To be published in *Behavioral and Brain Sciences* (in press)

Precis of: Foundations of Language: Brain, Meaning, Grammar, Evolution

Oxford and New York: Oxford University Press, 2002 (506 pp.)

Ray Jackendoff
Program in Linguistics, MS 013
Brandeis University
Waltham, MA 02454 USA
jackendoff@brandeis.edu
tel: 617-484-5394
fax :617-484-0164

Short abstract: The goal of this study to reintegrate the theory of generative grammar into the cognitive sciences. Generative grammar was correct to focus on the child's acquisition of language as its central problem, leading to the hypothesis of an innate Universal Grammar. However, generative grammar was mistaken to assume that the syntactic component is the sole course of combinatoriality, and that everything else is "interpretive." The proper approach is a parallel architecture, in which phonology, syntax, and semantics are autonomous generative systems, linked by interface components. The parallel architecture leads to an integration within linguistics, and to a far better integration with the rest of cognitive neuroscience. It fits naturally into the larger architecture of the mind/brain and permits a properly mentalistic theory of semantics. It leads to a view of linguistic performance in which the rules of grammar are directly involved in processing. Finally, it leads to a natural account of the incremental evolution of the language capacity.

1. Introduction

In the 1960s, when I became a graduate student in linguistics, generative grammar was *the* hot new topic. Everyone from philosophers to psychologists to anthropologists to educators to literary theorists was reading about transformational grammar. But by the late 1970s, the bloom was off the rose, although most linguists didn't realize it; and by the 1990s, linguistics was arguably far on the periphery of the action in cognitive science. To some extent, of course, such a decline in fortune was simply a matter of fashion and the arrival of new methodologies such as connectionism and brain imaging. However, there are deeper reasons for linguistics' loss of prestige, some historical and some scientific.

The basic questions I want to take up here, then, are:

What was *right* about generative grammar in the 1960s, such that it held out such promise?
What was *wrong* about it, such that it didn't fulfill its promise?
How can we fix it, so as to restore its value to the other cognitive sciences?

The goal is to integrate linguistics with the other cognitive sciences, not to eliminate the insights achieved by any of them. To understand language and the brain, we need all the tools we can get. But everyone will have to give a little in order for the pieces to fit together properly.

The position developed in *Foundations of Language* is that the overall program of generative grammar was correct, as was the way this program was intended to fit in with psychology and biology. However, a basic technical mistake at the heart of the formal implementation, concerning the overall role of syntax in the grammar, led to the theory being unable to make the proper connections both within linguistic theory and with neighboring fields. *Foundations of Language* develops an alternative, the *parallel architecture*, which offers far richer opportunities for integration of the field. In order to understand the motivation for the parallel architecture, it is necessary to go through some history.

2. Three founding themes of generative grammar

The remarkable first chapter of Noam Chomsky's *Aspects of the Theory of Syntax* (1965) set the agenda for everything that has happened in generative linguistics since. Three theoretical pillars support the enterprise: *mentalism, combinatoriality*, and *acquisition*.

Mentalism. Before *Aspects*, the predominant view among linguists – if it was even discussed – was that language is something that exists either as an abstraction, or in texts, or in some sense "in the community" (the latter being the influential view of Saussure (1915), for example). Chomsky urged the view that the appropriate object of study is the linguistic system in the mind/brain of the individual speaker. According to this stance, a community has a common language by virtue of all speakers in the community having essentially the same linguistic system in their minds/brains.[1]

The term most often used for this linguistic system is "knowledge", perhaps an unfortunate choice. However, within the theoretical discourse of the time, the alternative was thinking of language as an ability, a "knowing how" in the sense of Ryle (1949), which carried overtones of behaviorism and stimulus-response learning, a sense from which Chomsky with good reason wished to distance himself. It must be stressed, though, that whatever term is used, the linguistic system in a speaker's mind/brain is deeply unconscious and largely unavailable to introspection, in the same way that our processing of visual signals is deeply unconscious. Thus language is a kind of mind/brain property hard to associate with the term "knowledge", which commonly implies accessibility to introspection. *Foundations of Language* compromises with tradition by systematically using the term *f-knowledge* ('functional knowledge') to describe whatever is in speakers' heads that enables them to speak and understand their native language(s).

There still are linguists, especially those edging off toward semiotics and hermeneutics, who reject the mentalist stance and assert that the only sensible way to study language is in terms of the communication between individuals (a random example is Dufva and Lähteenmäki 1996). But on the whole, the mentalistic outlook of generative grammar has continued to be hugely influential throughout linguistics and cognitive neuroscience.

More controversial has been an important distinction made in *Aspects* between the study of *competence* – a speaker's f-knowledge of language – and *performance*, the actual processes (viewed computationally or neurally) taking place in the mind/brain that put this f-knowledge to use in speaking and understanding sentences. I think the original impulse behind the distinction was methodological convenience. A competence theory permits linguists to do what they have always done, namely study phenomena like Bulgarian case marking and Turkish vowel harmony, without worrying too much about how the brain actually processes them. Unfortunately, in response to criticism from many different quarters (especially in response to the collapse of the derivational theory of complexity as detailed in e.g. Fodor, Bever, and Garrett 19xx), linguists have tended to harden the distinction into a firewall: competence theories were taken to be immune to evidence from performance. And so began a gulf between linguistics and the rest of cognitive science that has persisted until the present.

Foundations does not abandon the competence-performance distinction, but does return it to its original status as a methodological rather than ideological distinction. Although the innovations in *Foundations* are largely in the realm of competence theory, one of their important consequences is that there is a far closer connection to theories of processing, as well as the possibility of a two-way dialogue between competence and performance theories. We return to this issue in section 9.3.

Combinatoriality: The earliest published work in generative grammar, Chomsky's *Syntactic Structures* (1957), began with the observation that a language contains an arbitrarily large number of sentences. Therefore, in addition to the finite list of words, a characterization of a language must contain a set of *rules* (or a *grammar*) that collectively describes or "generates" the sentences of the language. *Syntactic Structures* showed that the rules of natural language cannot be characterized in terms of a finite-state Markov process, nor in terms of a context-free phrase structure grammar. Chomsky proposed that the appropriate form for the rules of a natural language is a context-free phrase structure grammar supplemented by transformational rules. Not all subsequent traditions of generative grammar (e.g. Head-Driven Phrase Structure Grammar (Pollard and Sag 1994) and Lexical-Functional Grammar (Bresnan 1982, 2001)) have maintained the device of transformational rules; but they all contain machinery designed to overcome the shortcomings of context-free grammars pointed out in 1957.[2]

Transferred into the mentalistic framework of 1965, the consequence of combinatoriality is that speakers of the language must have rules of language (or mental grammars) in their heads as part of their f-knowledge. Again there is a certain amount of controversy arising from the term "rules". Rules of grammar in the sense of generative grammar are not like any of the sorts of rules or laws in ordinary life: rules of etiquette, rules of chess, traffic laws, or laws of physics. They are unconscious principles that play a role in the production and understanding of sentences. Again, to ward off improper analogies, *Foundations* uses the term *f-rules* for whatever the combinatorial principles in the head may be. Generative linguistics leaves open how directly the f-rules are involved in processing, but, as suggested above, the unfortunate tendency among linguists has been not to care. The theory in *Foundations*, though, makes it possible to regard the rules as playing a direct role in processing; see again section 9.3.

An important reason for the spectacular reception of early generative grammar was that it went beyond merely claiming that language needs rules: it offered rigorous formal techniques for characterizing the rules, based on approaches to the foundations of mathematics and computability developed earlier in the century. The technology suddenly made it possible to say lots of interesting things about language and ask lots of interesting questions. For the first time ever it was possible to provide detailed descriptions of the syntax of natural languages (not only English but German, French, Turkish, Mohawk, Hidatsa, and Japanese were studied early on). In addition, generative phonology took off rapidly, adapting elements of Prague School phonology of the 1930s to the new techniques. With Chomsky and Halle's 1968 *Sound Pattern of English* as its flagship, generative phonology quickly supplanted the phonological theory of the American structuralist tradition.

Acquisition: Mentalism and combinatoriality together lead to the crucial question: How do children get the f-rules into their heads? Given that the f-rules are unconscious, parents and peers cannot verbalize them; and even if they could, children would not understand, since they don't know language yet. The best the environment can do for a language learner is provide examples of the language in a context. From there on it is up to the language learner to construct the principles on his or her own – unconsciously of course.

Chomsky asked the prescient question: what does the child have to "(f-)know in advance" in order to accomplish this feat? He phrased the problem in terms of the "poverty of the stimulus": many different generalizations are consistent with the data presented to the child, but the child somehow comes up with the "right" one, i.e. the one that puts him or her in tune with the generalizations of the language community. I like to put the problem a bit more starkly: The whole community of linguists, working together for decades with all sorts of crosslinguistic and psycholinguistic data unavailable to children, has still been unable to come up with a complete characterization of the grammar of a single natural language. Yet every child does it by the age of ten or so. Children don't have to make the choices we do: for instance they don't have to decide whether the "right" choice of grammar is in the style of transformational grammar, the Minimalist Program, Optimality Theory, Role and Reference Grammar, Tree-Adjoining Grammar, Cognitive Grammar, connectionist networks, or some as yet unarticulated alternative. They already f-know it in advance.

One of the goals of linguistic theory, then, is to solve this "Paradox of Language Acquisition" by discovering what aspects of linguistic f-knowledge are *not* learned, but rather *form the basis* for the child's learning. The standard term for the unlearned component is *Universal Grammar* or *UG*, a term that again perhaps carries too much unwanted baggage. In particular, UG should not be confused with universals of language: it is rather what shapes the acquisition of language. I prefer to think of it as a toolkit for constructing language, out of which the child (or better, the child's brain) f-selects tools appropriate to the job at hand. If the language in the environment happens to have a case system (like German), UG will help shape the child's acquisition of case; if it has a tone system (like Mandarin), UG will help shape the child's acquisition of tone. But if the language in the environment happens to be English, which lacks case and tone, these parts of UG will simply be silent.

What then is the source of language universals? Some of them will indeed be determined by UG, for instance the overall "architecture" of the grammatical system: the parts of the mental grammar and their relations (of which much more below). Other universals, especially what are often called "statistical" or "implicational" universals, may be the result of biases imposed by UG. For instance, UG may say that if a language has a case system, the simplest such systems are thus-and-so; these will be widespread systems crosslinguistically; they will be acquired earlier by children; and systems may tend to change toward them over historical time. Other universals may be a consequence of the functional properties of any relatively efficient communication system: for instance, the most frequently used signals tend to be short. UG doesn't have to say anything about these universals at all; they will come about through the dynamics of language use in the community (a process which of course is not very well understood).

If UG is not learned, how does the child acquire it? The only alternative is through the structure of the brain, which is determined through a combination of genetic inheritance and the biological processes resulting from expression of the genes, the latter in turn determined by some combination of inherent structure and environmental input. Here contemporary science is pretty much at an impasse. We know little about how genes determine brain structure and nothing about how the details of brain structure determine anything about language structure, even aspects of language as simple as speech sounds. Filling out this part of the picture is a long-term challenge for cognitive neuroscience. It is premature to reject the hypothesis of Universal Grammar, as some have (e.g. Elman et al. 1996 and Deacon 1997), arguing that we don't know how genes could code for language acquisition. After all, we don't know how genes code for birdsong or sexual behavior or sneezing either, but we don't deny that there is a genetic basis behind these.

There next arises the question of how much of UG is a human cognitive specialization for language and how much is a consequence of more general capacities. The question has often been oversimplified to a binary decision between language being entirely special or entirely general, with a strong bias inside generative linguistics towards the former and outside generative linguistics towards the latter. The truth of the matter undoubtedly lies somewhere in between. To be sure, many people (including myself) would find it satisfying if a substantial part of language acquisition were a consequence of general human cognitive factors; but the possibility of some specialization overlaying the general factors must not be discounted. My view is that we cannot determine what is general and what is special until we have comparable theories of other cognitive capacities, including other *learned* cognitive capacities. To claim that language is parasitic on, say, motor control, perhaps because both have hierarchical and temporal structure (this seems to be the essence of Corballis's (1991) position) – but without stating a theory of the f-knowledge involved in motor control – is to coarsen the fabric of linguistic theory to the point of unrecognizability. The closest approach to a comparable theory is the music theory of Lerdahl and Jackendoff 1983, which displays some striking parallels and some striking differences with language.

