PREFINAL DRAFT: Kirby, S. (1998). Fitness and the selective adaptation of language. In Hurford, J. R., S.-K. M. and C., K., editors, Approaches to the Evolution of Language: Social and Cognitive Bases. Cambridge University Press, Cambridge.

Fitness and the selective adaptation of language

Simon Kirby

September 4, 1996

1 Introduction

The question that is at the centre of this paper is how can we go about explaining the observed constraints on variation across languages — in other words, language universals.[1] What makes many of these constraints interesting is that they appear to have 'evolved' in that they are adaptive; a point that is made repeatedly in the functionalist literature. What I will argue here, however, is that we should not rush into a biological evolutionary explanation for such universals. Where language is concerned, adaptation may be the result of several different interacting dynamic systems.

The order of presentation will be as follows. The next section is an introduction to what it means for something to be adaptive. The following two sections will illustrate this with two processing and typological asymmetries connected with word order and relative clauses. A potential biological explanation will be put forward and rejected in favour of a historical explanation in the next two sections. This alternative explanation will be tested using a computational model of interacting agents using language over several generations. Finally, the implications of this approach for the phylogenetic evolution of language will be sketched in the final section.

The novel contribution of this paper is the suggestion that languages evolve *historically* to be optimal communicative systems, and that the innately specified human language learning mechanisms have evolved in order to learn these systems more efficiently. In other words, human language has *not* evolved (in the phylogenetic timescale) directly in response to communicative needs, even though the cross-linguistic evidence is that language is well adapted for communication.

2 The fitness of universals

I have suggested that many of the constraints on cross linguistic variation found in the typological literature (see e.g. Greenberg 1963; Comrie 1981; Hawkins 1988; Croft 1990) are adaptive. By this I mean that they have a striking *appearance of design*. Cziko (1995) refers to this feature of systems as the "puzzle of fit". Many attempts at explaining universals have pointed out the fit of universals to the communicative functions of language. Although these observations are important and insightful, I believe they leave the real mystery unsolved. Rather than explaining the origin of universals, this fit is in itself puzzling. Assuming that

---

[1] Much of the material for this paper is covered in more detail in Kirby (1996c), although the analysis of the relation between glossogenetic and phylogenetic adaptation is new.

languages don't have a designer how do they end up looking like they have been crafted with a particular purpose in mind? Why is it that, as we shall see, the constraints on relative clauses give an optimal balance between speakers' and hearers' needs in communication, for example?

Before answering these questions, we need to review briefly what we mean by "fitness". In general, the term can be used to describe the appearance of design in some structure whose formal properties correlate with its function. So, an umbrella's fitness derives from the fact that its canopy shields its user from rain and its handle and stem are arranged so that it can be easily carried. The formal and functional elements of the umbrella so neatly dovetail that we are led to expect that it has been designed.

Similarly, the term 'typological fitness' can be used to describe the property of universals whose formal structure dovetails with the functions of language. In this paper I will concentrate on one aspect of this fitness: the match of cross-linguistic asymmetries with asymmetries in language processing. Just as umbrellas cannot weigh more than a few pounds in order to be carried by their users, languages appear to be constrained cross-linguistically to be efficiently processed in real time.

There are many explanations for language universals given in the literature that make reference to this kind of fitness. For example, Cutler *et al.* (1985) aim to explain the cross-linguistic preference for suffixes (as opposed to prefixes) in terms of the way in which language is processed by hearers in real time. The crucial feature of this processing is that it is constrained by the left-to-right, serial nature of speech. The start of a word is clearly received by the processor before the end, and the assumption is that work starts on processing input as soon as it arrives. Simplifying their argument somewhat, Cutler *et al.* point out that early lexical access is preferred by hearers because placing of salient information early in the word aids processing. If lexical access is stem-based — as they argue from experimental evidence — then the tendency for languages to be suffixal matches the preference of the processor.

3 Processing asymmetries

For the examples in this paper I will be using the general theory of processing complexity put forward in Hawkins (1994a). This is useful since Hawkins is primarily concerned with using his theory to explain language universals, and so it provides us with an ideal opportunity to explore the origin of typological fitness.

Hawkins' theory is based on the idea of defining a measure of tree-complexity associated with a particular node in a constituent that is relative to a particular psycholinguistic operation. Any processing operation involves a subset of the tree structure of an utterance that is relevant to the operation — a *structural domain* in Hawkins' terminology. The assumption is that processing mechanisms that have to operate as quickly as possible prefer *minimal structural domains*. Two applications of this idea are discussed here: minimal constituent recognition domains, and minimal relativisation domains.

3.1 Immediate constituent attachment

One of the most important jobs of the parser is to attach constituents to mother nodes in the left-to-right processing of an utterance. Hawkins (1990) and the bulk of Hawkins (1994a) are concerned with the implications of assuming that the parser will prefer constituent orders whose immediate constituents (ICs) are attached as rapidly as possible for a particular mother once attachment starts.[2] Importantly, it is possible to attach all ICs to a mother based on a subset of the words dominated by that mother. This structural domain is the set of words that starts from the point in the string where the first IC dominated by the mother is constructed and ends where the last IC is constructed.

The relevant feature of the string of words presented to the parser thus becomes the order of the categories that construct immediate constituents. One type of category (and the most important for Hawkins) that is involved in construction is the mother-node constructing category (MNCC). These are categories which *uniquely determine* a mother node. These MNCCs are similar to heads in traditional syntactic theory, but may also include some closed-class function words such as determiners which uniquely construct noun phrases. So, for example, in the verb phrase *watched the cricket*, *watched* can construct VP, and *the* and *cricket* can both construct NP. A simple demonstration of the impact of reordering such categories can be had from some examples of Particle Movement in English. In the sentences below, there are three MNCCs which construct categories that are immediately dominated by the verb phrase: *looked*, *the* and *up*. As the distance between the first of the MNCCs and the last increases so too does the awkwardness of the construction. The last sentence is the most acceptable since the three MNCCs are adjacent.

(1) John **looked the** number **up**.

(2) John **looked the** number of the pub **up**.

(3) John **looked the** number of the pub on Rose Street **up**.

(4) John **looked up the** number of the pub on Rose Street.

This example gives a flavour of Hawkins' approach, although the details of his metric are rather more complex than there is room for here. We will return to IC attachment later when we compare this processing asymmetry with the cross-linguistic data.

