allow a full-fledged linguistic structure to be related to its core argument structure as the latter's highly symmetric state is inferred from the former's less symmetric state, much as the two shapes in Figure 1 above or in Leyton's dented can thought experiment. The mind infers the past symmetric state of argument structure from the present asymmetric state of a full functional structure.

The idea that argument structure is the core symmetry of linguistic structure was explicitly encoded in Chomsky's (1981) Projection Principle which (among other things) imposes that the requirements of argument structure be projected into the syntax and represented at every level of syntactic structure. To state things in the present terms, the Projection Principle ensures that the past-state, central symmetry of syntax (argument structure) will always be recoverable from the output (present state) of a syntactic derivation. In other words, the Projection Principle simply formalizes the task for linguistic perception as it is understood here: Recovery of past symmetry from present asymmetry.

The particular past states of the syntactic object are also encoded in our present theories within the formalizations of the two basic transformations employed by syntax: movement and projection. For movement, this is encoded explicitly in the system of copies or traces that relate the surface position of an object to its original position (the latter dictated by argument structure). For projection, things are less explicit, but formalizations can be found in the idea of the extended projection (Grimshaw 1991) or the morphological word (Brody 2000), both encoding the idea that functional structure is a projection of its lexical base. Even in systems in which functional heads are not formally related to their lexical counterparts (such as those employing Merge as a basic structure-building operation), the history of projection is encoded in the hierarchical order of functional and lexical projection in the syntactic tree: those lower in the tree and thus closer to the lexical projection are projected earlier than those higher in the tree. The most explicit encoding in the Merge systems is in the so-called 'cartographic' approach which takes the number and ordering of functional heads in the hierarchical structure to be universal (e.g. Cinque 2002, Belletti 2004).

Regardless of the precise formalization, what is important here is that modern theories of syntax have formal ways of encoding the reconstruction of the past symmetric state (argument structure) from the present asymmetries

[13] By the 'perception' of linguistic structure, I am referring to the process by which LF representations are interpreted by the conceptual-intensional system. As I attempt to show below, my intention is not to supplant the LF representations of any particular theory, but rather to illustrate that the technology employed in such representations of the Chomskyan tradition encodes the fact that the reconstruction of past symmetries I discuss here is a primary function of interpretation at LF.

imposed by the syntactic transformations of movement and projection. It is thus an insight of linguistic theory that the principles being discussed here are crucial for the human linguistic system.

Within this line of thinking, it is useful to consider exactly what kinds of transformations projection and movement represent and how, specifically, past states must be inferred from them. Projection, as it is commonly understood, iteratively expands the structural space of the core symmetry of argument structure. Functionally, this is identical to the geometric transformation of dilation, a transformation that relates an image to an identical image that is proportionally larger or smaller.[14] A good example is a system of concentric circles. Each circle is related to the next smallest or largest circle by dilation. A system of such circles is reduced to its smallest member by beginning with the outermost circle and iteratively applying dilation, removing one outer circle at a time until only the smallest, most central circle remains. Projection is the same. Given an object to which projection has applied, past states of the object must be inferred by removing these dilational expansions one-by-one. To illustrate using conventional tree structures, the leftmost tree in Figure 7 would be inferred to have the previous states to its right. Here H0 is understood as a lexical element while H1 and H2 are functional heads projected from it.

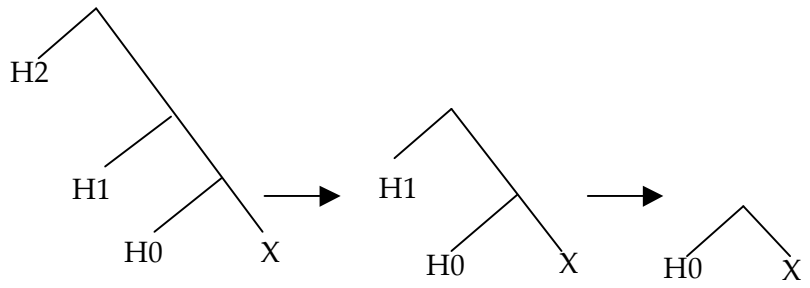

Figure 5: Removing functional projection to restore symmetries of argument structure

The movement transformation, on the other hand, involves displacing a lexical element. This is directly analogous to the geometric transformation of translation, discussed in section 2 above. Inferring a past state to which movement has applied therefore involves reconstructing a moved element to its original position. This is illustrated in Figure 6 and is in fact commonly referred to in the linguistic literature as 'reconstruction'.[15,16]

---

[14] The analogy of dilation was explicitly discussed by Cedric Boeckx in an earlier 2004 version of what became Boeckx (2008). The discussion of dilation is absent from the latter work, though the basic idea of projection as an iterative expansion of functional structure projected directly from the lexical item remains.

[15] Reconstruction may be implemented in many ways. A reviewer asks how the perceptual reconstruction discussed here compares to the understanding of reconstruction associated with the copy theory of movement. On the view developed here, however, the fact that a moved element leaves a copy of itself in its original position is simply an implementation of the fact that linguistic perception requires reconstructing moved elements to their positions of origin. A system based in trace theory would work just as well for my purposes.

[16] As a reviewer points out, these two transformations differ in the groups they belong to.

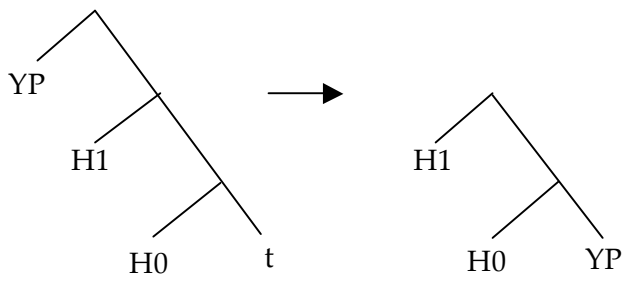

Figure 6: Reconstructing a moved element

Finally, it is also important to note that in most of the literature projection and movement do not apply randomly and without cause. Rather, they occur as the results of distinct causal relationships built into the theory. The concept of causation also has a central role in the discussion of shape history in Leyton (1992). Leyton postulates that present asymmetries are not only used to infer a past symmetric state, but are also assumed by the mind to be the result of an outside causal entity acting upon that state. Not only do we infer the undented can from the dented one, but we also infer that there must have been some outside agent that caused the dent. Interestingly, this principle of causality is also encoded in present syntactic theories in systems of feature checking and (some version of) Full Interpretation (Chomsky 1986). Features are outside causal entities that induce transformations.[17] Functional features induce the projection transformation in order to create representational space for their expression. Movement-related features (EPP features, 'strong' features, criterial features, etc., depending upon the framework), on the other hand, induce the movement transformation in order to satisfy the morphological and/or syntactic requirements of a language (in many systems, by in some sense 'checking' the features that projection has created space for). These two kinds of features (or at least two kinds of ways of talking about them) encode the fact that projection and movement are distinct types of transformations with distinct purposes in the computational system (creating representational space and fulfilling morphological requirements, respectively). Since the two have distinct causes, the present view suggests that collapsing projection and movement into one general transformation as some have proposed (Starke 2001, Boeckx 2008, among others) might be misguided (though certainly not ruled out in principle).[18]

Since it only affects the property of size, images related by dilation do not differ in their symmetry groups while images related by translation often do. Therefore, removing layers of projection as in Figure 5 does not by itself increase the symmetry group of the structure (though I will suggest in section 4.3 that, indirectly, it does).

17 A reviewer points out that features are methodological tools of a theory rather than empirically observable facts and so are unlike real-world causes of real-world asymmetries. However, the commonality pointed to here is that in both kinds of perception there is an innate assumption that present asymmetries are brought about by outside causal entities. Thus the reason we assume someone dented the can is the same reason we assume features. Whether or not features are the best methodological tool to encode this is a separate question I will not deal with.

18 Again, I refer the reader to Moro (2000) for a different view of movement-as-symmetry-breaking that is not based upon a feature-based theory of movement. Note, however, that

Crucially, however, while I am claiming projection and movement should be considered distinct and independent operations in the imposition of asymmetries on linguistic structure, this says nothing about the process by which those asymmetries are removed from the system when linguistic structure is perceived. That is, we must differentiate between the production side of the linguistic computational system and its perception side.[19] The distinction is motivated by the primary requirements of these two sub-systems: While the production sub-system is chiefly concerned with satisfying well-formedness requirements (checking features, satisfying Full Interpretation, etc.), the perception sub-system is chiefly concerned with reconstructing the past symmetric states of system in the most economic way possible. Thus, while we expect both sub-systems to make use of the same basic principles in producing and perceiving linguistic structure, they may apply them in subtly different ways, specifically when the perception system is able to remove asymmetries in a way that is more economic than the production system was able to impose them. I return to this point below.

To recap the present section, projection (dilation) and movement (translation) are simplex transformations that apply to the symmetric state of argument structure (typically iteratively). The task for linguistic perception is to reconstruction the past symmetric states of the derivation given the asymmetries that projection and movement have imposed upon it. Note that, given the discussion in section 3 above, we expect that the particular past states inferred for a present state should be related via single simplex transformations where possible. In the next section, I attempt to show that this prediction yields important results for our understanding of successive-cyclic movement.

4.2. Successive-Cyclic Movement and Its Discontents

Successive-cyclic movement is the simple idea that movement in syntax that is sufficiently long-distance does not take place in one fell-swoop, but requires a successive series of shorter movements. Though originally proposed as a solution to a theoretical problem in Chomsky (1973), a variety of empirical evidence from every sub-discipline of linguistics has since converged on the idea that some version of S-C movement exists. To present just two examples of semantic evidence for S-C movement, consider the following data from Fox (2000), an instance of topicalization.

---

even in Moro's approach, movement and projection have distinct causes, with movement caused by linearization and projection caused by feature requirements.

[19] To reiterate fn 13, by 'perception,' I refer to a (sub-part of) interpretation that is a part of the linguistic computational system and not an active, online perceptual process. I stand agnostic about how these perceptive principles interact with models of language processing, though it is possible, as a reviewer suggests, that just as apparent motion is an effect brought about by the processing of certain visual stimuli, the sorts of effects I examine here could be taken as the results of processing certain linguistic stimuli, or of more general aspects of processing such as its top-down nature.

(2) [The papers that he_i wrote for Ms. $Brown_j$], every $student_i$ asked her_j to grade ___.

(2) is interesting since standard binding consideration prevent the moved NP from being interpreted in either its surface position or its original position. Variable binding requires that *he* be c-commanded by the quantifier phrase *every student* in order to yield the proper interpretation, while Condition C of the binding theory requires that the R-expression *Ms. Brown* not be c-commanded by the coreferential pronoun *her*. In (2) the only place the NP can be interpreted to satisfy both conditions is in an intermediate position somewhere between *every student* and *her*. Of course, in order to be interpreted in this position, the constituent must have moved through it in the course of its movement from its original position to its surface position, thus providing evidence for S-C.

Another widely-known piece of semantic evidence for S-C goes back to Barss (1986) and concerns anaphor interpretation. According to Condition A of the binding theory, anaphors must be bound by a local antecedent. In (3a), for instance, *himself* can only refer to Bill and not John. In (3b), however, the NP containing *himself* has undergone movement. Note that here both interpretations are possible.

(3) a. $John_i$ thinks that $Bill_j$ hates the picture of himself $_{*i/j}$.
b. [the picture of himself $_{i/j}$] that $John_i$ thinks that $Bill_j$ hates ___.

While the coindexation with *Bill* follows from the NP being interpreted in its base position, coindexation with *John* is only possible if the NP containing *himself* enters a local binding relation with *John* in the course of the derivation. In other words, the NP must pass through an intermediate position somewhere below *John*, but above *Bill* in (3b), another argument for S-C.

While the phenomenon of S-C movement is well-established, its exact nature continues to be a subject of great debate. At the core of the discussion are two questions, one empirical and one theoretical. The empirical question is, precisely what intermediate positions does movement target? As both Abels (2003) and Boeckx (2007) have discussed, there are two general options available. Either S-C movement targets particular positions (what Abels calls the *punctuated paths* possibility) or S-C movement targets every possible position between its first and final landing sites (Abels' *uniform paths* possibility). While a number of recent proposals have come down in favor of a uniform approach (Fox 2000, Richards 2002, and Bošković 2002, 2007, among others; see Boeckx 2007 for an overview), the majority of the work on S-C movement has assumed a punctuated path approach. These works assume that certain nodes are bounding nodes (or barriers or phase level categories, depending upon the specific system assumed) that are barriers to movement. In order for movement to be considered legitimate by the computation system, it must proceed through the specifier positions of the barrier nodes. A version of each of these systems is illustrated below using the sentence *What did Mary think John bought?* (4a) represents a (version of a) punctuated understanding of S-C wherein the XP moves through only intermediate [Spec,CP]'s on its way to is final landing site. (4b) represents the

same structure in a uniform path framework. Here the XP is adjoined to every projection between its original and final positions.

(4) a. [$_{CP}$ ***what did*** [$_{TP}$ ***Mary*** *<did>* [$_{VP}$ *<Mary>* ***think*** [$_{CP}$ *<what>* C [$_{TP}$ ***John*** T [$_{VP}$ *<John>* [$_{V'}$ ***bought*** *<what>*]]]]]]]

b. [$_{CP}$ ***what did*** [$_{TP}$ *<what>* [$_{TP}$ ***Mary*** *<did>* [$_{VP}$ *<what>* [$_{VP}$ *<Mary>* ***think*** [$_{CP}$ *<what>* C [$_{TP}$ *<what>* [$_{TP}$ ***John*** T [$_{VP}$ *<what>* [$_{VP}$ *<John>* [$_{V'}$ ***bought*** *<what>*]]]]]]]]]]]

The central empirical question, then, is whether (4a) or (4b) is a more accurate description of the derivation of such sentences. I suggest that in fact both are accurate, in a way to be made clear below.

Another concern in the decision between the uniform and punctuated paths approaches, however, is more conceptual in nature. Namely, if one adopts a punctuated path view, then why does S-C movement target some nodes and not others? While there have been many systems proposed for deriving a punctuated system of S-C movement, none of them has achieved the sort of natural implementation that the minimalist program seeks from a theory. No natural system of bounding nodes or barriers was ever achieved, and in the current phase system of Chomsky (2001, 2008) there does not seem to be a natural connection between available EPP positions and phase category, even if the latter are taken to be natural propositional chunks of structure (see Boeckx & Grohmann 2007 for an overall critique of the phase system). The punctuated approach to S-C movement, then, has proven extremely difficult to motivate and implement theoretically. The uniform approach, however, does not suffer from this problem since it takes S-C movement to simply be a property of movement itself: Movement must be local, targeting every projection between its first and last position. There is no need to single out particular projections as special with regard to movement.

Even under the uniform approach, however, a conceptual question arises: If movement really is so local, why is it so local? Bošković (2007) and Boeckx (2007) contend that movement essentially comes for free as a derivational option that the grammar allows in order to allow syntactic objects to check their uninterpretable features as efficiently as possible. As long as some element has an unchecked feature, it continues to move up the tree (merging with each project in Boeckx's case) until all of the heads required for it to check all of its features have been introduced. Then movement stops in that position. Going further, Boeckx (2008) suggests that movement appears to target every projection due to the fact that, though typically formalized separately, movement and projection are really one and the same operation. An element moves from its initial position to its final landing site not separately from the projection operations that articulate the sentence's structure, but along with this projection. As each level of projection is iteratively added to the structure, movement 'piggybacks' onto it, in this way percolating up through the clause until its final landing site is reached.

While these approaches certainly provide us with a natural way of thinking about S-C movement, it is not clear that they are the best at capturing the facts typically taken as clear evidence for S-C. Furthermore, the idea that movement

comes 'for free' with projection only seems natural to me if movement and projection are indeed fully collapsed to a single operation as Boeckx (2008) argued. Yet, as discussed above, projection and movement are unique transformations (dilation vs. translation) and, more importantly, have unique causal origins. Therefore, I do not believe collapsing them to a single operation is the best approach. Despite this, however, I do believe that a partial conflation is possible if we take seriously the idea that the derivation of a structure may employ principles in subtly different ways than the perception of the same structure as suggested above. I expand on this below.

4.3. Uniform Successive Cyclicity as a Perceptual Phenomena

Recall the conclusions drawn from the field of visual perception discussed above. There I suggested that cognitive judgments involve using the simplest possible principles of symmetry to reconstruction past states of a present object with perceived asymmetries. I have also suggested, however, that the simplex transformations that are inferred to relate a present state to a (more symmetric) past state may be different from the transformational operations that caused the relevant present state to come into being in the first place. In particular, though multiple complex transformations may change a highly symmetric state to a less symmetric one, the cognitive faculty will, given only the present state, infer the change to have resulted from single, simplex transformations when possible, as in cases of apparent motion perception. Thus there may be differences between the temporal construction of an asymmetric state on the one hand and the inference of its past symmetric state on the other: While both are constrained by general concerns of economy and simplicity, those concerns may be manifested in different ways depending upon the particular requirements of the system.

I would like to suggest that this difference is relevant for the derivation/construction and perception/inference of linguistic structure as well. In particular, I am suggesting that while in the derivation of linguistic structure asymmetries are introduced by two distinct transformations with distinct causal relations (projection and movement), in the perception of linguistic structure, the past symmetric state of this structure (argument structure) is inferred to have resulted from a single, simplex transformation (projection only).

To illustrate what I mean, consider a structure to which both movement and projection have applied. Reconstructing the past symmetric state of the structure must involve removing the asymmetries imposed by both of these transformations. Crucially, this needn't be accomplished by re-applying both movement and projection. Rather, the effect of re-applying projection alone is enough to undo the effects of both movement and projection. This is illustrated in Figure 9. The leftmost object has been derived via projection of functional structure from the lexical item H0 and movement of the complement of H0 (YP) to a position higher in the structure. In each state represented to its right, one layer of functional projection has been removed under an iterative application of the projection transformation. Note that as a side effect of this process the moved constituent YP also gets closer and closer to its original position in the argument structure. As each layer of functional structure is removed, YP becomes adjoined

to the next lowest functional projection until there are no more and it is reconstructed to its original position in the lexical item's argument structure.

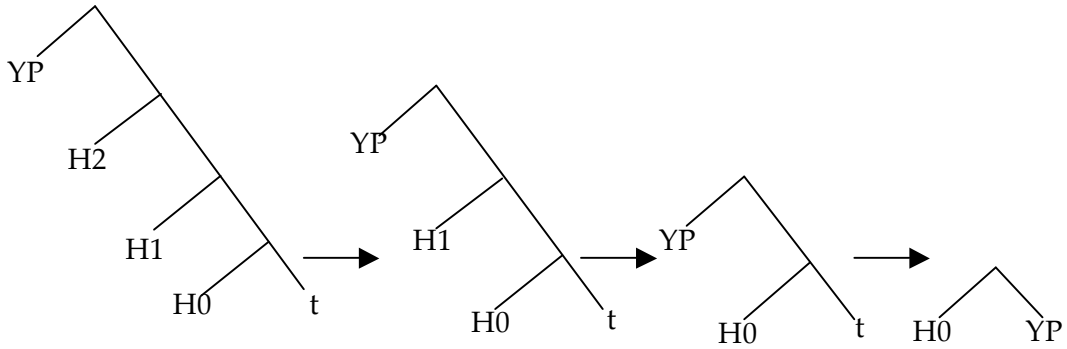

Figure 9: Removing effects of projection also removes effects of movement

I propose that this conflation of the effects of movement and projection in the perception of linguistic structure is responsible for the interpretative effects of successive-cyclic movement. As a side effect of inferring the past symmetries of argument structure via pure projection, rather than via the combination of projection and movement that resulted in the obfuscation of those symmetries, we interpret a moved XP as being adjoined to every functional head in the clausal architecture.

Note that adopting this idea has led us to a position very similar to the uniform paths hypothesis: The moved element adjoins to every head between its original and final positions. However, in the present understanding this adjunction takes place at the level of perception and not in the construction of the syntactic representations. Crucially, the conclusion that perception requires moved elements to reconstruct in this iterative fashion says nothing about the way that the movement asymmetries were imposed in the first place. That is, though movement is reconstructed to its original position via the projection transformation, we may still maintain that a unique movement transformation exists independently of projection and that this transformation is responsible for displacement. To put things in familiar syntactic terms, we may maintain that movement only occurs when it is triggered by an appropriate matching feature (say, a [wh] feature in an English question). Movement occurs in order to check this feature. With regard to syntactic derivation, this is all one needs to say. There is no need to posit intermediate movement positions between the original and final landing place of the moved constituent (but see below). Rather, movement can take place in one fell swoop. It is only in the perception of linguistic structure, when asymmetries are removed from the system, that the cyclic effects of movement are derived.

