The Development of Language

Acquisition, Change, and Evolution

David Lightfoot

University of Maryland

BLACKWELL
Publishers

3

Grammars and Language Acquisition

In this chapter we consider what it means to know a language. Individuals grow grammars, which are biological entities represented in people's brains and which characterize their linguistic knowledge. We shall move away now from talk of languages and focus instead on these biological grammars. Cutting some corners, we can say that grammars grow in accordance with genetic or developmental principles, and we can discover what those principles are.

Grammars vary from one person to another, and they may change in groups of people from one generation to the next. This happens if the initial conditions change somewhat, if people are exposed to different childhood experiences. We'll go on to argue in the next chapter that if we observe a language undergoing some significant shift over the course of time, what really is occurring is that grammars are changing in certain individuals and that grammatical change then spreads through a population of speakers. In this way, we adopt Jon Elster's "methodological individualism," mentioned in the last chapter, and study individual behavior as the basic building block of aggregate social phenomena, explaining the complex by the simple. For now we shall deal with the prior matter and argue that a person's mind/brain contains a grammar, which characterizes his or her linguistic capacity and interacts with other mental components.

3.1 We Know More than we Learn

A striking property of language acquisition is that children attain infinitely more than they experience – literally so, we shall see. They attain a productive system, a grammar, on the basis of very little experience. So there is more – much more – to language acquisition than mimicking what we hear in childhood; and there is more to it than the simple transmission of a set of words and sentences from one generation of speakers to the next.

To understand some fundamental aspects of language acquisition, let us go straight to a few simple but specific phenomena. I want to convince you that you know far more than you are consciously aware of, that you're smarter than you think you are.

Consider some subtleties that people are not consciously aware of. The verb *is* may be used in its full form or its reduced form: people say *Kim is happy* or *Kim's happy*. However, certain instances of *is* never reduce: for example, the underlined items in *Kim is happier than Tim is* or *I wonder where the concert is on Wednesday*. Most people are not aware of this, but we all know subconsciously not to use the reduced form here. How did we come to know this? As children, we were not instructed to avoid the reduced form in certain places. Yet, all children typically attain the ability to use the forms in the adult fashion, and the ability is quite independent of intelligence level or educational background. Children attain this ability early in their linguistic development. More significantly, children do not try out the nonoccurring forms as if testing a hypothesis, in the way that they "experiment" by using forms like *goed* and *taked*. The ability emerges perfectly, as if by magic.

Another example: pronouns like *she*, *her*, *he*, *him*, *his* sometimes refer back to a noun previously mentioned in a sentence (1a–c). However, one can understand (1d) only as referring to two men, Jay and somebody else; here the pronoun may not refer to Jay, unlike in (1a–c).

1 (a) Jay hurt his nose.
  (b) Jay's brother hurt him.
  (c) Jay said he hurt Ray.
  (d) Jay hurt him.

As adults we generalize that a pronoun may refer to a preceding noun *except* under very precise conditions (1d). But then, how did we all acquire the right generalization, particularly knowledge of the exception?

Recall the nature of our childhood experience: we were exposed to a haphazard set of linguistic expressions. We heard various sentences containing pronouns; sometimes the pronoun referred to another noun in the same sentence, sometimes to a person not mentioned there. Problem: because we were not informed about what cannot occur, our childhood experience provided no evidence for the "except" clause, that pronouns sometimes do not co-refer. That is, we had evidence for generalizations like "*is* may be pronounced *z*" and "pronouns may refer to a preceding noun," but no evidence for where these generalizations break down.

As children, we came to know the generalizations and their exceptions, and we came to this knowledge quickly and uniformly. Yet our linguistic experience was not rich enough to determine the limits to the generalizations. We call this the problem of the "poverty of the stimulus." Children have no data which show them that *is* may not be reduced in some contexts, and they have no data showing that *him* may not refer to Jay in (1d). These two small illustrations are examples of the form that the poverty-of-stimulus problem takes in language. It may look as if children are behaving magically, but there is no magician, and magic is no answer.

There are two "easy" solutions to the poverty-of-stimulus problem, but neither is adequate. One is to say that children do not overgeneralize, because they are reliable imitators. That is, children do not produce the reduced *is* in the wrong place or use a pronoun in (1d) wrongly to refer to Jay, because they never hear language being used in this way. In other words, children acquire their native language simply by imitating the speech of their elders. We know that this approach is not tenable, because everybody constantly says things they have never heard. We express thoughts with no conscious or subconscious consideration of whether we are imitating somebody else's use of language. This is true of the most trivial speech: in saying *I wanna catch the 3:25 p.m. bus, which leaves from outside Border's bookstore*, I am using a sentence that I have almost certainly not heard.

A variant on this approach is that children learn not to say the deviant forms because they are corrected by their elders. This view offers no better insight for several reasons. First, it would take an acute observer to detect and correct the error. Second, where linguistic correction is offered, young children are highly resistant and just don't get the correction. Third, in the examples discussed, children do not overgeneralize, and therefore parents have nothing to correct; this will become clearer towards the end of this chapter, when we discuss experimental work on young children.

So the first "easy" solution to the poverty-of-stimulus problem is to deny that it exists, and to hold that the environment is rich enough to provide evidence for where the generalizations break down. The problem is real, but the "solution" does not address the problem.

The second "easy" answer also denies that there is a problem, but it denies that there is anything to be learned and holds that a person's language is fully determined by genetic properties. Yet this answer also cannot be right, because people speak differently, and many of the differences are environmentally induced. There is nothing about my genetic inheritance that makes me a speaker of English; if I had been raised in a Dutch home, I would have become a speaker of Dutch.