3.2 Relativisation

Of course, IC attachment is not the only problem posed to the parser in attempting to build a representation of the structure of an input string. One of the other problems that Hawkins' looks at is that of associating a trace or resumptive pronoun with the head noun in a relative clause. Just as IC attachment involves a subset of nodes in the tree, the size of the structural domain for relativisation will vary depending on the type of relative clause and will typically be less than the whole of the clause.

Hawkins' (1994a:28–31) definitions for relativisation domains are given below:

Structural complexity of relative clause The structural complexity is calculated by counting the nodes in the *relativisation domain*.

Relativisation domain The relativisation domain consists of that subset of nodes within the NP dominating relative clause that structurally integrate the trace or pronoun.

---

[2] Attachment of ICs to a mother does not necessarily happen straight away. A constituent may be recognised on its right boundary, for example, in which case a look-ahead buffer for non-attached constituents is used. For the details of Hawkins' approach refer to one of the cited texts.

Structural integration of a node X in C The set of nodes which structurally integrate X in C are:

- all nodes dominating X within C (including C itself)
- all sisters of X[3]
- all sisters of the nodes dominating X within C

These definitions capture the idea that relating a trace (or pronoun) with a head noun becomes more complex the more that trace (or pronoun) is embedded within the subordinate clause. Figures 1 and 2 show how the relativisation domain is larger in an object relative than in a subject relative, since the trace is more deeply embedded.

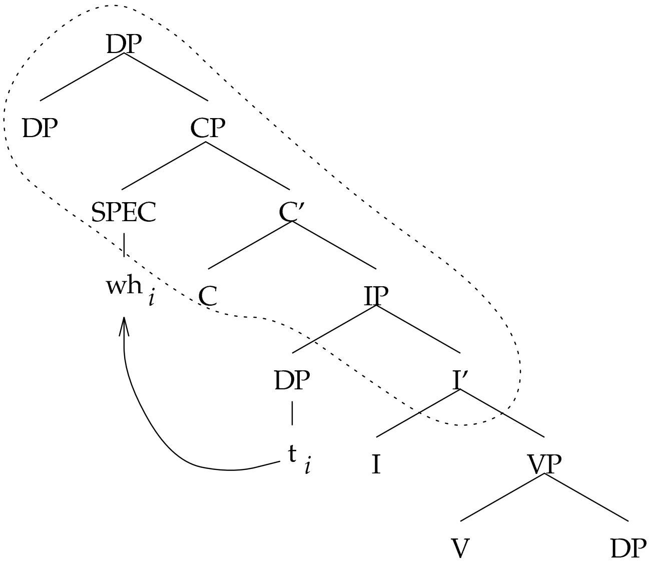

Figure 1: Subject relative with domain circled.

Just as the complexity predictions for the order of MNCCs are reflected in examples such as Particle Movement, the prediction that subject relatives should be easier to process than object relatives is born out by the psycholinguistic literature on the subject (see, e.g. Keenan & Hawkins 1987).

4 Typological asymmetries

In the previous section two processing asymmetries have been reviewed. The evidence that universals are fit comes when we compare processing with the cross-linguistic data. For both the processing asymmetries above (and many others) there are corresponding typological asymmetries.

---

[3] In fact, some sisters may be excluded if the language has a flatter configurational structure. See Hawkins (1994a:27–28) for discussion.

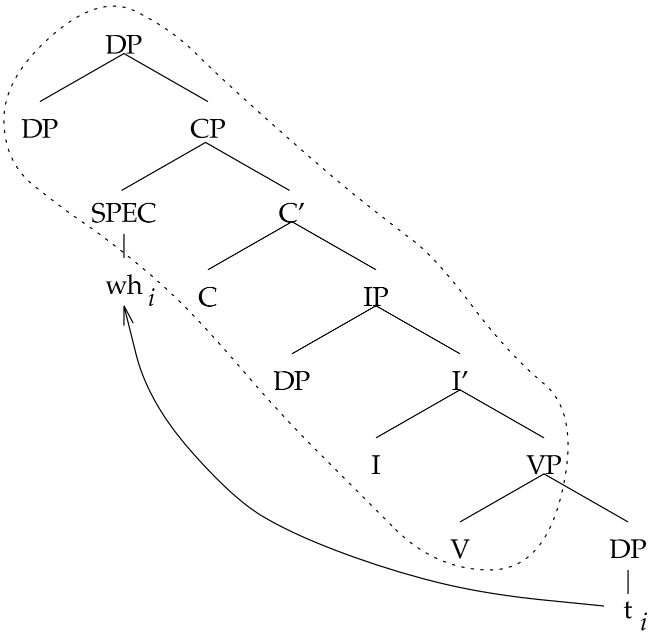

Figure 2: Object relative with domain circled.

4.1 Consistent branching direction

On the basis of a large sample of languages and using a method intended to compensate for biases in the sample due to historical relatedness of languages and areal groupings, Dryer (1992) formulates the following universal:

> "...a pair of elements X and Y will employ the order XY significantly more often among OV languages if and only if X is a nonphrasal category and Y is a phrasal category." (Dryer 1992:89)

One of the features of this *Branching Direction Theory* is that it predicts that heads will tend to order consistently on one side or another of non-heads cross-linguistically — a general tendency that has long been recognised.

One of the clearest predictions that Dryer's universal makes and that was recognised earlier by Greenberg (1963) is the correlation of the order of verb and object and the order of adposition and NP:

$$VO \leftrightarrow PrNP$$

Figure 3 shows the basic structures of VP internal PPs in the four possible language types defined by this typology. The nodes that have been circled are the ones involved in immediate constituent attachment for the mother VP. Just as in the Particle Movement examples, the parser prefers the MNCCs to be close together. In this case, the relevant MNCCs are the verb and adposition. Clearly the structures that are easier to process are the ones that are more common cross-linguistically. A similar argument can be made for other predictions of Dryer's universal if non-branching categories tend to be MNCCs.[4]

---

[4] This will not always be the case, which is why Dryer's statement goes further than consistent head ordering. For example, Dryer notes that the order of adjective and noun does *not* correlate with VO/OV order. These

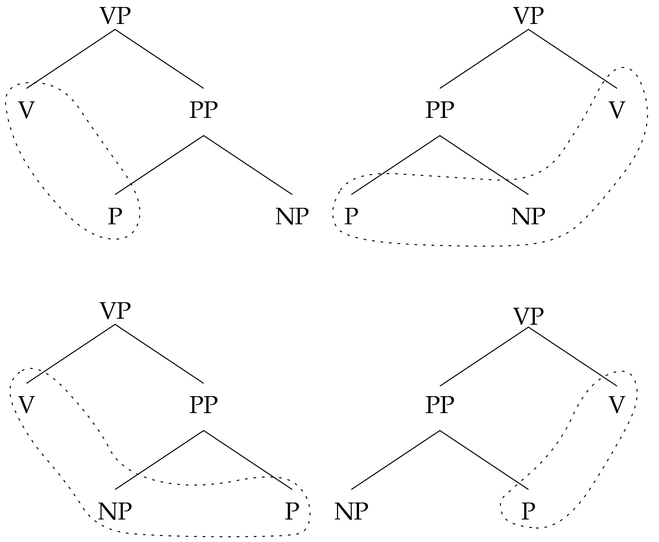

Figure 3: Four VP internal PPs.