Of course, this understanding of successive-cyclic movement makes an important prediction about the sorts of effects movement should produce.[20] Since

[20] A reviewer inquires about the implications of the present approach for island effects. The implications are not obvious as island effects are a challenge for any theory of locality. However, the present approach at least suggests that one way to think of (some) island effects may be as a failure to infer past states. We might therefore ask what sorts of factors

I have claimed that the cyclic component of movement occurs in linguistic perception and not in production, the only effects of cyclic movement should be interpretative effects. That is, we shouldn't see any phonological or morphological effects of successive-cyclic movement (or even syntactic effects of a particular kind). While it is true that syntactico-semantic arguments are the most prominent in work that argues for the uniform paths approach to S-C movement, in fact, there are a wide variety of claims for morpho-phonological effects of S-C movement in the literature. Unfortunately I haven't the space to review them all here (see Boeckx 2007: chap. 2 for an overview); instead, I will simply note two general observations that would seem to be compatible with approach developed here. First, as Boeckx (2007) notes, the phonological and morphological evidence for successive cyclic movement is much weaker than the syntactic and semantic evidence, chiefly because it may be interpreted in a variety of ways. For instance, so-called wh-agreement effects seen in languages like Chamorro (Chung 1994, 1998), Kinande (Schneider–Zioga 2007), and Irish (McCloskey 2002) have been argued to constitute phonological evidence for intermediate wh-movement wherein intermediate verbs or complementizers agree with a wh-word undergoing S-C movement on its way to its final landing site. However, these effects could alternatively be analyzed as a series of agreement relations between features of verbs or complementizer and the wh word before movement of the latter takes place (see, e.g., Schneider–Zioga's 2006 analysis of the Kinande facts). In other words, these agreement effects do not necessarily provide evidence for successive-cyclic movement.

Second, I would like to point out that the vast majority of (non-semantic) evidence for successive-cyclic movement is really only evidence for movement through intermediate [Spec,CP] positions. Complementizer agreement facts like those mentioned for Irish and Kinande as well as subject-auxiliary inversion in embedded clauses in French (Kayne & Pollock 1978) and even subject alternations in Ewe (Collins 1993) are all effects that, if they are evidence for intermediate movement, can only be related to intermediate movement of the moved element through intermediate [Spec,CP] positions. The same is true for the most convincing evidence for successive-cyclic wh-movement, namely the wh-copying that occurs in languages like Afrikaans (du Plessis 1977, Felser 2004). In such languages, a wh-word fronted to the beginning of the clause is sometimes repeated in intermediate [Spec,CP] positions. An example appears below:

(5) **Waarvoor** dink julle **waarvoor** werk ons? *Afrikaans*
wherefore think you wherefore work we
'What do you think we are working for?' (du Plessis 1977: 725)

In this case, it is indeed difficult to think of an analysis that does not require intermediate movement through these positions in narrow syntax, that is, on the production rather than just the perceptual side of things. However, note that if

might prevent the reconstruction of moved elements via pure projection in the manner discussed here. Exploring this question fully is beyond the present paper, so I leave it for future work.

this is granted, it only reintroduces successive cyclic movement into the derivation in a limited and conceptually-justified way, namely as movement through [Spec,CP] positions. Importantly, of all the nodes in the syntactic architecture through which punctuated S-C movement might be required to pass through, [Spec,CP] is surely the least stipulative since it corresponds to a natural barrier of syntax, namely the topmost level of the extended projection (the clause level). If indeed this is true, as wh-copying phenomena seem to suggest, then we are forced to keep successive-cyclic movement, but only in a very limited and punctuated form that derives from a natural locality imposed by the size of the clause's extended projection. We still predict that movement should not have intermediate phonological or morphological effects that derive unambiguously from successive-cyclicity at any level other than the CP level.[21] Rather, only interpretative effects for these intermediate movements should be found.

Taken together, the present view has the surprising conclusion that in the debate between punctuated and uniform paths approach for S-C movement, both turn out to be correct, though in different domains. While the punctuated paths approach, restricted to the natural barrier of the CP level, characterizes the production of syntactic structure, the uniform paths approach characterizes its perception. As discussed above, this makes an important prediction: Morpho-phonological effects of S-C movement should be limited to evidence for movement through intermediate [Spec,CP] positions while interpretative effects of S-C movement should be unrestricted, giving evidence for movement through all intermediate positions. To the extent the data bear this out, the present approach is superior to the pure production-oriented views of S-C movement discussed above since those approaches must explain why some positions are more highly privileged than others: The punctuated approach must explain why elements move to some positions and not others, while the uniform approach must explain why morpho-phonological traces of S-C movement are limited to intermediate CP positions.

5. Conclusions

In this paper, I have argued that just as in visual computation, principles of symmetry and economy play a strong role in the computation of linguistic structure. It is crucial to note, however, that these two computational systems differ in an important respect: While the visual system is almost wholly concerned with perception alone (it is unclear what the 'product' of the visual system

[21] As reviewers point out, the conclusion that [Spec,CP] is a 'natural' barrier in this way reinstates the notion of phases, at least lending support the idea that C is a phase head. I would not go so far, but clearly I have at least reintroduced the significance of a notion like Chomsky's (1986) 'Complete Functional Complex,' the level of structure at which all functional roles are satisfied. My position is that such a notion has significance because (i) there is overt syntactic evidence that this position is relevant for intermediate movement, and (ii) it is a conceptually natural syntactic object. However, whether the notion has significance beyond this (e.g., for a general theory of phases or extraction effects such as islands) is a separate research question that I will not go into here.

might be), the linguistic system is both a perceptual and productive system, and I have argued here that though both sides of the system make use of the same sorts of principles of symmetry and economy, exactly how those principles apply to production and/or perception may subtly differ. It is therefore possible that certain phenomena for which it is difficult to find a natural implementation in the productive system might find a more natural account in the perceptive system (and possibly vice versa). I have argued that successive cyclic movement is such a phenomena: While movement through intermediate [Spec,CP] positions can be naturally implemented in the productive system, evidence for movement through other intermediate positions should be interpreted as a by-product of the perceptual system alone. In large part, this lines up with the available semantic and morpho-phonological evidence for intermediate movement.

Another purpose of this paper was to suggest that some of the principles responsible for visual perception have been exapted for the perception and production of linguistic structure. In particular, I have argued for a central role for the dilation and translation transformations in the building and perceiving of linguistic structure. Given the basic character of language as a system that combines lexical items to form larger structures, it is easy to see why these two transformations in particular would be useful as exaptions: Dilation expands the representational space of linguistic structure and translation allows lexical items to be rearranged within that structure. Whether or not other components of basic group theory that seem active in visual perception (e.g., rotation, reflection) might also be active in the perception and creation of linguistic structure is an open question that requires attention.[22] Other questions also arise, in particular why is it that human minds employ such principles at all? Shepard has controversially addressed the latter question, arguing that the human mind makes use of principles of kinematic geometry because these principles have been extracted from environmental experience and 'internalized' in the course of human evolution. However, many have pointed out that principles of pure kinematics are seldom observed in the natural world where motion and shape tend to be messy and highly asymmetric. It is therefore hard to see why such internalizations would be favored by natural selection; that is, how such principles could be seen as adaptations to properties of the natural environment. It may be more likely that explanations will be found in more purely mathematical models of common organizational rules, such as those found in group theory, as Foster (2001) suggests. If that is the case, then evolutionary explanations in terms of the internalization of external regularities seem even more difficult to maintain. Rather, it may suggest that the various modules of the mind (language, perception, etc.) share common emergent organizational principles that can be described in mathematical terms, inviting more extensive and explanatory psychological descriptions along mathematical lines.

---

[22] Reflection comes to mind, for instances, when considering the structures involved in Parallel Merge (Citko 2005) and object sharing (Hiraiwa & Bodomo 2008) as well as in the general binary nature of hierarchical structure (see brief discussion of Kayne 1994 and Moro 2000 above).

References

Abels, Klaus. 2003. Successive cyclicity, anti-locality and adposition stranding. Storrs, CT: University of Connecticut dissertation.

Barss, Andrew. 1986. Chains and anaphoric dependence: On reconstruction and its implications. Cambridge, MA: MIT Dissertation.

Belletti, Adriana (ed.) 2004. *The Cartography of Syntactic Structures*, vol. 3: *Structures and Beyond* (Oxford Studies in Comparative Syntax). New York: Oxford.

Boeckx, Cedric & Kleanthes Grohmann. 2007. Putting phases in perspective. *Syntax* 10, 204–222.

Boeckx, Cedric. 2007. *Understanding Minimalist Syntax: Lessons from Locality in Long-Distance Dependencies*. Oxford: Blackwell.

Boeckx, Cedric. 2008. *Bare Syntax*. Oxford: Oxford University Press.

Bošković, Željko. 2002. A-movement and the EPP. *Syntax* 5, 167–218.

Bošković, Željko. 2007. On the locality and motivation of Move and Agree: An even more minimal theory. *Linguistic Inquiry* 38, 589–644.

Brody, Michael. 2000. Mirror theory: Syntactic representation in perfect syntax. *Linguistic Inquiry* 31, 29–56.

Chomsky, Noam. 1973. Conditions on transformations. In Stephen R. Anderson & Paul Kiparsky (eds.), *A Festschrift for Morris Halle*, 232–286. New York: Holt, Rinehart and Winston.

Chomsky, Noam. 1981. *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris.

Chomsky, Noam. 1986. *Knowledge of Language: Its Nature, Origin, and Use* (Convergence). New York: Praeger.

Chomsky, Noam. 2001. Derivation by phase. In Michael Kenstowicz (ed.), *Ken Hale: A Life in Language* (Current Studies in Linguistics 36), 1–52. Cambridge, MA: MIT Press.

Chomsky, Noam. 2005. Three factors in language design. *Linguistic Inquiry* 36, 1–22.

Chomsky, Noam. 2008. On phases. In Robert Freidin, Carlos P. Otero & Maria Luisa Zubizarreta (eds.), *Foundational Issues in Linguistic Theory: Essays in Honor of Jean-Roger Vergnaud* (Current Studies in Linguistics 47), 133-166. Cambridge, MA: MIT Press.

Chung, Sandra. 1994. Wh-agreement and 'referentiality' in Chamorro. *Linguistic Inquiry* 25, 1–44.

Chung, Sandra. 1998. *The Design of Agreement: Evidence from Chamorro.* Chicago: University of Chicago Press.

Cinque, Guglielmo (ed.). 2002. *The Cartography of Syntactic Structures,* vol. 1: *Functional Structure in DP and IP* (Oxford Studies in Comparative Syntax). New York: Oxford University Press.

Citko, Barbara. 2005. On the nature of merge: external merge, internal merge, and parallel merge. *Linguistic Inquiry* 36: 475–497.

Collins, Chris. 1993. Topics in Ewe syntax. Cambridge, MA: MIT Dissertation.

du Plessis, Hans. 1977. Wh-movement in Afrikaans. *Linguistic Inquiry* 8, 723–726.

Epstein, Samuel David & T. Daniel Seely (eds.). 2002. *Derivation and Explanation in*

the Minimalist Program (Generative Syntax 6). Malden, MA: Blackwell.
Felser, Claudia. 2004. Wh-copying, phases, and successive cyclicity. *Lingua* 114, 543–574.
Foster, David. 2001. Natural groups of transformations underlying apparent motion and perceived object shape and color. *Behavioral and Brain Sciences* 24, 665–668.
Fox. Danny. 2000. *Economy and Semantic Interpretation* (Linguistic Inquiry Monograph 35). Cambridge, MA: MIT Press.
Grimshaw, Jane. 1991. Extended projection. Ms., Brandeis University, Waltham, MA.
Hahn, Ulrike & Nick Chater. 1997. Concepts and similarity. In Koen Lamberts & David Shanks (eds.), *Knowledge, Concepts and Categories*. Hove: Psychology Press, 43–92.
Hahn, Ulrike, Nick Chater & Lucky Richardson. 2003. Similarity as transformation. *Cognition* 87, 1–32.
Hiraiwa, Ken & Adams Bodomo. 2008. Object sharing as symmetric sharing: Evidence from Dagaare. *WCCFL* 26, 243–251.
Imai, S. 1977. Pattern similarity and cognitive transformations. *Acta Psychologica* 41, 433–447.
Kayne, Richard S. & Jean-Yves Pollock. 1978. Stylistic inversion, successive cyclicity, and Move NP in French. *Linguistic Inquiry* 9, 595–621.
Kayne, Richard S. 1994. *The Antisymmetry of Syntax* (Linguistic Inquiry Monograph 25). Cambridge, MA: MIT Press.
Kemp, Charles, Aaron Bernstein & Joshua B. Tenenbaum. 2005. A generative theory of similarity. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*, 1132–1137.
Koffka. Kurt. 1935. *Principles of Gestalt Psychology.* New York: Harcourt, Brace & World.
Leyton, Michael. 1992. *Symmetry, Causality, Mind.* Cambridge, MA: MIT Press.
Marr, David. 1982. *Vision*: *A Computational Investigation into the Human Representation and Processing of Visual Information.* San Francisco, CA: W.H. Freeman.
McCloskey, James. 2002. Resumption, successive cyclicity, and the locality of operations. In Epstein & Seely (eds.), 184–218.
Moro, Andrea. 2000. *Dynamic Antisymmetry* (Linguistic Inquiry Monograph 38). Cambridge, MA: MIT Press.
Richards, Norvin. 2002. Very local A-bar movement in a root-first derivation. In Epstein & Seely (eds.), 227–248.
Schneider–Zioga, Patricia. 2006. Bounded unbounded movement. Paper presented at the *29th Annual Colloquium of Generative Linguistics in the Old World (GLOW 29)*, Barcelona. [6–8 April, 2006]
Schneider–Zioga, Patricia. 2007. Anti-agreement, anti-locality and minimality: The syntax of dislocated subjects. *Natural Language and Linguistic Theory* 25, 403–446.
Shepard, Roger. 1984. Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review* 91, 417–47.

Shepard, Roger. 1994. Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin and Review* 1, 2–28.

Shepard, Roger. 2001. Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences* 24, 581–601.

Starke, Michal. 2001. Move dissolves into merge: A theory of locality. Geneva: Université de Genève dissertation.

Todorovič, Dejan. 2001. Is kinematic geometry an internalized regularity? *Behavioral and Brain Sciences*, 24, 641–651.

Tversky, Amos. 1977. Features of similarity. *Psychological Review* 84, 327–52.

Wertheimer. Max. 1912. Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie* 61, 161–265. [Translated in part in Shipley, Thorne (ed.). 1961. *Classics in Psychology*. New York: Philosophical Library.]

Brent Henderson
University of Florida
Linguistics Program
4131 Turlington Hall
PO Box 115454
Gainesville, FL 32611
USA
bhendrsn@ufl.edu

Minimalist Meaning, Internalist Interpretation

Paul M. Pietroski

This article offers a conception of semantics, and of what makes the human language faculty distinctive, based on five theses: Meanings are instructions to build concepts; concatenation calls for conjunction of monadic concepts; grammatical relations invoke certain thematic relations and a kind of existential closure; lexicalization is a partly creative process of abstraction; and meanings are internalistic properties of expressions. Each of these claims is defended elsewhere. The aim here is to connect them explicitly, and compare the result with alternatives, in the hope of providing a plausible conception of natural language meaning that coheres with Chomsky's minimalist program.

Keywords: concepts, internalism, lexicalization, meanings, semantics

1. Proposal: Meanings as Conjunctive *Begriffsplans*

This paper combines five theses, each discussed in more detail elsewhere (see Pietroski 2002 *et seq.*).

(1) Meanings are instructions to build concepts.
(2) Concatenation calls for conjunction of monadic concepts.
(3) Grammatical relations invoke certain thematic relations and a restricted operation of existential closure (∃-closure).
(4) Lexicalization is a partly creative process of abstraction.
(5) Meanings are internalistic properties of expressions.

These claims, explained below, are logically independent. But they cohere, inviting comparison with alternative quintets. The net result is a version of semantic internalism: Open-class lexical items are instructions to fetch monadic concepts that may have been abstracted in the course of acquisition; and the meaning of a phrase is an instruction to build a conjunctive monadic concept from fetchable elements, given a few relational/thematic concepts and an operation of existential closure. The position identified owes much to Chomsky's (1977, 1995a, 2000a) discussion of meaning and the general program of biolinguistics.

For helpful comments and discussion, my thanks to Cedric Boeckx, Susan Dwyer, Norbert Hornstein, Terje Lohndal, Dennis Ott, and two anonymous referees.

biolinguistics
© The Author (2008)
ISSN 1450–3417
Biolinguistics 2.4: 317–341, 2008
http://www.biolinguistics.eu

1.1. Meanings are Instructions to Build Concepts

For these purposes, I take it as given that humans have a language faculty that is in some respects distinctively human, and that the languages human children naturally acquire can be described as biologically implemented *procedures* that generate expressions. Chomsky (1986) calls these procedures, which correspond to stable states of the faculty, *I-languages*. By contrast, E-languages are sets of expressions; even for languages with endlessly many expressions, a single E-language might be determined by two or more generative procedures. The 'I'/'E' distinction connotes the contrast between intensions (procedures, algorithms) and the extensions we characterize by appeal to procedures, as when we specify sets without listing their elements; cf. Church (1941).

Using this terminology, we can describe the Human Faculty of Language (HFL) as a biologically implemented capacity to acquire and use one or more I-language*s* that associate phonological instructions to articulatory/perceptual systems with semantic instructions to conceptual/intentional systems, by means of a constrained syntax; cf. Chomsky (1995b, 2000b). Abstracting from phonology, expressions of an I-language *are* semantic instructions.[2] These expressions have various grammatical properties of interest to syntacticians. But whatever additional properties they have, expressions of an I-language (henceforth, *i*-expressions) are presumably instructions to access certain mental representations and generate others, in accord with certain principles of lexicalization and composition. If only for simplicity, let's say that the representations accessed and generated are concepts — or more specifically, '*i*-concepts'. I assume that concepts are *composable* mental representations; see e.g. Fodor (1987, 2003, 2008). But this is compatible with many proposals about how *i*-concepts are related to the representations that children lexicalize and the full range of representations available to human thinkers.

If some of the representations that children lexicalize are not conceptual — think of mental images, maps, and prototypes — then lexicalization must somehow associate these pre-lexical representations with concepts. But following Fodor and others, I suspect that children typically lexicalize concepts, many of which are not uniquely human. This allows for the hypothesis that lexicalization is fundamentally a matter of concept labeling, and that *i*-concepts just *are* lexicalized concepts; cf. Bloom (2000). Alternatively, one can hypothesize that lexicalized concepts are linked to formally distinct *i*-concepts that *are abstracted* in the course of lexicalization, and that *i*-concepts form a special subset of the concepts available to humans; see section 1.4 below. But whatever the details, the broad idea is one that many theorists should be able to accept: Lexical items are instructions to *fetch* concepts that meet certain conditions, while phrases are instructions to *combine* fetchable concepts in certain ways.

From this perspective, *i*-expressions are concept-construction-instructions (*Begriffsplans*) that reflect principles governing combination of *i*-expressions and

[2] Likewise, abstracting from semantics, expressions of an I-language are phonological instructions. The point is not to privilege semantics. And even if instructions to articulatory/perceptible systems appeared later, in terms of HFL's evolutionary history, ontogeny may not recapitulate phylogeny in this respect.

'interfaces' between HFL and other aspects of human cognition. Correlatively, *i*-concepts are concepts that can interface with HFL: *i*-concepts are results of executing instructions (*i*-expressions) that are generated by I-languages; and for a given context, an *i*-concept can be expressed with an *i*-expression whose execution would create the *i*-concept.

Thesis (1) is to be understood, accordingly, as a claim about *i*-expressions.

(1) Meanings are instructions to build concepts.