Of course, if UG – the ability to learn language – is in part a human cognitive specialization, it must be determined by some specifically human genes, which in turn had to have come into existence sometime since the hominid line separated from the other great apes. One would therefore like to be able to tell some reasonable story about how it could be shaped by natural selection or other evolutionary processes. We return to this issue in section 9.4.

This approach to the acquisition of language has given rise to a flourishing tradition of developmental research (references far too numerous to mention) and a small but persistent tradition in learnability theory (e.g. Wexler and Culicover 1980, Baker and McCarthy 1981). And certainly, even if the jury is not yet in on the degree to which language acquisition is a cognitive specialization, there have been all manner of phenomena investigated that bear on the issue, for instance:

- the sensitive period for language learning and the consequences for first and second language acquisition at a later age (Lenneberg 1967, Curtiss 1977, Newport 1990, Flynn and O'Neill 1988, Klein and Perdue 1997)

- the limited ability of apes to acquire even rudimentary versions of human language, even with extensive training (Premack 1976, Seidenberg and Petitto 1978, Terrace 1979, Savage-Rumbaugh et al. 1998)

- the characteristic brain localization of language functions, resulting in characteristic aphasias (Zurif 1990)

- the grammatical parallels between spoken and signed languages and the parallels in acquisition and aphasia (Klima and Bellugi 1979, Bellugi, Poizner, and Klima 1989, Fischer and Siple 1990)

- the existence of characteristic language deficits associated with various genetic conditions (Bellugi et al. 1994, Gopnik 1999, Clahsen and Almazan 1998)

- the creation of creole languages by communities of pidgin-speaking children (Bickerton 1981, DeGraff 1999)

- most strikingly, the creation of a signed language de novo by a newly assembled community of deaf children in Nicaragua (Kegl et al. 1999)

My impression is that, while there are questions about all of these cases, en masse they offer an overwhelming case for some degree of genetic specialization for language learning in humans.

These three foundational issues of generative grammar – mentalism, combinatoriality, and acquisition – have stood the test of time; if anything they have become even more important over the years in the rest of cognitive science. It is these three issues that connect linguistics intimately with psychology, brain science, and genetics. Much of the promise of generative linguistics arose from this new and exciting potential for scientific unification.

3. The broken promise: Deep Structure would be the key to the mind

A fourth major point of *Aspects*, and the one that seeped most deeply into the awareness of the wider public, concerned the notion of Deep Structure. A basic claim of the 1965 version of generative grammar was that in addition to the surface form of sentences, i.e. the form we hear, there is another level of syntactic structure, called Deep Structure, which expresses underlying syntactic regularities of sentences. For instance, a passive sentence like (1a) has a Deep Structure in which the noun phrases are in the order of the corresponding active (1b).

(1) a. The bear was chased by the lion.
    b. The lion chased the bear.

Similarly, a question such as (2a) has a Deep Structure closely resembling that of the corresponding declarative (2b).

(2) a. Which martini did Harry drink?
    b. Harry drank that martini.

In the years preceding *Aspects*, the question arose of how syntactic structure is connected to meaning. Following a hypothesis first proposed by Katz and Postal (1964), *Aspects* made the striking claim that the relevant level of syntax for determining meaning is Deep Structure.

In its weakest version, this claim was only that regularities of meaning are most directly encoded in Deep Structure, and this can be seen in (1) and (2). However, the claim was sometimes taken to imply much more: that Deep Structure *IS* meaning, an interpretation that Chomsky did not at first discourage.[3] And this was the part of generative linguistics that got everyone really excited. For if the techniques of transformational grammar lead us to meaning, we can uncover the nature of human thought. Moreover, if Deep Structure is innate – being dictated by Universal Grammar – then linguistic theory gives us unparalleled access to the essence of human nature. No wonder everyone wanted to learn linguistics.

What happened next was that a group of generative linguists, notably George Lakoff, John Robert Ross, James McCawley, and Paul Postal, pushed very hard on the idea that Deep Structure should directly encode meaning. The outcome, the theory of Generative Semantics (e.g. McCawley 1968, Postal 1970, Lakoff 1971), increased the "abstractness" and complexity of Deep Structure, to the point that the example *Floyd broke the glass* was famously posited to have eight underlying clauses, each corresponding to some feature of the semantics. All the people who admired *Aspects* for what it said about meaning loved Generative Semantics, and it swept the country. But Chomsky himself reacted negatively, and with the aid of his then-current students (full disclosure: present author included), argued vigorously against Generative Semantics. When the dust of the ensuing"Linguistics Wars" cleared around 1973 (Newmeyer 1980, Harris 1993, Huck and Goldsmith 1995), Chomsky had won – but with a twist: he no longer claimed that Deep Structure was the sole level that determines meaning (Chomsky 1972). Then, having won the battle, he turned his attention, not to meaning, but to relatively technical constraints on movement transformations (e.g. Chomsky 1973, 1977).

The reaction in the larger community was shock: for one thing, at the fact that the linguists had behaved so badly; but more substantively, at the sense that there had been a "bait and switch." Chomsky had promised Meaning with a capital M and then had withdrawn the offer. Many researchers, both inside and outside linguistics, turned away from generative grammar with disgust, rejecting not only Deep Structure but mentalism, innateness, and sometimes even combinatoriality. And when, later in the 1970s, Chomsky started talking about meaning again, in terms of a syntactic level of Logical Form (e.g. Chomsky 1981), it was too late: the damage had been done. From this point on, the increasingly abstract technical apparatus of generative grammar was of no interest to more than a tiny minority of cognitive scientists, much less the general public.

Meanwhile, various nonChomskyan traditions of generative grammar developed, most notably Relational Grammar (Perlmutter 1983), Head-Driven Phrase Structure Grammar (Pollard and Sag 1987, 1994), Lexical-Functional Grammar (Bresnan 1982, 2001), Formal Semantics (Partee 1976, Heim and Kratzer 19xx), Optimality Theory (Prince and Smolensky 1993), Construction Grammar (Fillmore et al. 1988, Goldberg 1995), and Cognitive Grammar (Lakoff 1987, Langacker 1987, Talmy 2000). On the whole, these approaches to linguistics (with the possible exception of Cognitive Grammar) have made even less contact with philosophy, psychology, and neuroscience than the recent Chomskyan tradition. My impression is that many linguists have simply returned to the traditional concerns of the field: describing languages, with as little theoretical and cognitive baggage as possible. While this is perfectly fine – particularly since issues of innateness don't play too big a role when you're trying to record an endangered language before its speakers all die – the sense of excitement and danger that comes from participating in the integration of fields has become attenuated.

4. The scientific mistake: syntactocentrism

So much for pure intellectual history. We now turn to what I think was an important mistake at the core of generative grammar, one that in retrospect lies behind much of the alienation of linguistic theory from the cognitive sciences. Chomsky did demonstrate that language requires a generative system that makes possible an infinite variety of sentences. However, he explicitly assumed, without argument (1965: 16, 17, 75, 198), that generativity is localized in the syntactic component of the grammar – the construction of phrases from words – and that phonology (the organization of speech sounds) and semantics (the organization of meaning) are purely "interpretive", that is, that their combinatorial properties are derived strictly from the combinatoriality of syntax.

In 1965 this was a perfectly reasonable view. The important issue at that time was to show that *something* in language is generative. Generative syntax had provided powerful new tools, which were yielding copious and striking results. At the time, it looked as though phonology could be treated as a sort of low-level derivative of syntax: the syntax gets the words in the right order, then phonology massages their pronunciation to adjust them to their local environment. As for semantics, virtually nothing was known: the only things on the table were the rudimentary proposals of Katz and Fodor (1963) and some promising work by people such as Bierwisch (1967, 1969) and Weinreich (1966). So the state of the theory offered no reason to question the assumption that all combinatorial complexity arises from syntax.

Subsequent shifts in mainstream generative linguistics stressed major differences in outlook. But one thing that remained unchanged was the assumption that syntax is the sole source of combinatoriality. Figure 1 diagrams the architecture of components in three major stages of Chomskyan syntactic theory: the *Aspects* theory, Principles and Parameters (or Government-Binding) Theory (Chomsky 1981), and the Minimalist Program (Chomsky 1995). The arrows denote direction of derivation.

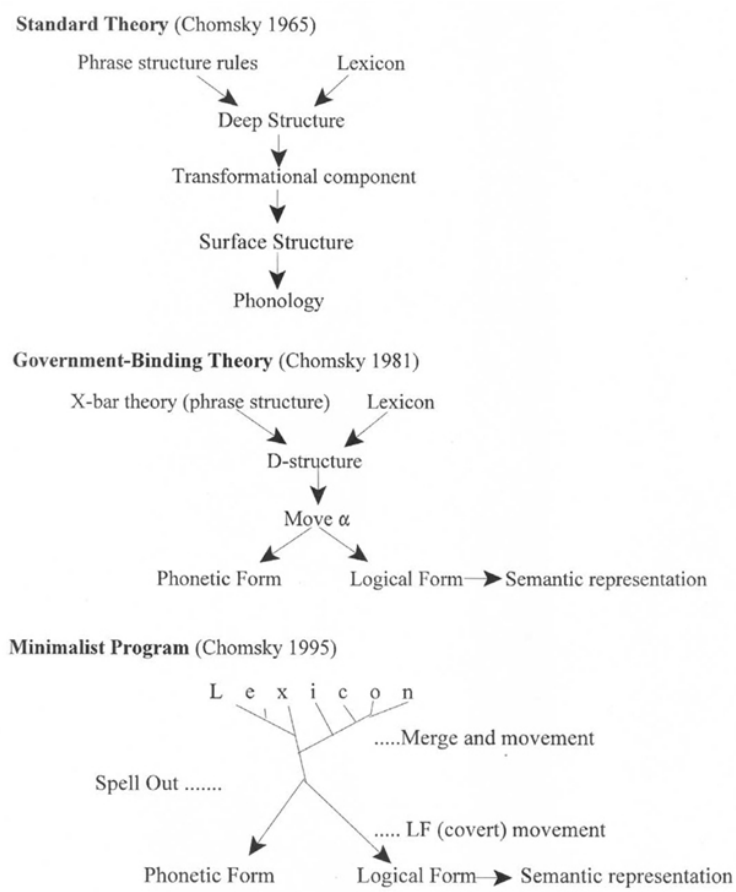

Figure 1. Architecture of Chomsky's theories over the years.

These shifts alter the components of syntax and their relation to sound and meaning. What remains constant throughout, though, is that (a) there is an initial stage of derivation in which words or morphemes are combined into syntactic structures; (b) these structures are then massaged by various syntactic operations; and (c) certain syntactic structures are shipped off to phonology/phonetics to be pronounced and other syntactic structures are shipped off to “semantic interpretation” to be understood. In short, syntax is the source of all linguistic organization.

I believe that this assumption of "syntactocentrism" – which, I repeat, was never explicitly grounded – was an important mistake at the heart of the field.[4] The correct approach is to regard linguistic structure to be the product of a number of parallel but interacting generative capacities – at the very least, one each for phonology, syntax, and semantics. As we will see, elements of such a "parallel architecture" have been implicit in practice in the field for years. What is novel in the present work is bringing these practices out into the open, stating them as a foundational principle of linguistic organization, and exploring the large-scale consequences.