This, then, is an example of typological fitness. The branching direction theory, a statement about the distribution of languages in the world, dovetails with the minimal attachment domain theory, a statement of the processing complexity of particular utterances.

4.2 Object relativisation

A similar fit of universal and function can be found if we look at the cross-linguistic distribution of relative clause types. The well-known implicational hierarchy of Keenan & Comrie (1977) puts noun-phrase types in order of accessibility to relativisation:

Subject>Direct Object>Indirect Object>Oblique>Genitive>Object of Comparison

Accessibility Hierarchy Constraint "If a language can relativise any position on the AH with a primary strategy [a strategy for forming relatives that is used for subjects], then it can relativise all higher positions with that strategy." (Comrie & Keenan 1979:653)

This hierarchy unpacks as a series of chained implicational statements:

$$\begin{gathered}(Object\ of\ Comparison \rightarrow Genitive)\& \\ (Genitive \rightarrow Oblique)\& \\ (Oblique \rightarrow Indirect\ Object)\& \\ (Indirect\ Object \rightarrow Direct\ Object)\& \\ (Direct\ Object \rightarrow Subject)\end{gathered}$$

---

discrepancies between Hawkins' processing theory and the cross-linguistic data may be the result of a problem with either the processing theory or the sampling method. Another possibility, however, is that match between typology and processing may not always be one-to-one for some other reason (see Kirby 1996b for discussion of problems like this). I will not go into these mismatches here, although it is important that they be a focus of future typological and psycholinguistic research.

For this paper we will concentrate on the fourth of these: if a language relativises indirect objects with a primary strategy, then it will relativise direct objects with a primary strategy.

Hawkins argues that this typological asymmetry reflects a psycholinguistic asymmetry in terms of the minimal structural domain for relativisation.

> "The minimal SD [structural domain] of a DO [direct object] is properly included in that of an IO [indirect object], since a clause with an IO generally requires an accompanying DO and also a SU [subject], whereas a DO can occur both with and without an IO, i.e. Min SD(DO)$\subset$Min SD(IO)." (Hawkins 1994a:39)

This is also backed up by the psycholinguistic literature as noted earlier (e.g. Keenan & Hawkins 1987 in their experiments on a repetition task with English-speaking children).

Once again, there is a clear fit of the typological facts and the preference of the parser as predicted by Hawkins' theory and backed-up by psycholinguistic experiments. The explanation for this fit is the focus of the remainder of this paper.

5 A phylogenetic explanation?

The appearance of design is a characteristic of biological structures whose evolution is mediated by natural selection. Since language is a product of a biological organism — part of our *extended phenotype* in Dawkins (1982) terms — then a phylogenetic explanation for the adaptive nature of its universals seems the obvious place to look.

5.1 Natural selection and the LAD

It is widely believed that the human ability to acquire language is at least in some part innately given, and that the Chomskyan Universal Grammar is embodied in this ability. Since we may assume that on the whole this biological endowment is shared by all members of our species, then we have a ready explanation for universals — they are simply the result of idiosyncratic properties of the innate mechanism for acquisition (see, e.g. Hoekstra & Kooij 1988).

Various authors (e.g. Pinker & Bloom 1990, Hurford 1989, Hurford 1991) have argued that a complex innately specified language acquisition device (LAD) must have evolved through natural selection.

> "All we have argued is that human language, like other specialized biological systems, evolved by natural selection. Our conclusion is based on two facts that we would think would be entirely uncontroversial: Language shows signs of complex design for the communication of propositional structures, and the only explanation for the origins of organs with complex design is the process of natural selection." (Pinker & Bloom 1990:726)

In particular the LAD is claimed to have evolved through selection *for the function it now fulfils*. This relies on the assumption that human language confers a survival or reproductive advantage on the organisms that have it. This assumption is to be fairly well accepted, although when we get to specific features of UG there seems to be greater unease. Lightfoot (1991:69), for example, pours scorn on this argument suggesting "the Subjacency Condition has many virtues, but I am not sure that it could have increased the chances of having fruitful sex." We shall return to the Subjacency and Sex Problem in the concluding section, but for the moment let us continue following the logic of the phylogenetic argument.

5.2 Selection for parsability

In order to solve our problem of the origin of typological fitness all we need to do now is argue that different possible LAD's could be selected for on the basis of the parsability of the output of the grammars they allow to be acquired. Newmeyer (1991) has made precisely this argument with respect to features of the LAD specific to the Government and Binding theory of syntax (see, e.g. Haegeman 1991).

> "It is quite plausible that the design of the grammatical model as a whole or *some particular grammatical principle* might have become encoded in our genes by virtue of its being successful in facilitating communication that the survival and reproductive possibilities of those possessing it were enhanced." (Newmeyer 1991:7) (emphasis my own)

One such feature of UG that Newmeyer approaches in this fashion is the Subjacency Condition (the object of Lightfoot's scorn as mentioned above). Briefly, the Subjacency Condition is a universal that constrains the structural distance between two elements that are related by some rule of the grammar (Riemsdijk & Williams 1986, 62, cited in Newmeyer 1991, 12):

Subjacency Condition No rule can relate X, Y in the structure

$$\ldots X \ldots [_\alpha \ldots [_\beta \ldots Y \ldots$$

or

$$\ldots Y \ldots]_\beta \ldots]_\alpha \ldots X \ldots$$

where α and β are bounding nodes.