If *i*-expressions pair phonological instructions with semantic instructions, then *i*-expression meanings can be described as instructions (generable via HFL) to fetch and/or combine *i*-concepts in certain ways. One can stipulate that meanings are not semantic instructions in this sense. But then it is hardly obvious that there are theories of meaning for I-languages, much less that such theories perspicuously describe the natural phenomenon of human linguistic understanding. Put another way, one can stipulate that *i*-expressions pair phonological instructions with *i*-meanings and hypothesize that theories of the natural phenomenon are theories of *i*-meanings. On this view, endorsed here, understanding an *i*-expression (perceiving its meaning) is a matter of recognizing that expression *as* a certain concept–construction–instruction. This proposal may be wrong. But it is less tendentious than it might initially appear. So let me distinguish (1) from some alternatives in the vicinity.

It is an ancient idea that lexical items *label* concepts, and that if a lexical item λ labels a concept C for a speaker S, then S can use λ to talk about whatever S thinks about by using C. I prefer to say that lexical items *are instructions to fetch* concepts, if only because (i) a polysemous item like *book* or *set* may be lexically linked to more than one fetchable concept, and (ii) as we'll see in section 1.4, lexical items may fetch concepts that are formally distinct from the concepts initially lexicalized. Though given enough caveats, one can say that lexical *i*-expressions (as used in contexts) indicate the concepts they normally fetch.

Some theorists prefer to say that words like *rabbit* and *table* indicate things like rabbits and tables, as opposed to concepts. In many cases, this is harmless. Predicates like *unicorn* and names like *Vulcan* present familiar worries. But one can maintain that words are 'normally' used to indicate 'real' things, given a suitably restricted notion of normal use, so long as one isn't too demanding about what counts as real; cf. Chomsky (2000a), Pietroski (2005a), Hinzen (2007). In saying that *i*-expressions are *Begriffsplans*, my point is not to deny that speakers often use words to talk about language-independent things. On the contrary, my suggestion is that *i*-expressions are often used in this way *because* they call for construction of *i*-concepts, which are often constituents of thoughts that are (somehow) about mind-independent things. Since *i*-expressions are also associated with phonological instructions, they can be used to communicate thoughts. So in suitably controlled contexts, intuitions about the truth or falsity of a thought communicated with an *i*-expression can serve as useful data points for theories of *i*-meanings.[3]

[3] Likewise, I grant that speakers can use *i*-expressions to make assertions whose contents can

As these concessive remarks suggest, (1) is fully compatible with psychologized versions of Truth Conditional Semantics, according to which *i*-expressions are instructions to construct concepts that have Tarski-style satisfaction conditions.[4] If semantic instructions are individuated in this externalistic fashion — if the instructions require construction of concepts with certain *truth-theoretic* properties — then thesis (5) is false.

(5) Meanings are internalistic properties of expressions.

But (1) can be combined with the hypothesis that each *i*-expression not only has a Tarskian satisfaction condition, it has that satisfaction condition essentially.

That said, (1) is also compatible with a claim that is more friendly to internalism about meaning/understanding: *i*-expressions are instructions to construct *i*-concepts, some of which can be refined and used (in contexts where truth matters) to form thoughts that count as true or false because they are (modulo some complications like vagueness) sufficiently like 'Ideal Thoughts' whose constituent concepts really do have Tarski-style satisfaction conditions. This claim, which remains agnostic about the nature of semantic instructions and concepts constructed, is compatible with (5). So while (1) does not presuppose (5), (1) may be offered as part of a package that includes (5).

In subsequent sections, it will be important to be clear about what (1) does and does not imply. So let me stress these points now. Thesis (1) is incompatible with 'spare' theories according to which *i*-expressions have *no* semantic properties other than Tarskian satisfaction conditions.[5] But accepting (1), and describing

---

be characterized in terms of sets of possible worlds. When a suitably clear thought is asserted, we can talk about the worlds in which that thought is true; and for any suitably clear concept, we can talk about its extension at a world. But it does not follow that *i*-expressions have intensions, much less that the meaning of an *i*-expression is its intension. One can say instead that *i*-expressions are used to construct concepts that have (context-dependent) intensions; where these concepts can be combined to form thoughts whose contents are determined by the intensions of the constituent concepts.

4 See e.g. Larson & Segal (1995). This is not what Davidson (1967), Montague (1974), or Lewis (1972) proposed. But it is in the spirit of Harman (1970) and Partee (2004). A variant proposal is that *i*-expressions are instructions to construct concepts, and that the concepts constructed have satisfaction conditions, even though the instructions do not require this.

5 If we invent a Tarskian language in which 'Rx_1' is satisfied by a sequence σ iff σ(1) is a rabbit, then 'Rx_1' can be used to talk about rabbits, even by speakers who do not associate 'Rx_1' with any concept (much less a concept of rabbits as such). We might invent such a language, and stipulate that 'Rx_1' has no other semantic properties, in order to guarantee that 'Rx_1' has the same (concept-independent) interpretation for all users of the language. We can likewise stipulate that the logical constant 'p' is satisfied by σ iff σ(p) = Phosphorus, where σ(p) is the element of σ corresponding to the pth logical constant, and that 'h' is satisfied by σ iff σ(h) = Hesperus. Given the usual ancillary apparatus, it will follow that 'Rp' is true iff Phosphorus is a rabbit, and that 'Rh' is true iff Hesperus is a rabbit.

In general, expressions of the invented language can have (and be understood as having) truth-theoretic properties as their only semantic properties. Thesis (1) implies that I-languages do not have this spare Tarskian character. But to *define* a language whose semantics is concept-independent in this way, humans may need a prior I-language whose expressions are used as devices to fetch concepts that can be combined to form truth-evaluable thoughts.

the meanings of *i*-expressions as semantic instructions, leaves room for many conceptions of the relevant fulfillment conditions.

For example, a neo-Davidsonian might regard the untensed verb phrase *stab Caesar* as the following complex instruction: Fetch a singular concept of the individual Caesar; fetch a concept satisfied by ordered triples <e, x, y> such that e is a stabbing by x of y; and saturate (the most internal variable of) the latter with the former, thereby forming a concept satisfied by ordered pairs <e, x> such that e is a stabbing by x of Caesar. On this view, the instruction *stab Caesar* is fulfilled — and in that sense, satisfied — by constructing any concept with the specified truth-theoretic profile. In which case, one might say that *stab Caesar* is itself satisfied by <e, x> iff e is a stabbing by x of Caesar. Likewise, one might say that *dog* is an instruction to fetch a concept satisfied by x iff x is a dog, and that the *i*-expression *dog* inherits this satisfaction condition.

There are, however, other coherent conceptions of semantic instructions and fulfillment. For each speaker, each lexical item might be an instruction to fetch a concept that has a certain 'address' in mental space — viz. the address of the concept that was lexicalized with the relevant phonological form (PF). Suppose that at least typically, a lexical item is acquired via some process in which: A PF is linked to a pre-lexical concept, perhaps initially by mere association, but eventually by pairing the concept with an address A such that (i) HFL can generate a lexical *i*-expression that links the PF to an instruction to fetch a concept paired with A, and (ii) the resulting *i*-expression can have certain additional features corresponding to grammatical idiosyncracies. If the concept lexicalized remains the only concept at address A, then fulfilling the instruction is always a matter of fetching that concept. But suppose that one or more additional concepts get assigned to that address. Perhaps a formally new concept is defined in terms of the concept lexicalized, and HFL does not assign a new address to a concept so formed. Or perhaps the cognitive processes underlying polysemy can result in a family of concepts having the same address, so far as HFL is concerned. In such cases, there may well be more than one way of fetching a concept from address A.[6]

Of course, the PF of *dog* could have been lexically linked to a concept of cats, while the PF of *cat* was linked to a concept of dogs. And we can imagine two individuals that differ only this respect: In Oscar, the PFs are linked to concepts in this 'inverted' way; while in Felix, the PFs are linked to concepts in a 'proper-English' way. We can go on to say that Oscar's I-language does not count as an idiolect of English. And it may well be that the normal way of acquiring an (I-language that counts as an) idiolect of English involves using the PF of *dog* to lexicalize a concept with which the lexicalizer can think about dogs. But even if the concept lexicalized has a doggish satisfaction condition, it does not follow

[6] Suppose you tell me to fetch a box from a certain room. If I find four boxes in the room, I can fulfill your instruction by bringing back any one of the four. But if you also tell me to fetch a ball from another room, and put it in the box I fetch, then the sizes of the boxes/balls will impose constraints on which choices permit fulfillment of your tripartite instruction. In this context, see Pietroski (2005a) on the difference between *France is a hexagonal republic* and *France is a hexagonal, and France is a republic*; cf. Chomsky's (1995a) discussion of *London* as a device for referring to different 'things' on different occasions of use.

that *dog* is an instruction to fetch a concept with that satisfaction condition, not even if the concept lexicalized is the concept fetched. An instruction to *fetch a concept from a given address* is not an instruction to *fetch a concept with a certain truth-theoretic property* — not even if every concept at the address has that truth-theoretic property.[7]

This last point will be important when we turn to thesis (5). But for now, the point is just that thesis (1) makes it possible to reject externalist conceptions of meaning while accepting many externalistic claims about concepts. Readers already inclined to agree, or not care, can skim through to section 1.2 below.

Let's grant that if a lexical item λ counts as an instance of the English word *dog*, then λ is lexically linked to a concept of dogs — i.e. a concept with which one can think about dogs, and think of them as such (cf. Fodor 2003). I am more skeptical of the following claim: λ is not an instance of the English word *dog* if λ is lexically linked to a concept that would apply to things of a kind biologically distinct from but superficially similar to the actual dogs. Yet even given this much externalism, along with the related idea that a concept of dogs needs to be a concept that is essentially tied to the biology of actual dogs, a lexical *i*-expression can count as an instance of the English word *dog* without being an instruction to fetch a concept of dogs. For an *i*-expression can be an instruction to fetch a concept from a certain address — an address that actually links the PF of *dog* to a concept of dogs, where being a dog is a biologically demanding condition — without being an instruction to fetch a concept of dogs, or even an instruction to fetch a concept that applies to dogs.

In terms of the much discussed 'twin earth' thought experiments, doppelgangers might execute the same address-focused instruction — fetch a concept from address A — yet thereby fetch concepts with different contents, at least if doppelgangers in different environments can have concepts with different contents.[8] For just as we can imagine Oscar and Felix linking different concepts to the PF of *dog*, we can imagine Felix and Twin-Felix linking different concepts to the PF of *dog*, at least if Felix and Twin-Felix do indeed have different concepts. If the difference between dogs and cats is salient to Oscar and Felix, they may not link the PF of *dog* to the same address in mental space; in which case, Oscar and Felix may not have the very same I-language, even given an address-focused conception of semantic instructions. But if any differences between dogs and twin-dogs would go undetected by Felix/Twin-Felix, then presumably, these duplicates are type-identical with regard to how they link the PF of *dog* to an address in mental space. If this is correct, then so far as their shared HFL is concerned, the PF of *dog* is linked to the same address. In which case, Felix and Twin-Felix do have the same I-language, given an address-focused conception of

---

[7] An instruction to fetch a box from a certain room differs from an instruction to fetch a red box, even if every box is both in that room and red.

[8] For these purposes, we needn't worry about the differences between Putnam's (1975) scenarios — which need not correspond to possible worlds in Kripke's (1980) sense — and Burge's (1979) counterfactual situations that often hold the 'non-linguistic environment fixed. Extending Kaplan's (1988) notion of character to thought, Fodor (1987) famously appealed to 'narrow' contents; see Segal (2000) for helpful discussion. But whatever one says about concepts and *their* fulfillment conditions, the issues here concern *i*-expressions (hypothesized instructions to fetch and combine concepts).

semantic instructions.

Note that an address-focused instruction differs from any instruction to fetch a concept with a certain 'narrow' content, even if any concept fetched from that address indicates the same mapping from contexts to extensions. But my aim is not to eschew appeals to contents of any breadth. For all that I have said here, semantic instructions may have wide (but mind-dependent) contents that are individuated by the relevant addresses, or narrow contents, or both. More importantly, just as twins might execute the same lexical instruction, so they might execute the same phrasal instruction. If twins share the addresses A1 and A2, along with a binary operator '•', each twin might fulfill the following instruction: Form an instance of 'C1(_) • C2(_)' by replacing 'C1' with a concept found at A1 and replacing 'C2' with a concept found at A2. If '•' is a conjunction operator, which need not be invoked under a truth-theoretic guise, then twins who execute the instruction might form conjunctive concepts with different contents.[9]

On this view, two individuals can share an I-language and competently use the same lexical items to fetch concepts from the same addresses, even if (for whatever reason) the concepts fetched have *distinct* contents. My suspicion is that such cases are common, even without twin earth. It seems all too likely that speakers in the same conversation often share an *i*-expression (e.g. *water*), while severally using the expression to fetch concepts that differ extensionally in ways that are often but not always irrelevant — at least if we suppose that the concepts fetched via lexical items really do have extensions in contexts. Given two concepts that apply to the stuff in Lake Michigan, only one may apply to the stuff coming out of a certain tap; cf. Chomsky (2000a).

More generally, one shouldn't insist that if two otherwise linguistically similar speakers each use a lexical item to fetch concept, a mere extensional difference in the concepts fetched guarantees that either (i) the lexical *instructions differ*, or (ii) each instruction is associated with the *same function* from contexts to extensions, given a substantive notion of context — e.g., a sequence of potential values for conceptual variables, as opposed to entire possible worlds or their 'extension-determining' aspects. Another option, considered but not developed by Putnam (1975), is to deny that meanings determine extensions in contexts.

The meaning of an *i*-expression may determine the general 'shape' of any concept constructed in accord with that semantic instruction, without determining a characteristic function (from contexts to contents) for any concept so constructed. One can invent a *Begriffsschrift* that employs lots of indices, and thereby confines the content variation in concepts constructed to values of variables; cf. Stanley (2000). But I-languages may care less about conceptual content. One can say, if one likes, that Twin-Felix fails to have an I-language that

---

[9] This assumes that twin-concepts can have the same address, and hence that addresses can be individuated without reference to the satisfiers of addressed concepts. But if this is a substantive assumption, so is its negation. One can still say that each twin executes a demonstrative instruction, akin to "Fetch one of those," with a distinctive 'wide' fulfillment condition (if each twin has his own concepts). My claim is not that it is incorrect to describe twin-earth cases externalistically, but rather, that theorists can — and perhaps should — describe such cases internalistically.

counts as an idiolect of English (and that his I-language fails to employ an English 'water'-address), while Felix fails to have an I-language that counts as an idiolect of Twinglish. But there may be no theoretically interesting distinction between an I-language that counts as an idiolect of English and an I-language that counts as idiolects of Twinglish. If the difference is simply a matter of which environment the implemented procedure inhabits, then for scientific purposes, there may be no difference English procedures and Twinglish procedures.

There certainly are substantive issues about concepts, and their relation to normative notions like truth, in the vicinity; see Burge (2005). And with or without such issues in mind, one can hypothesize that duplicates can have different I-languages, taking the relevant procedures (intensions) to be individuated externalistically. Chomsky (1986) did not stipulate that I-languages are individuated internalistically; his was a proposal about how to count natural languages for scientific purposes. In any case, Ludlow (forthcoming) offers a notion of 'Ψ-language' that is neutral in this respect, while preserving the idea that Ψ-languages are like I-languages in being psychologically implemented algorithms — as opposed to E-languages, which are sets (extensions) of expressions. Externalists can thus eschew appeal to E-languages, yet allow for duplicates who implement different Ψ-languages.

If the *only* constraints on semantic instructions are truth-theoretic, then *many* distinctions will be irrelevant to fulfillment. If the *i*-expression *groundhogs who like coriander* is satisfied by construction of any concept with the right truth-theoretic properties, and similarly for *woodchucks fond of cilantro*, then individuals who implement *very* different psychological processes may fulfill the same instruction. But one can individuate instructions finely — saying *both* that truth-theoretically identical phrases can be semantically distinct (internal constraints matter), *and* that twins can execute distinct instructions (external constraints matter). One cannot, however, stipulate that meanings are instructions, so individuated. For one cannot stipulate the nature of *i*-expressions.[10]

One can stipulate that instructions are *semantic* only if they are truth-theoretic. But then internalists can say that *i*-expressions are psemantic instructions, with silent 'p' for 'psycho' (cf. Katz & Fodor 1963), leaving it open whether or not psemantic instructions are individuated truth-theoretically. Expressions of an I-language may be meaningful, in the sense of being psemantic instructions, without being semantic instructions in the stipulated sense. If it is *also* stipulated that expressions are meaningful only if they are semantic instructions, one can regard it is an open question whether *i*-expressions are meaningful in this doubly stipulated sense. The substantive debates concern the nature of the instructions. And to repeat, thesis (1)

(1) Meanings are instructions to build concepts.

---

[10] One can introduce a technical notion of 'Meaning' according to which Meanings determine extensions, and describe a purpose P such that the following normative claim is plausible: other things equal, expressions of a language used for purpose P *ought* to have Meanings. But that is a very different project.

provides a relatively neutral way of framing these debates. But (1) can be combined with independently confirmable theses that are, when taken together, less neutral.

1.2. Concatenation Calls for Conjunction of Monadic Concepts

At the heart of compositional semantics lies a simple question: What is the significance of combining expressions? For invented languages, answers can be stipulated. But for I-languages (cf. Higginbotham [1985]), the questions are empirical: What is the significance of combining *i*-expressions? Or put another way, taking *i*-expressions to be *Begriffsplans*, what operations of concept combination do complex *i*-expressions invoke?

Any proposal has immediate consequences for the significance of combinable expressions, since combinables must be of the right types for purposes of combination. In part because of the justified influence of Frege (1879, 1892) and Montague (1974) — despite their own non-psychologistic (though not anti-psychological) projects — it became standard to assume that combination signifies function-application. In which case, for any pair of combined expressions, one must signify an element in the domain of the function signified by the other. This made appeals to type-shifting unavoidable, given that I-languages allow for complex predicates like *red ball*, along with simple predicates like *red* and *ball*. On this view, combination always signifies function-application, but with the caveat that combination may *also* signify a further operation of type adjustment for adjuncts. Put another way, the familiar idea was that while concatenation is a univocal instruction to apply a function to an argument, combination may have significance that goes beyond that of mere concatenation of expressions.

More recently, the trend has been towards views according to which *some but not all* cases of combination signify function-application (or equivalently, predicate-saturation), even within the tradition of transformational grammar.[11] Elsewhere, I have argued for a stronger hypothesis: Combination always signifies predicate-conjunction, but with the caveat that combination may *also* signify a further operation of type adjustment for arguments. In my view, while concatenation is a univocal instruction to conjoin monadic concepts, combination may have significance that goes beyond that of mere concatenation.

To clarify, it may help to consider some proposals about expression formation. We can characterize an operation COMBINE, such that if α and β are combinable expressions of an I-language, then COMBINE(α, β) is the result of combining them. This operation is somehow *asymmetric*, since phrases are headed; for example, combining a verb with a noun yields a verb phrase, not a mere verb–noun concatenation. But the operation COMBINE(α, β), so characterized, may or may not be *primitive*. It might be basic, like Chomsky's (1995b,

[11] See e.g. Higginbotham (1985), Larson & Segal (1995), and Heim & Kratzer (1998) — and in the tradition of categorial grammar, Steedman (2000). Of course, Davidson (1967b) had proposed logical forms involving both saturation and conjunction, at least for certain adverbially modified verb phrases.

2000b) operation MERGE(α, β). And there are several ways of being non-basic. The operation COMBINE(α, β) might turn out to be a disjunction of basic operations — e.g., ADJOIN(α, β) or SELECT(α, β) — reflecting two or more asymmetric ways of forming phrases. But another possibility is that while phrases are always formed in the same way, COMBINE(α, β) is a complex operation. In particular, phrase formation may be a process of *concatenating* expressions *and labeling* the result; see Hornstein (in press) and Hornstein & Pietroski (2008a, 2008b). So one can adopt the hypothesis below,

COMBINE(α, β) = LABEL[CONCATENATE(α, β)]
CONCATENATE(α, β) = α^β
LABEL[α^β] = $[\alpha\ \beta]_{\alpha/\beta}$

with CONCATENATE(α, β) and LABEL[α^β] as the posited basic operations invoked by I-languages.

The subscript on $[\alpha\ \beta]_{\alpha/\beta}$ indicates that one of the concatenates is itself the label of the structured object, which can be identified with an ordered pair like {{α, β}, α}; cf. Chomsky (1995b). The source of phrasal asymmetry is then confined to the labeling operation, whose inputs are symmetric concatenations. Given any such view, one can go on to say that concatenation is an instruction to perform a certain binary operation on concepts, and that labeling is an instruction to perform some further operation on the concept associated with the non-dominant concatenate. Let me first illustrate this point from a traditional perspective, according to which concatenation signifies saturation, before turning to the idea that concatenation signifies conjunction.