5. Phonology as an exemplar of the parallel architecture

An unnoticed crack in the assumption of syntactocentrism appeared in the middle to late 1970s, when the theory of phonology underwent a major seachange. Before then, the sound system of language had been regarded essentially as a sequence of speech sounds. Any further structure, such as the division into words, was thought of as simply inherited from syntax. However, beginning with work such as Goldsmith (1979) and Liberman and Prince (1977), phonology rapidly came to be thought of as having its own autonomous structure, in fact multiple structures or *tiers*. Figure 2 provides a sample, the structure of the phrase *the big apple*. The phonological segments appear at the bottom, as terminal elements of the syllabic tree.

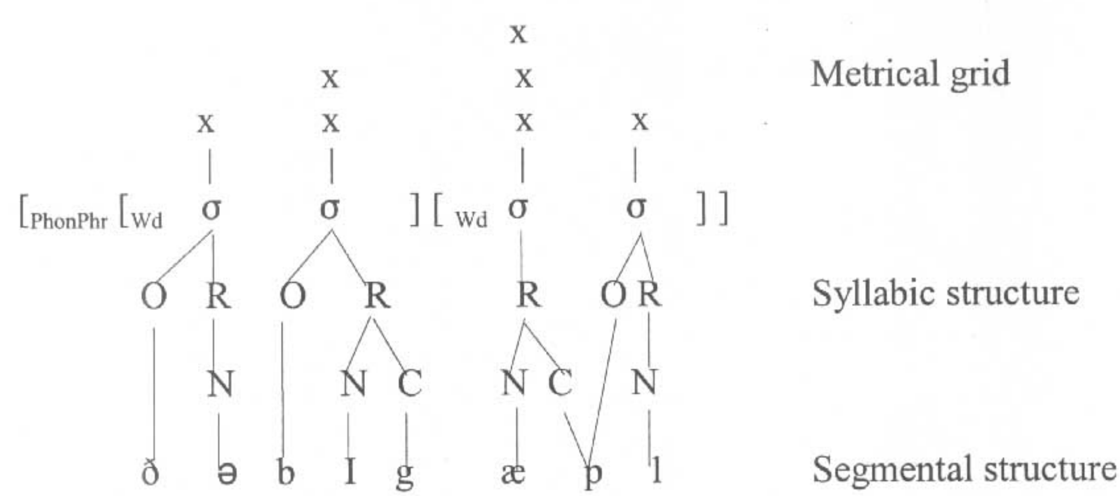

Figure 2. Phonological structure of *the big apple*.

There are several innovations here. First, syllabic structure is seen as hierarchically organized. At the center of the syllable (notated as σ) is a syllabic nucleus (notated *N*), which is usually a vowel but sometimes a syllabic consonant such as the *l* in *apple*. The material following the nucleus is the syllabic coda (notated *C*); this groups with the nucleus to form the rhyme (notated *R*), the part involved in rhyming. In turn, the rhyme groups with the syllabic onset (notated *O*) to form the entire syllable. Syllables are grouped together into larger units such as feet and phonological words (here, the bracketing subscripted *Wd*). Notice that in Figure 1, the word *the* does not constitute a phonological word on its own; it is attached (or cliticized) onto the word *big*. Finally, phonological words group into larger units such as phonological phrases. Languages differ in their repertoire of admissible nuclei, onsets, and codas, but the basic hierarchical organization and the principles by which strings of segments are divided into syllables are universal. (It should also be mentioned that signed languages have parallel syllabic organization, except that the syllables are built out of manual rather than vocal constituents (Klima and Bellugi 1979, Fischer and Siple 1990).)

These hierarchical structures are not built out of syntactic primitives such as nouns, verbs, and determiners; their units are intrinsically phonological. In addition, the structures, though hierarchical, are not recursive.[5] Thus the principles governing these structures are not derivable from syntactic structures; they are an autonomous system of generative rules.

Next consider the metrical grid in Figure 2. Its units are *beats*, notated as columns of *x*s. A column with only one *x* is a weak beat, and more *x*s in a column indicate a relatively stronger beat. Each beat is associated with a syllable; the strength of a beat indicates the relative stress on that syllable, so that for example in Figure 2 the first syllable of *apple* receives maximum stress. The basic principles of metrical grids are in part autonomous of language: they also appear, for instance, in music (Lerdahl and Jackendoff 1983), where they are associated with notes instead of syllables. Metrical grids place a high priority on rhythmicity: an optimum grid presents an alternation of strong and weak beats, as is found in music and in much poetry. On the other hand, the structure of syllables exerts an influence on the associated metrical grid: syllables with heavy rhymes (i.e. containing a coda or a long vowel) "want" to be associated with relatively heavy stress. The stress rules of a language concern the way syllabic structure comes to be associated with a metrical grid; languages differ in ways that are now quite well understood (e.g. Kager 1995, Halle and Idsardi 1995).

Again, metrical grids are built of nonsyntactic units. As they are to some degree independent of syllabic structure, they turn out to be a further autonomous "tier" of phonological structure.

At a larger scale of phonological organization we find prosodic units over which intonation contours are defined. These are comparable in size to syntactic phrases but do not coincide with them. Here are two examples.

(3) *Syntactic bracketing:*
[Sesame Street] [is [a production [of [the Children's Television Workshop]]]]
Prosodic bracketing (two pronunciations):
a. [Sesame Street is a production of] [the Children's Television Workshop]
b. [Sesame Street] [is a production] [of the Children's Television Workshop]

(4) *Syntactic bracketing*
[This] [is [the cat [that chased [the rat [that ate [the cheese]]]]]]
Prosodic bracketing:
[This is the cat] [that chased the rat] [that ate the cheese]

The two pronunciations of (3) are both acceptable, and other prosodic bracketings are also possible. However, the choice of prosodic bracketing is not entirely free, since for instance

[Sesame] [Street is a production of the] [Children's Television Workshop] is an impossible phrasing. Now notice that the first constituent of (3a) and the second constituent of (3b) do not correspond to any syntactic constituent. We would be hard pressed to know what syntactic label to give to *Sesame Street is a production of*. But as an *intonational* constituent it is perfectly fine. Similarly in (4), the syntax is relentlessly right-embedded, but the prosody is flat and perfectly balanced into three parts. Again, the first two constituents of the prosody do not correspond to syntactic constituents of the sentence.

The proper way to deal with this lack of correspondence is to posit a phonological category of Intonational Phrase, which plays a role in the assignment of intonation contours and the distribution of stress (Beckman and Pierrehumbert 1986, Ladd 1996). Intonation Phrases are to some degree correlated with syntax; their boundaries tend to be at the beginning of major syntactic constituents; but their *ends* do not necessarily correlate with the ends of the corresponding syntactic constituents. At the same time, intonational phrases have their own autonomous constraints, in particular a strong preference for rhythmicity and parallelism (as evinced in (2) for example), and a preference for saving the longest prosodic constituent for the end of the sentence.[6]

Another example of mismatch between syntax and phonology comes from contractions such as *I'm* and *Lisa's* (as in *Lisa's a doctor*). These are clearly phonological words, but what is their syntactic category? It is implausible to see them either as noun phrases that incidentally contain a verb or to see them as verbs that incidentally contain a noun. Keeping phonological and syntactic structure separate allows us to say the natural thing: they are phonological words that correspond to two separate syntactic constituents.

(5) *Syntactic structure*: $[_{NP}$ I] $[_{V}$ (a)m] $[_{NP}$ Lisa] $[_{V}$ (i)s]
Phonological structure: $[_{Wd}$ I'm] $[_{Wd}$ Lisa's]

Since every different sentence of the language has a different phonological structure, and since phonological structures cannot be derived from syntax, the usual arguments for combinatoriality lead us to the conclusion that phonological structure is generative. However, in addition to the generative principles that describe these structures, it is necessary to introduce a new kind of principle into the grammar, what might be called "correspondence rules" or "interface rules." These rules (I revert to the standard term "rules" rather than being obsessive about "f-rules") regulate the way the independent structures correspond with each other. For instance, the relation between syllable weight and metrical weight is regulated by an interface rule between syllabic and metrical structure; the relation between syntactic and intonational constituents is regulated by an interface rule between syntactic and prosodic structure.

An important property of interface rules is they don't "see" every aspect of the structures they are connecting. For instance, the rules that connect syllabic content to metrical grids are totally insensitive to syllable onset: universally, stress rules care only about what happens in the rhyme. Similarly, although the connection between syntax and phonology "sees" certain syntactic boundaries, it is insensitive to the depth of syntactic embedding, Moreover, syntactic structure is totally insensitive to the segmental content of the words it is arranging (e.g. there is no syntactic rule that applies only to words that begin with *b*). Thus interface rules implement not isomorphisms between the structures they relate, but rather only partial homomorphisms.

This is not to say that we should think of speakers as thinking up phonological and syntactic structures independently in the hope they can be matched up by the interfaces. That would be the same sort of mistake as thinking that speakers start with the symbol *S* and generate a syntactic tree, finally putting in words so they know what the sentence is about. At the moment we are not thinking in terms of production; rather we are stating the principles (of “competence”) in terms of which sentences are well-formed. We will get back to how this is related to processing in section 9.3.

Now the main point of this section. This view of phonological structure, developed in the late 1970s and almost immediately adopted as standard, is deeply subversive of the syntactocentric assumption that all linguistic combinatoriality originates in syntax. According to this view, phonological structure is not just a passive hand-me-down derived from low-level syntax: it has its own role in shaping the totality of linguistic structure. But at the time of these changes, no great commotion was made about this most radical aspect of the new phonology. Phonologists for the most part were happy to get on with exploring this exciting way of doing things, and for them, the consequences for syntax didn’t matter. Syntacticians, for their part, simply found phonology irrelevant to their concerns of constraining movement rules and the like, especially since phonology had now developed its own arcane technical machinery. So neither subdiscipline really took notice; and as the technologies diverged, the relation between syntax and phonology became a no-man’s-land (or perhaps only a very-few-man’s-land). Tellingly, as far as I can determine, in all of Chomsky’s frequent writings on the character of the human language capacity, there is no reference at all to post-1975 phonology – much less to the challenge that it presents to his overall syntactocentric view of language.

6. The syntax-semantics interface

I have treated the developments in phonology first because it is less controversial. But in fact the same thing happened in semantics. Over the course of the 1970s and 1980s, several radically different approaches to semantics developed: within linguistics, at least Formal Semantics (growing out of formal logic)(Partee 1976, Heim and Kratzer 19xx), Cognitive Grammar (Lakoff 1987, Langacker 1987, Talmy 2000), and Conceptual Semantics (Jackendoff 1983, 1990, Pinker 1989, Pustejovsky 1995), plus approaches within computational linguistics and cognitive psychology. Whatever their differences, all these approaches take meaning to be deeply combinatorial. None of them take the units of semantic structure to be syntactic units such as NPs and VPs; rather, the units are intrinsically semantic entities like objects, events, actions, properties, and quantifiers.[7] Therefore, whichever semantic theory we choose, it is necessary to grant semantics an independent generative organization, and it is necessary to include in the theory of grammar an interface component that correlates semantic structures with syntactic and phonological structures. In other words, the relation of syntax to semantics is qualitatively parallel to the relation of syntax to phonology. However, apparently no one pointed out the challenge to syntactocentrism – except the Cognitive Grammarians, who mostly went to the other extreme and denied syntax *any* independent role, and who have been steadfastly ignored by mainstream generative linguistics.

The organization of phonological structure into semi-independent tiers finds a parallel in semantics as well. Linguistic meaning can be readily partialled into two independent aspects. On one hand there is what might be called "propositional structure": who did what to whom and so on. For instance, in *The bear chased the lion*, there is an event of chasing in which the bear is the chaser and the lion is "chasee". On the other hand, there is also what is now called "information structure": the partitioning of the message into old vs. new information, topic vs. comment, presupposition vs. focus, and so forth. We can leave the propositional structure of a sentence intact but change its information structure, by using stress (6a-c) or various focusing constructions (6d-f):

(6) a. The BEAR chased the lion.
    b. The bear chased the LION.
    c. The bear CHASED the lion.
    d. It was the bear that chased the lion.
    e. What the bear did was chase the lion.
    f. What happened to the lion was the bear chased it.