In English, the bounding nodes are IP and NP, hence the ungrammaticality of the sentences below, where *who* has moved over two bounding nodes (with no intermediate 'landing site'[5]):

(5) a. *I met the fan who$_i$ we played $_{NP}[$ the song which$_j$ $_{IP}[$ t_i liked $t_j]]$
b. *Who$_i$ did $_{IP}[$ Matt tell you when$_j$ $_{IP}[$ he had met $t_i t_j]]$

The standard assumption is that the subjacency condition is one of a set of constraints on the application of the transformation move-α that form part of our innate knowledge of language. Although there is some cross-linguistic variability in the inventory of bounding nodes, the constraint can, in principle, be applied to any language. How can the existence of this constraint be explained? Berwick & Weinberg (1984) point out that the subjacency condition tends to rule out sentences in which the distance between the wh-element and its co-indexed gap is long. As already discussed above, there is a pressure from the parser to keep this distance to a minimum (although clearly there is not a direct parallel between the formulation of the subjacency condition and minimal relativisation domains). Newmeyer's argument is that this parsing pressure led to the biological selection of a language acquisition device that had some way of eliminating the worst wh-extractions from the language.

---

[5] See, for example, Haegeman 1991, §6.2 for further details of the applicability of the subjacency condition.

6 A glossogenetic explanation

Although the argument presented in the previous section is persuasive, there is an alternative that should be considered — one which will allow us to approach the specific examples of typological fitness that were reviewed earlier. It will be argued in this section that languages *themselves* adapt to aid their own "survival" in the transmission from speaker to speaker. In this view, selection for parsability acts on a historical rather than biological timescale, or in Hurford's (1990) terms, this is a *glossogenetic* as opposed to *phylogenetic* explanation. A glossogenetic explanation for fitness relies on a theory of linguistic selection. Such a theory is set out in more detail elsewhere (e.g. Kirby 1996a, Kirby 1996c, Kirby 1994), but a sketch of its main components will be given here.

Selection takes place whenever there is differential replication of information through a dynamic system. Gell-Mann's (1992) treatment of complex adaptive systems treats the objects over which selection is made as *schemata*. In his view, schemata "unfold" in the environment to produce effects or behaviour which have consequences for the viability of the schema in terms of survival to a later generation. There is therefore a cycle of:

1. unfolding of schemata to produce environmental effect,
2. differential selection of schemata based on that effect,
3. compression of information in environment to produce next generation of schemata.

Notice that this cycle involves two different types of object: the schema itself, and its "external" effects. It is the transformations between these objects that gives rise to the selective effect. Lewontin (1974) (cited in Sober 1984) talks about biological evolution in exactly these terms, highlighting the importance of transformation between genotype and phenotype in biology. Kirby (1996a) adapts Lewontin's map of the transformations for linguistics, replacing genotypes with *I-language* and phenotypes with *E-language* (figure 4).

For us the important transformations on this diagram are the ones that map between the I- and E-language domains: production (**T1**) and parsing (**T3**). If we assume that the transformations are not perfect mappings from competence to performance (and vice-versa) then we have selection. To put it another way; as long as the only way for language to be transmitted from generation to generation is through the being repeatedly produced and parsed, then the form of the language that survives this process will naturally be adapted to being produced and parsed. This is what I mean by saying that languages adapt to aid their *own* survival.

7 Computer simulations

Now that there is a basic framework in which to think about linguistic selection, we need some way of seeing what effect various possible processing pressures have on the eventual distribution of languages. The methodology employed here is to use computational simulations of populations of language users. This allows us to test exactly what effect changing the selection criteria over utterances has on the languages spoken. In this way, the computer provides us an with excellent tool for testing the behaviour of complex adaptive systems whose behaviour is the *emergent* result of the interactions of many individuals.

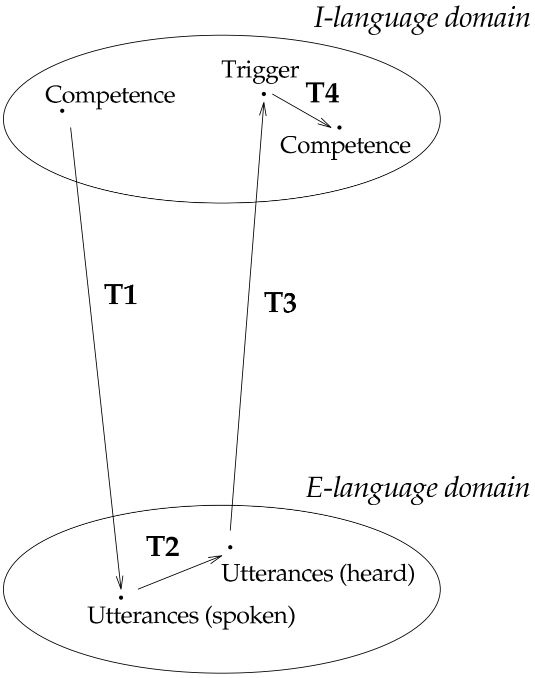

Figure 4: Transformations within and between I- and E-language domains

7.1 General set up of the simulations

The simulations are made up of three central components, each of which consists of a two-dimensional toroidal surface or array of elements arranged so that each element has four neighbours (above, below, to the left, and to the right). The first array of elements represents speakers, the second represents utterances and the third, acquirers. What form each of these kinds of element can take depends on the particular processing metric being tested.

Each generation, all the speakers randomly produce utterances in line with their particular grammars. At this point, the choice of which utterances to produce may be biased according to some measure of production complexity — this is the kind of selection associated with transformation **T1** in figure 4. After the production phase has taken place, all the acquirers take a sample of the utterances produced by the speakers closest to their position in the array (the speaker in the same position, and the four neighbours). Again, this sample may be biased according to some measure of parsing complexity — the selection in **T3** in figure 4. The resultant sample of utterances is used to set the acquirer's grammar. Finally, the acquirers become the speakers of the next iteration of the simulation (see figure 5).

7.2 The S-shaped curve

Kroch (1994) refers to situations where languages change their relative frequency of variants as *grammar competition*. Under his formulation, two or more broadly syntactic doublets behave in the same way as morphological doublets in competition for a paradigm slot. This view follows from a move in syntactic analysis to treat cross-linguistic variation as a reflection of variation in the properties and inventories of functional heads.