One can hypothesize that the phrase $[\text{red ball}]_{\text{ball}}$ is the following instruction: Fetch a concept linked to *ball*, say BALL(_), with '_' as a variable ranging over individuals; fetch a concept linked to *red*, say RED(_), and perform a 'lifting' operation to obtain the corresponding concept λX.X(X) & RED(_); then saturate one of these concepts with the other, thereby obtaining a concept like BALL(_) & RED(_). One can describe the lifting operation as a reflex of the labeling: The label on $[\text{red ball}]_{\text{ball}}$ serves as an instruction to type-adjust the concept fetched via the other constituent; cf. Parsons (1970) and Kamp (1975). On this view, only *some* labels trigger this reflex. For the leading idea is that a phrase like $[\text{stab Caesar}]_{\text{stab}}$, in which a verb takes an argument, is an instruction to fetch concepts linked to the lexical constituents and then simply saturate one concept with the other (without any type-shifting). But asymmetry in the semantic effects of labeling is easily encoded.

Suppose that phrasal instructions are issued by *pairs* of expressions that have been concatenated *and labeled as a unit*, as opposed to the two concatenates taken separately. Think of $[\alpha\ \beta]_{\alpha}$ as a complex instruction that includes two sub-instructions that correspond to *label-relativized* concatenates: β relative to the label α, and α relative to the label α; where one of these sub-instructions is simply α relative to itself. If phrasal instructions are issued by labeled concatenations, and CONCAT is the operation signified by concatenation, then one can say that the phrasal expression $[\alpha\ \beta]_{\alpha}$ is an instruction to apply CONCAT to the concepts obtained by executing two sub-instructions: β relative to α, and α relative to α;

where relativizing an expression/instruction to itself makes no semantic difference.[12] The idea, which is independent of any particular hypothesis about CONCAT, is that the non-dominant concatenate β may need to be adjusted for purposes of combining with α. The *kind* of adjustment called for, if any, will depend on CONCAT and β. But one can hypothesize that CONCAT is the operation of saturation, and that predicate–adjunct combination (unlike predicate–argument combination) invokes type-lifting. For one can say that combining a predicate with a grammatical argument makes no difference to the instruction associated with the argument, while combining a predicate with an adjunct is an instruction to type-lift the adjunctive instruction.

If only for simplicity, let's assume that one way or another *Caesar* is marked as an argument, while *stab*, *ball*, and *red* are marked as predicates. And for ease of notation, let's underline arguments. (From a traditional perspective, one might think of underlining as an instruction to ignore phrasal relativization, thus precluding any substantive 'shifting' of the constituent semantic instruction.) Then the phrase [stab <u>Caesar</u>]$_{stab}$ can be an instruction to perform saturation on concepts obtained by executing two vacuously relativized expressions/instructions: *stab* relative to *stab*, and <u>*Caesar*</u> relative to *stab*. That is, despite the phrasal label, [stab <u>Caesar</u>]$_{stab}$ can be an instruction to perform saturation on concepts obtained by simply executing *stab* and *Caesar*. By contrast, [red ball]$_{ball}$ can be an instruction to perform saturation on concepts obtained by executing two relativized expressions/instructions, one of which is *non*-vacuously relativized: *red* relative to *ball*, and *ball* relative to *ball*. That is, given the phrasal label, [red ball]$_{ball}$ can be an instruction to perform saturation on (i) a concept obtained by executing *red* and type-lifting, and (ii) a concept obtained by simply executing *ball*.

Alternatively, one can invert the traditional perspective and adopt the following hypothesis: CONCAT is an operation of conjunction — more specifically, conjunction of monadic concepts; and predicate-argument combination, unlike predicate-adjunct combination, invokes a kind of type-adjustment. I discuss my specific neo-Davidsonian proposal, in terms of thematic roles, in section 1.3 below. But if combining a predicate with an adjunct makes no difference to the instruction associated with the adjunct, then [red ball]$_{ball}$ can be an instruction to conjoin two monadic concepts, obtained by executing two vacuously relativized instructions: *red* relative to *ball*, and *ball* relative to *ball*. That is, despite the phrasal label, [red ball]$_{ball}$ can be an instruction to conjoin concepts obtained by simply executing *red* and *ball*. And we can say that combining a predicate α with a grammatical argument β is an instruction to type-adjust the instruction issue by β itself. For given the phrasal label, [stab <u>Caesar</u>]$_{stab}$ can be an instruction to conjoin a pair of monadic concepts obtained by executing two relativized expressions/instructions, one of which is non-vacuously relativized: *stab* relative to *stab*, and *Caesar* relative to *stab*. That is, [stab <u>Caesar</u>]$_{stab}$ can be an instruction to conjoin concepts obtained by (i) simply executing *stab*, and (ii) executing *Caesar* and type-adjusting.

---

[12] This is a restricted version of the relativization that Higginbotham (1986) employs. For these purposes, I treat $ball_N$ and $stab_V$ as primitives. But see Hornstein & Pietroski (forthcoming) for more detailed discussion, and some implications of treating labels as formatives that can be combined with unlabeled lexical roots.

If it helps, think of underlining on this nontraditional view as an instruction to relativize *non*-vacuously: The 'default' effect of label-relativization is null, with a restricted kind of (thematic) relativization as the marked case. I cannot here defend this proposal in any detail. But in I-languages, adjunction seems to be open-ended (especially given relative clauses), while predicates seem to combine with at most three grammatical arguments.[13] *Prima facie*, this casts doubt on the semantic tradition of treating adjuncts as the marked cases that call for special treatment. And if adjunction invokes monadic concept conjunction, as a (biologically implemented) recursive combination operation, one might wonder if I-languages *also* invoke a recursive operation of saturation in which a concept of adicity *n* combines with a semantic argument to yield a concept of adicity *n–1*. But the more important point here is that we can (at least conceptually) decompose the semantic effect of combining expressions into the semantic effects of concatenating and labeling. And since labeling need not be semantically vacuous, one can distinguish the semantic effect of combining expressions from the semantic effect of concatenating expressions — even if in some cases, this distinction is not semantically significant.

In short, each expression of the form $[\alpha\ \beta]_\alpha$ can be a 'macro' instruction: execute the sub-instructions α and β, obtaining concepts C1 and C2, respectively; then form the concept CONCAT(C1, LAB[C2, α]), where LAB[C2, α] is the concept formed by subjecting C2 to the (perhaps vacuous) operation induced by the label α. Again, this general idea is compatible with various proposals about which operations CONCAT and LAB are. But suppose CONCAT is an operation of monadic concept conjunction, signified with '•'. Then $[\alpha\ \beta]_\alpha$ is an instruction to build a monadic concept of the following form: C1(_) • LAB[C2(_), α]; where LAB[C2(_), α] may be a complex thematic concept of the form ∃χ[C(χ) • Θ(_, χ)]. Potential values of the variable '_' include events, which can have individual participants, and individuals (which can be event participants). So concatenation can be an instruction to conjoin monadic concepts — and in this sense, conjunction can be the basic I-language mode of semantic combination — without each expression of the form $[\alpha\ \beta]_\alpha$ being an instruction to conjoin concepts formed by executing the sub-instructions α and β.

Strictly speaking, '•' is a little more permissive than an operator that can be flanked *only* by monadic concepts. The idea is that (the conceptual operation indicated with) '•' can be used to combine a monadic concept C(χ) with a formally dyadic concept Θ(_, χ) — with 'χ' ranging over potential participants of values of '_' — so long as one variable in the resulting concept, C(χ) • Θ(_, χ), is immediately closed to create a formally monadic concept. But one can hypothesize that this *slightly* relaxed operation of monadic concept conjunction is the one invoked by I-languages.[14]

---

[13] In this context, see Pietroski (2005a) for an argument that *that*-clauses are semantically like adjuncts, and a review of the trend towards (i) analyzing ditransitive constructions in terms of a covert preposition (suggesting a maximum of *two* arguments) and (ii) analyzing many transitive constructions in terms of a covert light verb.

[14] Here, I am also ignoring the difference between verbish instructions (which tend to be tensable) and nounish instructions (which tend to be indexable); see Hornstein & Pietroski (forthcoming) for discussion drawing on Baker (2005).

1.3. Grammatical Relations Invoke Thematic Relations and ∃-Closure

For present purposes, I take it as given that in I-languages, predicate-argument relations are at least often associated with thematic relations. Correspondingly, I assume, it is not *ad hoc* to say that an argument like *Caesar* in *stab Caesar* is somehow associated with a thematic concept like PATIENT(_, CAESAR); where this complex monadic concept applies to 'events' that have Caesar as their patient. If we can seriously entertain the more traditional idea that $[\text{red ball}]_{\text{ball}}$ is an instruction to form a concept like λX.X(_) & RED(_), which gets saturated with a concept like BALL(_), then we can seriously entertain a structurally similar idea: $[\text{stab } \underline{\text{Caesar}}]_{\text{stab}}$ is an instruction to form a concept like PATIENT(_, CAESAR) and conjoin it with another monadic concept of events like STAB(_); likewise, $[\underline{\text{Brutus}}\ [\text{stab } \underline{\text{Caesar}}]_{\text{stab}}]_{\text{stab}}$ is an instruction to add a conjunct like AGENT(_, BRUTUS).

I also assume that I-languages invoke a cognitive operation that is like existential closure in two respects. First, when the operation is applied to a monadic concept C(_), it yields a complete thought of the form ∃_[C(_)]; where a thought of this form is correct iff C(_) applies to one or more things. More briefly, the operation converts C(_) into a thought that is true iff C(_) is not empty. Second, the operation can convert a formally dyadic concept of the form 'C(X) • Θ(_, X)' — with values 'X' being potential participants of values of '_' — into a monadic concept of form '∃X[C(X) • Θ(_, X)]'; where a concept of this form applies to one or more potential values of '_' iff they are related, in the right thematic way, to one or more things that fall under the relevant monadic concept. But I do not assume that the closure operation invoked by I-languages, signified here with '∃', has the full power of existential closure to convert any concept of adicity *n* into a concept of adicity *n-1* by binding any conceptual variable. For example, I do not assume that this operation can convert a tetradic concept of the form 'R(X, Y) & S(W, X, Z) ⊃ T(Y, Z, W)]' into a triadic concept of the form '∃X[R(X, Y) & S(W, X, Z) ⊃ T(Y, Z, W)]'. Correlatively, the cognitive effect of the posited closure operation need not be characterized (*à la* Tarski) in terms of the satisfaction conditions imposed by open sentences with arbitrarily many variables on sequences of arbitrary many potential values of variables.

In Pietroski (2005b, 2006, 2008a, 2008b), I show how these relatively modest resources can provide a compositional semantics that accommodates a wide range of constructions — including causative and ditransitive constructions, plural noun phrases, prepositions, negation, and quantificational determiners like *every* that take tensed clauses as their external arguments. This is not the place for an explicit fragment of a semantic theory. But to illustrate, suppose that internal and external arguments of predicates are relativized as such.

For any grammatical argument ..., let *int*-... be the instruction issued by ... when it appears as the internal argument of a predicate, and let *ext*-... be the instruction issued by ... when it appears as the external argument of a predicate. (If each predicate takes at most two grammatical arguments, '*ext*' can be replaced with '~*int*', treating each external argument as the non-internal argument of its predicate.) And suppose that each argument is, by itself, an instruction to form a monadic concept. This concept may be complex and context sensitive, reflecting a

complex expression that includes a lexical noun along with a covert determiner and/or index. But for now, we can idealize, taking the relevant concepts to be atomic: CAESARIZER(_) and BRUTUSIZER(_); cf. Quine (1963), though see section 1.4 below. If grammatical arguments are instructions to build monadic concepts, it is easy to provide composition principles according to which the relativized expressions '*int*-Caesar' and '*ext*-Brutus' are instructions to build concepts of things with internal/external participants: ∃X[CAESARIZER(X) • INTERNAL(_, X)], ∃X[BRUTUSIZER(X) • EXTERNAL(_, X)]; where in each case, the variable introduced by the proper noun is ∃-closed.

Thus, [Brutus [stab Caesar]$_{stab}$]$_{stab}$ can be an instruction to construct a concept like the following: ∃X[BRUTUSIZER(X) • EXTERNAL(_, X)] • [STAB(_) • ∃X[CAESARIZER(X) • INTERNAL(_, X)]]. And if *stab* is understood as an instruction to fetch a concept of things (actions) whose internal/external arguments are the patients/agents of those things, the formalistic notions INTERNAL(_, X) and EXTERNAL(_, X) can be replaced with more specific thematic contents: ∃X[BRUTUSIZER(X) • AGENT(_, X)] • [STAB(_) • ∃X[CAESARIZER(X) • PATIENT(_, X)]].[15]

Adding tense and a final existential closure can yield a complete thought, corresponding to *Brutus stabbed Caesar*: ∃_[PAST(_) • [∃X[BRUTUSIZER(X) • AGENT(_, X)] • [STAB(_) • ∃X[CAESARIZER(X) • PATIENT(_, X)]]]]. Alternatively, the untensed clause can itself be the internal argument of a larger verb phrase as in *see Brutus stab Caesar*, as shown:

SEE(_) • ∃X[INTERNAL(_, X) • [∃X'[BRUTUSIZER(X') & AGENT(X, X')] •
[STAB(X) • ∃X[CAESARIZER(X') • PATIENT(X, X')]]]]

And this construction can be treated on a par with *see a tree*:

SEE(_) • ∃X[INTERNAL(_, X) • TREE(X)]

Of course, one wants to see the details for other construction types. But for present purposes, I take it as premise that the requisite homework can be done, along the lines suggested in Pietroski (2005b, 2006, 2008a, 2008b): The Conjunctivist idea outlined in section 1.2 can be turned into detailed theory, by hypothesizing that (i) certain grammatical relations serve as instructions to introduce certain dyadic/thematic relations, and (ii) each cycle or 'phase' directs construction of a conjunctive concept whose main variable, introduced by the syntactic head, can be ∃-closed; where this may leave open exactly one thematically introduced variable, thereby allowing for subsequent conjunction with other monadic concepts. A more radical suggestion, not explored here, is that the requisite ∃-closures reflect interface conditions between HFL — which always generates instructions to construct monadic concepts, as opposed to complete sentences of type <t> — and cognitive systems whose representations have *correctness* conditions along with at least some *dyadic/thematic* constituents.

---

[15] One can add that the internal/external participants of stabs are their agents/patients: ∀E[STAB(E) ⊃ ∀x[AGENT(E, X) ≡ EXTERNAL(E, X)] & ∀X[PATIENT(E, X) ≡ INTERNAL(E, X)]].

1.4. Lexicalization is a Partly Creative Process of Abstraction

I have been assuming that lexical items can be instructions to fetch monadic concepts, even with regard to words like *Brutus* and *stab*, which were presumably introduced in the course of lexicalizing non-monadic concepts like BRUTUS and STAB(X, Y) — or STAB(X, Y, Z), with 'z' as a variable for instruments, or STAB(X, Y, E), or STAB(X, Y, Z, E). I readily grant that humans and other animals have singular and polyadic concepts, and that such concepts are often lexicalized. But the concept lexicalized with an *i*-expression need not be the concept subsequently fetched with that expression. For lexicalization can be a process in which non-monadic concepts are paired with monadic analogs, even if the monadic analogs have to be abstracted from the concepts lexicalized.

Imagining mechanisms for such abstraction is not difficult. For illustration, suppose the concept lexicalized with *stab* is dyadic, with no event variable. Given STAB(X, Y), and enough logical apparatus to define new concepts, lexicalizers might be able to introduce a triadic concept STAB(_, X, Y): ∀X∀Y{STAB(X, Y) ≡ ∃_[STAB(_, X, Y)]}; cf. Davidson (1967b). Then STAB(_) might be introduced via thematic notions: ∀X∀Y∀_{STAB(_, X, Y) ≡ AGENT(_, X) & STAB(_) & PATIENT(_, X)}; cf. Castañeda (1967) and Davidson (1985). Then *stab* can be an instruction to fetch STAB(_), with thematic notions introduced via grammatical arguments and/or prepositional phrases, which may be lexically mandatory or optional; see Pietroski (2008a, 2008b) for further details.

This kind of abstraction may require (non-recursive) cognitive resources that are not needed for Conjunctivist *composition* of *i*-concepts. But lexicalization is one thing, composition another. Or perhaps lexicalizers can use STAB(X, Y) to introduce STAB(_) more directly and in more restricted terms:

$$\exists X \exists Y[\text{STAB}(X, Y)] \equiv \exists_\{\exists X[\text{AGENT}(_, X)] \bullet [\text{STAB}(_) \bullet \exists X[\text{PATIENT}(_, X)]]\}$$

But whatever the details, the idea is that children can abstract a monadic concept of stabs from the presumably polyadic concept lexicalized with *stab*. And given an independent grip on the relevant thematic notions, one can allow for patientless (and/or agentless) stabs, at least as conceptual possibilities; cf. Parsons (1990). For even if each actual stab is a stabbing of something by someone, there may be no contradiction in the thought that some stab that lacks a patient also lacks an agent.

In general, a lexicalizable *n*-adic concept can be used to define a monadic analog, given *n* thematic relations. And there are various ways of reanalyzing singular concepts as monadic, even without appealing to predicates like *Brutusizer*, which apply to at most one thing. Suppose the name *Brutus* is introduced in the course of lexicalizing an atomic singular concept. Even if the pre-lexical concept is a simple 'mental tag' of type <e>, the name can be a complex expression used to fetch and combine two concepts — one demonstrative, and one metalinguistic (cf. Burge 1973 or Katz 1994). In many languages, proper nouns can and often must appear with an overt determiner or demonstrative, as in *the/that/our Brutus*, suggesting that the proper noun is (used to fetch a concept that is) of type <e, t>; see Longobardi (1994), Giannakidou & Stavrou (1999). If

names in English are structurally similar, with a covert functional element, the lexical proper noun *Brutus* can be analyzed as an instruction to fetch a concept of things called (with the sound of) *Brutus*. Independent evidence in favor of some such analysis is independent evidence that singular concepts are indeed paired with monadic analogs in the course of lexicalization; see Pietroski (2007), drawing on many others, for further discussion. But let me summarize, to this point, by returning to a more theoretically neutral idea.

Whatever one says about the significance of the COMBINE(α, β), expressions formed via this grammatical operation can be viewed as instructions to construct concepts of certain types. Unless one thinks that grammatical combination is effectively unrestricted, there will be many possible conceptual types such that *i*-expressions are never instructions to construct concepts of those types. On the assumption that COMBINE(α, β) signifies at most a few operations, and that lexical types are subject to nontrivial constraints of some kind, there will be some range of types such that each *i*-expression is an instruction to construct a concept of one of those types. All *i*-concepts will belong to this class of concepts.

Conjunctivism is a very restrictive thesis according to which all *i*-concepts are monadic. Theorists are free to adopt more permissive semantic theories, according to which *i*-concepts exhibit many types: <e>, <t>, <e, t >, <e, <e, t>>; <t, t>, <t, <t, t>>; <<e, t>, t>, <<e, t>, <<e, t>, t>>; <<e, t>, <e, t>>, and perhaps others. But I know of no good reason for positing *i*-concepts of type <<t, e>, e>, and likewise for endlessly many other possible concepts definable in terms of <e> and <t>. Ideally, one wants a characterization of the possible types for *i*-concepts that is empirically adequate without overgenerating. In this respect, Conjunctivists try to err on the side of positing too little: The idea is to start by supposing that all *i*-concepts are type <e, t>, and then find out which facts can/cannot be accommodated in this fashion.[16] But whatever one says about the space of *i*-concepts, the concepts fetched via lexical items have to be in this space, since COMBINE(α, β) only operates on combinable expressions that are used to fetch or construct i-concepts.

By contrast, there is no requirement that all concepts be *i*-concepts. Humans and other animals may have concepts that cannot be combined via the operation(s) signified via COMBINE(α, β). For example, concepts are often saturated. But it does not follow that some *i*-concepts saturate others. And on any view, some of the concepts that a child lexicalizes may fail to be of the right form for purposes of I-language combination. In which case, some and perhaps many pre-lexical concepts will have to paired with *i*-concept analogs that can be fetched via the resulting words.