The two "easy" answers attribute everything either to the environment or to the genetic inheritance. Neither position is tenable. Instead, language emerges through an interaction between our genetic inheritance and the linguistic environment to which we happen to be exposed. English-speaking children learn from their environment that the verb *is* may be pronounced *iz* or *z*, and native principles prevent the reduced form from being used in the wrong places. Likewise, children learn from their environment that *he*, *his*, etc. are pronouns, and native principles dictate where pronouns may not refer to a preceding noun. The interaction of the environmental information and the native principles accounts for how the relevant properties emerge in an English-speaking child.

We'll sketch the relevant principles in a moment. It is worth pointing out two things. First, we are doing a kind of Mendelian genetics here. Working in the mid-nineteenth century, Mendel postulated genetic "factors" to explain the variable characteristics of his pea plants, without the slightest idea of how these factors might be instantiated biologically. Similarly, linguists seek to identify information which must be available independently of experience, in order for a grammar to emerge in a child. We have no idea whether this information is encoded directly in the genome or whether it results from epigenetic, developmental properties of the organism; it is, in any case, native. As a shorthand device for these native properties, I shall write of the "linguistic genotype," that part of our genetic endowment which is relevant to our linguistic development. Each individual's genotype determines the potential range of functional adaptations to the environment (Dobzhansky 1970: 36), and I assume that the linguistic genotype – what linguists call "Universal Grammar," or "UG" – is uniform across the species (short of pathological cases). That is, linguistically we all have the same potential for functional adaptations, and any of us may grow up to be a speaker of Catalan or Hungarian, depending entirely on our circumstances and not at all on variation in our genetic makeup.

Second, our ideas about how this a priori information is represented are very tentative. There are different models of UG, and those models all change frequently as researchers learn more about language variation, acquisition, and change. I shall invoke the technicalities of particular models as little as possible. Where I do need theoretical ideas, I pick them from what are known as Government-Binding or Minimalist models, but I am not concerned here with outlining a coherent model of UG and shall simply pick ideas which help us understand the phenomena under discussion. I shall be eclectic and opportunistic in this regard, and not much concerned about consistency. There may be equally useful ideas in other models. I am interested in arguments to the effect that property *p* needs to be postulated at the level of UG, and am less concerned about the form that property *p* might take or how property *p* is biologically encoded.

Since children are capable of acquiring any language to which they happen to be exposed between infancy and puberty, the same set of genetic principles which accounts for the emergence of English (using "genetic" now in the extended sense I have indicated) must also account for the emergence of Dutch, Vietnamese, Hopi, or any other of the thousands of languages spoken by human beings. This plasticity imposes a strong empirical demand on hypotheses about the linguistic genotype; the principles postulated must be open enough to account for the variation among the world's languages. The fact that people develop different linguistic capacities, depending on whether they are brought up in Togo, Tokyo, or Toronto, provides a delicate tool with which to refine claims about the nature of the native component.

So there is a biological entity, a finite mental organ, which develops in children along one of a number of paths. The paths are determined in advance of any childhood experience. The language organ that emerges, the grammar, is represented in the brain and plays a central role in the person's use of language, whether for speaking, listening, writing poems, or solving crossword puzzles. We have gained some insight into the nature of people's language organs by considering a wide range of phenomena: the developmental stages that young children go through, the way language breaks down in the event of brain damage, the manner in which people analyze incoming speech signals, and more. At the center is the biological notion of a language organ, a grammar.

3.2 The Nature of Grammars

So children acquire a productive system, a grammar, in accordance with the requirements of the genotype. If asked to say quite generally what is now known about the linguistic genotype, I would say that it permits finite grammars, because they are represented in the finite space of the brain, but that they range over infinity. Finite grammars consist of a limited, precise set of operations which allow for infinite variation in the expressions that are generated. The genotype is plastic, consistent with speaking Japanese or Quechua. It is modular and uniquely computational.

By "modular" I mean that the genotype consists of separate subcomponents each of which has its own distinctive properties, which interact to yield the properties of the whole. These modules are, in many cases, specific to language. Research has undermined the notion that the mind possesses only general principles of "intelligence" which hold of all kinds of mental activity. One module of innate linguistic capacity contains abstract structures which are compositional (consisting of units made up of smaller units) and which fit a narrow range of possibilities. Another module encompasses the ability to relate one position to another within these structures by movement, and those movement relationships are narrowly defined.

To see the kind of compositionality involved, consider how words combine. Words are members of categories like noun (N), verb (V), preposition (P), adjective/adverb (A), etc. If two words combine, then the grammatical properties of the resulting phrase are determined by one of the two words, which we call the "head." So, if we combine the verb *visit* with the noun *Chicago*, the resulting phrase *visit Chicago* has verbal, not nominal, properties. It occurs where verbs occur and not where nouns occur: *I want to visit Chicago*, but not **the visit Chicago*, or **We discussed visit Chicago*. So the expression *visit Chicago* is a verb phrase (VP), where the verb *visit* is the head projecting to VP. This can be represented as a labeled bracketing (2a) or as a tree diagram (2b). The verb is the head of the VP, and the noun is the "complement." (Here I adapt the novel, bottom-up approach to phrase structure of Chomsky 1995: 244ff.)