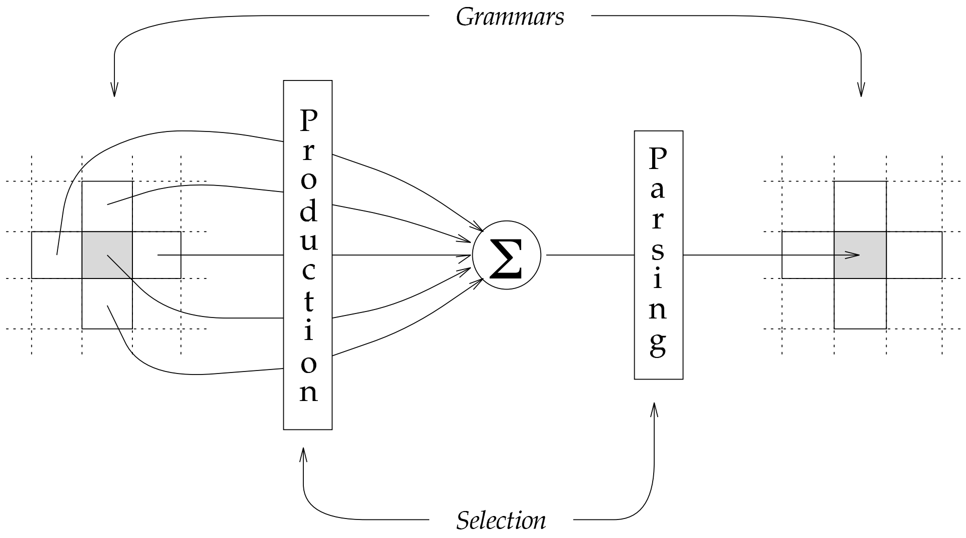

Figure 5: Computer simulation of glossogenetic adaptation.

> "If we take this view seriously, we are led to the conclusion that syntactic variation should be governed by the same principles as variation in morphology, since the locus of the variability in the two cases is the same — the formative. Just as morphological variants which are not functionally distinguished are disallowed, so we should not expect to find variation between semantically non-distinct syntactic heads. To the extent that such variability is found, it poses the same theoretical problem as the appearance of doublets does in morphology." (Kroch 1994:5)

Kroch points out that the "blocking effect" in morphology (whereby the presence of an irregular form in a paradigm slot blocks the occupation of that slot by a regular form) is a central tenet of modern morphology. Doublets *are* in fact often observed in languages; however, if the doublets are functionally equivalent, speakers "learn either one or the other form in the course of basic language acquisition, but not both" (p. 6). Later on the same speakers may recognise the existence of the variant form, which "for them has the status of a foreign accent" (p. 6). Finally, one of these two doublets will tend to win out in a particular community or the two forms will become functionally differentiated.

The first example of the simulation in action looks at the time course of change where there are two such variants in competition. This is the simplest possible example, with only two possible grammars and two possible utterances. Furthermore, for the run displayed here, the spatial organisation of the population described in the last section is "switched off" so that all speakers can potentially contribute to all acquirers' trigger experience.

Imagine a situation where there is a language with basic VO order and postpositions. This type of language is sub-optimal with regard to rapid IC attachment in constructions such as $_{VP}[V\ _{PP}[NP\ P]]$ since the structural domain for attachment of the ICs to the VP stretches from the initial verb to the postposition (over the noun-phrase). Now, if prepositions were to be introduced as a minor variant, perhaps during language contact, then this variant will be selected more frequently by the parser to trigger acquisition than the postpositional variant.

In the simulation of this scenario, the initial speech community is made up of 450 speakers with the postpositional variant, and 50 speakers with prepositions. Each utterance is assigned a "parsability score" which is based on Hawkins' suggested metric for VO languages (Hawkins 1990:238).[6] In general, for an individual acquirer, the probability of a particular variant v being acquired is:

$$P(v) = \frac{\sum_{i \in I_v} w_i}{\sum_{j \in (I_v \cup I_{v'})} w_j}$$

where I_v is the set of utterances in the input that can trigger the acquisition of v, w_i is the parsing complexity of an utterance, and v' is the competitor for v. The result of this simulation run is given in figure 6.

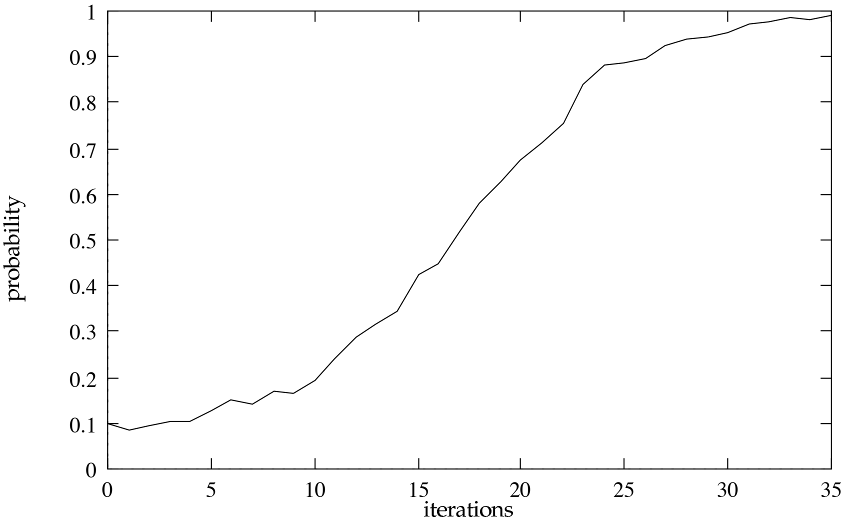

Figure 6: The S-shaped curve.

The striking feature of this result is its similarity to the S-shaped curve that Kroch (1989a, 1989b) finds in his data for the rise of periphrastic *do* in English — a case of the replacement by competition of one variant by another. This result is also found by Niyogi & Berwick (1995) in their analysis of the dynamics of parameterised grammatical change, although in their case, the complexity of parsing is not considered. Research is currently underway to analyse more carefully the dynamic properties of parametrically specified grammars under the influence of parsing pressures (Clark 1996). For the moment, however, the closeness of the time course of change resulting from the simulation to that observed in the historical data is certainly encouraging.

7.3 Emergence of branching direction

The previous example simply demonstrated one form taking over another in a community. In order to show how typological fitness can emerge, the simulation should give rise to

---

[6] This means that $w_{prepositions} = 1$ and $w_{postpositions} = 0.79$ in the following equation.

more than one "winning" language type in a multi-dimensional typology. Consider the case of consistent branching-direction; the minimum we need to show this emerging is two variable orthogonal typological parameters, such as verb/object order and adposition/noun order. In the previous example, only adposition order was variable. If both are variable, then we should see the emergence of a correlation between VO and prepositions *and* OV and postpositions.