---

[16] If all *i*-concepts are number neutral (see Schein 1993, 2002, 2006 and Pietroski 2005b, 2008a, 2008b), perhaps we should say that *i*-concepts are uniformly of type <_, t>; where judgments are of type <t>, and a concept is of type <_, t> if it applies to *one-or-more* things, allowing for concepts that are neither essentially singular (in that they apply to two or more things only distributively) nor essentially plural (in that only apply, and only apply non-distributively, to two or more things). In which case, it may be that *no* pre-lexical concepts are *i*-concepts. For it may be that all pre-lexical concepts are either essentially singular or essentially plural. But let's set this complication aside.

Imposing an I-language on pre-lexical thought may thus require some conceptual 'reformatting'. Frege (1879, 1884, 1892) envisioned a process of imposing a *Begriffsschrift* on pre-scientific thought, recognizing that this would require considerable reformatting of our natural ideas. By inventing modern logic, and offering some 'fruitful definitions' of key arithmetic notions, Frege provided a model of how such reformatting might proceed — at least in principle — for the special case of imposing a now familiar hierarchy of types, which correspond to concepts that can be combined via function-application. (See Horty (2007) for extended discussion.) Frege was not offering hypotheses about HFL; he was concerned with an idealized language of scientific thought, designed to reflect the mind-independent world, not the I-languages that children so readily acquire. But one can use the basic types <e> and <t> to define a space of 'F-concepts', and then characterize sub-regions of this space in terms of constraints corresponding to possible operations for combining F-concepts.

Frege's operation of saturation imposes no constraint; for any F-concept, there is another such that they can be combined via saturation. But other operations, like monadic concept conjunction, only permit certain combinations of F-concepts. More generally, given one or more composition operations, let's say that an F-concept of type <α> is *licensed* iff for some type <β>, an F-concept of type <α> can be combined with an F-concept of type <β> given the operation(s). Correlatively, given one or more composition operations, we can at least imagine minds that go through a critical period in which many mental representations that are *not* licensed by those operations get used to abstract formally distinct but analytically related concepts that *are* licensed by those operations. One can hypothesize that human children go through some such critical period, because they have a language faculty, and that we humans thereby acquire a Fregean 'second nature' in the following sense: We acquire certain *i*-concepts, and thereby acquire some F-concepts we did not already have; where these formally new concepts can be combined via operations invoked by I-languages. Conjunctivists adopt the view that all *i*-concepts are monadic. In which case, the combination operations can be rather restrictive, but lexicalization needs to be a little creative, in ways that dovetail with the special role that certain dyadic/thematic concepts play in allowing I-languages to interface with pre-lexical thought.

1.5. Meanings are Internalistic Properties of Expressions

As already noted, one can say that meanings are instructions to create *i*-concepts, while maintaining that the semantic properties of expressions are individuated externalistically. Likewise, Conjunctivism is logically compatible with truth conditional semantics. One can say that theories of meaning for I-languages will take the form of Tarski-style theories of truth, with concatenation signifying conjunction. But truth conditional semantics faces considerable difficulties, even if one ignores biologically imposed constraints on composition. And in my view, it is especially implausible that *i*-expressions have Tarskian satisfaction conditions that compose in the fashion sketched above. In short, (1)–(4) may together make (5) plausible.

(1) Meanings are instructions to build concepts.
(2) Concatenation signifies predicate-conjunction.
(3) Grammatical relations invoke thematic relations and ∃-closure.
(4) Lexicalization is a partly creative process of abstraction.
(5) Meanings are internalistic properties of expressions.

For if Conjunctivism is even roughly correct, it seems that the nature of semantic composition is determined by relatively simple computations, which permit construction of representations with formal properties that make them suited for 'interfacing' between HFL (the human faculty of language) and other aspects of human cognition. It would be amazing — in ways suggesting cosmic benevolence — if these simple computations also let us generate representations that reflect the language-independent world well enough to have compositionally determined satisfaction conditions. Frege and Tarski showed us, among other things, just how hard it is to *design* a language that has a truth-theoretic semantics, even when one can *stipulate* the operative composition principles. Especially given the much discussed paradoxes — involving, for example, self-reference and vagueness — a plausible externalism about *truth* may impose limits on the kinds of languages for which a truth-theoretic semantics can be given. One cannot just declare that an expression like *set of bald linguists who often use heterological words* has a compositionally determined satisfaction condition. But even setting such considerations aside, constraints on I-languages seem to be at best orthogonal to the requirements of compositional truth theories, and often at odds with the idea that there are Conjunctivist truth theories for I-languages.

To take a much discussed example, consider *The sky is blue*, which seems to be on a grammatical par with *The pie is round*. *Prima facie*, a truth conditional semantics will need to treat *sky* as a predicate satisfied by skies, one of which can be (in a context c) the relevant thing *x* such that the truth or falsity of the sentence (relative to c) depends on whether or not *x* is blue (in c). For now, set aside worries about what contexts need to be, and what it is be blue (in a context). What is *x* supposed to be? Tarski could stipulate that the satisfiers of a given predicate were, for example, natural numbers; where these abstracta were antecedently well defined. Correspondingly, the right hand sides of metalinguistic claims like 'Pb is true iff 2 is prime' are couched in a theoretical idiom that is antecedently understood in terms of an explicit model of certain aspects of reality. Given certain notational conventions, the number two can be identified with a certain set; and we know, in a scientific way, what it is for a number to prime. But what are skies, and what is it to identify one of them as the one said to be blue? Absent answers to these questions, it is hard to even begin evaluating the hypothesis that an *i*-expression like *The sky is blue* has a compositionally determined truth condition.

I don't deny that we humans have one or more sky-*concepts*, along with various color-concepts. On the contrary, I think speakers can and do access such concepts in response to semantic instructions like *The sky is blue*. But even if one speculates that these concepts have satisfaction conditions, one need not burden semantic theorizing with the further speculation that *i*-expressions inherit these

satisfaction conditions. Recalling section 1.1, in any given I-language, the word *sky* may be an instruction to fetch a concept from a certain lexical address — as opposed to an instruction to fetch a concept with a certain satisfaction condition. If a child acquires an I-language that counts as an idiolect of English, then presumably, the child uses lexical *i*-expressions to fetch concepts that have been paired with phonological forms in ways that respect substantive constraints. But it doesn't follow that the child's *i*-expressions are (or that the child somehow takes her *i*-expressions to be) instructions to fetch concepts that meet these constraints. The instructions may be address-focused, even if more demanding conditions are met, *de facto*.

To be sure, theorists can inscribe 'axioms' and 'theorems' like the following: *x* satisfies *Brutus* iff *x* is a sky; *x* satisfies *blue* iff *x* is blue; *x* satisfies *blue sky* iff *x* is a sky and *x* is blue; etc. But absent a specific proposal about the potential values of the variable, this seems like an overly technical (and therefore misleading) way of saying that *blue* and *Brutus* are devices for fetching monadic concepts, and that *blue sky* is an instruction to conjoin concepts fetched via *blue* and *Brutus*. Of course, any one example is just that. But especially if concatenation uniformly signifies a simple operation like conjunction, truth conditional semanticists will continually face uncomfortable questions about values of the relevant variables.

One can be lulled into thinking that at least event variables are friendly to truth conditional semantics, since there are indeed events. But as Davidson (1985) himself noted, difficult questions of event individuation arise as soon as one takes such variables seriously in the context of a truth theory. (Are events individuated in terms of their spatiotemporal location, their causes/effects, their participants, or still other factors? Under what conditions are counterfactual claims about an event true?) And if we take events to be the values of event variables in theories of meaning for I-languages, without *stipulating* that these variable-values are also language-independent spatiotemporal particulars, then appealing to thematic relations quickly leads to the conclusion that events are distinguished in ways that reflect grammatical distinctions; see e.g. Schein (1993, 2002) and Tenny (1995). For example, no event of you facing me can be identical with an event of me facing you: Otherwise, a facing whose agent is *you* would be a facing whose agent is *me*; in which case, you'd be me. Or suppose there was an event of Jim Higginbotham drinking a (pint of) beer in a minute. It seems that qua value of an event variable, any such event must be distinct from any simultaneous event of Jim drinking beer for a minute, at least if *in a minute* and *for a minute* are truth-theoretic conjuncts of the relevant event descriptions. For if there is just *one* event of drinking that satisfies *both* conjuncts, then *prima facie*, *Jim drank beer in a minute* should be just as *true* as *Jim drank beer for a minute*.

To take another kind of example, *France is a hexagonal republic* is somehow weird in way that *France is hexagonal and France is a republic* is not. Prima facie, this asymmetry is at odds with the idea that a sentence of the form 'France is Φ' is true iff the satisfier of 'France' satisfies 'Φ', and that *hexagonal republic* is satisfied by *x* iff *x* satisfies both *hexagonal* and *republic*. I don't expect these simple illustrations to convince. Arguing against truth conditional semantics requires

discussion of many specific constructions, and the many potential replies.[17] But it isn't hard to get a feel for the general Cartesian worry, to which any nativist should be sensitive — viz., that *i*-expressions are concept–construction-instructions whose basic architecture is determined by endogenous constraints, not the language-independent world; cf. Chomsky (2000a).

Humans are lucky to have HFL, a faculty that lets us acquire I-languages in conditions of limited experience, and then use *i*-expressions that are compositional in *some* sense that is compatible with both the nature of HFL and the concepts we lexicalize. If *i*-expressions are also compositional in a truth-theoretic sense, then the demands of HFL and infant psychology somehow conspire to yield I-languages that are relevantly like the languages that Frege and Tarski *designed* to be truth-theoretically compositional. That sounds, to my ear, like wishful thinking. I'll return to this point after a brief detour.

2. Speculation: Lexicalization as a Human Cognitive Tool

If my proposed account of semantic composition is roughly on the right track, it invites a non-standard conjecture about what makes HFL distinctively human.

Consider two theses, at least one of which is presumably true.

(L) Humans have a special capacity to lexicalize mental representations.
(C) Humans have a special capacity to combine mental representations.

Perhaps both are true. But other things equal, one doesn't want to posit two distinctively human capacities that somehow manage to interact in the right ways. Thesis (C) has a distinguished heritage; see Hauser, Chomsky & Fitch (2002) for a recent discussion. It can seem obvious that recursion is the key to human language. But as Chomsky (1957) famously discussed, there's recursion, and then there's recursion. A mere concatenater of atomic expressions can generate arbitrarily many complex expressions of the form α^β; and non-human animals can surely concatenate at least some representations.

I suspect that other animals also have the basic capacities required to treat concatenations as Conjunctivist instructions: Do this, do that, and connect the results with something like an AND-gate. It seems obvious that nonhuman animals have many concepts — or if you prefer, pre-lexical mental representations — that are at least candidates for lexicalization. In terms of nonlinguistic cognitive capacities, our evolutionary cousins are quite impressive, especially compared with human infants. And it seems quite plausible that other social

[17] See e.g Stanley (2000, 2002) and Schein (forthcoming); cf. Pietroski (2005b, 2006b). This is *not* to deny the enormous value of extant work done within the framework of truth-conditional semantics. Nor is to suggest that a Conjunctivist theory will itself be adequate to capture the meanings of *i*-expressions. On the contrary, a Conjunctivist theory may capture only a semantic 'primal sketch' that is supplemented by other processes that are not compositional in the same way. As with vision, it is hard to say where one faculty/module ends, and interfaces to 'general cognition' begin. But in my view, talk of truth conditions reflects a mix of linguistic factors and other cognitive (or metaphysical) factors that are orthogonal to the study of HFL.

primates can (when suitably stimulated) come to have some thematic concepts corresponding to participation relations exhibited by events and agents/patients/places/times/etc.[18] But humans have still more, on at least two fronts, even if we set aside the possibility that the concepts humans lexicalize are somehow especially lexicalizable.

First, as already noted, grammatical combination is asymmetric in a way that mere concatenation is not. Following Hornstein (in press), I have suggested that labeling is the source of this asymmetry; see also Boeckx (2008). Given labeled phrases, grammatical relations like 'being the internal/external argument of' can serve as devices for invoking thematic relations, and thereby distinguishing the significance of mere *concatenation* from the significance of I-language *combination*. But one might wonder: If labeling is all it takes to unleash the power of HFL from antecedently available animal capacities, why don't many species have analogs of HFL?

This leads to the second point, which is that the cognitive value of I-languages is limited if there are only a few lexical *i*-concepts. If only a few atomic concepts can be fetched, for purposes of creating complex concepts, recursively generating concept construction instructions won't do much good. So especially if concatenation signifies an elementary cognitive operation, which by itself only allows for simple cases of concept composition, it may be that lexicalization is the key new linguistic trick. Though correlatively, the utility of abstraction mechanisms that pair lexicalizable concepts with *i*-concepts may depend on a capacity to efficiently combine *i*-concepts via independently available operations. Perhaps some ancestral primates fortunately *connected* a capacity to label concatenations (or otherwise introduce some such source of grammatical asymmetry) with a capacity for formal abstraction that allowed for lexicalization (i.e. the creation of *i*-concepts, not mere pairs of signals with pre-lexical concepts).

Since this already seems like a lot to posit, in terms of distinguishing humans from other animals, I am wary of also positing a capacity to employ an operation like function-application (concept-saturation) recursively. By contrast, conjunction appears to be a simple and ubiquitous computation. Moreover, whatever we say about the significance of combination, we are faced with the empirical fact that human children naturally lexicalize with a vengeance; whereas other animals can 'only' learn to pair signals with concepts to a certain extent, given lots of explicit training. So returning to theses (L) and (C), we have independent evidence in favor of (L), making it undesirable to *also* posit (C). And other things equal, I would rather not say that other primates have a (perhaps unused) capacity to construe concatenations as recursive instructions to fetch and saturate polyadic concepts of arbitrary adicity.

Given that Conjunctivism seems to have a chance of yielding descriptive adequacy, I conclude that we should explore the possibility that lexicalization is a large part of the uniquely human aspect of HFL — and that *i*-concepts are monadic *because* conjoining monadic concepts is a common cognitive capacity that can be used in novel and fruitful ways, once existing concepts are paired with (monadic) *i*-concepts. From this perspective, I-languages impose a common

[18] See Hurford (2007) for relevant and congenial discussion of various evolutionary issues.

and simple format on animal concepts, with the result that human *i*-concepts can be combined via simple operations. This may have real value in terms of computational efficiency.

It may also be occasionally useful for purposes of representing the world. But on the whole, one would predict what one actually sees when looking at I-languages in any detail: Many features (as in expressions like *bald linguists who like blue skies*) that seem like quirks, and often serious design flaws, if one assumes that expressions have compositionally determined Tarskian satisfaction conditions; and many features, like invocation of thematic roles, that make sense if one assumes that expressions are compositional in a more simple-minded, less world-directed way.

Again, I don't expect this compressed argument to convince. But the idea is that if Conjunctivism turns out to be the best account of semantic composition — and of how it came to pass that the biological world includes semantically compositional I-languages — then that is itself an argument against truth conditional semantics, given the plausible assumption that there are no *Conjunctivist* theories of *truth* for I-languages.

More generally, we should evaluate semantic externalism and other high-level theses about meaning in light of our best proposals concerning (i) the significance of the operation COMBINE, and (ii) how humans came to have a faculty in which this operation has that significance. Even if my proposal is entirely wrong in detail, I hope to have illustrated how a cluster of theses like (1)–(5) can hang together.

(1) Meanings are instructions to build concepts.
(2) Concatenation signifies predicate-conjunction.
(3) Grammatical relations invoke thematic relations and ∃-closure.
(4) Lexicalization is a partly creative process of abstraction.
(5) Meanings are internalistic properties of expressions.

For we can and should ask how alternative clusters compare, if only to loosen the current grip of various semantic dogmas on our collective theoretical imagination.

References

Bloom, Paul. 2000. *How Children Learn the Meanings of Words* (Learning, Development, and Conceptual Change). Cambridge, MA: MIT Press.
Boeckx, Cedric. 2008. *Bare Syntax*. Oxford: Oxford University Press.
Burge, Tyler. 1973. Reference and proper names. *Journal of Philosophy* 70, 425–39.
Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies* 4, 73–121.
Burge, Tyler. 2005. *Truth, Thought, Reason: Essays on Frege*. Oxford: Oxford University Press.

Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.
Chomsky, Noam. 1977. *Essays on Form and Interpretation*. New York: North Holland.
Chomsky, Noam. 1986. *Knowledge of Language*. New York: Praeger.
Chomsky, Noam. 1995a. Language and nature. *Mind* 104, 1–61.
Chomsky, Noam. 1995b. *The Minimalist Program*. Cambridge, MA: MIT Press.
Chomsky, Noam. 2000a. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
Chomsky, Noam 2000b. Minimalist inquiries: The framework. In Roger Martin, David Michaels & Juan Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, 89–155. Cambridge, MA: MIT Press.
Church, Alonzo 1941. *The Calculi of Lambda-Conversion*. Princeton, NJ: Princeton University Press.
Castañeda, Hector-Neri, 1967. Comments. In Rescher (1967), 104–112.
Davidson, Donald. 1967a. Truth and meaning. *Synthese* 17, 304-23.
Davidson, Donald. 1967b. The logical form of action sentences. In Rescher (1967), 81–95.
Davidson, Donald. 1985. Adverbs of action. In Bruce Vermazen & Merril B. Hintikka (eds.), *Essays on Davidson: Actions and Events*, 230–241. Oxford: Clarendon Press.
Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Explorations in Cognitive Science). Cambridge, MA: MIT Press.
Fodor, Jerry A. 2003. *Hume Variations* (Lines of Thought). Cambridge: Blackwell.
Fodor, Jerry A. 2008. *LOT II*. Oxford: Oxford University Press.
Frege, Gottlob. 1879. *Begriffsschrift*. Halle: Louis Nebert. [English translation in van Heijenoort, Jean (ed.). 1967. *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931* (Source Books in the History of the Sciences). Cambridge, MA: Harvard University Press.]
Frege, Gottlob. 1884. *Die Grundlagen der Arithmetik*. Breslau: Wilhelm Koebner. [English translation in Frege, Gottlob. 1974. *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number* [transl. by J.L. Austin]. Oxford: Basil Blackwell.]
Frege, Gottlob. 1892. Function and concept. In Peter Geach & Max Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*, 21–41. Oxford: Blackwell, 1980.
Giannakidou, Aanastasia & Melita Stavrou. 1999. Nominalization and ellipsis in the Greek DP. *The Linguistic Review* 16, 295–331.
Harman, Gilbert. 1970. Deep structure as logical form. *Synthese* 21, 275–97.
Hauser, Marc, Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298, 1569–1579.
Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar* (Blackwell Textbooks in Linguistics). Oxford: Blackwell.
Higginbotham, James. 1985. On semantics. *Linguistic Inquiry* 16, 547–93.
Higginbotham, James. 1986. Linguistic theory and Davidson's program. In Ernest Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, 29–48. Oxford: Basil Blackwell.

Hinzen, Wolfram 2007. *An Essay on Names and Truth*. Oxford: Oxford University Press.

Hornstein, Norbert. In press. *A Theory of Syntax: Minimal Operations and Universal Grammar*. Cambridge: Cambridge University Press.

Hornstein, Norbert & Paul Pietroski. Forthcoming. Basic operations. *Catalan Journal of Linguistics*.

Horty, John. 2007. *Frege on Definitions: A Case Study of Semantic Content*. Oxford: Oxford University Press.

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

Kamp, Hans. 1975. Two theories about adjectives. In Edward L. Keenan (ed.), *Formal Semantics of Natural Languages*, 123–155. Cambridge: Cambridge University Press.

Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry & Howard K. Wettstein (eds.), *Themes from Kaplan*, 481–564. New York: Oxford University Press.

Katz, Jerrold. 1994. Names without bearers. *Philosophical Review* 103, 1–39.

Katz, Jerrold & Jerry A. Fodor 1963. The structure of a semantic theory. *Language* 39, 170–210.

Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Larson, Richard K. & Gabriel Segal. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

Lewis, David. 1972. General semantics. In Donald Davidson & Gilbert Harman (eds.), *Semantics of Natural Language*, 169–218. Dordrecht: Reidel.

Longobardi, Giuseppe. 1994. Reference and proper names. *Linguistic Inquiry* 25, 609–665.

Ludlow, Peter. Forthcoming. *The Philosophy of Generative Linguistics*. New York: Oxford University Press.

Montague, Richard. 1974. *Formal Philosophy*. New Haven: Yale University Press.

Parsons, Terence. 1970. Some problems concerning the logic of grammatical modifiers. *Synthese* 21, 320–334.