Thus the propositional structure and the information structure are orthogonal dimensions of meaning, and can profitably be regarded as autonomous tiers. (*Foundations* proposes a further split of propositional structure into *descriptive* and *referential* tiers, an issue too complex for the present context.)

Like the interface between syntax and phonology, that between syntax and semantics is not an isomorphism. Some aspects of syntax make no difference in semantics. For instance, the semantic structure of a language is the same whether or not the syntax marks subject-verb agreement, verb-object agreement, or nominative and accusative case. The semantic structure of a language does not care whether the syntax calls for the verb to be after the subject (as in English), at the end of the clause (as in Japanese), or second in a main clause and final in a subordinate clause (as in German). As these aspects of syntax are not correlated with or derivable from semantics, the interface component disregards them.

Similarly, some aspects of semantics have little if any systematic effect in syntax. Here are a few well-known examples.

! The syntactic form of a question can be used to elicit information (7a), test someone's knowledge (7b), request an action (7c), or express sarcasm (7d). Thus these choices of illocutionary force are not mapped into syntactic structure.

(7) a. Where is my hat?
    b. (Now, Billy:) What's the capital of New York?
    c. Would you please open the window?
    d. Is the Pope Catholic?

- In example (8a), the interpretation is that Jill jumped multiple times. This aspect of interpretation does not arise from any single word in the sentence, nor from the syntactic structure. If we change the verb to *sleep*, as in (8b), we don't interpret the sentence as implying multiple acts of sleeping. If we change *until* to *when*, as in (8c), only a single jump is entailed.

(8) a. Jill jumped until the alarm went off.
 b. Jill slept until the alarm went off.
 c. Jill jumped when the alarm went off.

The standard account of this contrast (Talmy 2000, Verkuyl 1993, Pustejovsky 1995, Jackendoff 1997) is that the meaning of *until* is to set a temporal bound on an ongoing process. When the verb phrase already denotes an ongoing process, such as sleeping, all is well. But when the verb phrase denotes an action that has a natural temporal ending, such as jumping, then its interpretation is "coerced" into *repeated* action – a sort of ongoing process – which in turn can have a temporal bound set on it by *until*. For present purposes, the point is that the sense of repetition arises from semantic combination, without any direct syntactic reflex. (On the other hand, there are languages such as American Sign Language that have a grammatical marker of iteration; this will have to be used in the translation of (8a).)

- In the examples in (9), the "understood" subject of the sentence is not the entity normally denoted by the actual subject (Nunberg 1979).

(9) a. [One waitress says to another]:
 The ham sandwich wants another cup of coffee.
 [Interpretation: '*the person who ordered/is eating* the ham sandwich...']
 b. Chomsky is on the top shelf next to Plato.
 [Interpretation: '*the book by* Chomsky ...']

Such cases of "reference transfer" contain no syntactic reflex of the italicized parts of the interpretation. One might be tempted to dismiss these phenomena as "mere pragmatics", hence outside the grammatical system. But this proves impossible, because reference transfer can have indirect grammatical effects. A clear example involves imagining that Richard Nixon went to see the opera *Nixon in China* (yes, a real opera!), and what happened was that:

(10) Nixon was astonished to see himself sing a foolish duet with Pat.

The singer of the duet, of course, is the *actor playing* Nixon; thus the interpretation of *himself* involves a reference transfer. However, we cannot felicitously say that what happened next was that:

(11) *(Up on stage,) Nixon was astonished to see himself get up and walk out.

That is, a reflexive pronoun referring to the acted character can have the real person as antecedent, but not vice versa (Fauconnier 1985, Jackendoff 1992). Since the use of reflexive pronouns is central to grammar, reference transfer cannot be seen as "extragrammatical."

! One of the most controversial issues within generative grammar has been the syntactic status of quantifier scope. Consider the two interpretations of (12).

(12) Everyone in this room knows at least two languages.
   a. 'John knows English and French; Sue knows Hebrew and Hausa; ...."
   b. '... namely, Mandarin and Navajo.'

Should there be two different syntactic structures associated with these two interpretations? Chomsky 1957 said no; Chomsky 1981 said yes; Generative Semantics said yes; I am inclined to say no (Jackendoff 1996, *Foundations* chapter 12). The problem with finding two different syntactic structures is that it requires systematic and drastic distortions of the syntactic tree that never show up in the surface syntax of any language. The problem with having only one syntactic structure is that it makes the syntax-semantics interface more complex. The point to be made here is that the scope of quantification may well be a further example of the "dirtiness" of the interface between syntax and semantics; this continues to be an important issue in linguistic theory.

In each of these cases, a syntactocentric theory is forced to derive the semantic distinctions from syntactic distinctions. Hence it is forced into artificial solutions such as empty syntactic structure and elaborate movement, which have no independent motivation beyond providing grist for the semantics. On the other hand, if the semantics is treated as independent from syntax but correlated with it, it is possible to permit a less than perfect correlation; it is then an empirical issue to determine how close the match is.

If we abandon syntactocentrism, it is logically possible that there are aspects of semantics that have no impact on syntax but *do* have an effect on phonology through a direct phonology-semantics interface. Such a treatment is attractive for the correlation between prosody and information structure. For instance, the differences among (6a-c) do not show up in syntax at all – only in the stress and intonation in phonology, and in the focus-presupposition relations in semantics. In a syntactocentric theory, one is forced to generate these sentences with a dummy syntactic element [+Focus], which serves only to correlate phonology and meaning and does not affect word order or inflection. (Such was the approach in Jackendoff (1972), for instance.) But this element does no work in syntax per se; it is only present in order to account for the correlation between phonology and semantics. By introducing a direct phonology-to-semantics interface sensitive to this correlation, we can account for it with minimal extra machinery; but of course this requires us to abandon syntactocentrism.

7. The outcome: Parallel architecture

The argument so far has been that theoretical thinking in both phonology and semantics has proceeded in practice as though their structures are due to independent generative capacities. What has attracted far less notice among syntacticians, phonologists, and semanticists alike is that such an organization logically requires the grammar to contain interface components that correlate the independent structures. Carrying this observation through the entire architecture of grammar, we arrive at an overall picture like Figure 3: the grammar contains multiple sets of formation rules (the "generative" components), each determining its own characteristic type of structure, and the structures are linked or correlated by interface components.

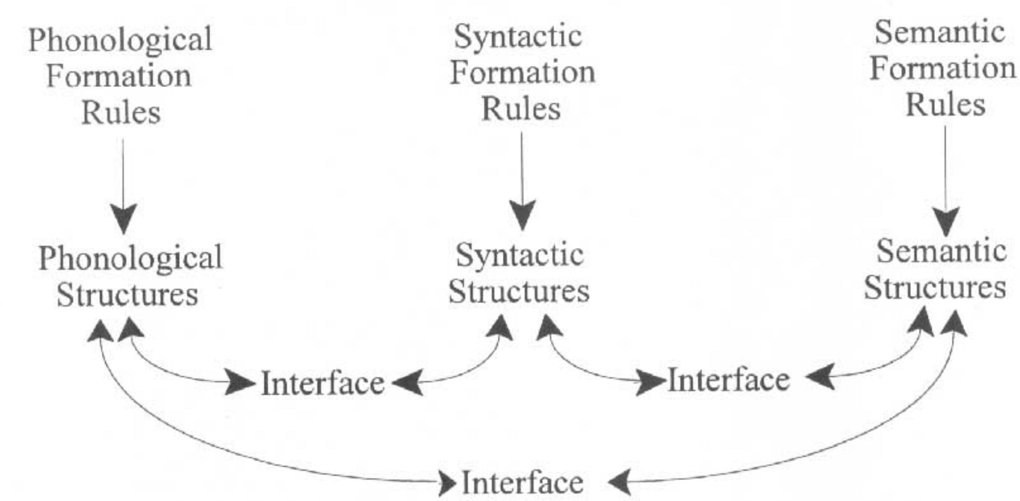

Figure 3. The parallel archetecture.

In the syntactocentric architecture, a sentence is well-formed when its initial syntactic tree is well-formed and all the steps of derivation from this to phonology and semantics are well-formed. In the parallel architecture, a sentence is well-formed when all three of its structures – phonological, syntactic, and semantic – are independently well-formed and a well-formed correspondence among them has been established by the interfaces.

One of the primary interface rules between phonology and syntax is that the linear order of units in phonology corresponds to the linear order of the corresponding units in syntax. One of the primary interface rules between syntax and semantics is that a syntactic head (such as a verb, noun, adjective, or preposition) corresponds to a semantic function, that the syntactic arguments of the head (subject, object, etc.) correspond to the arguments of the semantic function. The consequence of these two primary interface principles is that for the most part, syntax has the linear order of phonology but the embedding structure of semantics.

An illustration of some of these properties of the parallel architecture appears in Figure 4, the structure of the phrase *the cats*. The three independent structures are displayed side by side;[8] the subscripting indicates the connections established by the interfaces between the parts of the three structures. For example, the clitic pronounced $\xi\chi$ is coindexed with the determiner in the syntax and the with definiteness feature in semantics. Notice that the lowest nodes in the syntactic tree are syntactic features, not the customary notation *the cat-s*. The reasons for this are explained in the next section.

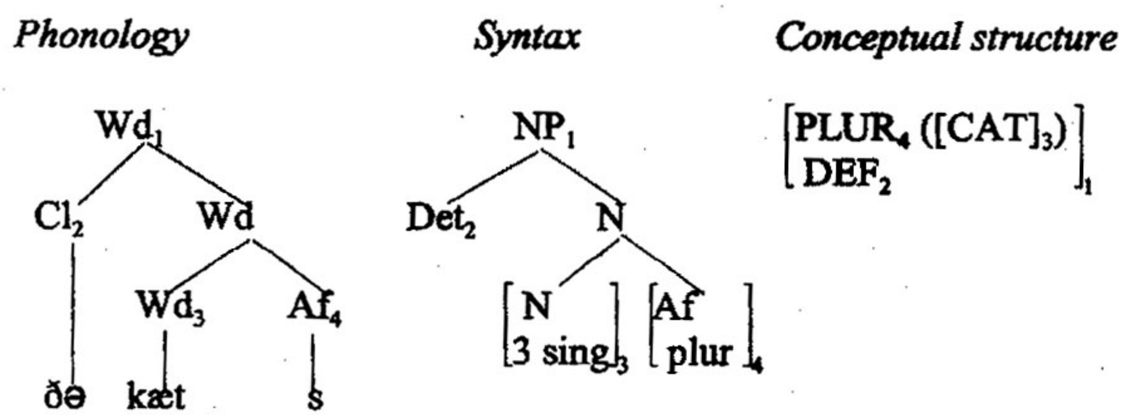

Figure 4. The structure of *the cats* in the parallel architecture.

The overall architecture laid out in Figure 3 provides a model within which many different theories of grammar can be embedded and compared. For instance, Figure 3 does not dictate whether the syntactic formation rules are along the lines of transformational grammar, the Minimalist Program, Head-Driven Phrase Structure Grammar, or many other alternatives.
Moreover, the syntactocentric framework is a version of Figure 3 in which the phonological and semantic formation rules are null, so that everything in phonological and semantic structures is determined only by their interfaces with syntax. The framework favored by many in Cognitive Linguistics minimizes or even eliminates the syntactic formation rules, so that syntax is determined entirely by meaning.

The organization into parallel generative components is not new here. In addition to the innovations in phonology discussed in section 3, Lexical-Functional Grammar divides syntax into two tiers, c-structure and f-structure; Autolexical Syntax (Sadock 1991) has a different division into morphosyntax and phrasal syntax; Role and Reference Grammar (Van Valin and LaPolla 1997) has, in addition to a morphosyntax/phrasal syntax division, the propositional/ information tier division in semantics, with interfaces going every which way among the tiers. In other words, various elements of this architecture are widely present in the literature. What is novel here is recognizing that this organization runs through the entire grammar, from phonology through semantics (and further, as we will see in section 9).[9]

It might well be argued that the standard syntactocentric framework has served the field well for forty years. Why should anyone want to give it up? A reply might come in five parts. First, no one has ever argued for the syntactocentric model. In *Aspects*, it was explicitly only an assumption, which quickly hardened into dogma and then became part of the unstated background. By contrast, the parallel architecture now *has* been argued for, in part based on well-established results in phonology and semantics, which have never played a role in syntactocentric argumentation.