2 (a) $_{VP}$[$_{V}$[visit] $_{N}$[Chicago]]

(b)

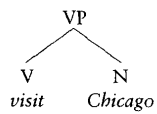

In general, two categories *merge* to form a new category. So an "inflectional" element like *will* might merge with the VP *visit Chicago*, to yield the more complex expression *will visit Chicago*, with a structure $_{IP}$[$_{I}$[will] $_{VP}$[visit Chicago]]. The inflectional *will* heads the new phrase and projects to a phrasal category IP. This means that *visit Chicago* is a unit (VP), which acts as the complement of *will*, but *will visit* is not a unit; that is, there is no single node which dominates *will visit* and nothing else in this example. This is how grammars assign structures to phrases.

A full sentence like *The student from Denver will visit Chicago* is formed by successive merger, yielding the structure of (3). *The student from Denver* is a unit which acts as the subject of the clause. It is formed by successive merger. *From* merges with *Denver* to yield a preposition phrase (PP), which merges with *student* to form a NP, which merges with *The* to form a constituent called, in recent work, a "determiner phrase" (DP). This DP, in turn, is the *specifier* of the inflectional head *will*, acting as the subject of the clause.[1]

3

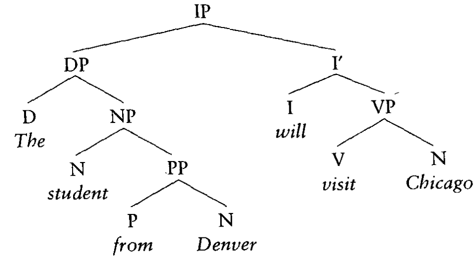

This means that there are two kinds of phrase, one consisting of a head and its complement, *visit Chicago*, in (3), and the other, a more complex phrase consisting also of a specifier, like the IP in (3).

The units defined by these trees are the items which the computational operations manipulate; they are the items which move and can be deleted and which receive indices. Non-units are not subject to these operations.

One of the computational operations involved is that of overt movement to account for the displacement of elements. So an expression like *What city will the student visit?* (where *What city* is understood as the complement of *visit*) has a structure along the lines of (4). Here we need more structure to enable *will* to merge with the IP, and then to enable *What city* to merge with the rest of the clause. *Will* is a head (labeled C for "complementizer"), which merges with the IP, and the determiner phrase (DP) *What city* is a specifier of that head. I indicate the positions from which these elements have moved.

4

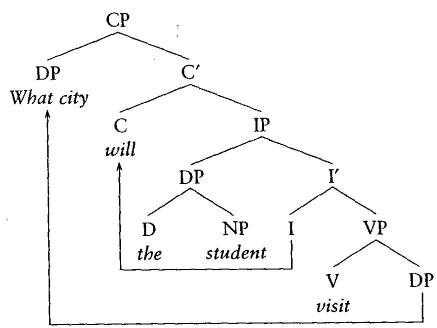

Let us return now to the problem of the reduced *is*, which we discussed earlier. A computational operation attaches the reduced *is* not to the preceding word but to the *following* word, contrary to what the apostrophe *'s* spelling convention suggests. That makes it one of very many instances where orthography does not reflect grammatical reality. Orthography reflects the direction of *phonetic* assimilation: the reduced *'s* takes on phonetic characteristics of the preceding word (compare *Pat's here*, where *'s* is voiceless, and *The dog's here*, where *'s* is voiced, pronounced like *z*). However, rightward attachment at the syntactic level explains why, if there is no following word, the reduced form does not occur, because there is nothing to the right for it to attach to: *Kim is happier than Tim is*. The same holds for *'ve*, *'re*, and the reduced forms of *am*, *will*, *would*, *shall*, and *should*. Poets make linguistic jokes from these principles: the Gershwins were famous for contraction jokes, and in *Girl Crazy* (1930) a chorus begins "I'm bidin' my time / 'Cause that's the kind of guy I'm."

An intervening understood element also blocks attachment to the right, so that there is no reduction: *I wonder where*$_x$ *the concert is* x *on Wednesday*. Here *where* has moved from the position indicated and is understood in the position x. This shows that the reduced *is* has to be in a certain kind of structural relationship with the following word: it has to *govern* it. *Is* does not govern *on Wednesday* in this example. The distinction is illustrated in (5). In (5a) *is* governs the following word, *happy*, and therefore may be reduced: *happy* is the complement of *is*. Similarly with the first *is* of (5b). However, the underlined form in (5b) does not govern the following word, *Tim*, and therefore may not reduce; in this case, the subject *Tim* and the copula verb *is* have permuted. Intuitively, a head like *is* governs a following word which is its complement (in due course, we shall need to extend this definition slightly). So now we have an answer to the problem sketched at the beginning of this chapter: a reduced *is* is attached to a word to its immediate right which it governs.

5 (a) Kim's happy.
  (b) Kim's happier than is Tim.

Consider another reduction process, whereby *want to* is often pronounced *wanna*: *I want to be happy* or *I wanna be happy*. Here, the *to* may be attached to the preceding *want* and then reduced to give the *wanna* form. But an intervening understood element blocks the reduction process. For example, *Who do you want to see?* has the logical structure of (6a) (corresponding to the statement *You want to see Jill*), and it may be pronounced with the *wanna* form. *Who* is understood as the direct object (complement) of the verb *see*, indicated here by x.

6 (a) Who$_x$ do you want $_{IP}[_{I}$to see x]
  (b) Who$_x$ do you want [x to go]

However, when an understood element intervenes between *want* and *to*, *wanna* does not occur in the speech of most people. So an expression *Who do you want to go?* has the logical structure of (6b) where *who* is understood in the position x as the subject of the verb *go* (corresponding to the statement *I want Jill to go*). Logical structures like (6b) do not have a corresponding phonological structure where *want to* is reduced to *wanna*.