Parsons, Terence. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics* (Current Studies in Linguistics 19). Cambridge, MA: MIT Press.

Partee, Barbara H. 2004. *Compositionality in Formal Semantics*. Oxford: Blackwell.

Pietroski, Paul. 2002. Function and concatenation. In Gerhard Preyer & Georg Peters (eds.), *Logical Form and Language*, 91–117. Oxford: Oxford University Press.

Pietroski, Paul. 2003. Quantification and second-order monadicity. *Philosophical Perspectives* 17, 259–298.

Pietroski, Paul. 2005a. Meaning before truth. In Gerhard Preyer & Georg Peters (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 253–300. Oxford: Oxford University Press.

Pietroski, Paul. 2005b. *Events and Semantic Architecture*. Oxford: Oxford University Press.

Pietroski, Paul. 2006a. Interpreting concatenation and concatenates. *Philosophical Issues* 16, 221–45.

Pietroski, Paul. 2006b. Character before content. In Judith Thomson & Alex Byrne

(eds.), *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, 34–60. Oxford: Oxford University Press.

Pietroski, Paul. 2007. Systematicity via monadicity. *Croatian Journal of Philosophy* 7, 343–374.

Pietroski, Paul. 2008a. Lexicalizing and Combining. Ms., University of Maryland, College Park.

Pietroski, Paul. 2008b. Minimal Semantic Instructions. Ms., University of Maryland, College Park.

Poeppel, David & David Embick. 2005. Defining the relation between linguistics and neuroscience. In Anne Cutler (ed.), *Twenty-First Century Psycholinguistics: Four Cornerstones*, 103–118. Mahwah, NJ: Lawrence Erlbaum.

Putnam. Hilary. 1975. The meaning of 'Meaning'. In Keith Gunderson (ed.), *Language, Mind and Knowledge*, 131–193. Minneapolis: University of Minnesota Press.

Rescher, Nicholas (ed.). 1967. *The Logic of Decision and Action*. Pittsburgh, PA: University of Pittsburgh Press.

Schein, Barry. 1993. *Plurals and Events* (Current Studies in Linguistics 23). Cambridge, MA: MIT Press.

Schein, Barry. 2002. Events and the semantic content of thematic relations. In Gerhard Preyer & Georg Peters (eds.), *Logical Form and Language*, 263–344. Oxford: Oxford University Press.

Schein, Barry. 2006. Plurals. In Ernest Lepore & Barry Smith (eds.), *The Oxford Handbook in the Philosophy of Language*, 716–767. Oxford: Oxford University Press.

Schein, Barry. Forthcoming. *Conjunction Reduction Redux*. Cambridge, MA: MIT Press.

Segal, Gabriel. 2000. *A Slim Book about Narrow Content*. Cambridge, MA: MIT Press.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.

Stalnaker, Robert. 1999. *Context and Content.* Oxford: Oxford University Press.

Stalnaker, Robert. 2003. *Ways a World Might Be.* Oxford: Oxford University Press.

Stanley, Jason. 2000. Context and logical form. *Linguistics and Philosophy* 23, 391–424.

Stanley, Jason. 2002. Making it articulated. *Mind and Language* 17, 149–68.

Stanley, Jason & Zoltán Gendler Szabó. 2000. On quantifier domain restriction. *Mind and Language* 15, 219–261.

Steedman, Mark. 2000. *The Syntactic Process* (Language, Speech, and Communication). Cambridge, Mass: MIT Press.

Tenny, Carol L. 1994. *Aspectual Roles and the Syntax–Semantics Interface* (Studies in Linguistics and Philosophy 52). Dordrecht: Kluwer.

Paul Pietroski
University of Maryland
Department of Linguistics
1401 Marie Mount Hall
College Park, MD 20742
USA
pietro@umd.edu

Not the Only Word

Kenneally, Christine. 2007. *The First Word: The Search for the Origins of Language*. New York: Penguin Books.

By Lyle Jenkins

Christine Kenneally's *The First Word* is a review of work in the area of language evolution intended as an introductory overview for the general reader and, as such, is a valuable resource for pointers to work in progress on a wide range of evolutionary topics; these include primate calls, birdsong, categorical perception, gene research, computer simulation studies, to name a few. This is the subject of the second and third parts of the book. These parts contain the most valuable information on language evolution research. However, the first part of the book is devoted to interviews with Noam Chomsky, Sue Savage–Rumbaugh, Stephen Pinker and Paul Bloom, and Philip Lieberman. Here the stage is set for a kind of 'linguistics wars' on language evolution, with Chomsky on one side of a 'debate' and just about everybody else on the other side. This is the least convincing part of the book, as we will see.

According to Kenneally, the study of the evolution of language can be divided into several phases — one starting in 1866, when "the Société de Linguistique of Paris declared a moratorium on the topic" (p.7), and another phase when "the official ban developed fairly seamlessly into a virtual ban" (p.79) which was maintained, it is claimed, until the publication of a paper by Pinker and Bloom around 1990 (but see below). The "virtual ban" on the study of evolution of language seems to be ascribed by Kenneally almost solely to Chomsky. Unfortunately, the historical record, which Kenneally examines somewhat superficially, does not bear out her claim of any "ban", virtual or otherwise.

Kenneally reports that the "academic censorship" lasted for more than a century and language evolution remained a "disreputable pursuit", until several conferences on the origins of language were organized in the 1970s. She neglects to mention that at the earliest conferences on the biological foundations of language during this period, Chomsky was a key participant and speaker, apparently in violation of his own ban. These include the interdisciplinary meeting on language and biology at Dedham, Massachusetts, in 1974, sponsored by the Royaumont Center for a Science of Man, as well as a conference on *Ontogenetic and Phylogenetic Models of Cognitive Development* in Paris in 1975, both organized by Massimo Piattelli–Palmarini (see also Boeckx & Grohmann 2007, in their editorial to the first issue of *Biolinguistics*).

biolinguistics
© The Author (2008)
ISSN 1450–3417
Biolinguistics 2.4: 342–347, 2008
http://www.biolinguistics.eu

Moreover, Kenneally doesn't mention that at one of the conferences she lists, *Origins and Evolution of Language and Speech*, organized by the New York Academy of Sciences in 1976, one of the participants and speakers was Noam Chomsky. Far from being regarded as someone out to squelch discussion of language evolution, one of the conference organizers, Stevan Harnad, noted in his introductory remarks that "the revolution of linguistics due to Noam Chomsky has provided a very different idea of what the nature of the 'target' for the evolutionary process might actually be". That is to say, to attain a deep understanding of language evolution, one must first understand what the 'target' of evolution is (the quotes indicate that 'target' is not to be understood in a teleological sense). This was one of the reasons that Eric Lenneberg (1967) chose to include an appendix on generative grammar by Chomsky in his classic *Biological Foundations of Language* as early as 1967.

Or, as Chomsky framed it at the Lenneberg symposium at around the same time — and on later occasions — the study of biolinguistics was concerned with standard questions about mechanisms, development, and evolution that any biologist would ask about a biological system:

(1) What is knowledge of language?

(2) How does language develop in the child?

(3) How does language evolve in the species?

Although these three questions can be studied in any order, or even in parallel, the better answers we have to question (1), i.e. have an understanding of the mechanisms and structure of the language faculty, the better we can answer question (2), how these mechanisms and structure unfold in the child — and, as recent work on evolutionary developmental biology (evo-devo) has shown, the better we understand the answers to questions (1) and (2), the better equipped we are to study what Harnad above called the "target" of language evolution. For example, recursion is one the important syntactic mechanisms available to the language faculty. This insight from theoretical linguistics in turn opened the cross-species investigation of recursion in other species (as, for example, in tamarins and starlings).

As another example, at a conference on *Maturational Factors in Cognitive Development and the Biology of Language* held in 1978, there is a discussion between Chomsky and the neurologist Norman Geschwind about a number of questions concerning evolution of language; among others, about the cerebral asymmetries in the great apes, auditory tasks in the left temporal lobe of the monkey, asymmetry for emotional behavior in the brain, the recognition of species-specific cries in the left hemisphere of Japanese monkeys, male-female differences in left-right asymmetry for areas involved in bird song, and so on. (For the full exchange, see Jenkins 2000.) This interchange provides a flavor of the discussions taking place among linguists and neurologists at the time. In any case, this does not seem to be the behavior of someone out to pose a "virtual ban" on the topic of evolution.

Kenneally tends to sprinkle her factual reporting with unsupported

generalizations: "For example, scientists assumed for a long time that the parts of the brain that have to do with language must be wholly new, recently evolved additions that we do not share with nonhumans". In asserting this, she totally overlooks a whole generation of work on brain and language; one looks in vain for any mention of Geschwind and Galaburda's work on cerebral dominance and asymmetry (e.g. Geschwind & Galaburda 1986) or LeMay and Geschwind's work on the morphological asymmetries of the brains and skulls of nonhuman primates (e.g. LeMay & Geschwind 1975), to mention only a few examples. Nor was work on evolution of lateralization limited to the language areas; consider, for example, the work by Denenberg and colleagues on functional asymmetries in the rat (e.g. Denenberg 1981) as well as the numerous asymmetries in lower organisms documented by Corballis and Morgan (e.g. Corballis & Morgan 1978; Morgan & Corballis 1978).

In 1998, when Gannon and colleagues reported their findings of an asymmetry in the left planum temporale area of chimpanzee brains (Gannon *et al*. 1998), this was heralded in the press as "challenging cherished notions of how language evolved in humans and why apes cannot talk" (*New York Times*). However, left-right asymmetries in non-human primates, including the planum temporale area, had been long known (see e.g. Cunningham 1892, Fischer 1921, Yeni–Komshian & Benson 1976). The German magazine *Der Spiegel* claimed that until the study of Gannon *et al.*, it had been thought that the left and right sides of the brains of non-human primates were absolutely equal, although this had been shown twenty years earlier not to be the case by the study of Beheim–Schwarzbach (1975), who had compared the temporal regions in humans, chimpanzee, and the orangutan.

Moreover, in 1979, the Linguistics Society of America held its first Summer Institute abroad at the Joint Linguistic Society of America and University of Salzburg Summer Linguistics Institute at the University of Salzburg, Austria with the theme of *Linguistics and Biology*, which included courses, seminars, and other presentations as well as discussions on linguistics and biology of language, including neurology and the evolution of language. Although many scholars from Europe and around the world attended this Linguistics Institute, no mention was made of the "virtual ban" that had supposedly been imposed by Chomsky. Also around this time there were many fruitful contacts between the ethologist and evolutionary biologist Konrad Lorenz and his colleagues in Austria and at the Max Planck Institute in Germany, and generative linguists at the universities of Vienna and Salzburg. In 1976, Lorenz and his colleagues participated in a symposium on language and biology at the Salzburg Summer School of Linguistics (apparently in violation of the "ban").

Much is made by Kenneally of a paper by Pinker and Bloom, in which the central thesis is: "In one sense our goal is terribly boring […]. All we argue is that […] the only way to explain the origins of such abilities [such as language, vision, etc. — LJ] is through the theory of natural selection" (Pinker & Bloom 1990). However, the idea that natural selection and adaptation play a role in language evolution was hardly controversial long before the Pinker and Bloom paper. In fact, the application of these standard biological ideas may be seen in the popular writings of the biologists and Nobel Laureates Monod, Jacob, and Luria in the

early 1970s in their discussion of the ideas of Chomskyan generative grammar. In fact, Kenneally notes: "He [Chomsky — LJ] reiterated that there were factors in evolution other than natural selection, which were as likely to be significant. And in this regard, Chomsky, Pinker and Bloom were essentially in agreement, their debate arising more from differing emphases than actual discord". So the reader is left scratching his or her head as to what the fuss is all about, and why people make such fantastical claims, such as that Chomsky believes that language is not a product of evolution (Plotkin 1998) or that he is trying to enforce a ban on language evolution studies.

Kenneally (p. 200) states that after the discovery of the language effects of FOXP2, "the possibilities (were hailed) for a new science" and that it was called "neurogenetics" by Vargha–Khadem and her colleagues (who worked on the phenotype of the FOXP2 system; cf. Vargha–Khadem 1995). However, the field of neurogenetics has been around a lot longer than since 2001, and its scope is considerably broader than language, referring to the genetics of the nervous system in general. For example, nearly 30 years ago, Xandra Breakefield edited a volume on the subject. This collection also included an article by Kenneth K. Kidd & Mary Ann Records entitled 'Genetic methodologies for the study of speech' (Kidd & Records 1979). Although the term was even in use much earlier than this, this shows that the genetics of language was understood to be a subfield of the much broader field of neurogenetics. In addition, the *Journal of Neurogenetics* launched its first issue in 1983.

Kenneally notes that "it's been pointed out that the rules of phonology contradict Chomsky's notion of the poverty of stimulus — the idea that there is not enough information in the language a child hears for it to learn language" (p. 155). However, the work cited only argues that certain data from phonology don't require appeal to poverty of stimulus. Moreover, poverty of stimulus doesn't require that all language data be part of man's genetic endowment. Clearly, then you would not be able to explain how an infant born in a Japanese-speaking environment,but moved at birth to an English-speaking environment comes to learn English. It has always been understood that language learning requires a complex interplay between internal and environmental factors. So, for example, to argue that poverty of stimulus is superfluous, you would have to show that the phenomena explained by it; e.g., structure-dependence of syntactic rules such as question inversion in English can be derived from the data available to the child. There have been serious efforts to do this, e.g. Pullum & Scholz (2002). However, Legate & Yang (2002) examined the child language corpus used and showed that this particular attempt failed. Even if the attempt had succeeded, one would still have to go on to provide an alternative explanation for all of the numerous phenomena that have been discovered over the years and explained by poverty of stimulus (island conditions and others) to rule it out. In fact, however, researchers in biolinguistics now accept that there is a genetic endowment underlying the language faculty. Most research has long since shifted to other questions such as what is genetically specified and what environmental input is required, and whether this genetic endowment is in part or in whole domain-specific or not, or even species-specific or not.

In fact, an alternative picture to Kenneally's 'linguistics wars' view of the

study of language evolution is that biolinguistics has moved gradually through several phases in recent decades, corresponding to the questions posed earlier about (1) knowledge of language, (2) acquisition of language, and (3) evolution of language. The first phase was primarily concerned with the construction of generative grammars that represented the knowledge of language (language faculty). The second phase extended the findings across many languages and attempted to account for both the universal properties of language as well as their variation (e.g., in the Principles–and–Parameters model). The third phase, building on the results of the first two phases, is increasingly concerned with questions of function, design, and evolution (e.g., the Minimalist Program).

Kenneally cites the famous maxim of Dobzhansky (1964) that "nothing in biology makes sense except in the light of evolution". Modern work on evolution ('evo-devo') has turned that maxim on its ear, as Pennisi (2002) has noted, so that one might say that higher stages of evolution make no sense except in light of (developmental) biology. Accordingly, if we want to understand the evolution of language, we must also understand the (developmental) mechanisms underlying one of evolution's magnificent achievements, the human language faculty.

References

Beheim–Schwarzbach, Dorothée. 1975. Further studies of the cytoarchitectonic division in the dorsal surface of the 1st temporal gyrus of a linguistic genius and 2 anthropoids. *Zeitschrift für mikroskopisch-anatomische Forschung* 89, 759–776.

Boeckx, Cedric & Kleanthes K. Grohmann. 2007. The *Biolinguistics* manifesto. *Biolinguistics* 1, 1–8.

Corballis, Michael C. & Michael J. Morgan. 1978. On the biological basis of human laterality: I. Evidence for a maturational left-right gradient. *Behavioral and Brain Sciences* 1, 261–269.

Cunningham, Daniel J. 1892. *Contribution to the Surface Anatomy of the Cerebral Hemispheres* (Cunningham Memoirs). Dublin: Royal Irish Academy of Science.

Denenberg, Victor H. 1981. Hemispheric laterality in animals and the effects of early experience. *Behavioral and Brain Sciences* 4, 1–49.

Dobzhansky, Theodosius. 1964. Biology, Molecular and Organismic. *American Zoologist* 4, 443–452.

Fischer, E. 1921. Über die Variationen der Hirnfurchen des Schimpansen. *Anatomischer Anzeiger* 54, 48–54.

Gannon, Patrick J., Ralph L. Holloway, Douglas C. Broadfield & Allen R. Braun. 1998. Asymmetry of chimpanzee planum temporale: Humanlike pattern of Wernicke's brain language area homolog. *Science* 279, 220–222.

Geschwind, Norman & Albert M. Galaburda. 1986. *Cerebral Lateralization: Biological Mechanisms, Associations, and Pathology*. Cambridge, MA: MIT Press.

Kidd, Kenneth K. & Mary Ann Records. 1979. Genetic methodologies for the study of speech. In Xandra O. Breakefield (ed.), *Neurogenetics: Genetic Approaches to the Nervous System*, 311–343. New York: Elsevier.

Jenkins, Lyle. 2000. *Biolinguistics: Exploring the Biology of Language*. Cambridge: Cambridge University Press.

LeMay, Marjorie & Norman Geschwind. 1975. Hemispheric differences in the brains of great apes. *Brain, Behavior, and Evolution* 11, 48–52.

Legate, Julie Anne & Charles D. Yang. 2002. Empirical re-assessment of stimulus poverty arguments. *The Linguistic Review* 19, 151–162.

Lenneberg, Eric H. 1967. *Biological Foundations of Language*. New York: John Wiley.

Morgan, Michael J. & Michael C. Corballis. On the biological basis of human laterality: II. The mechanism of inheritance. *Behavioral and Brain Sciences* 1, 270–277.

Pennisi, Elizabeth. 2002. Evo-devo enthusiasts get down to details. *Science* 298, 953–955.

Pinker, Steven & Paul Bloom. 1990. Natural language and natural selection. *Behavioral and Brain Sciences* 13, 707–784.

Plotkin, Henry. 1998. *Evolution in Mind: An Introduction to Evolutionary Psychology.* Cambridge, MA: Harvard University Press.

Pullum, Geoffrey K. & Barbara C. Scholz. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 19, 9–50.

Vargha-Khadem, Faraneh, Kate Watkins, Katie Alcock, Paul Fletcher & Richard Passingham. 1995. Praxic and nonverbal cognitive deficits in a large family with a genetically transmitted speech and language disorder. *Proceedings of the National Academy of Sciences USA* 92, 930–933.

Yeni–Komshian, Grace H. & Dennis A. Benson. 1976. Anatomical study of cerebral asymmetry in the temporal lobe of humans, chimpanzees, and rhesus monkeys. *Science* 192, 387–389.

Lyle Jenkins
Biolinguistics Institute
103 River St., Suite 3
Cambridge, MA 02139
USA
ljenkins2@comcast.net

★ FORUM ★

Prospects for an Explanatory Theory of Semantics

Wolfram Hinzen

I

This essay reports and reflects on a workshop on *Biosemantics — Semantics within the Biolinguistic Program*, organized by Boban Arsenijević and myself, which took place at Leiden University from 10 to 12 September 2008 and was funded in the frame of the 'Origins of Truth' project by the Dutch NWO. The reader can hardly expect that I will now review the state of the art on semantics within the biolinguistic program: Truth be told, there is no field here with a state of the art to report. Do we know even the rough outlines of what an explanatory research program on the origins of human semantics looks like? Everyone present at an informal pre-workshop meeting agreed, in particular, that one aspect of the problem was beyond the scope of current research: the fact that there is semanticity at all — the problem of 'intentionality', as philosophers have dubbed it. Syntax allows us to generate complex descriptions of the world (or possible worlds) — but the very term 'description' is an intentional one. Syntactic structures actually *mean* something — we can use them to refer to real or possible worlds in highly systematic ways. This is something that syntax as we know it does not explain. This, indeed, is the business of semantics (jointly with pragmatics, whatever the difference between these two may amount to). We can *formalize* semantics, to be sure. Yet, this won't tell us why it exists or why it does what it does.

That said, I think it is fair to say that the workshop did isolate some questions that we can put on our research agendas. I will summarize these in the final section. Before that, I will report on some of the talks and the discussions that took place. I will embed this within some stage-setting on what I am calling an explanatory theory of semantics. I should warn upfront that my account is necessarily selective and coloured by my perceptions, if not simply by what I understood.