In this more complex example, the parsability of a particular variant is dependent on the proportions of the orthogonal variant in the input. This is modelled by the following complexity functions:

$$
\begin{aligned}
w_{i \in I_{prep}} &= \alpha |I_{vo}| + (1-\alpha)|I_{ov}| \\
w_{i \in I_{postp}} &= \alpha |I_{ov}| + (1-\alpha)|I_{vo}| \\
w_{i \in I_{vo}} &= \alpha |I_{prep}| + (1-\alpha)|I_{postp}| \\
w_{i \in I_{ov}} &= \alpha |I_{postp}| + (1-\alpha)|I_{prep}|
\end{aligned}
$$

where $|I_v|$ is the number of utterances in the input that trigger the acquisition of v, and α is a measure of relatedness of verb/object order and adposition order varying from 0 to 1. With α less than 0.5, VO and postpositions are correlated in terms of complexity, with α greater than 0.5, VO and prepositions are correlated. For the run shown here, α is 0.6.

For this simulation, a 40x40 grid of speakers was set up so that each speaker initially had a random grammar corresponding to either VO&Prep, VO&Postp, OV&Prep, or OV&Postp. This initial condition is shown in figure 7. Each shade of square corresponds to a different grammar. After 30 iterations of the simulation, consistent branching direction emerges (figure 8). The language types have organised themselves into two major groups spatially, corresponding to left- and right-branching. These groups shift slowly over time, both in size and position. The inconsistently branching languages remain in the minority, occurring as transitional types as a particular area switches from one order to another where two language groups are in contact.

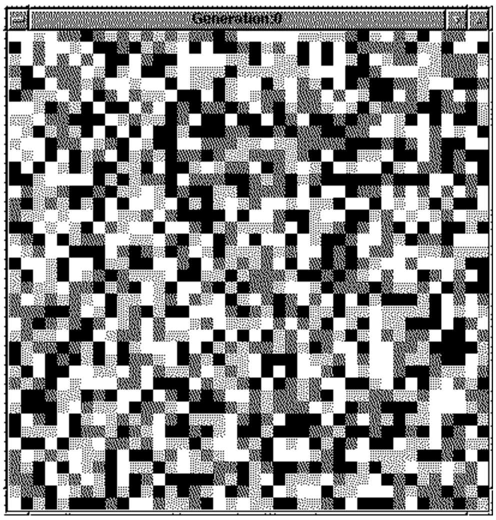

Figure 7: Initial state of the simulation

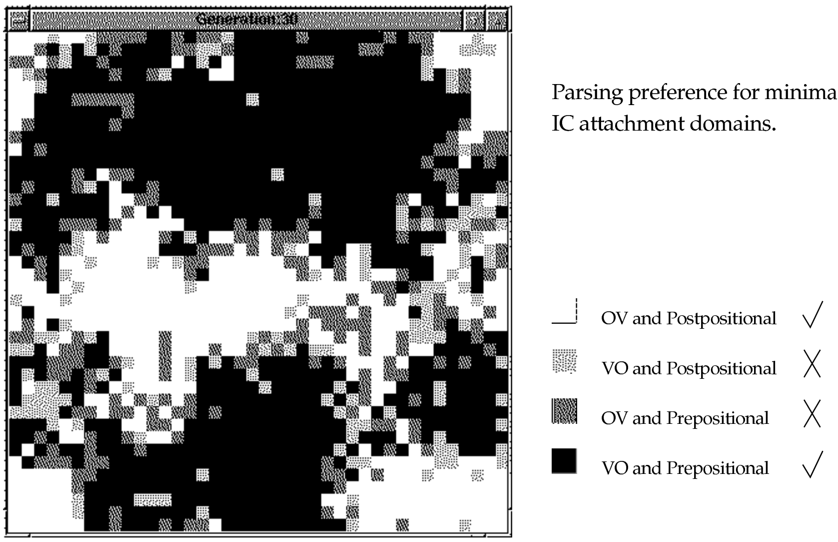

Figure 8: Emergence of branching-direction

This simulation shows that large sub-groupings in a speech community can emerge and remain stable even though only *individuals* are being modelled. More importantly, typological fitness also emerges. The language groups organise themselves so that they allow a particular variant only if it makes utterances easy to parse, given the other variants in the language. This global property appears to be inevitable given the large number of individuals interacting on a purely local basis, as long as this interaction is differentially successful for particular variants.

7.4 Emergence of the object relative universal

The final example shows how an implicational hierarchy may emerge given two types of selection in conflict. Consider the object relative universal presented earlier:

$$IndirectObject \rightarrow DirectObject$$

In order to test if this emerges from the simulation we need to set up the complexity of indirect and direct object relatives:

$$0 < w_{IO} < w_{DO} < 0.5$$

$$w_{v'} = (1 - w_v)$$

This means that indirect object relatives are more complex to parse than direct object relatives (as shown in section 3.2), and both are more complex than a competing variant that does not involve object relativisation. Such a non-object relative variant might be a subject relative in which the object has been promoted by a passive, for example. See Keenan & Hawkins (1987) for some discussion of the complexity of promoted and non-promoted relatives.

The problem with this set up is that, given any set of values of w_{IO} and w_{DO} that conform to the inequality above, the end result is always solely languages with no object relativisation. However, what we find cross-linguistically are

1. languages with no object relatives,
2. languages with only direct object relatives, and
3. languages with both direct and indirect object relatives.

What seems to be happening is that languages are adapting to fit the only processing pressures that are about in the simulation — in other words, the pressures relating to the complexity of object relatives. The inevitable result of this is that these relatives will be avoided altogether.

On the other hand, it is possible object relatives are not rejected wholesale since *the alternative is also complex*. The sentences below show an object relative, and its derived-subject variant:

(6) The batsman who Atherton dropped got a century.

(7) The batsman who was dropped by Atherton got a century.

The second example has a lower parsing complexity because it relativises the subject of the subordinate.[7] Counteracting this it involves extra morphemes to support the passivisation of *the embedded clause*. Kirby (1996a) argues that the speaker (as opposed to the acquirer) will select from variant forms on the basis of the morphological (as opposed to parsing) complexity. This means that there is a *competing motivation* involved in the selection of relative clause types in the cycle of acquisition and use.