The government restriction also holds for the reduced *wanna* form, even though *to* attaches to its left: the reduced form occurs only if *want* governs *to*. *Want* governs *to* in (6a) and therefore may be reduced. However, *want* does not govern the following *to* in the sentences of (7), and it may not be reduced.[2] I give some partial structure in (7a′, b′), enough to show the lack of government between *want* and the immediately following *to*. In (7a′) *want* does not govern the IP *to win games*, because the IP is not its complement; it acts as the subject of the next IP, which, in turn, is the complement of the verb *want*. Similarly, *to vote for the incumbent* is not the complement of *want* in (7b′) and therefore is not governed by it.

7 (a) They don't want to win games to be their only goal.
  (a′) . . . want $_{IP}[_{IP}[_{I}$to win games] $_{I}$to be their only goal]
  (b) They expect people who continue to want to vote for the incumbent.
  (b′) They expect $_{IP}[_{NP}$[people who continue to want] $_{I}$to vote for the incumbent]

The restrictions on reduced forms, as we have them now, are general: reduced items must be in a government relation with a lexical element. One productive approach is to treat reduced *to* and reduced *is*, etc. as *clitics*. Clitics are little words which occur in many, perhaps all, languages and have the property of not being able to stand alone. There are rightward clitics, which attach to a word to the right, like reduced *is*, and leftward clitics, which attach leftwards, like reduced *to*. Part of what a child developing a grammar needs to do is to determine the clitics in his or her linguistic environment, knowing in advance of any experience that these are small, unstressed items left- or right-attached to an adjacent element under a structural relationship of government, with no other (phonetic or "understood") element intervening. This predetermined knowledge is contributed by the linguistic genotype and is what the child brings to language acquisition. So hearing a reduced form like *'s cold in here* and knowing that it is equivalent in some way to *It is cold in here* suffices to show the child that *'s* is a rightward clitic.

On this approach, the child is faced with a chaotic environment and scans it, looking for rightward and leftward clitics – among many other things, of course. This is my solution to the problem with which I opened this chapter, and it is an answer of the right shape. It makes a general claim at the genetic level (clitics and their behavior are predefined) and postulates that the child arrives at a plausible analysis on exposure to a few simple expressions like *'s cold in here*. The analysis that the child arrives at predicts no reduction for the underlined *is* in *Kim is happier than Tim is*, *I wonder where the concert is on Wednesday*, and countless other cases, and the child needs no correction to arrive at this system. The very fact that *'s* is a clitic, defined in advance of any experience, dictates that it may not occur in certain contexts. It is for this reason that the generalization that *is* may be pronounced as *'s* breaks down at certain points and does not hold across the board.

The mature grammar treats *'s* as a rightward clitic and *-na* as a leftward clitic. This is a feature of an English speaker's phenotype, the mature expression of the potentiality defined by the genotype. By pointing to a plausible triggering experience, linguists show how the English speaker arrives at the relevant analysis, given specific claims at the genotypical level. Linguists spend a lot of time analyzing a wide range of languages, sensitive always to the poverty-of-stimulus problems which each grammar throws up and seeking a single body of information, the invariant linguistic genotype, which solves the poverty-of-stimulus problems in all grammars.

Consider now the second problem: the reference of pronouns. They taught us in school that pronouns refer to a preceding noun, but the data of (1) show that this just isn't right. As we saw, in (1d) *him* may not refer to Jay; in (1b) *him* may refer to Jay but not to Jay's brother. The best account of this complex phenomenon seems to be to invoke a native principle which says that pronouns may *not* refer back to a local nominal element, where "local" means contained in the same clause or in the same DP. This is known as Principle B of the Binding Theory, which deals with the indexical properties of NP/DPs and whose effects have been studied in many languages.[3]

In (8) I give the relevant structure for the corresponding sentences of (1). In (8b) the DP *Jay's brother* is local to *him*, so *him* may not refer back to that DP – we express this by indexing them differently. On the other hand, *Jay* is contained within the DP and therefore is not available to *him* for indexing purposes, so those two nouns do not need to be indexed differently – they may refer to the same person, and they may be co-indexed. Again

we see the constituent structure illustrated in (2), (3), and (4) playing a central role in the way in which the computations of the Binding Theory are carried out. In (8d) *Jay* is local to *him*, and so the two elements may not be co-indexed; they do not refer to the same person. In (8c) *Jay* is not local to *he*, because the two items are not contained in the same clause, and Principle B does not block co-indexing: *Jay* and *he* may refer to the same person. In (8a) *his* is contained within a DP and may not be co-indexed with anything else within that DP; what happens outside the DP is irrelevant to Principle B; so *his* and *Jay* may co-refer and do not need to be indexed differently.

8 (a) $_{IP}$[Jay$_{i}$ hurt $_{DP}$[his$_{i/j}$ nose]]
(b) $_{IP}$[$_{DP}$[Jay$_{i}$'s brother]$_{k}$ hurt him$_{i/j}$]
(c) $_{IP}$[Jay$_{i}$ said $_{IP}$[he$_{i/j}$ hurt Ray]]
(d) $_{IP}$[Jay$_{i}$ hurt him$_{j}$]

We could have illustrated this principle equally well with data from French or Dutch, because the principle applies quite generally, to pronouns in all languages. If we assume Principle B as a native principle, available to the child independently of any actual experience, language acquisition is greatly simplified. Now the child does not need to "learn" why the pronoun may refer to Jay in (8a) or (8b, c) but not in (8d). Rather, the child raised in an English-speaking setting has only to learn that *he*, *his*, *him* are pronouns – that is, elements subject to Principle B. This can be learned by exposure to a simple sentence like (1d) (8d), uttered in a context where *him* refers to somebody other than Jay.[4]

One way of thinking about the contribution of the linguistic genotype is to view it as providing invariant principles and option points, or "parameters." There are invariant principles to the effect that clitics attach to elements in a government relation and that pronouns are not locally co-indexed. Meanwhile, there are options such that direct objects may precede the verb in some grammars and follow it in others, that some clitics attach to the right and some to the left. These are parameters of variation, and the child sets these parameters one way or another in light of her particular linguistic experience. As a result, a grammar emerges in the child, part of the linguistic phenotype. The child has learned that *'s* is a clitic and that *her* is a pronoun; the genotype ensures that *'s* is attached to a lexical category in a government relation and that *her* is never used in a context where it refers to a local nominal.