Second, an advocate might argue that the syntactocentric model is a priori simpler: why should we admit so many different components into the grammar? The reply would be that the choice among theories must be determined by empirical adequacy as well as a priori simplicity. If the parallel architecture allows a more perspicuous account of, say, intonation contours or the relation of focus to stress, these are arguments in its favor.

A third point concerns the relation of syntax and semantics. Since syntax is now not responsible for determining every semantic distinction, it is to some degree liberated from semantics and can therefore be considerably simplified. However, some compensating complexity must be introduced into the syntax-semantics interface, so it doesn't disappear from the grammar entirely. It now becomes an empirical question how to parcel the complexity out, and this question can be addressed; it is not just an issue of opinion or preference (see e.g. Jackendoff 1992, Culicover 1999, Culicover and Jackendoff 1995, 1997, 1999). At the same time, syntax does not go away entirely (as opponents of Chomsky would often like). The syntax of a language still has to say where the verb goes, whether the verb agrees with the subject, how to form relative clauses and questions, and so on. The differences among languages in these respects are not predictable from semantics, and children have to learn them.

A fourth point concerns the nature of Universal Grammar. In the parallel architecture, the issues of acquisition and innateness don't go away, they are exactly the same, namely, How does the child acquire the grammar of its native language on the basis of environmental evidence? However, as just suggested, the sorts of questions that most often arise concern the balance of power among components. We don't find ourselves invariably asking: what do we have to add to syntax to account for such-and-such a phenomenon? Rather, we find ourselves asking: in which component does this phenomenon belong? Is it a fact of syntax, of semantics, or of the interfaces? And to what extent is it realistic to attribute such a bias to the child learning the language?

A final point concerns not linguistic structure itself but its connection to the rest of the theory of the brain/mind. On the face of it (at least in my opinion), one should favor approaches that permit theoretical integration. Section 9 will show four ways that the parallel architecture invites such integration but the syntactocentric theory does not.

8. The lexicon and the words vs. rules controversy

Every theory of language has to take a word to be a stored complex of phonological, syntactic, and semantic features or structure; commonly the store of words is called the *lexicon*. . However, theories differ in the role of the lexicon in the construction of sentences. In all of the syntactocentric architectures shown in Figure 1, words are inserted into syntactic trees at the beginning of a syntactic derivation, at the point when syntactic trees are being built and before they begin to be manipulated and fed to phonology and semantics. Thus the traditional notation for trees in Figure 5a is actually intended as an abbreviation of Figure 5b, in which the lexical items are spelled out in full. The consequence is that the syntax is carrying around with it all the phonological and semantic features of words, which are totally invisible to syntactic rules and are of use to the grammar only when handed over and “interpreted” by the proper component.

Figure 5a. Traditional notion for *the cat*. b. What the traditional notation abbreviates.

The parallel architecture, by contrast, insists that each kind of feature belongs only in its own structure. Thus the traditional syntactic notation in Figure 5 is formally incoherent, because it has phonological and semantic features in a syntactic structure. Thus it is formally impossible to insert full lexical items into syntactic structure. How then do words get into linguistic structures? The answer is that each of the three structures making up a word is inserted into its proper structure, and each of them carries with it an index that connects it to the others. So, for example, the word *cat* is notated as Figure 6; its contribution to the larger structure in Figure 4 should be evident.

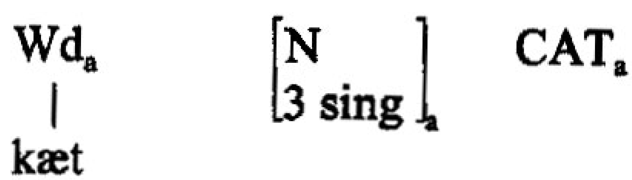

Figure 6. The structure of the word *cat*.

Thus a word is best regarded as a type of interface rule that establishes partial correspondences among pieces of phonological, syntactic, and semantic structure (each piece in turn conforming to the formation rules of its own component). In other words, the language does not consist of a lexicon *plus* rules of grammar; rather, lexical items are *among* the rules of grammar – very particular rules to be sure, but rules nonetheless.

This treatment of the lexicon offers an attractive account for a number of previously troublesome phenomena. For instance, consider an idiom such as *kick the bucket*. This can be treated as a lexically listed VP that is coindexed with phonology in the normal way, but which lacks indexes connecting the individual words to semantics: instead, the VP as a whole is coindexed with the semantic structure *DIE*. As a consequence, the individual words *kick, the*, and *bucket* do not contribute individually to meaning. This is precisely what an idiom is supposed to be: a stored unit in which the words do not have their normal meaning.

A sort of converse is found in irregular morphology. Consider something like the irregular plural *feet*. It has to be listed syntactically as a plural noun, and the two syntactic parts are coindexed in the normal way to semantics: the word denotes multiple entities of the type *FOOT*. However, the syntactic parts are not connected in normal fashion to phonology; rather the whole syntactic complex is coindexed with the undifferentiated lump *feet* in phonology.

Notice by contrast how the *regular* plural is coded in Figure 4. The regular plural consists of a piece of meaning, namely plurality, plus a piece of syntax, namely an affix attached to nouns, plus a piece of phonology, namely a suffix *s* or *z* or *χz*, the choice determined contextually. That is, the regular plural has all the same parts as a word, and it determines a connection between them. We can notate this as a lexical item along the lines of Figure 7. (The italicized bits denote contextual features that determine how this item is combined with its environment.) The contribution of this item to the overall structure in Figure 4 is entirely parallel to the contribution of the word *cat*.

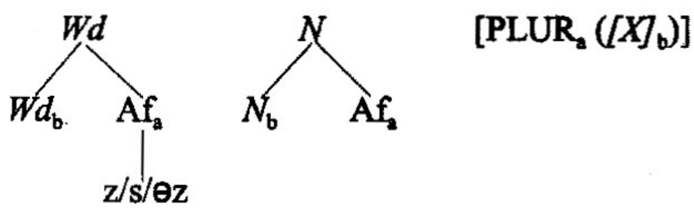

Figure 7. The English regular plural as a lexical item.

This view of regular morphology puts a new and unexpected spin on the by now hoary "words vs. rules" controversy (e.g. Rumelhart and McClelland 1986, Elman et al. 1996, Pinker 1999). Traditionally, the issue is taken to be:

- Everyone agrees that irregular plural nouns like *feet* have to be listed in the lexicon. Are regular plurals all listed as well, or is there a separate *rule* for the regular cases that says "To form the plural of a noun, add *-z*"? And, therefore, when children learn to form regular plurals, are they learning something qualitatively different from learning the rough-and-ready generalizations among irregular plurals?

On the present view, words *are* rules – interface rules that help connect phonological, syntactic, and semantic structures. And Figure 7, the "rule" for the regular English plural affix, is qualitatively no different. Its contextual features are qualitatively not unlike those of, say, transitive verbs. It combines with nouns the same way a transitive verb combines with its object. Thus the formation of regular plurals is an instance of ordinary combinatoriality. In this approach, the issue comes to be restated like this:

- Are regular plurals all listed, or is there a separate *lexical item* that encodes the regular affix, which combines with any singular noun to form a plural noun? And, therefore, when children learn to form regular plurals, are they learning this new lexical item by extracting it as a regularity from the contexts in which it appears – in the same way that they extract verbs from the phrasal contexts in which they appear?

I submit that even to the most committed of connectionists, this latter way of framing the question can hardly be objectionable.

In fact, the advocates of rules, such as Pinker, have not made the case nearly as strong as it can be. The connectionist argument has been to the effect: We can make a device that learns all English past tenses without making use of a rule, and we can find evidence from acquisition and processing that supports this account. The best version of the anti-connectionist argument has been: Connectionist modeling offers important innovations over standard models of language in dealing with case-by-case learning and analogy for the irregular past tenses. But -- you still need rule learning to account for children's acquisition of regular past tense, and we can find evidence from acquisition and processing that supports *this* account. The problem is that the debate has often been framed as though only the past tense were at issue, while the subtext behind the connectionist position is that if *this* can be learned without rules, then it is a good bet that the rest of language can be too.

But not only the past tense is at stake. To deal with the whole of language, it is necessary to account for the creative formation of things like verb phrases and relative clauses – linguistic entities that cannot be listed in the lexicon. On the present view, the way that the regular past tense affix combines with verbs is exactly like the way nouns combine with relative clauses and the way noun phrases combine with verbs and the way subordinate clauses combine with main clauses – it is just another case of free combinatoriality. In the decade and a half since the original connectionist past tense model, there has been no demonstration that the model scales up to acquisition of the full free combinatoriality of language – the issue that grounds generative linguistics.

At the same time, notice that within the parallel architecture, the terms of the dispute become far less contentious. The regular past tense is no longer a qualitatively different phenomenon from words: words are a type of rule, and the posited regular past tense morpheme, Figure 7, is in the relevant respects just like a word. (It differs only in that it is grammatically smaller and it requires a word as its grammatical host.) So the issue is only whether there is such a separate lexical item, not whether there are two wholly different kinds of linguistic animal, namely words and rules. Thus in the end the fate of the past tense doesn't seem like such a big deal.

9. Four ways the parallel architecture helps integrate linguistics with cognitive neuroscience

The parallel architecture may be an intriguing technical alternative to the Chomskyan orthodoxy in linguistics, but is there any reason why it should be of interest to anyone other than linguists? The end of the previous section may have begun to offer some hints. This section will sketch out a little more fully some ways in which the parallel architecture offers opportunities to unify linguistics with the other cognitive sciences.

9.1. The place of the parallel framework in the larger architecture of the mind/brain.

To sum up the larger picture: The parallel architecture claims that language is organized into a number of semi-independent combinatorial systems, each of which has its own organizing principles. These systems are linked by systems of interface principles. Interface principles establish a correlation between pieces of structure in two (or more) of the combinatorial systems. Some interface principles deal with large-scale and general correspondences such as the parallel between linear order in syntax and in phonology. On the other hand, some of them are extremely specialized, for instance individual words, idioms, and regular affixes. The interface principles as a whole do not implement an isomorphism between the structures they connect. Rather they implement a partial homomorphism, a "dirty" correspondence in which not all parts of the structures in question are correlated, and in which many-to-many mappings are altogether common.

This conception of the interfaces within language is perfectly in tune with the way linguistic structures connect to the rest of the mind. Consider how phonology interacts with the auditory system in speech perception and with the motor system in speech production. As is well known (to the dismay of fifty years of computer scientists working on automated speech recognition), the mapping between a frequency analysis of the speech signal and the phonological structure of an utterance is frighteningly complex. In particular, some aspects of the speech signal play no role in phonological structure and must be factored out, for instance the individual timbre of the speaker's voice, the speaker's tone of voice, and the speed of production, not to mention ambient noise. These aspects of the speech signal are put to use for other cognitive purposes, but not for speech. Moreover, having factored all these things out from the acoustic signal, still not every part of the phonological structure is predictable from what is left: most prominently, word boundaries are not present as pauses in the signal. Thus the auditory-to-phonological mapping has the same general characteristics as the interfaces inside language: it establishes a "dirty" correspondence between certain aspects of two disparate mental structures.

Speech production has similar properties. Not every aspect of phonological structure corresponds to an aspect of the motor control involved in operating the vocal tract. In particular, word boundaries do not correspond at all consistently to pauses in production. And not every aspect of motor control is controlled by phonological structure. For instance, one can talk intelligibly with a pipe in one's mouth, which hugely distorts the motor commands involved in speech without changing the phonological structure a bit. And of course the same muscles in the vocal tract are used for chewing, swallowing, and so on. Without going into more detail, it should be clear that again the same sort of interface is in play here.