If this is the case, we need to include a distinction in the simulation between the two types of complexity:

p-complexity The complexity of the utterance from the point of view of the acquirer. (parsing-complexity)

$$0 < w^p_{IO} < w^p_{DO} < 0.5$$

m-complexity The complexity of the utterance from the point of view of the speaker. (morphological-complexity)

$$0 < w^m_{IO'} = w^m_{DO'} < 0.5$$

The result of running the simulation with these complexity considerations is shown in figure 9. The actual settings of the various complexity measures is rather more complex than in previous simulations since the relative prominence of p- and m-complexity in the selection of utterances varies dependent on other features in the language, which are assumed to be random with respect to relativisation possibilities. For further details see Kirby (1996a).

The result shows clearly that the implicational universal emerges out of the conflict between hearer and speaker related processing considerations. This, then, is a paradigm example of the emergence of typological fitness. An individual language may "decide" to impose zero, one, or two object relatives upon its hearers (in response to the needs of speakers), but the relatives it allows will always be the ones which hearers find easier to parse. All this emerges once again solely from individual language users interacting on a purely local basis without any teleological motivation.

---

[7] This is not exactly what Hawkins (1994b) predicts, however it seems to be consonant with the data from repetition tasks in Keenan & Hawkins (1987).

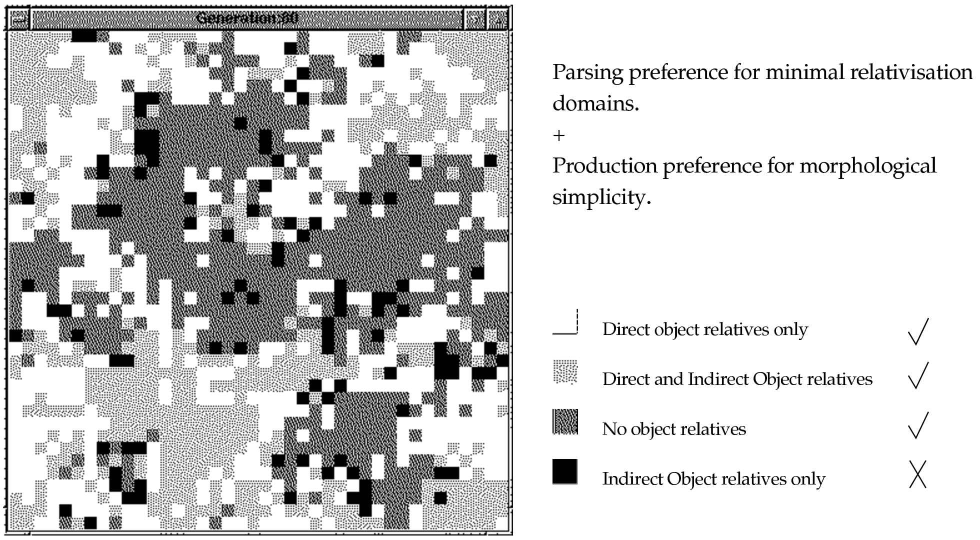

Parsing preference for minimal relativisation domains.
+
Production preference for morphological simplicity.

Figure 9: The emergence of an implicational universal

8 Conclusion: language evolution as the interaction of adaptive systems

In this paper I have shown that the occurrence of typological fitness does not necessarily mean we should immediately conclude that there is a phylogenetic explanation for that fitness. The other possibility is that languages may adapt themselves on a glossogenetic timescale simply by virtue of their need to propagate themselves through the medium of communication. The computer simulation results show that, even with very simple assumptions, complex adaptive behaviour emerges, with stable end states which show a characteristic "appearance of design". None of this rules out the phylogenetic explanation, however, and in this final section I will conclude with some speculative remarks about how this glossogenetic perspective may solve some problems that beset any potential account for the evolution of the human language faculty.

The problems with the phylogenetic account outlined in section 5 are to do with the need to assume that a cross-linguistic constraint could evolve by giving an individual a survival advantage who has a mutated LAD that gives rise to that constraint. Even if that particular constraint will in the long run give rise to languages that have constructions that are easy to parse, then it is difficult to see how the individual's survival chances are increased. The classic problems (which Pinker & Bloom (1990) attempt to solve) are:

1. How is the selfish individual better off if her output is easier to parse?
2. How is an innovation in a single individual any use when language is a shared activity?
3. Even if the previous two points are not problems for evolution given enough time, does that mean language could not have evolved as quickly as it seems to have?
4. Does the subjacency condition really increase the chances of fruitful sex by making sentences easier to parse?

In view of the glossogenetic approach, we can say that these questions become irrelevant. Languages evolve orders of magnitude faster than their users' species, and they appear to do so in order to increase the survival chances of the rules with which they are made up. This just happens to make things easier for us as speakers and hearers.

Where does this leave innateness? Well, we can simply say that the branching direction universal and the Accessibility Hierarchy are not innate. In other words, they are not hard-wired into whatever part of our brain carries out language acquisition. But what about other universals, such as the Subjacency Condition, that have been argued by many — and for reasons other than their usefulness for parsing — to be part of our innate endowment?

In the view set out in this paper, there is absolutely no need to reject innateness. Instead we need to realise that the LAD has evolved for the specific (and totally selfish from the individual's point of view) purpose of *learning language quickly*. Even if language were to prove to be acquirable using some kind of general purpose learning mechanism, there is no particular reason why it should be acquired this way. If certain features of the languages that humans need to learn show up again and again, it makes sense for these features to be part of our innate endowment so that we are not forced to learn them from scratch. The Baldwin Effect (see, e.g. Baldwin 1896; Hinton & Nowlan 1987; Kirby & Hurford 1996) predicts that just this sort of nativisation will arise wherever learning and evolution interact.

If languages have adapted glossogenetically to consistently avoid some type of construction because it is hard to parse, then an individual will be at an advantage who does not need to learn this fact and instead has it given innately. In other words, the Subjacency Condition may be an innate "fossil" of a glossogenetic adaptation to avoid long-distance dependencies. In this way, the LAD does not evolve to constrain languages to be better communicative tools — this is simply a bi-product of the evolution of faster and better ways of learning a cultural artifact that is itself an adaptive system.