Here we have looked at a couple of specific acquisition problems and considered what ingredients are needed for their solution. Now let us stand back and think about these matters more abstractly.

3.3 The Acquisition Problem: The Poverty of the Stimulus

The child acquires a finite, generative system – a grammar – which generates structures which correspond more or less to utterances of various kinds. Children acquire these grammars despite a poverty of stimulus on three levels.

First, the child hears speech from adults, peers, and older children. This stimulus does not consist uniformly of complete, well-formed utterances; it also includes sentence fragments, slips of the tongue, incomplete thoughts, ill-formed utterances from people who do not know the child's language well, and even sentences artificially simplified, supposedly for the benefit of children. Even if only 5 percent of the expressions the child hears are of this latter type, there will be a significant problem in generalizing to the set of grammatical sentences of the language, because the pseudo-sentences do not come labeled as defective.

Second, a child encounters only a finite range of expressions, but she comes to be able to produce and understand an infinite range of novel sentences, going far beyond the sentences heard in childhood. We know this at the intuitive level when we recognize that we constantly encounter novel sentences that we have not heard or used before. Consider the sentences on this page; it is unlikely that the reader has encountered any of them before in the precise form in which they occur here.

More formally, to understand that there is an infinite number of English sentences, one has only to realize that, in principle, any given sentence may be of indefinite length. Three iterative devices permit this and they may occur in various combinations:

Relativization: This is the dog that chased the cat that killed the rat that caught the mouse that nibbled the cheese that . . .
Complementation: I think that Jim asked me to tell Kim that Tim thought that I said that . . .
Co-ordination: Jay went to the movies and to the concert, and Ray and Kay went out for dinner, and Fay stayed at home, and . . .

If a sentence may be of indefinite length, then it follows that people have a capacity to use and express an indefinite number of sentences. Since a person's experience is finite and the mature capacity ranges over infinity, the stimulus alone cannot fully determine the mature capacity.

Third, people come to know things subconsciously about their language, things for which no direct evidence is available in the data to which they are exposed as children learning to speak. People eventually understand and utter complex, ambiguous sentences, identify paraphrases, and distinguish sentences that may occur in their language from ones that may not. We have just spent a few pages illustrating this kind of thing, and the subconscious, mature capacity that ordinary people have involves language properties made explicit by linguists. However, the crucial properties lie outside the primary linguistic data available to young children. Children are not systematically informed that hypothetical sentences do not in fact occur (e.g. **Who do you wanna go?*, **Kim's happier than Tim's*), that a given sentence is ambiguous, or that certain sets of sentences are paraphrases of each other. Also, many legitimate, acceptable sentence types may never occur in a child's linguistic experience. Such data are not available to preschool children and are not part of their verbal experience. The distinction between the range of data known to the linguist and the much more limited data available to the beginning speaker is of vital importance for the biological view of language development.

This third deficiency is quite crucial. The first two, the imperfection and finiteness of the stimuli, are not decisive kinds of data deficiencies. They do not deny that relevant experience for language learning is available; they simply assert that the experience is "degenerate," hard to sort out. The fundamental deficiency is the third, which says not that relevant experience is degenerate but that in certain areas it does not exist at all. This deficiency shapes hypotheses about the linguistic genotype.

Some structural principle prevents forms like **Who do you wanna go?* from occurring in the speech of English speakers, as we have seen. Children are not exposed to pseudo-sentences like this and informed systematically that they are not said. Speakers come to know subconsciously that they cannot be said, and this knowledge emerges somehow, even though it is not part of the input to the child's development. Furthermore, it is hard to imagine how the inventory of sentences and sentence fragments that constitute the child's linguistic environment could provide even indirect evidence that such sentences do not occur. It is not enough to say that people do not utter such forms because they never hear them. This argument is insufficient, because people say many things that they have not heard, as we have noted. Language is not learned simply by imitating or repeating what has been heard.

This third deficiency of the stimulus is of particular importance in defining our approach to language acquisition. A good deal of evidence exists that the contrast between the child's experience and the range of data available to the linguist is quite substantial. Over the last 40 years, much of the linguistic literature has focused on areas where the best description cannot be derived directly from the data to which the child has access, or is underdetermined by those data, as in the examples with the clitics *to* and *'s* and the pronouns. If the child's linguistic experience does not provide the basis for establishing a particular aspect of linguistic knowledge, another source must exist for that knowledge. That aspect must be known a priori, in advance, in the sense that it is available independently of linguistic experience. We tentatively assume that it is available through genetic prescriptions, while not ignoring other possibilities – for example, that it arises as a consequence of other, nonlinguistic experience.