Next consider the visual system. Beyond the very early levels of vision, there is little detailed theory of the f-knowledge involved in vision – the necessary levels of representation and so on (I take Marr (1983) to have been attempting to lay out such a theory, but the enterprise has been largely abandoned since his death). On the other hand, the neuroscience of vision reveals a qualitatively similar picture: numerous independent brain areas, each specializing in a particular aspect of vision such as shape, motion, color, and spatial relations, each interacting with certain others by dedicated pathways, and no area where "it all comes together" to form a full representation of the visual field. This has precisely the flavor of the parallel architecture in linguistics, where the notion of a "sentence" or "phrase" is distributed among several structures, communicating with each other via dedicated interfaces, as shown in Figure 4.[10]

A syntactocentric architecture, by comparison, shows no resemblance to the rest of the mind/brain. A master "computational system" that generates syntactic structures, which in turn determine phonological structures and meanings, simply has no known parallel in the brain. Even the connection of language to speech is markedly different from the connections among the components inside of language.

9.2. The role of semantics

Another important advantage of the parallel architecture is the connection of semantics to the rest of the mind/brain. *Foundations of Language* (chapters 9 and 10) advocates that if generative grammar is to truly adopt the mentalist stance, this stance must be applied to meaning as well. According to this stance, the basic function of language is to convert thoughts into communicable form; the virtue of human language over other natural communication systems is that it is so broad in the range of messages it can convey. Each of the indefinitely many sentences of a language conveys a different thought. Since not all these thoughts can be stored in a single head, it is necessary that thoughts be constructed combinatorially. Therefore a goal for semantic theory is to uncover the combinatorial system underlying human concepts. Such a goal converges with important trends in psychology and philosophy.

However, another influential strain in semantics (and the predominant one in Anglo-American philosophy, dating back to Frege (1892) and shared by people as different as David Lewis (1972), Hilary Putnam (1976), John Searle (1980), and Jerry Fodor (1987)) takes it that semantics is the study of the connection of language to the *world*. On this view, a proper semantics has to be concerned above all with how the noise *kΦt* is connected with cats. How language users make that connection is quite a different issue (and to many semanticists, not of interest). There is no room here to begin the critique of this view; chapters 9 and 10 of *Foundations* take up the argument in detail. My overall conclusion is that even if it is worthwhile undertaking such a "realist" semantics, the enterprise of discovering how language users do it is also worthwhile. I don't care whether you call the latter enterprise semantics or shmenatics or whatever: it's this enterprise whose central issues intercalate naturally with those of generative linguistics, cognitive psychology, and neuroscience. Just to be clear, I will call this enterprise *conceptualist semantics*.

Conceptualist semantics requires us to rethink the traditional issue of reference, which takes as its starting point the unshakeable intuition that the phrase *my cat* does indeed pick out an individual in the world. In a mentalist linguistic theory, the language user's linguistic system connects the phonological string /maykΦt/ to the concept of a feline animal, and to the concept of this feline animal being possessed by the speaker of the phrase. How then does the language user get from there to the actual individual out there in the world? The answer in brief is that it isn't just language users who have to connect something in their head to a sense of individuals in the world: any organism with a visual system about like ours (e.g. babies and apes) has precisely the same problem. The environment acting on the visual system produces some set of activations in the brain, resulting in the organism experiencing real objects out there. In other words, conceptualist semantics allows us to recognize that the problem of reference is not a problem about *language*, it's at bottom a problem about *perception* and *cognition* which has to be solved by psychology and neuroscience. By contrast, conventional realist theories of reference divorce reference from the mind and make no contact whatsoever with research on perception.

In order for the system of meaning to be influenced by perception, of course, there has to be an interface between conceptual/semantic structure and the "upper end" of the perceptual systems, where "the world" (i.e. the perceiver's conceptualization of the physical world) is organized in terms of stable three-dimensional objects that are located in space with respect to the perceiver and each other. This interface too can be shown to have the standard characteristics: it is a partial homomorphism between the quasi-algebraic format in which linguistic meanings are encoded and the quasi-geometric/topological format(s) in which spatial understanding is encoded. Thus at the semantic end of the language faculty, just as at the phonological end, the relation between language and the rest of the mind is of the same general character as the interfaces within the language faculty itself.

Studying the conceptual system as a combinatorial system leads to the same questions about acquisition as studying syntax. How does the child learning language acquire the meanings of all those thousands of words on the basis of experience, both perceptual and linguistic? What perceptual biases and innate structures does the child bring to the task of interpreting the world? Here conceptualist semantics makes contact with a rich literature on word and concept learning and its innate bases (Bloom 2000, Carey 1985, Spelke et al. 1994, Baillargeon 1986, Gleitman and Landau 1994, Keil 1989, to mention only a few parochial examples). Moreover, since humans doubtless share with monkeys and apes at least the parts of the conceptual system dealing with physical space and perhaps some of the parts dealing with social relations and other minds, conceptualist semantics further makes contact with research on primate cognition (Premack 1976, Tomasello 2000, Hauser 2000, Povinelli 2001, Cheney and Seyfarth 1990, Köhler 1927).

Again, these are issues that conventional realist semantics cannot address. Nor are they particularly accessible to semantics studied in a syntactocentric linguistic theory. For if the combinatorial properties of semantics were completely attributable to the combinatorial properties of syntax, then it would be impossible for nonlinguistic organisms to have combinatorial thoughts. There are of course important strains of philosophy that have embraced this view, identifying the capability for thought with the capability for overt language (Descartes comes to mind, for instance). But I think contemporary cognitive neuroscience has outgrown such a view, and linguistics ought to be able to follow suit gracefully.

9.3. The relation of grammar to processing

A theory of linguistic competence is supposed to simply define the permissible structures in the language, without saying how those structures are produced in real time. However, as pointed out in section 1, a competence theory ought to lend itself to being embedded in a theory of processing: we ought to be able to say how the f-knowledge that constitutes the competence theory is actually put to use.

There turns out to be an inherent structural reason why competence has to be isolated from performance in the syntactocentric view of language. If we flatten out and simplify all the syntactocentric architectures in Figure 1, they all have a logical directionality proceeding outward from syntax in the middle:

(13) *Logical directionality of syntactocentric architecture*
sound Z phonology Z syntax Ψ meaning

What I mean by logical directionality is that the possible phonological structures and meanings cannot be determined without first determining syntactic structures. Syntacticians may insist that they are being "metaphorical" when they talk about things happening "before" and "after" other things in a derivation; but the *logical* dependence is there nevertheless. Now contrast this to the logical directionality of language processing: language perception goes consistently from left to right, and language production from right to left.

(14) a. *Logical directionality of language perception*
sound Ψ phonology Ψ syntax Ψ meaning

b. *Logical directionality of language production*
sound Z phonology Z syntax Z meaning

Hence there is no way that the logical directionality in (13) can serve the purposes of both perception and production. Going from syntax to phonology in (13) seems inherently like production – but only *part* of production; going from syntax to semantics in (13) seems inherently like perception – but only part of it.

The parallel architecture, by contrast, is inherently *non*directional. The "information flow" between sound and meaning is through the sequence of interfaces, each of which is a system of correlations between two structures, not a derivation of one from the other. The correlations can be used in either direction (which is why they are drawn with double arrows in Figure 3). This makes it possible to think of speech perception as a process where structures are activated first at the auditory end of the chain of structures, "clamped" by the environmental input. The interfaces propagate activation rightward through the chain, each interface principle creating a partial resonance between the structures it connects. Eventually the structured activation reaches semantic structure, at which point it can interact with the hearer's understanding of the context to produce the understanding of the heard utterance. Similarly, in speech production, the speaker begins with a thought to convey, i.e. meaning is "clamped" by the speaker's communicative intent. Then the interface principles propagate activation leftward through the chain, eventually activating motor control of the vocal tract and producing speech. Crucially, except for the auditory and vocal parts of the chain, the very same structures and the very same interface principles are invoked in perception and production, just in opposite directions.

There is no need in this system for all of one level to be totally processed before activation of the next level sets in. Any activation of a level, no matter how incomplete, if it can be detected by the next interface, will start to propagate to the next level in the chain. Thus processing can be thought of as "incremental" or "opportunistic" rather than rigidly regulated. In addition, since the interfaces are trying to achieve "resonance", i.e. optimal mapping between levels, there is ample room in the processing theory for feedback in processing – semantics affecting syntactic processing in perception, and vice versa in production.

A crucial tenet of this theory, though, is that *the rules of grammar are the only source of information flow in language processing*. For example, knowledge of context cannot directly affect phonological processing, because there are no interface rules that directly relate contextual understanding to phonological structure. On the other hand, context can *indirectly* affect phonological processing – via the interfaces linking them through semantics and syntax. The prediction is that such feedback will take effect some time after constraints directly from phonology, because it has to go up the chain of interfaces and down again. On the whole such a prediction seems consistent with the experimental literature (Cutler and Clifton 1999, Levelt 1989); *Foundations* works out many details, in particular the relation of long-term memory to working memory during language processing.

The role of the lexicon in the processing theory is entirely parallel to that in the competence theory. Recall that words are little interface rules, providing partial routes for mapping between sound and meaning. Now consider the logic of language perception. The auditory system and the interface from audition to phonology produce some string of speech sounds in the hearer's head, and this activates a call to the lexicon: "Do any of you guys in there sound like *this*?" And various items raise their hand, i.e. get activated. At this point the processor has no way of knowing which of these items is semantically appropriate, because no contact has yet been made with semantics. However, each item over time activates a connection to potential syntactic and semantic structures, which can be integrated with previous words and with context to determine which candidate word makes most sense in context. This scenario corresponds precisely to the results in lexical access experiments (Swinney 1979, Tanenhaus et al. 1979), in which every possible sense of a given phonological string is activated at first, later to be pruned down by semantic context.

A parallel story can be told for speech production. The speaker has activated some conceptual structure that s/he wishes to communicate. The first step is to call the lexicon: "Do any of you guys in there mean *this*?" And various items raise their hand. All the lexical retrieval and speech error literature now comes to bear in showing us the flow of information from this point to actual vocal production; for the most part it proves to correspond nicely to the options made possible by the components of the parallel framework (Levelt 1989, 1999).

It is significant that the parallel architecture accords words a very active role in determining the structure of sentences, in concurrence with evidence from the psycholinguistic literature. By contrast, the syntactocentric architecture views words as much more passive: they simply ride around at the bottom of syntactic trees, while the derivational rules of syntax do all the interesting work. Thus again, in the area of the lexicon, the syntactocentric framework again makes it hard to connect competence and performance.

The conclusion here is that the parallel architecture permits a far closer relation between competence and performance theories. The rules of the language, including the words, are posited to be precisely what the processing system uses in constructing mappings between sound and meaning. This opens the door for a two-way dialogue between linguistics and psycholinguistics. Linguistics has always dictated the structures that psycholinguistics should be investigating. But now there is the possibility that psycholinguistic experiments may help determine what component of the grammar is responsible for a particular phenomenon. For instance, PiZango et al. (1999) have shown that aspectual coercion (e.g. the sense of repetition in *Jill jumped until the alarm went off*, example (8a) above) causes a processing load at a point in time that is appropriate to semantic, not syntactic processing. This result conforms to the theoretical claim that aspectual coercion is a matter of adjusting semantic well-formedness, not a matter of syntactic deletion of an iterative morpheme. In short, the idealization of a competence theory is not a rigid abstraction, it is rather a convenient methodological move, to be bridged freely when the occasion arises. (More detail on the relation of rules to processing is offered in *Foundations*, chapter 7.)