Future research into the evolution of the language faculty must take this interaction of adaptive systems seriously. We can imagine a type of bootstrapping process where a particular stage in the phylogenetic evolution of language licences certain kinds of glossogenetic adaptation. If the fitness of individuals is related to their ability to acquire the language around them, then any universals that emerge from this glossogenetic adaptation will in turn alter the fitness of the language faculties around at the time. If this leads to the evolution of those faculties, then the types of glossogenetic adaptation that are possible will also change. In this way, phylogenetic evolution may "track" glossogenetic evolution solely through selection pressure to learn quickly.

This view of the evolution of language raises many questions, and challenges our understanding of complex adaptive systems. In this paper I have not *shown* that glossogenetic adaptation interacts with phylogenetic evolution — simply that wherever there is adaptation, natural selection is not the only explanation. Further modelling work is needed to marry the type of simulation discussed in this paper and those that look at the evolution of learning (e.g. Ackley & Littman 1991). This will require an interesting synthesis of functionalist accounts of diachrony and universals on the one hand and generative/innatist accounts of language acquisition on the other. Perhaps the biggest challenge facing those studying the evolution of language is finding common ground between the two.

References

ACKLEY, DAVID, & MICHAEL LITTMAN. 1991. Interactions between learning and evolution. In *Artificial Life II*, ed. by C.G. Langton, C. Taylor, J.D. Farmer, & S. Ramussen, 487–509. Addison-Wesley.

BALDWIN, J.M. 1896. A new factor in evolution. *American Naturalist* 30.441–451.

BERWICK, R.C., & A.S. WEINBERG. 1984. *The Grammatical Basis of Linguistic Performance: Language Use and Acquisition*. MIT Press.

CLARK, ROBERT, 1996. Internal and external factors affecting language change: A computational model. Master's thesis, University of Edinburgh. In preparation.

COMRIE, BERNARD. 1981. *Language Universals and Linguistic Typology*. Basil Blackwell.

——, & EDWARD KEENAN. 1979. Noun phrase accessibility revisited. *Language* 55.649–664.

CROFT, WILLIAM. 1990. *Typology and universals*. Cambridge: Cambridge University Press.

CUTLER, A., J. HAWKINS, & G. GILLIGAN. 1985. The suffixing preference: a processing explanation. *Linguistics* 23.723–758.

CZIKO, GARY. 1995. *Without Miracles: Universal Selection Theory and the Second Darwinian Revolution*. MIT Press.

DAWKINS, RICHARD. 1982. *The Extended Phenotype*. Freeman.

DRYER, MATTHEW. 1992. The Greenbergian word order correlations. *Language* 68.81–138.

GELL-MANN, MURRAY. 1992. Complexity and complex adaptive systems. In *The Evolution of Human Languages*, ed. by J.A. Hawkins & M. Gell-Mann. Addison-Wesley.

GREENBERG, JOSEPH. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In *Universals of Language*, ed. by Joseph Greenberg, 73–113. MIT Press.

HAEGEMAN, LILIANE. 1991. *Introduction to Government and Binding Theory*. Blackwell.

HAWKINS, JOHN A. 1988. Explaining language universals. In *Explaining Language Universals*, ed. by John A. Hawkins. Basil Blackwell.

—— 1990. A parsing theory of word order universals. *Linguistic Inquiry* 21.223–261.

—— 1994a. *A performance theory of order and constituency*. Cambridge University Press.

——, 1994b. Some issues in a performance theory of word order. Manuscript, University of Southern California.

HINTON, G., & S. NOWLAN. 1987. How learning can guide evolution. *Complex Systems* 1.495–502.

HOEKSTRA, TEUN, & JAN G. KOOIJ. 1988. The innateness hypothesis. In *Explaining Language Universals*, ed. by John A. Hawkins. Blackwell.

HURFORD, JAMES. 1989. Biological evolution of the Saussurean sign as a component of the language acquisition device. *Lingua* 77.187–222.

——. 1990. Nativist and functional explanations in language acquisition. In *Logical Issues in Language Acquisition*, ed. by I.M. Roca, 85–136. Foris Publications.

——. 1991. The evolution of the critical period for language acquisition. *Cognition* 40.159–201.

KEENAN, EDWARD, & BERNARD COMRIE. 1977. Noun phrase accessibility and universal grammar. *Linguistic Inquiry* 8.63–99.

——, & SARAH HAWKINS. 1987. The psychological validity of the accessibility hierarchy. In *Universal Grammar: 15 essays*, ed. by Edward Keenan, 60–85. London: Croom Helm.

KIRBY, SIMON. 1994. Adaptive explanations for language universals: a model of Hawkins' performance theory. *Sprachtypologie und Universalienforschung* 47.186–210.

——. 1996a. Competing motivations and emergence: explaining implicational hierarchies. *Language Typology* . Forthcoming.

——, 1996b. Constraints on constraints, or the limits of functional adaptation. To appear in: *Functionalism and Formalism*.

——. 1996c. *The Emergence of Universals: Function, Selection and Innateness*. University of Edinburgh. PhD thesis.

——, & JAMES HURFORD, 1996. The evolution of incremental learning. In preparation.

KROCH, ANTHONY. 1989a. Function and grammar in the history of English. In *Language Change and Variation*, ed. by Ralph Fasold & Deborah Schiffrin, 133–172. .

——. 1989b. Reflexes of grammar in patterns of language change. *Language Variation and Change* 1.199–244.

——. 1994. Morphosyntactic variation. In *Papers from the 30th Regional Meeting of the Chicago Linguistics Society*, ed. by K. Beals.

LEWONTIN, R. 1974. *The Genetic Basis of Evolutionary Change*. Columbia University Press.

LIGHTFOOT, DAVID. 1991. Subjacency and sex. *Language and Communication* 11.3–28.

NEWMEYER, FREDERICK J. 1991. Functional explanation in linguistics and the origins of language. *Language and Communication* 11.3–28.

NIYOGI, PARTHA, & ROBERT BERWICK. 1995. The logical problem of language change. Technical Report AI Memo 1516 / CBCL Paper 115, MIT AI Laboratory and Center for Biological and Computational Learning, Department of Brain and Cognitive Sciences.

PINKER, STEVEN, & PAUL BLOOM. 1990. Natural language and natural selection. *Behavioral and Brain Sciences* 13.707–784.

RIEMSDIJK, H. VAN, & E. WILLIAMS. 1986. *Introduction to the Theory of Grammar*. MIT Press.

SOBER, ELLIOT. 1984. *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. MIT Press.