All this is not to say that imitation plays no role: just that it does not provide a sufficient explanation, given the third data deficiency. This is worth emphasizing, because antagonists sometimes caricature this approach to language acquisition as "denying the existence of learning," when in fact its adherents merely deny that learning is the whole story – a very different matter. The quotation is from a remarkable article in *Science* magazine, in which the authors assert that "Noam Chomsky, the founder of generative linguistics, has argued for 40 years that language is unlearnable," and that they, on the other hand, have "rediscovered" learning (Bates and Elman 1996)!

Caricatures of this type show up in the writing of people who claim that *all* information is derived from the environment, and that there is no domain-specific genetic component to language acquisition. These people deny the poverty-of-stimulus problems, claiming that children derive all relevant information from their linguistic environment. Bates and Elman provide a recent, particularly clear and striking instance of this line, claiming that artificial neural networks can learn linguistic regularities from imperfect but "huge computerized corpora of written and spoken language."[5]

Others have appealed to a structured input which allows children to circumvent the poverty-of-stimulus problems (Snow 1977). Parents and other people often adopt a simple, sometimes artificial style of speech when addressing children; but it is scarcely plausible that this "Motherese" provides sufficient structure for language acquisition to take place on a purely inductive basis. Children do not simply generalize patterns without the aid of genetically determined principles.

There are at least four reasons why this kind of pattern generalization is not the answer to how children acquire speech. First, although children no doubt register only part of their linguistic environment, there is no way of knowing exactly what any individual child registers. Therefore, there is no factual basis for the claim that children register only what is filtered for them through parents' deliberately simplified speech. Children have access to more than this, including defective utterances, sentence fragments, and all the imperfections that the world throws at us. Second, even supposing

that they register only perfectly formed expressions (and hence that the first data deficiency does not hold), this isn't enough to show that the child has a sufficient inductive base for language acquisition. The third data deficiency still holds, and the child would need to know that *wanna* fails to occur in certain contexts. If children learn by induction, we must ask why quite ordinary inductive generalizations like this – that *want to* may be pronounced *wanna* – break down. The artificial, simplified speech of the Motherese style does not show where inductive generalizations must stop. Third, if the child registered only the simplified, well-formed sentences of Motherese, the problem of language learning would be *more* difficult, because the child's information would be more limited. Fourth, careful studies of parents' speech to children (like Newport et al. 1977) show that an unusually high proportion consists of questions and imperatives, and that simple declarative sentences are much rarer than in ordinary speech. This suggests that there is very little correlation between the way the child's language emerges and what parents do in their speech directed to children. The existence of Motherese in no way eliminates the need for a genetic basis to language acquisition. The child is primarily responsible for the acquisition process, not parents or playmates.

Nobody denies that the child must extract information from her environment; it is no revelation that there is "learning" in that technical sense. My point is that there is more to language acquisition than this.

Children react to evidence in accordance with specific principles. It is not at all clear what role induction plays. Induction does not enable a child to determine what a well-formed sentence is; nor does it explain how children "learn" the meanings of even the simplest words. Children do not have sufficient evidence to induce the meaning of *house*, *book*, or *city*, or of more complex expressions, even if we grant everything to advocates of Motherese or those who argue that it's all data processing of huge corpora (see n. 5). The biggest problem with ordinary sentences like *Kim is too clever to catch* or *Everybody saw him* is not that they exist, but to characterize their meaning and the meanings of the individual words. There is another problem with the Motherese hypothesis, which is more trivial than this but nonetheless real: children typically acquire the language not of their parents, but of their older siblings and peers.

The problem demanding explanation is compounded by other factors. Despite variation in background and intelligence, people's mature linguistic capacity emerges in fairly uniform fashion, in just a few years, without much apparent effort, conscious thought, or difficulty; and it develops with only a narrow range of the logically possible "errors." Children do not test random hypotheses, gradually discarding those leading to "incorrect" results and provoking parental correction. In each language community the non-adult sentences formed by very young children seem to be few in number and quite uniform from one child to another, which falls well short of random hypotheses. Normal children attain a fairly rich system of linguistic knowledge by five or six years of age, and a mature system by puberty. This is impressive when compared with the laborious efforts of squads of adult linguists who have made it through graduate school and try with only moderate success to characterize explicitly what people know when they know Estonian, Hopi, or Bengali. In this regard, language is no different from, say, vision, except that vision is taken for granted, and ordinary people give more conscious thought to language.

These, then, are the salient facts about language acquisition or, more properly, language growth. The child masters a rich system of knowledge without significant instruction and despite a triple deficiency of experiential data. The process involves only a narrow range of "errors" or false hypotheses and takes place rapidly, even explosively, between two and three years of age. The main question is how children acquire so much more than they experience.

Given these facts, especially the third data deficiency, certain properties must be available to the organism independently of linguistic experience. These properties permit language growth to circumvent environmental deficiencies and thus to take place quickly and not solely by trial and error. The environmental stimulus is impoverished, unstructured, and fairly random; the child hears a haphazard selection of sentences and pseudo-sentences and receives no significant instruction. So the environmental stimulus is just a trigger. Much of the ability eventually attained is determined by genotypical principles, which are activated by environmental stimuli.

The further fact that children can master *any* human language to which they happen to be exposed in infancy imposes strong limitations on the principles which the scientist can attribute to the genotype. An answer to the problem of acquisition does not lie merely in the properties of the specific language a particular child acquires. This would amount to a claim that the specific properties of, say, Hindi are innately prescribed, which permits no explanation of how a language with a significantly different structure – Polish, say – is acquired. The native principles must therefore be somewhat abstract and not language-specific. The principles must meet strict empirical demands: for each child, the principles must provide a basis for attaining the eventual mature grammar given exposure only to the haphazard and unstructured experience available in the child's community. To require the native principles to give an account of how a child may master *any* natural grammar under the conditions noted is to ask a lot.