9.4. Evolution of language

Let us return to a point from section 2. If Universal Grammar is a human cognitive specialization, it has to be transmitted by genes that have emerged in the course of our evolutionary divergence from the chimpanzees. Of course the actual evidence for the evolution of the language faculty is practically nonexistent. There is some evidence about the evolution of the human vocal tract (Fitch 2000), but the ability to make speech sounds is only one element of language – and of course there are signed languages, which don't involve speech at all. In addition, it has begun to look like many of the mechanisms for auditory perception are already in place in other mammals (Hauser, Chomsky, and Fitch 2002). But the real issue is: how did the ability to systematically map combinations of concepts into sequences of speech sounds and back again develop in our species, and how did the ability to *learn* such systematic combinatorial mappings develop?

In the absence of evidence, we would like at least to be able to tell a plausible story about the emergence of Universal Grammar, an important aspect of which is the overall architecture of language. In particular, we would not like to have to explain language through miraculous emergence, given that (as argued by Pinker and Bloom 1990) it has the hallmarks of being shaped by natural selection. Pinker and Bloom, however, do not offer any concrete proposals as to how language evolved. As is well known, Chomsky himself has been notably evasive on the issue of the evolution of the language faculty, often seeming to cast aspersions on the theory of natural selection (Newmeyer 1998 collects representative quotes). Chomsky is correct that other factors besides natural selection play a role in evolution, for instance environmental pressure and the biochemistry of proteins. Nevertheless there is nothing in these other factors that provides any helpful hints on what brought about the emergence of language.

The logic of the syntactocentric architecture suggests a reason why such evasion has been necessary. The problem is in providing a route for incremental evolution, such that some primitive version of the faculty could still be useful to the organism. In the syntactocentric architecture, everything depends on syntax. Meaning cannot have evolved before syntax, because its structure is totally dependent on the syntactic structure from which it is derived. For the same reason, phonological structure cannot have evolved before syntax. Thus the complexity of syntax had to evolve before the complexity of the other components. But what would confer an adaptive advantage on a syntactic faculty that just generated meaningless and imperceptible syntactic structures? And what would enable children to acquire such syntactic structures if there were no perceptible output to which they could attach it? We quickly see that, at this very crude level at least, the syntactocentric theory is stuck: there is no logical way to build it incrementally, such that the earlier stages are useful.

The parallel architecture offers a better alternative. The system of concepts that language expresses is an independent generative component in the mind/brain. Since it is taken to exist, in some degree at least, in other primates as well, it also could have existed in our ancestors, prior to language. That is, our ancestors had interesting thoughts, but lacked any way to say them: meaning therefore would be the first generative component of language to emerge (a similar view is urged by Hauser 2000).

Most speculation on language evolution goes on to say that the earliest stage would have been the symbolic use of simple vocalization, without grammatical organization. Such a stage is theoretically impossible in the syntactocentric theory, since even single-word utterances have to arise from syntactic structure. But such a stage is quite natural in the parallel architecture: it consists of stored associations of vocalizations and concepts, i.e. a "paleo-lexicon." Lexical items that can serve on their own as utterances still exist in modern language, for instance *hello, oops, ouch,* and *gadzooks*. The provision for them in language might be viewed as an evolutionary relic of this earliest stage.

Assuming that there would be an adaptive advantage to a larger number of signals, a regimentation of vocalization along the lines of phonological structure would be the next generative component of language to emerge. Phonological organization in effect digitizes vocalizations, making a large vocabulary reliably discriminable and learnable. (Proto-)words at this point would be simply duples of phonological and semantic structure, without syntax.

A next innovation might be the provision of concatenating words into larger utterances. However, when words are concatenated, the issue arises of how the meanings of words in a string are related to each other semantically. In a string like *eat apple Fred*, it is pretty clear on pragmatic grounds that Fred is eating the apple and not the reverse. But pragmatics can only go so far: in *chase lion bear*, who is the chaser? Something as elaborate as English syntax is not entirely necessary to fix this. One can actually get considerable mileage from simple functional principles of linear word order. For example, the principle "Agent First" would tell us that the lion is chasing the bear and not the reverse. Such a principle is a straight phonology-to-semantics mapping, relating linear order to semantic function. And principles like this appear to be widespread in pidgin languages (Bickerton 1981) and the grammars of speakers who have acquired their languages late in life, after the sensitive period (Klein and Perdue 1997).

Finally, principles like Agent First have their limitations too. One can imagine the capacity for modern syntactic structure evolving *last*, as a way of making more complex semantic relations among the words of an utterances more precisely mappable to linear word order in phonology. That is, syntax comes along in evolution as a refinement, a "supercharger" of a preexisting interface between phonology and semantics. This seems exactly appropriate to its function within the parallel architecture.

In short, the parallel architecture cannot tell us exactly how language evolved – I don't think anything can ever tell us that. But it does afford a far more plausible hypothesis than the syntactocentric architecture (*Foundations*, chapter 8 develops this story in considerably more detail). Thus the parallel architecture opens the door for linguistics to participate far more fully in the mainstream of evolutionary psychology, yet another desirable connection.

10. Conclusions

Putting this all together, the parallel architecture makes it possible both to internally integrate linguistic theory, establishing the proper relation between phonology, syntax, semantics, and the lexicon, and also to integrate linguistic theory more comprehensively with the brain and with biology. In addition, by liberating semantics from its syntactic shackles, the parallel architecture makes it possible to develop a fully psychological theory of meaning and its relation to perception. As observed above, these points of connection were precisely what early generative grammar promised but ultimately couldn't bring off; I have tried to show here why syntactocentrism was a major reason behind this disappointment.

Of course, to propose a new architecture only begins the work. It opens major questions about exactly what components the grammar requires and what interfaces connect them. Vast numbers of phenomena have been studied in the context of the traditional architecture; to what extent can the analyses proposed there be duplicated or even improved upon? In particular, a thorough overhaul of syntactic theory is necessary, in order to overcome decades of accretions motivated solely by syntactocentric assumptions (Culicover and Jackendoff (forthcoming) begin to undertake this task). Perhaps the hardest part of all this will be maintaining a sense of global integration, keeping the subdomains of the field in closer touch than they have recently been.

But linguistics alone cannot sustain the weight of the inquiry. We need all the help we can get from every possible quarter. And in return, one would hope that linguistic theory might be a more fruitful source of evidence and puzzles for other fields. Above all, my aspiration for *Foundations* is that it can help encourage the necessary culture of collaboration.

References

Baillargeon, Renée (1986) Representing the existence and the location of hidden objects: Object permanence in 6- and 8-month-old infants. *Cognition* 23, 21-41.

Baker, C. L., and John J. McCarthy (eds.) (1981) *The logical problem of language acquisition*. Cambridge, MA: MIT Press.

Beckman, Mary, and Janet Pierrehumbert (1986) Intonational structure in English and Japanese. *Phonology* 3, 255-309. Bellugi et al. 1994

Bellugi, Ursula, Howard Poizner, and Edward S. Klima (1989) Language, modality, and the brain. *Trends in Neurosciences* 12, 380-388.

Bickerton, Derek (1981) *Roots of language*. Ann Arbor, MI: Karoma.

Bierwisch, Manfred (1967) Some semantic universals of German adjectivals. *Foundations of Language* 3, 1-36.

Bierwisch, Manfred (1969) On certain problems of semantic representation. *Foundations of Language* 5, 153-184.

Bloom, Paul (2000) *How children learn the meanings of words*. Cambridge, Mass.: MIT Press.

Bresnan, Joan W. (ed.) (1982) *The mental representation of grammatical relations*. Cambridge, MA: MIT Press.

Bresnan, Joan W. (2001) *Lexical-functional syntax*. Oxford: Blackwell.

Carey, Susan (1985) *Conceptual Change in Childhood*. MIT Press.

Cheney, Dorothy, and Robert Seyfarth (1990) *How monkeys see the world*. Chicago: University of Chicago Press.

Chomsky, Noam (1957) *Syntactic structures*. The Hague: Mouton.

(1965) *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

(1966) *Cartesian Linguistics*. New York: Harper & Row.

(1972) *Studies on Semantics in Generative Grammar*. The Hague: Mouton.

(1973) Constraints on transformations. In S. Anderson and P. Kiparsky (eds.), *A Festschrift for Morris Halle*, 232-286. New York: Holt, Rinehart & Winston.

(1977) On wh-movement. In P. Culicover, T. Wasow, and A. Akmajian (eds.), *Formal syntax*, 71-132. New York: Academic Press.

(1981) *Lectures on government and binding*. Dordrecht: Foris.

(1995) *The Minimalist Program*. Cambridge, MA: MIT Press.

Chomsky, Noam, and Morris Halle (1968) *The sound pattern of English*. New York, Harper & Row.

Clahsen, Harald, and Mayella Almazan (1998) Syntax and morphology in Williams syndrome. *Cognition* 68, 167-198.

Corballis, Michael C. (1991) *The lopsided ape*. Oxford: Oxford University Press,

Culicover, Peter W. (1999) *Syntactic nuts: Hard cases in syntax*. Oxford: Oxford University Press.

Culicover, Peter, and Ray Jackendoff (1995) *Something else* for the binding theory. *Linguistic Inquiry* 26, 249-275.

(1997) Semantic subordination despite syntactic coordination. *Linguistic Inquiry* 28, 195-217.

(1999) The view from the periphery: The English correlative conditional. *Linguistic Inquiry* 30, 543-571.

(forthcoming) *Minimal syntax*.

Curtiss, Susan (1977) *Genie: A linguistic study of a modern-day "wild child"*. New York, Academic Press.

Cutler, Anne, and Charles Clifton, Jr. (1999) Comprehending spoken language: A blueprint of the listener. In Colin M. Brown and Peter Hagoort (eds.), *The neurocognition of language*, 123-166. Oxford: Oxford University Press.

Deacon, Terrence W. (1997) *The Symbolic Species*. New York: Norton.

Degraff, Michel (ed.) (1999) *Language creation and language change: Creolization, diachrony, and development*. Cambridge, MA: MIT Press.

Dufva, H., and M. Lähteenmäki (1996) But who killed Harry? A dialogical approach to language and consciousness. *Pragmatics and Cognition* 4, 35-53.

Elman, Jeffrey (1990) Finding structure in time. *Cognitive Science* 14, 179-211.

Elman, Jeffrey, Elizabeth Bates, Mark Johnson, Annette Karmiloff-Smith, Domenico Parisi, and Kim Plunkett (1996) *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Fauconnier, Gilles (1985) *Mental spaces: Aspects of meaning constrution in natural language*. Cambridge, MA: MIT Press.

Fillmore, Charles, Paul Kay, and Mary Catherine O'Connor (1988) Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language* 64, 501-538.

Fischer, Susan D., and Patricia Siple (eds.) (1990) *Theoretical issues in sign language research 1*. Chicago: University of Chicago Press.

Fitch, W. Tecumseh (2000) The evolution of speech: A comparative review. *Trends in Cognitive Sciences* 4, 258-267.

Flynn, Suzanne, and Wayne O'Neill (eds.) (1988) *Linguistic theory in second language acquisition*. Dordrecht: Reidel.

Fodor, Jerry (1987) *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.

Fodor, Jerry, Thomas Bever, and Merrill Garrett (1974) *The Psychology of Language*. New York: McGraw-Hill.

Frege, Gottlob (1892) Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100, 25-50. English translation in Peter Geach and Max Black (eds.), *Translations from the philosophical writings of Gottlob Frege*, 56-78. Oxford: Blackwell, 1952.

Gleitman, Lila R., and Barbara Landau (eds.) (1994) *The acquisition of the lexicon*. Cambridge, MA: MIT Press.

Goldberg, Adele (1995) *Constructions: A Construction Grammar approach to argument structure*. Chicago: University of Chicago Press.

Goldsmith, John (1976) *Autosegmental phonology*. Doctoral dissertion, MIT. New York: Garland Press, 1979.

Gopnik, Myrna (1999) Some evidence for impaired grammars. In R. Jackendoff, P. Bloom, and K. Wynn (eds.), *Language, Logic, and Concepts*, 263-283. Cambridge, Mass.: MIT Press.

Halle, Morris, and William Idsardi (1995) Stress and metrical structure. In J. Goldsmith (ed.), *Handbook of theoretical phonology*, 403-443. Oxford: Blackwell.