II

Under the broad umbrella of the biolinguistic program, syntax has clearly held centre-stage. This is, presumably, because syntax as a domain of inquiry fits into

I am very grateful to Boban Arsenijević and Dora Achourioti for perceptive comments on an earlier version of this review.

biolinguistics
© The Author (2008)

ISSN 1450–3417

Biolinguistics 2.4: 348–363, 2008
http://www.biolinguistics.eu

the mould of questions that we know how to tackle. It's the 'easy' part. In regards to syntax, in any case, the biolinguistic program now appears relatively well-defined and the basic agenda can be summarized as follows: Minimize appeals to a genetically defined Universal Grammar, don't proliferate 'modules', maximize appeals to domain-general principles in linguistic organization, from learning algorithms of a non-linguistic sort to general principles of design economy and geometry (e.g., symmetry breaking, prominent in Boeckx's and Di Sciullo's presentations). In short, the basic parameters of inquiry are set, and daily work can take the form of relatively technical, problem-solving routines, as in analyzing derivations using the apparatus of Merge, Move, feature-checking, and, most prominently, phases: This is the equivalent, in linguistics, to Kuhnian 'normal science'.

That said, what *exactly* the domain of syntax involves is quite unclear. As Boban Arsenijević pointed out in his talk, the minimalist operation Merge, given its austerity and utmost generality, may be a bad candidate for what is specific to language or grammar: Rather, grammar may arise as a specific kind of *restriction* on the operations of a basic domain-general computational system based on Merge that subserves cognition more generally. This important question resonated in Boeckx's and Roeper's talks. How do we characterize syntax — as unrestricted Merge or as a particular kind of restriction on it, due to demands of the interfaces?

It also resonated in Zeijlstra's presentation, in whose model syntax clearly isn't the universal cognitive machine, but the real, restricted, and linguistic syntax whose inherent job is to mediate between sound and meaning. We therefore expect its elegance and minimality to be compromised by having to serve two interfaces, which may impose conflicting demands, as well as possible. Here, syntax is *not* a general capacity to generate structures that go with a systematic form of semantics, but a much more idiosyncratic system that has to be rationalized as a solution that meets a very specific design problem in the context of the externalization of cognitive and conceptual structure in a physical medium. If, on the other hand, one thinks of syntax as the very format of human thought (Leiss, in press), one privileges the semantic interface over the phonological one, and one's account of syntax will tend to be different.

A more specific aspect of this broader issue of the interface between language and cognition transpired in Uriagereka's presentation, which asked: How much of linguistically familiar semantic territory does syntax carve out (let's call this 'Uriagereka's Question'; see Uriagereka 2008)? For example, does syntax have to be organized in such a way that it derives — or architecturally reflects — the dimensional layering of concepts as characterized, say, in Vendler's work on verbs? Are Vendler's distinctions, intrinsically, syntactic cuts, as well as, or over and above, semantic ones?

As Hinzen pointed out in his opening remarks on the 'prehistory of semantics', the philosophical and linguistic answer to Uriagereka's Question for nearly 700 years has arguably been: by and large, nothing. Ever since the nominalists under Occam's lead demolished Modistic Universal Grammar in the early 14th century, language was largely regarded as an arbitrary means to express thought: It is deprived of any intrinsic relation to either the structure of

the thoughts conveyed or, for that matter, the external world out there, which language can be used to describe. It is a tool for conveying thought, not its cause. If so, language is deprived of the role it was given by the Modists: that of an instrument of knowledge, a format in which systematic knowledge of the world is possible (see Leiss, in press for a reconstruction of the history).

When Cartesian rationalists resurrected the Universal Grammar project, the nominalist view of language as an arbitrary device for externalizing independently constituted 'thoughts' was kept: Language is still not the format of a specifically human form of knowledge. Neither did this axiomatics change when Chomsky resurrected Cartesian linguistics and syntax became 'autonomous'. As many have understood this claim of autonomy, syntax is an arbitrary formal system, unmotivated in terms of semantics or thought. Even the latest minimalist architectures of grammar reiterate a version of this basic axiom, when syntax is explained as an optimal solution to the 'conditions imposed by thought', a view that entails the independent constitution of thought, which thus can impose 'conditions' to which language is somehow answerable. Syntax, again, therefore, is not instrumental in formatting thought: It is still merely a way of expressing it (though, now, optimally).

The same axiom is found in the opponents of generative syntax, such as cognitive linguists taking the notion of a 'communicative intention' as a starting point of linguistic inquiry, or in the 'language of thought' tradition (Fodor 1975, Pinker 2007). Here thought is again an independent system. Despite the term 'language of thought' it is not structured by natural language, or the basic forms and mechanisms that Universal Grammar provides our minds with (Fodor 2001). Michiel van Lambalgen, at the workshop, provided evidence for a related view, arguing that non-linguistic logical processing entering the construction of a discourse-model is prior to the evolution of language. This makes linguistic organization an instance of cognitive structuring rather than the format of a specifically human form of knowledge.

Taking another viewpoint, Hinzen suggested turning the tables against the nominalists and seeing the computational system of language as the engine of specifically human, propositional thought. Syntax provides the forms of a possible human thought (cf. Hinzen 2006). Effectively, syntax therefore *is* a theory of semantics. If syntax does not constrain the space of semantics — if it does not define the concept of a possible meaning, in the same way in which Universal Grammar defined the concept of a possible language — what could? How, why, and by what, would semantics be constrained in the first place? (It had better *be* constrained, if there is to be an explanatory theory of it.) Where the axiomatic foundations for semantics are broadly those of Fregean philosophy, this is hard to even conceive. What would it mean to say that the abstract realm of Fregean 'propositions' (the realm of 'content'), would be empirically constrained? Propositions are not meant to be mind-dependent objects, and neither are they meant to be subject to constraints coming from the physical world in some broader sense. Some might want to insist on the archaic character of these views, but then, it is unclear what restricts the format of human thought on other views. Human semantics obeys a highly specific format, including the 'duality' of semantics (thematic structure vs. discourse structure) or the duality of

semantic evaluation (reference vs. truth): From which general, non-linguistic constraints might these aspects derive?

The guiding force behind 20th century philosophy was naturalism, involving an opposition to Frege's metaphysics. For semantics this meant, in one version of the naturalist story, that Fregean propositions were stipulated to be necessarily mediated by 'mental representations' (Fodor 1975). Unlike propositions, the latter are definitely internal to heads and brains, hence would seem subject to physiological and physical constraints. Yet, this is hardly how they have been looked at. The 'representational' theory of mind was embedded in the doctrine of functionalism, which, in its characterization of mind, purposefully abstracted from the way in which minds actually depend on brains (see Hinzen 2006). Mental representations therefore were not studied empirically. Their syntax was a logical syntax, their semantics a stipulated, Tarskian recursive definition of truth. The mental representations were thus not empirically but functionally motivated, as a device for representing the propositions or contents that were inherited from the 19th century. A representation in this sense is what it is, precisely because it relates to something external to itself, its meaning or 'content', of which there is no constrained theory at all. A categorical syntax-semantics separation, unproblematic in the logical tradition and arguably even necessary for purposes of logic, was thus kept as a foundational axiom in the analysis of mind. Having been evicted from the 'Third Realm', Fregean propositions were alive and well. Their fate was to be re-located in the domain of 'denotations'. 'Semantic values' is how they were now called.[1]

None of this is easy to square with the development of our understanding of syntax in linguistics. We have gone a long way from the 'autonomy of syntax' as understood by some in the 1950s and 1960s. Now syntax is seen much more as a system that intrinsically subserves the purposes of semantics. 'Full Interpretation' at the semantic interface, for example, is a prime principle of Minimalist Syntax. Positing a new head in the hierarchy of the clause is constrained by its contribution to semantic composition. The most basic architectural elements of grammar, such as phases, characterize units of computation as having an identity that is syntactic as much as it is semantic (Chomsky 2007). As Marantz (2000) points out, domains of 'special meaning' — where meaning, prior to becoming compositional, has some wiggle-room and allows for idiomaticity — are syntactically defined (e.g., the 'internal domain of *v*'). Syntactic definitions of so-called lexical categories (e.g. Baker 2003) will naturally incorporate semantic information. Event structure can largely be read off syntactic configurations (see e.g. Hale & Keyser 2002 and Marantz 2007). Quantification and scope are mediated syntactically (Huang 1995). A categorical syntax–semantics separation on purely conceptual or philosophical grounds, in the light of even these few instances of recent theorizing, makes little sense.

One may even make the stronger case, as Chomsky (2007) does, that the

---

[1] As the movement towards 'embodied cognition' gained momentum in the philosophy of mind, the idea of propositional mental representation in the relational sense above was abolished, but so were the genuine insights in the generative tradition on the highly systematic and structured character of human linguistic cognition.

very reason that operator-variable interpretations of *wh*-expressions arise, is by virtue of the interaction of syntactic principles and configurations: Within *v*P, the *wh*-element is interpreted thematically; given an appropriately minimal theory of movement, the copy theory, plus an account of phasing, another configuration will then necessarily arise in which two copies of the same lexical item appear in two phases, with the lower copy interpreted as a variable and the higher as an operator. This, it needs to be admitted, is a mere correlation between logic (semantics) and syntax. Yet, engaging in such correlation research may be our best bet at the moment, if we are interested in the origin of a logical mind capable of (binary) quantification.

This would solve a problem of ontology, which we face in semantics, but not (or less so) in syntax: We would make progress, as Uriagereka described it at the workshop, on the 'naturalization of meaning'. The ontological problem in question is simple enough to state. What *is* meaning? Linguistic *form*, while highly abstract and relatively removed from the visible surface of language, is manifestly there. Semantics, by contrast, to the extent that it doesn't trivially reproduce aspects of linguistic form, is entirely abstract. Linguists tend to have intuitions as to when something is a 'semantic question' or a 'syntactic' one, but this doesn't answer the ontological question just posed — a metaphysical question, essentially. To be sure, meaning is 'real' enough, an inherent part of the scheme of things, an aspect of nature. Yet it is intangible in a way syntax is not.

What proves this intangibility is the persistent attempt of philosophers throughout the 20th century and even today to get rid of it — say, by redefining semantics so as to reduce meaning to causal relations between words and objects (the 'causal theory of reference', coherent with the doctrine of physicalism, and first proposed by Skinner in 1955), or by denying the very reality of meaning, be it through elimination (Quine 1960) or claims of radical indeterminacy (Davidson 1984). No matter which of these paths we take, a 'science' of meaning will be a conceptually impossible enterprise. The validity of the enterprise of a 'science of syntax' in Chomsky's sense has been routinely questioned on similar grounds. Yet, in this instance, the arguments seem somewhat easier to refute (Chomsky 2000: chap. 3).

Syntax as a theoretical form of inquiry arises with the insight that human grammar, surprisingly, has a formal structure to it, which can be empirically studied and has a contingent character: It does not follow from any logical necessity. Hence there is such a thing as a theoretical 'science' of syntax. A theoretical 'science of semantics' in a similar sense is not yet in view, I claim — which perhaps is why semantics, unlike syntax, has primarily been a sub-chapter of logic, and is taught as such in just about any philosophy of language course around the world — courses which, by contrast, often don't as much as touch upon syntax.

III

When late medieval grammarians around 1300 first tried their luck on a 'scientific', and hence universal grammar, their focus of attention was on the parts of speech qua conditions for the possibility of syntax. This was for a good

reason: Consider, as they did, the difference between *curr–ere* ([*v*–RUN]) and *cur–sus* ([*n*–RUN]), or better, between *dol–or* ([*n*–PAIN], *dol–eo* ([*I* [*v–feel* PAIN]], *dol–enter* ([*a*–PAIN], 'painful'), and *heu!* ('Pain!'). Clearly, while there is an element that is identical throughout these sequences, namely the roots √RUN and √PAIN, respectively, the semantics of these roots — good candidates for conceptual 'atoms' in Fodor's (1998) sense — does not explain or predict the existence or function of any of the above parts of speech into which these roots enter. Nor would they predict or explain that these parts of speech enter into syntactic configurations in systematic ways. Presumably, non-linguistic beings could have a concept like PRUN, while being incapable of grammaticalizing it in a way that they can systematically distinguish between *run–s*, *a run*, or *run(n)–ing*, and deliberately switch between these three systematically and categorically different perspectives on what is in some sense the same thing, and in the very same external or environmental circumstances.

In short, the principle that allows what the Modists called *constructio*, i.e. syntax, and thereby systematic forms of knowledge and inference, is not meaning, concepts, or denotations. Syntax is the way to warp us out of the immediacy of our experience with the outer denotations of the words we use. Quite a different kind of meaning, call it grammatical meaning, kicks in the moment we categorize concepts and insert them into configurations. At this stage, the rules of combination are not sensitive any more to the semantics of lexical roots: It is categorial information that drives them. There seems to be no way to track this kind of meaning denotationally — there is no outer physical correlate of verb phrases or propositional configurations, for all we know on empirical grounds. This is part of the problem with a 'referential semantics' as Chomsky (2000) has exposed it: To whatever extent meaning is internally determined, by principles of grammar, external relations between words and things viewed as mind-independent won't illuminate it. Yet, the way in which meaning follows systematically from the compositional rules of language — I-meaning, in short — presumably is what a science of meaning is primarily about.

IV

Cedric Boeckx' presentation, on 'I-semantics', addressed precisely this issue. Can we, once we see syntax constrained in particular ways, somehow see semantics falling into place? To sketch his account, say we start from a pure 'theory of form': All there is, let us suppose, is unconstrained Merge, triggered by a lexical item's having an edge-feature ('Wild-type Merge'). The operations of this function disregard anything else: configurations as much as theta-theoretic restrictions, what 'makes sense', etc. The system's austere beauty consists in its libertarian tolerance of, in fact blindness to, nonsense. "The rule in the Library", as Boeckx quotes Borges, "is not 'sense', but 'non-sense', and 'rationality' (even humble, pure coherence) is an almost miraculous exception". Such a beginning seems maximally unsuited to tell a story about the origins of semantics, until the moment that we see syntactic derivations to be phased. Phases, by definition, introduce asymmetries that acquire interpretive significance: For example, the asymmetry between the phasal head, which remains in the derivation and retains

its edge-feature, and its complement, which is transferred to the interpretive components when the phase boundary is reached.

The specific compositionality of semantics, as Hinzen and Roeper noted at the workshop, maybe is a consequence of phasing in this very sense: The meaning of compositionality is that there are constituents in the derivation which have an independent interpretation, which is not changed as the derivation proceeds. Phasal transfer has this consequence: encapsulating a unit of interpretation. A spelled-out item loses its edge-feature — hence the syntax has to treat it as an unanalyzable ('frozen') unit. Phasal transfer, however, must not be too early: For example, not before a head is merged that can serve as a landing site for an argument; and it must not be a transfer of the whole phase but only of its head's complement. The right solution, Boeckx suggested, is that Merge must apply externally twice; only this creates a rhythm of alternating phasal heads and non-heads that corresponds to the relevant interpretive asymmetries: argumenthood (the D-phase), thematic structure (the *v*-phase), and propositionality/truth (the C-phase). And those indeed, are the basic cuts that we see correspondences of when we look at neo-Davidsonian (Pietroskian) event-representations (Pietroski 2005). In short, as the syntax develops a rhythm of periodic phase-non-phase alternations, semantics as we know it swings into place.

The austere purity of Boeckx' architectural model of syntax raises fascinating questions, not least about how we account for syntactic variation: not at all (because there isn't any, in syntax), or by exploiting the rhythm of phasal heads and non-heads just described, and compromising the rigidity of this alternation somewhat, so as to allow for variation in how many non-phase heads can be inserted between two phasal ones. This, as Ángel Gallego argued at the workshop, takes care of some of the comparative data to be captured (see also Uriagereka, in press).

V

In a distinctly different tone, Michiel van Lambalgen formulated an argument against the whole idea of turning to syntax for explaining human semantics. Lambalgen presented EEG data supporting the view that semantics is not fully dependent on syntax, if it is so dependent at all: It is, in particular, richer, in that semantics systematically depends on non-linguistic sources of knowledge and involves processes of non-monotonic inference and causal reasoning that the syntax as such does not support. The view has become prominent in cognitive neuroscience, given the existence of neuroimaging data demonstrating the fast and seamless integration of non-linguistic with linguistic information in linguistic comprehension. Saying *I am an Aristocrat* in a Cockney accent will not easily lead a hearer to incorporate a new truth into his discourse model, for example (see Hagoort *et al*. 2007). On this view, the very notion of semantics, in the traditional sense of a syntax-driven 'interpretive' system in which compositional information is read off the syntactic representation of an expression, has come into question. Rethinking tradition here proves valuable in accounting for the origins of semantics.

Like everything else in evolution, any 'semantics module' needs precursors. Thinking of language comprehension as the process of updating a discourse model (rather than merely as a module for interpreting syntax and mapping it to truth conditions), van Lambalgen argued that a precursor to this can be found in the systems of executive control that both human and non-human animals exercise when planning a sequence of actions so as to achieve a particular goal in possibly adverse circumstances. This predicts the co-occurrence of deficits in planning and in discourse processing, and it is what we find, in both autism and ADHD. If making a plan is analogous to constructing a discourse model, and language comprehension is that, too, a biological substrate is found, in a non-communicative domain, from which language can develop. A planner doesn't necessarily talk, yet any model of a discourse that we construct will consist of events and some relations between them — as do plans. The hypothesis can then be that language in the sense of communicative medium evolved through the externalization of planning driven by a need for interactive as opposed to individualistic planning.

Let's consider one of van Lambalgen's examples, *The girl was writing a letter*, which, on the view under discussion, leads to the computation of a discourse model in which there exists a completed letter. If this is part of the semantics of this sentence, however, then semantics will include interpretive processes that may *or may not* co-occur with a particular linguistic expression. Clearly, all that the Tense of this sentence specifies is that the relevant writing event was in the past with respect to the point of speech, and all that the progressive Aspect adds to this is that the event was ongoing at this time. That letters are more often completed when being written may or may not be a statistically accurate observation. The compositional core of the meaning of our sentence, its 'grammatical meaning', doesn't seem to be affected by this (any more than the compositional meaning of *I am an Aristocrat*, above, is affected by the statistical fact, more credible in this instance, that Aristocrats are only very rarely accomplished speakers of Cockney).[2] Syntax-driven semantic effects are *not* defeasible or default expectations, however, in the way that the completion of a letter is. That the subject of 'brought' in the sentence *Who did Mary say brought her the chips* is controlled by *who*, not by *Mary*, say, is an interpretive consequence of sentence-internal syntax which has nothing to do with discourse or pragmatics, for example. The question therefore is whether the enriched conception of semantics misses a distinction between interpretive processes, if not a whole area of inquiry revolving around the question of which semantic effects are determined by linguistic form, and how.

Here is what's needed to refute a hypothesis of syntax-semantics alignment for those forms of semantics that *are* fully systematic: evidence of a semantic process that is systematic, monotonic, and *un*supported by syntactic laws.

---

[2] The following discourse felicitously describes a possible situation: 'The girl was writing a letter when her friend spilled coffee on the paper. The ink proved coffee-resistant and the letter was completed in no time.' We can't tell, from language, how resistant ink is, and therefore we need world knowledge as encoded in a situation to complement the information encoded in the syntax alone; yet, that won't affect either syntax or its compositional interpretive effects.

Semantic coercion, as in the forcing of an event interpretation of the verbal complement in *John began the book*, has been thought to be such a candidate; yet, whether this is so, is not clear from either the linguistic or the neuroimaging data available (see Pylkkänen *et al*. 2008). In the meantime, the appropriate way to proceed seems to be to approach the syntax-discourse from both sides: asking how syntax subserves the embedding of propositional information in discourse, and how van Lambalgenian processes of discourse model construction tune in. As Valentina Papa argued at the workshop, for example, a thorough analysis of elliptic comparatives reveals that relevant missing constituents in syntactic form may need to be retrieved from context, or be added inferentially. Compositionality in the syntax-semantics mapping is accordingly compromised. Clearly, the task ahead is a principled account of just when, and when not, in the construction of meaning, context can or even must be accessed. The extreme answers, 'never' and 'always', seem equally implausible. Some meaning is determined by linguistic form, some not.