The tight empirical demands make language particularly useful as a probe into the intrinsic properties of the human mind/brain.

A grammar represents what a speaker comes to know, subconsciously for the most part, about his or her native language. It represents the fully developed linguistic capacity, and is therefore part of an individual's phenotype. It is one expression of the potential defined by the genotype. Speakers know that certain sentences (in fact, an infinite number) may occur in their speech and that others may not; they know what the occurring sentences mean and the various ways in which they can be pronounced and rephrased. Most of this largely subconscious knowledge is represented in a person's grammar. The grammar may be used for various purposes, from everyday functions like expressing ideas, communicating, or listening to other people, to more contrived functions like writing elegant prose or lyric poetry, or compiling and solving crossword puzzles, or writing a book about language acquisition and change.

Universal Grammar (UG) represents the genetic equipment that makes language growth possible under the conditions assumed here (therefore part of the genotype) and delimits the linguistic knowledge that may eventually be attained – that is, the form and functioning of the grammar. The genotypical principles and parameters can be viewed as a theory of grammar. On this view, the theory of grammar is, in Chomsky's words,

> a common human attribute, genetically determined, one component of the human mind. Through interaction with the environment, this faculty of mind becomes articulated and refined, emerging in the mature person as a system of knowledge of language. To discover the character of this mental faculty, we will try to isolate those properties of attained linguistic competence that hold by necessity rather than as a result of accidental experience, where by "necessity" I of course mean biological rather than logical necessity. . . . The commitment to formulate a restrictive theory of UG is nothing other than a commitment to discover the biological endowment that makes language acquisition possible and to determine its particular manifestations. . . . we can explain some property of attained linguistic competence by showing that this property necessarily results from the interplay of the genetically-determined language faculty, specified by UG, and the person's (accidental) experience. (Chomsky 1977: 164)

It need hardly be pointed out that there is nothing necessary or God-given about this research goal; nor do I want to give the impression that all linguists adopt it. Crucially for our purposes, the vast majority of people who have worked on language change have not adopted this research goal. In fact, people have studied language with quite different goals in mind, ranging from the highly specific (to describe Dutch in such a way that it can be learned easily by speakers of Indonesian), to the more general, such as showing how a language may differ from one historical stage to another (comparing, for example, Chaucerian and present-day English). However, this is the goal I adopt, and this is the sense in which I shall construe a grammar, seeing it as a biological object; and it is important to keep this idea straight and not conflate it with other, more traditional notions of what a grammar is.

3.4 The Analytical Triplet

A grammar, for us, is a psychological entity, part of the psychological state of somebody who knows a language. For any aspect of linguistic knowledge, three intimately related items are included in the account. First, there is a formal, explicit characterization of what a mature speaker knows; this is the *grammar*, which is part of that speaker's phenotype. It is an internal system, what Chomsky 1986 called the I-language, as distinct from the external linguistic production, the E-language. Since the grammar is represented in the mind/brain, it must be a finite system, which can relate sound and meaning for an infinite number of sentences.

Second, also specified are the relevant principles and parameters common to the species and part of the initial state of the organism; these principles and parameters make up part of the *theory of grammar*, or *Universal Grammar*, and belong to the genotype.

The third item is the *trigger* experience, which varies from person to person and consists of an unorganized, fairly haphazard set of utterances, of the kind that any child hears (the notion of a trigger stems from ethologists' work on the emergence of behavioral patterns in young animals). The universal theory of grammar and the variable trigger together form the basis for attaining a grammar; grammars are attained on the basis of a certain trigger and the genotype.

In (9) I give the explanatory schema, using general biological terminology in (9a) and the corresponding linguistic terms in (9b). The triggering experience causes the genotype to develop into a phenotype; exposure to a range of utterances in, say, English allows the UG capacity to develop into a particular mature grammar. One may think of the theory of grammar as making available a set of choices; the choices are taken in the light of the trigger experience, or the primary linguistic data (PLD), and a grammar emerges when the relevant options are selected. A child develops a grammar by setting the parameters of UG in the light of her particular experience.

9 (a) linguistic triggering experience (genotype → phenotype)
  (b) primary linguistic data (Universal Grammar → grammar)

Each of the items in the triplet – trigger, UG, and grammar – must meet various demands. The trigger, or PLD, must consist only of the kinds of things that children routinely experience and includes only simple structures.[6] The theory of grammar, or UG, is the one constant and must hold universally, so that any person's grammar can be attained on the basis of naturally available trigger experiences. The mature grammar must define an infinite number of expressions as well-formed, and for each of these it must specify at least the sound and the meaning. A description always involves these three items, which are closely related; changing a claim about one of the items usually involves changing claims about the other two. This tight, ambitious system of description must meet many empirical demands. It is hard to imagine that we might have to choose between two or more descriptions which meet all empirical demands, hard to imagine problems of indeterminacy of the kind that plague the natural historian or the taxonomist who does not take a psychological view of grammars.