Harris, Randy Allen (1993) *The Linguistics Wars*. New York: Oxford University Press.

Hauser, Marc D. (2000) *Wild Minds: What Animals Really Think*. New York: Henry Holt.

Hauser, Marc D., Noam Chomsky, and Tecumseh Fitch (2002) *Science*

Heim, Irena, and Angelica Kratzer (199x)

Huck, Geoffrey, and John Goldsmith (1995) *Ideology and linguistic theory*. Chicago: University of Chicago Press.

Jackendoff, Ray (1972) *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.

(1983) *Semantics and cognition.* Cambridge, MA: MIT Press.

(1992) Mme. Tussaud meets the binding theory. *Natural Language and Linguistic Theory* 10.1, 1-31.

(1990) *Semantic structures*. Cambridge, MA: MIT Press.

(1996) The proper treatment of measuring out, telicity, and possibly even quantification in English. *Natural Language and Linguistic Theory* 14, 305-354.

(1997) *The architecture of the language faculty*. Cambridge, MA: MIT Press.

Kager, René (1995) The metrical theory of word stress. In J. Goldsmith (ed.), *Handbook of theoretical phonology*, 367-402. Oxford: Blackwell.

Katz, Jerrold, and Jerry Fodor (1963) The structure of a semantic theory. *Language* 39, 170-210.

Katz, Jerrold, and Paul M. Postal (1964) *An integrated theory of linguistic descriptions*. Cambridge, MA: MIT Press.

Kegl, Judy, Ann Senghas, and Marie Coppola (1999) Creations through contact: Sign language emergence and sign language change in Nicaragua. In DeGraff 1999, 179-237.

Keil, Frank C. (1989) *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Klein, Wolfgang, and Clive Perdue (1997) The Basic Variety, or: Couldn't language be much simpler? *Second Language Research* 13, 301-347.

Klima, Edward S., and Ursula Bellugi (1979) *The signs of language*. Cambridge, MA: Harvard University Press.

Köhler, Wolfgang (1927) *The mentality of apes*, Routledge & Kegan Paul.

Ladd, D. Robert (1996) *Intonational phonology*. Cambridge: Cambridge University Press.

Lakoff, George (1971) On generative semantics. In D. Steinberg and L. Jakobovits (eds.), *Semantics: An interdisciplinary reader in philosophy, linguistics, and psychology*, 232-296. New York: Cambridge University Press.

Lakoff, George (1987) *Women, fire, and dangerous things*. Chicago: University of Chicago Press.

Lamb, Sydney (1966) *Outline of Stratificational Grammar*. Washington: Georgetown University Press.

Langacker, Ronald (1987) *Foundations of cognitive grammar, Vol. 1*. Stanford: Stanford University Press.

Lenneberg, Eric (1967) *Biological foundations of language*. New York: John Wiley & Sons.

Lerdahl, Fred, and Ray Jackendoff (1983) *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.

Levelt, W. J. M. (1989) *Speaking*. Cambridge: MIT Press.

(1999) Producing spoken language: A blueprint of the speaker. In Colin M. Brown and Peter Hagoort (eds.), *The neurocognition of language*, 83-122. Oxford: Oxford University Press.

Liberman, Mark, and Alan Prince (1977) On stress and linguistic rhythm. *Linguistic Inquiry* 8, 249-336.

Marcus, Gary (2001) *The algebraic mind*. Cambridge, MA: MIT Press.

Marr, David (1982) *Vision*. San Francisco: Freeman.

McCawley, James D. (1968) Lexical insertion in a transformational grammar without deep structure. In B. Darden, C.-J. N. Bailey, and A. Davison (eds.) *Papers from the fourth meeting of the Chicago Linguistic Society*. Chicago: University of Chicago Department of Linguistics.

Newmeyer, Frederick J. (1980) *Linguistic theory in America: The first quarter-century of transformational-generative grammar*. New York: Academic Press.

(1998) On the supposed 'counterfunctionality' of Universal Grammar: some evolutionary implications, in James Hurford, Michael Studdert-Kennedy, and Chris Knight (eds.) *Approaches to the evolution of language,* 305-319. Cambridge: Cambridge University Press.

Newport, Elissa (1990) Maturational constraints on language learning. *Cognitive Science* 14, 11-28.

Nunberg, Geoffrey (1979) The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy* 3, 143-184.

Partee, Barbara, ed. (1976) *Montague Grammar*. New York, Academic Press.

Perlmutter, David M. (ed.) (1983) *Studies in Relational Grammar 1*. Chicago: University of Chicago Press.

PiZango, Maria Mercedes, Edgar Zurif, and Ray Jackendoff (1999) Real-time processing implications of enriched composition at the syntax-semantics interface. *Journal of Psycholinguistic Research*, 28, 395-414.

Pinker, Steven (1989) *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press

(1999) *Words and rules*. New York: Basic Books.

Pinker, Steven, and Paul Bloom (1990) Natural language and natural selection, *Behavioral and Brain Sciences* 13, 707-726

Pollard, Carl, and Ivan Sag (1987) *Information-based syntax and semantics*. Stanford: Center for the Study of Language and Information.

(1994) *Head-Driven Phrase Structure Grammar*. Chicago: University of Chicago Press.

Postal, Paul M. (1970) On the surface verb "remind." *Linguistic Inquiry* 1, 37-120.

Povinelli, Daniel (2002) *Folk Physics for Apes*. Oxford: Oxford University Press.

Premack, David (1976) *Intelligence in ape and man*. Hillsdale, NJ: Erlbaum.

Prince, Alan, and Paul Smolensky (1993). *Optimality Theory: Constraint Interaction in Generative Grammar*. Rutgers University Center for Cognitive Science, Piscataway, NJ.

Pustejovsky, James (1995) *The generative lexicon*. Cambridge, MA: MIT Press.

Putnam, Hilary (1975) The meaning of 'meaning.' In Keith Gunderson (ed.), *Language, mind, and knowledge*, 131-193. Minneapolis: University of Minnesota Press.

Rumelhart, David, and James McClelland (1986) On learning the past tense of English verbs. In J. McClelland, D. Rumelhart, and the PDP Research Group, *Parallel distributed processing*, Vol. 2, 216-271. Cambridge, MA: MIT Press.

Ryle, Gilbert (1949) *The concept of mind*. Chicago: University of Chicago Press.

Sadock, Jerrold (1991) *Autolexical syntax*. Chicago: University of Chicago Press.

Saussure, Ferdinand de (1915) *Cours de linguistique générale*, edited by Charles Bally and Albert Sechehaye. English translation: *Course in General Linguistics*. New York: Philosophical Library, 1959.

Savage-Rumbaugh, Sue, Stuart Shanker, and Talbot Taylor (1998) *Apes, language, and the human mind*. Oxford University Press.

Searle, John (1980) Minds, brains, and programs. *Behavioral and Brain Sciences* 3, 417-424.

Seidenberg, Mark S., and Laura Petitto (1978) Signing behavior in apes: A critical review. *Cognition* l7, 177-215.

Spelke, Elizabeth, Gary Katz, Susan Purcell, Sheryl Ehrlich, and Karen Breinlinger (1994) Early knowledge of object motion: Continuity and inertia. *Cognition* 51, 131-176.

Swinney, David (1979) Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18, 645-659.

Talmy, Leonard (2000) *Toward a Cognitive Semantics*. Cambridge, Mass.: MIT Press.

Tanenhaus, M. K., M. J. Spivey-Knowlton, K. M. Eberhard, and J. C. Sedivy (1995). Integration of Visual and Linguistic Information in Spoken Language Comprehension. *Science* 268, 1632-4.

Terrace, H. (1979) *Nim*. New York, Knopf.

Tomasello, Michael (ed.) (2000) *Primate cognition* (special issue). *Cognitive Science* 24.3.

Van Valin, Robert, and Randy LaPolla (1997) *Syntax: Structure, meaning and function*. Cambridge: Cambridge Unviersity Press.

Verkuyl, Henk (1993) *A theory of aspectuality: The interaction between temporal and atemporal structure*. Cambridge: Cambridge University Press.

Weinreich, Uriel (1966) Explorations in semantic theory. In T. Sebeok (ed.), *Current trends in linguistics*, vol. 3. The Hague: Mouton. Reprinted in U. Weinreich (1980), *On semantics*, 99-201. Philadelphia: University of Pennsylvania Press.

Wexler, Kenneth, and Peter Culicover (1980) *Formal principles of language acquisition*. Cambridge, MA: MIT Press.

Zurif, Edgar (1990) Language and the brain. In Dan Osherson and Howard Lasnik (eds.), *An introduction to cognitive science, Vol. 1: Language*, 177-198. Cambridge, MA: MIT Press.

Acknowledgment

Foundations of Language was written during 1999-2000, while I was a Fellow at the Wissenschaftskolleg zu Berlin, a year for which I am profoundly grateful. Both the book and this version of it were supported in part by NIH Grant 03660 to Brandeis University. I also must thank Marc Hauser and Dan Dennett for their help and encouragement on this version.

Notes

---

1. "Essentially the same" is a matter of perspective. When we are talking about "English speakers" as a whole we can treat them all as essentially the same. But if we're talking about dialect differences, dialect contact, or language change, we can just as easily switch to treating different speakers as having (slightly) different linguistic systems in their heads. And of course when we're talking about language acquisition, we take it for granted that the young child has a different system than the adults.

2. To some extent Chomsky's point has been lost on the larger cognitive neuroscience community. For instance, the widely cited connectionist parser of Elman (1990) is a variation of a finite-state Markov device, and is subject to some of the same objections raised by Chomsky in 1957. See Marcus 2001 and Pinker 1999 for extensive discussion.

3. For example: "The deep structure that expresses the meaning is common to all languages, so it is claimed [by the Port-Royal grammarians – who of course did not use the term "deep structure"], being a simple reflection of the forms of thought" (Chomsky 1966, 35).

4. Some opponents of generative grammar (for instance some Cognitive Grammarians) have rightly objected to syntactocentrism, but proposed instead that all properties of language are derivable from meaning. I take this to be equally misguided, for reasons that should be evident as we proceed.

5. A standard mark of recursivity is a constituent occurring within another constituent of the same type. For instance, a clause can appear within another clause: *The man* ***who comes from New York*** *is tall*; and a noun phrase can appear within a noun phrase: *the king of* ***the Cannibal Islands***. In phonology this sort of situation does not occur nearly so freely: in particular, a syllable cannot occur within another syllable.

6. Interestingly, Chomsky (1965) brings up an example like (4) and analyzes the prosody is a fact of performance: speakers don't pronounce the sentence in accordance with its syntactic structure. This is about the only way he can analyze it, given that he does not have independent principles of intonational constituency at his disposal. Contemporary theory allows us to say (correctly, I believe) that (4) is well-formed both syntactically and prosodically, with a well-formed correspondence between the two structures.

7. It is important to distinguish two interpretations of "syntactic" here. In the broader sense, every combinatorial system has a syntax: mathematics, computer languages, music, and even phonology and semantics. In the narrower sense of technical linguistics, "syntactic" denotes

the organization of units such as NPs, VPs, and prepositions. I am reserving “syntactic” for this narrower sense and using “combinatorial” for the broader sense.

8. For the semantics I have used the Conceptual Structure notation of Jackendoff (1983, 1990); readers invested in other frameworks should feel free to substitute their own notations.

9. Stratificational Grammar (Lamb 1966) also proposed a thoroughgoing organization into independent generative components linked by interfaces.

10. The lexicon, a large collection of learned arbitrary associations between very particular bits of structure, also has parallels in other domains of memory. For instance, it is an arbitrary fact that the sound *kΦt* means ‘feline animal’, and the child must learn it from the environment. Now consider the association between the appearance of foods and their tastes. It is similarly arbitrary (from the point of view of the organism) that something that looks like a cucumber tastes the way it does, and organisms learn probably hundreds or thousands of such associations. (There are even “ambiguous” looking foods: think of mashed potatoes and vanilla ice cream.)