VI

Formulating detailed mapping hypotheses between syntax and semantics that make this particular mapping come out as a motivated one is an instance of this broad task. In a perfect world, the semantic interface should at least be *transparent*. But syntactic structuring should be *meaningful* in regards to how it is interpreted beyond merely allowing for transparent compositional mappings. This basic intuition — that the semantic interface should be well-behaved in some ways — can and has been very fruitfully explored. Ideally, *nothing* should be *non*-interpretable on that side of the grammar, and Kayne (2007) has argued that this is indeed the case (though Zeijlstra's talk at the workshop entailed the opposite stance). That syntactic structures are not only mapped compositionally into semantic structures, but in a particularly 'strict' way, with heads only depending for their composition on their respective direct complements, has been argued by Larson & Segal (1995). As Uriagereka pointed out at the workshop, this aspect of linguistic architecture can be rationalized as a way of facilitating the learning of new lexical items that appear in such strictly compositional structures.

Using a different approach, Ronnie Wilbur showed that syntactic decompositions of event structures may be transparent not only in semantic respects but even phonologically, namely in sign languages, as per Wilbur's Event Visibility Hypothesis. Evidence was presented from American Sign Language that externalization may here reveal morphological structure unrecognizable in spoken languages. This is progress towards the 'embodiment' of semantics. Speaking on the same issue of transparency, Heather Burnett argued that there is no need to assign different semantic types to the same quantificational expression in cases of 'Quantification at a Distance' in Québec French, as long as one takes into account the *phase* within which they are merged. If so, semantic content and syntactic position align again. These scattered examples all appear to point us towards a deeper understanding of the semantic interface, according to which that side of the grammar is rather well-behaved indeed.

Uriagereka also developed a novel argument that compositionality in the syntax–to–semantics mapping is not in fact the point. It can't be, because one could apply the relevant composition functions top-down rather than from the bottom-up, reversing them (e.g., we could start by composing Agents with verbs, rather than Themes). However, there is a significant *asymmetry* between (for example) Agents and Themes, leading us to 'first-merge' Themes rather than Agents. In short, there is a derivational *directionality*, with Themes coming earlier, and Thematic Structure as such coming earlier than discourse-relevant structure, giving rise to crucial asymmetrical interpretive dependencies. These asymmetries, again, directly inform the child's learning, now not only in terms of what composes with what, but also in terms of what composes first. One way, therefore, of doing the compositions of meaning is syntactically *viable*, while another is not, and this speaks to the concern of 'naturalizing' semantics.

Proceeding on this course, Uriagereka showed that we can naturalize other aspects of semantics — in fact, we can do so in areas where one might have expected this least, such as the rigidity of reference, which Kripke (1980) still grounded in metaphysical and externalist considerations. Arguably, though, rigidity follows from no aspect of the external environment, or 'semantics', at all. Viewed as an aspect of linguistic form, and specifically of how names function, there needs to be some internal mechanism that achieves this effect. Crucial to Uriagereka's naturalization of it is an important distinction between atomism, in the sense of Fodor (1998), and rigidity. The former is due to the mechanism of lexicalization, which consists in the idiomatization of a complex syntactic structure and thus the creation of an encapsulated item that, unlike most sentences, is listed in the lexicon. Rigidity, by contrast — arising, unlike atomism, at the outset of language acquisition rather than its final stages — is embodied most crisply in pointing, and thus it needs to be due to something else. What is it about names that they have got rigidity in addition to atomism, while nouns don't, and how can sentences be rigid even though they lack atomicity? Enter phasal Transfer: Periodically, as with Boeckx' model, descriptive information contained in the derivation is discarded, or lost. The way to think about this, Uriagereka suggests, is as the elimination of parameters in an algebraic equation: Taking them away does not allow us to independently vary different descriptive factors any more, and the inflexible or 'flat' object resulting, complex as it may be (in the case of sentences), cannot refer but rigidly: It cannot *adapt* any more to the specific features of an object, as required when referring descriptively. In the case of names, this Transfer is induced by the very lexical nature of names.

VII

Remarkably, no one at the workshop even raised the issue of semantic variation, or parameterization. Meaning as configured in language may be so abstract that it simply cannot be subject to cross-linguistic variation, any more than, on some current minimalist views, the computational system of language is. But, if this is so, one wonders where semantics comes from. Again, if syntax is not its format and cause, what is? Picking up on the task of distilling semantic effects of linguistic form, Anna Maria Di Sciullo asked the bold and topical question of

why sentences have truth values whereas words do not. Her answer was a purely structural one: Intrinsically, a word is a unit of morphological organization that as such is structurally too restricted to support propositional information. We could imagine a creature that had words only, but no phrases or sentences. Such a creature would have Merge, it would have compositionality, it would understand scope relations, there would be uninterpretable feature checking, and its words would be structured by the three broad categories that structure the human clause: predicate-argument structure (e.g. *writ–er*), aspectual modifications (e.g. *re–write*), and operator-variable relations (e.g. *th–*, *wh–*, *–ed*). This creature would equally grasp asymmetry and its intrinsic role in semantic interpretation; in fact, as Di Sciullo argued, it would have it in a stricter form than syntax does, which has symmetries as well, as in predicational [XP YP] configurations.

Yet, for all that, our morphological creature would not refer to the world in the way we do, and it would have a discourse interface of a radically different sort. This we see from the fact that symmetries in syntax, where they arise, are resolved quickly through movement, as per Moro (2000). As movement correlates with discourse semantic properties as opposed to thematic structuring, that is not something that the domain of the word will therefore support. It also follows that words will exhibit a radically impoverished contribution to information structure. Neither can words be arguments or exhibit referential specificity: As noted already in Di Sciullo & Williams (1987), words remain 'generic' in meaning. They cannot locate a specific predicate in time (cf. **At six, John is a writer*; and neither can *bank robber* be used to talk about an act of *robbing a bank*), we cannot extract from a word (*John admires Nixon* allows for *who does John admire*, but *John is a Nixon-admirer* does not allow for **Who is John an admirer*), a form of referential opacity that also transpires in the inability of anaphoric pronouns to pick up on the external reference of parts of words (**Book-shelving I like, especially if they [the books, or the shelves] are not heavy*); see Di Sciullo (2005) for more discussion.

Morphology, in other words, is a possible language in its own right, but it lacks what philosophers at least since Russell have paradigmatically associated with words, namely reference. It takes the fully expanded structure of the clause, specifically [ForceP [EvalP [CP …]]] on Di Sciullo's model, for structures to qualify for a mapping to a truth value. This can be viewed as a formal specification of the pragmatic interface of the grammar; it is a structure that yields our 'sense of truth'. The asymmetry we see between words and sentences in regards to truth value evaluation reduces to a structural difference, an aspect of linguistic form. This is an internalist account of truth, in the sense of Hinzen (2006, 2007), also continuing the tradition of an internalism about meaning and the pragmatic determination of truth values, as characterized in Paul Pietroski's work (Pietroski 2005b).

While not addressing the issue of truth directly, Roeper and Hollebrande's take on recursion at the workshop may well stand in an intimate relation to it. Note that recursion as such is simply a formal property of human grammars that is just as 'autonomous' with regards to interpretive asymmetries as compositionality as such is: In fact, recursion as a mathematical property need

not as such intrinsically relate to semantics at all. And yet, as featuring in human language, it does, as Roeper showed: Recursion, in language, it turns out, is inherently a way of organizing and constraining semantic information (cf. also Hinzen 2006, 2007). Interpretive options are more open, if, rather than choosing for embedding one sentence in another, we let the two sentences form a discourse.

As Roeper and Hollebrandse amply documented, recursion is a construction-specific and language-specific phenomenon. Germanic languages have recursive nominal compounds, Romance languages lack them. English has recursive possessives, German, Swedish, and Dutch lack them. Bantu languages have recursive serial verbs, English doesn't. Most importantly, recursion, in whatever forms it does exist, is highly restricted. Most obviously, it is only particular domains of syntactic organization that productively 'recur' at all: Specifically, the cycle, or, in other terminology, the phase, is what can occur in itself. We don't get a recursive embedding of T in T, say, obtaining a hierarchy of the form V–*v*–T–T–T…C, or a recursion of *v*–V, as in V–*v*–V–*v*–V–*v*–T–C. Neither does C recur, directly. Any C–C–C sequence is in fact mediated by full phasal expansions in between: [C–T–*v*–V [C–T–*v*–V [C–T–*v*–V…]]] (Arsenijević & Hinzen 2007). In the nominal domain, as well, DP doesn't embed DP, unless something else intervenes (**Bill's knowledge Fred's knowledge John's knowledge is incorrect*), as Roeper pointed out. *Within* a phase or cycle, nothing recurs at all. So recursion depends on the cycle, and if a cycle is thought of as a phase, then it is not even the case, strictly, that a phase occurs within a phase: For, as a new phase starts off, a large part of the previous one is already transferred, hence it is never in fact the case, at any one derivational time, that a structure contains two phases embedded in one another.

Restrictions on recursion can also be of a theta-theoretic sort (*the city's destruction by the enemy* or *the enemy's destruction of the city*, but not **the city's enemy's destruction*), or they can reflect on the extent to which embedded propositions can be evaluated (cf. **John considered Bill to consider the food to be tasty*; **John knew Bill to know Fred to be a liar*). Clearly, the task here is a typology of recursive constructions, and to test the claim that recursion, wherever it happens, is 'indirect', occurring via a phase boundary, as hypothesized by Roeper and Hollebrandse.

The issue of constraints on recursion is interestingly connected to another issue raised by Andrea Moro in discussion after Hinzen's talk. Moro (1997: 198) noted a curious fact about English. Consider the pair of sentences *John is sad* and *it's that John is sad*. On Moro's analysis, the structure of the latter sentence is the following:

(1) $[_{IP}$ it$_i$ is $[_{SC}$ $[_{CP}$ that John is sad] t$_i$]]

That is, *it* is a raised pro-predicative element in what Moro terms an inverse copular construction, which in this case is interpreted as involving the assignment of a sentential predicate to the embedded proposition. Basically, this is the truth-predicate, given that *it's that John is sad* means *It's that it's true that John is sad* or *It's the fact that John is sad*. That said, it is interesting that the

construction in question does not recursively iterate: **It's that it's that John is sad*. This would be predicted if Hinzen was right that the evaluation of a sentence for truth is, in fact, the last thing that can happen to it: After that, recursion and compositionality stop. No truth-evaluated sentence (i.e. a declarative sentence whose truth is being asserted) can occur embedded: In *John believes Bill behaved suspiciously*, for example, the embedded clause must not be evaluated for truth, if the whole sentence's truth value is being determined.

As Roeper phrased this in discussion, truth may be like specificity. If something is referentially specific, it doesn't embed in something that's also referentially specific. If something has obtained full specificity, this is the end of it, and recursive structure-building stops. Unlike the *word* 'specific', which is a compositionally interpreted constituent of a larger structure, specificity is not, and neither is truth. Interestingly, if we say that it's true that it's true that John is pale, and then that it's true that it's true that it's true that John is pale, recursion is like an idle wheel: There is no semantic effect at all, virtually, and 'it's true' behaves like an identity function mapping one object back onto itself.

VIII

A final issue in any future 'biosemantic' enterprise, which Andrea Moro urged should be pursued relates to the question of what must a brain be like in order for it to process propositional, truth-evaluable information? This question depends on our ability to map propositions to brain space in the first space. New neuroimaging data that Moro presented (Tettamanti *et al*. 2008) suggest that we may be approaching a time where we can study the assignment of truth values to propositional units at a neurological level, hence a specifically propositional form of semantics as opposed to a lexical or conceptual one (which is rather standardly targeted in current neuroimaging work on 'semantics'; see again Pylkkänen *et al*. 2008). Arguably, negation is, inherently, a propositional operator, which reverses truth values. Using functional magnetic resonance imaging, brain activity was measured during passive listening of sentences characterized by a factorial combination of polarity (affirmative vs. negative) and concreteness (action-related vs. abstract). It turns out that sentential negation transiently de-activates a left-hemispheric action-representation system normally correlating with the processing of the negated, action-related sentence. If blocking access to the relevant area reflects the reversal of truth value, this study may provide for a fascinating inroad into caverns of the mind where propositionality and truth are housed. Indeed, on the assumption, defended by Moro, that negation is not available to non-linguistic species, this corner of our mind would be uniquely human as well.

IX

Form the above, it seems clear, at least philosophically, that there has been a major shift, away from locating 'semantics' in a relation between language and world, to locating it in the various *interfaces* that the grammar forms with extra-grammatical cognitive systems, particularly at the syntax-discourse interface. Within this frame, the following research questions seem particularly pressing:

(A) What exactly do the domains of syntax and semantics involve, and what is uniquely human about them? If on the syntactic side this is recursive Merge, what makes this operation so fundamentally restricted across languages? What, on the semantic side, are the basic combinatorial operations?

(B) How much of what is uniquely human about semantics depends on syntactic structuring? (Uriagereka's Question: How much semantic territory does syntax carve out?) For example, is the *word* a distinctive domain of semantic organization?

(C) On what path can the naturalization of semantics proceed? On what relations between the respective ontologies of semantics and of syntax does a natural science of semantics depend?

(D) Which aspects of meaning should (or shouldn't) involve access to extra-grammatical/mind-external factors like discourse contexts, inferences, reference to the external world or entities out there?

(E) Can we discard the idea of semantic parameters and variation?

(F) How can we make progress on the neuroscience of propositional and compositional semantic processing? What is it about the brain that makes it capable of propositional semantic processing?

Overall, it seems that these short remarks and reflections indicate a rich research landscape with an agenda that is relatively clearly defined and that calls for an open-minded community of interdisciplinary research.

References

Arsenijević, Boban & Wolfram Hinzen. 2007. Recursion as an epiphenomenon. Ms., Universiteit van Amsterdam.

Baker, Mark C. 2003. *Lexical Categories: Verbs, Nouns, and Adjectives* (Cambridge Studies in Linguistics 102). Cambridge: Cambridge University Press.

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

Chomsky, Noam. 2007. Approaching UG from below. In Uli Sauerland & Hans-Martin Gärtner (eds.), *Interfaces + Recursion = Language? Chomsky's Minimalism and the View from Syntax–Semantics* (Studies in Generative Grammar 89), 1–29. Berlin: Mouton de Gruyter.

Davidson, Donald. 2001. *Inquiries into Truth and Interpretation*, 2nd edn. (Philosophical Essays of Donald Davidson). New York: Oxford University Press.

Di Sciullo, Anna Maria. 2005. *Asymmetry in Morphology* (Linguistic Inquiry Monographs 46). Cambridge, MA: MIT Press.

Di Sciullo, Anna Maria & Edwin Williams. 1987. *On the Definition of Word* (Linguistic Inquiry Monographs 14). Cambridge. MA: MIT Press.

Fodor, Jerry A. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.

Fodor, Jerry A. 1998. *Concepts*. Oxford: Clarendon Press.

Fodor, Jerry A. 2001. Language, thought, and compositionality. *Mind and Language* 16, 1–15.

Hagoort, Peter & Jos J.A. van Berkum. 2007. Beyond the sentence given. *Philosophical Transactions of the Royal Society* B 362, 801–811.

Hale, Ken & Samuel Jay Keyser. 2002. *Prolegomenon to a Theory of Argument Structure* (Linguistic Inquiry Monographs 39). Cambridge, MA: MIT Press.

Halle, Morris & Alec Marantz. 1993. Distributed morphology and the pieces of inflection. In Ken Hale & Samuel Jay Keyser (eds.), *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger* (Current Studies in Linguistics 24), 111-176. Cambridge, MA: MIT Press.

Hinzen, Wolfram. 2006. *Mind Design and Minimal Syntax*. Oxford: Oxford University Press.

Hinzen, Wolfram. 2007. *An Essay on Names and Truth*. Oxford: Oxford University Press.

Huang, C.-T. James. 1995. Logical form. In Gert Webelhuth (ed.), *Government and Binding and the Minimalist Program*. Oxford: Blackwell, 125–175.

Kayne, Richard S. 2006. Notes on full interpretation. Paper presented at the *Inter-Phases* conference, Nicosia. [Castelliotissa Hall, 18–20 May 2006.]

van Lambalgen, Michiel & Fritz Hamm 2005. *The Proper Treatment of Events* (Explorations in Semantics). Oxford: Blackwell.

Larson, Richard & Gabriel Segal. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

Marantz, Alec. 2001. Words. Paper presented at *WCCFL XX*, Los Angeles. [University of Southern California, 23-25 February 2001.]

Marantz, Alec. 2007. Words and phases. Ms., New York University.

Moro, Andrea. 1997. *The Raising of Predicates: Predicative Noun Phrases and the Theory of Clause Structure* (Cambridge Studies in Linguistics 80). Cambridge University Press, Cambridge.

Moro, Andrea. 2000. *Dynamic Antisymmetry* (Linguistic Inquiry Monograph 38). Cambridge, MA: MIT Press.

Leiss, Elisabeth. In press. *Sprachphilosophie: Eine Einführung in die Axiomatik*. Berlin: Walter de Gruyter.

Pietroski, Paul. 2005a. *Events and Semantics Architecture*. Oxford: Oxford University Press.

Pietroski, Paul. 2005b. Meaning before truth. In Gerhard Preyer & Georg Peters (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.

Pinker, Stephen. 2007. *The Stuff of Thought*. New York: Penguin.

Pylkkänen, Liina, Rodolfo Llinás & Brian McElree. 2008. Distinct effects of semantic plausibility and semantic composition in MEG. Ms., New York University.

Quine, Willard Van Orman. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Skinner, B.F. 1955. *Verbal Behavior*. Acton, MA: Copley Publishing Group.

Tettamanti, Marco, Rosa Manenti, Pasquale A. Della Rosa, Andrea Falini, Daniela Perani, Stefano F. Cappa & Andrea Moro 2008. Negation in the brain: Modulating action representations. *Neuroimage* 43, 358–367.

Uriagereka, Juan. 2008. *Syntactic Anchors: On Semantic Structuring* (Cambridge Studies in Linguistics 118). Cambridge: Cambridge University Press.
Uriagereka, Juan. In press. Phase periodicity. In Jordi Fortuny (ed.), *Spelling Out Universal Grammar*. Barcelona: University of Barcelona Press.

Wolfram Hinzen
Durham University
Department of Philosophy
50 Old Elvet
Durham DH1 3HN
United Kingdom
wolfram.hinzen@durham.ac.uk

Notice

We would like to use this opportunity to thank all those involved in creating the first volume of *Biolinguistics* and all the work that went into preparing the entire endeavour over the past year and a half or so. For internal quality control, our gratitude goes to the reviewers that have served us in 2008, who are listed below. For everything else, we thank our supporters listed below as well as all active participation and feedback from the members of the *Biolinguistics* Advisory Board, the *Biolinguistics* Editorial Board, and the *Biolinguistics* Task Team that are not specifically mentioned by name.

Reviewers

David Adger
Boban Arsenijević
John Collins
Samuel David Epstein
Jordi Fortuny Andreu
Hans-Martin Gärtner
Ángel J. Gallego
Javier Guitérrez-Rexach
So-One Hwang
Johannes Krause
Winfried Lechner
Tommi Leung
Terje Lohndal
Gary Marcus
Hiroki Narita
Mcebisi Ntleki
Akira Omaki
David Ott
E. Phoevos Panagiotidis
Jeffrey K. Parrott
Charles Reiss
Bridget Samuels
Koji Sugisaki
Robert Truswell
Masaya Yoshida
Hedde Zeijlstra

Supporters

Noam Chomsky
John Davey
Anna Maria Di Sciullo
Norbert Hornstein
Terje Lohndal
Massimo Piattelli-Palmarini
Henk van Riemsdijk
Marinos Stylianou
Juan Uriagereka
Andreas Wadle
Andrew Winnard

and the University of Cyprus for financial support (especially the Department of English Studies)

ISSN 1450–3417
Biolinguistics 2.4: 364, 2008
http://www.biolinguistics.eu

Biolinguistics

Volume 2, Issue 4

Fall 2008

TABLE OF CONTENTS

260 The Asymmetry of Merge — Anna Maria Di Sciullo, *University of Quebec at Montreal*; Daniela Isac, *Concordia University*

291 Symmetry in Visual and Linguistic Perception — Brent Henderson, *University of Florida*

317 Minimalist Meaning, Internalist Interpretation — Paul Pietroski, *University of Maryland*

★ **REVIEWS** ★

342 Not the Only Word — Lyle Jenkins, *Biolinguistics Institute*

★ **FORUM** ★

348 Prospects for an Explanatory Theory of Semantics — Wolfram Hinzen, *Durham University*

364 Notice — *Biolinguistics* Editors

biolinguistics

ISSN 1450–3417

Biolinguistics 2.4: 260–364, 2008
http://www.biolinguistics.eu