The grammar is one sub-component of the mind, which interacts with other cognitive capacities or modules. Like the grammar, each of the other modules is likely to develop in time and to have distinct initial and mature states. So the visual system recognizes triangles, circles, and squares through the structure of the circuits that filter and recompose the retinal image (Hubel and Wiesel 1962). Certain nerve cells respond only to a straight line sloping downward from left to right, other nerve cells to lines sloped in different directions. The range of angles that an individual neuron can register is set by the genetic program, but experience is needed to fix the precise orientation. In the mid-sixties David Hubel, Torsten Wiesel, and their colleagues devised an ingenious technique to identify how individual neurons in an animal's visual system react to specific patterns in the visual field (including horizontal and vertical lines, moving spots, and sharp angles). They found that particular nerve cells were set within a few hours of birth to react only to certain visual stimuli, and, furthermore, that if a nerve cell is not stimulated within a few hours, it becomes totally inert in later life. In several experiments on kittens, it was shown that if a kitten spent its first few days in a deprived visual environment (a tall cylinder painted only with vertical stripes), only the neurons stimulated by that environment remained active; all other optical neurons became inactive, because the relevant synapses degenerated, and the kitten never learned to see horizontal lines or moving spots in the normal way. Thus learning is a selective process: parameters are provided by the genetic equipment, and relevant experience fixes those parameters. A certain mature cognitive structure emerges at the expense of other possible structures, which are lost irretrievably as the inactive synapses degenerate. The view that there is a narrowing down of possible connections from an overabundance of initially possible ones is now receiving more attention in the light of Hubel and Wiesel's Nobel Prize-winning success. At the moment, this seems to be a more likely means of fine tuning the nervous system as "learning" takes place, by contrast with the earlier view that there is an *increase* in the connections among nerve cells.

So human cognitive capacity is made up of identifiable properties that are genetically prescribed, each developing along one of various pre-established routes, depending on the particular experience encountered. These genetic prescriptions may be extremely specialized, as Hubel and Wiesel showed for the visual system. They assign some order to our experience. Experience elicits or triggers certain kinds of specific response, but it does not determine the basic form of the response.

This kind of modularity is very different from the view that the cognitive faculties are homogeneous and undifferentiated, and that the faculties develop through general problem-solving techniques. In the physical domain, nobody would suggest that the visual system and the system governing the circulation of the blood are determined by the same genetic regulatory mechanisms. Of course, the possibility should not be ruled out that the linguistic principles postulated here may eventually turn out to be special instances of principles holding over domains other than language; but before that can be established, more, much more, must be known about what kinds of principles are needed for language acquisition to take place under normal conditions, and similarly for other aspects of cognitive development. Only then can meaningful analogies be detected. Meanwhile,

> we are led to expect that each region of the central nervous system has its own special problems that require different solutions. In vision we are concerned with contours and directions and depth. With the auditory system, on the other hand, we can anticipate a galaxy of problems relating to temporal interactions of sounds of different frequencies, and it is difficult to imagine that the same neural apparatus deals with all of these phenomena . . . for the major aspects of the brain's operation no master solution is likely. (Hubel 1978: 28)

3.5 Real-Time Acquisition of Grammars

In the domain of language, some ingenious colleagues at the University of Maryland have shown that the sophisticated distinctions discussed at the beginning of this chapter do not result from learning, and that the hypothesized genetic constraints seem to be at work from the outset. The experimenters constructed situations in which the overriding temptation for children would be to violate the relevant constraints. The fact that children conformed to the hypothesized constraints, resisting the preferences they showed in other contexts, is taken to be evidence that they have the constraints under investigation, and have them at the earliest stage that they might be manifested (Crain 1991).

Stephen Crain and Rosalind Thornton developed an elicitation task that encouraged children to ask questions like **Who do you wanna go?* (cf. 6b), if these were compatible with their grammars. They hypothesized that children would generally show a preference for the reduced *wanna* form whenever this was consistent with their grammars. This preference would be revealed in a frequency count of legitimate forms, like *Who do you wanna see?* (cf. 6a). Comparing the frequency of the reduced forms in these contexts with non-adult reduced forms would indicate whether or not children's grammars contained the hypothetical genetic constraint. If the genetic constraint is at work, there should be a significant difference in frequency; otherwise, not.

In the simplest case, an experimenter asked for a child's help in obtaining information about rats. Help was sought because the rat (a puppet) was too timid to talk to grown-ups. The experimenter said, "The rat looks hungry. I bet he wants to eat something. Ask Ratty what he wants." And the children, who ranged in age from 2 years, 10 months, to 5 years, 5 months, typically would ask "What do you wanna/want to eat?" In this example, the wh- word is understood as the object of *eat*, and the reduced form could occur freely, as in (6a). In fact, the reduced form occurred 59 percent of the time in these examples, and the unreduced form occurred 18 percent of the time. So children do show a preference for the reduced form, and that is the extent of it.

Something very different happened when the wh- word had to be understood in the subject position, as in (6b), *Who do you want to go?* The protocol for this experiment was that the experimenter would ask: "There are three guys in this story: Cookie Monster, a dog, and this baby. One of them gets to take a walk, one gets to take a nap, and one gets to eat a cookie. And the rat gets to choose who does each thing. So, *one* gets to take a walk, right? Ask Ratty who he wants." And the child would typically reply: "Who do you want to take a walk?" Here *who* is understood as the subject of *take*, i.e. between *want* and *to*: who_x you want [x to take a walk]? In these contexts the frequency of the reduced form was quite different: the reduced forms occurred 4 percent of the time, the unreduced forms 67 percent. In fact, one child accounted for all three actual occurrences of the reduced form, which suggests that this particular child had some other analysis of *wanna* forms. So children prefer to use the reduced form in asking questions like (6a), but correctly resist this preference when they should. They use the reduced form in asking questions like (6a), but not in questions like (6b), so they manifest the hypothetical genetic constraint at a stage when their spontaneous production manifests very few instances of long-distance wh- movement. The ingenuity of the experiment shows that even at this stage the relevant principle is operating (Crain 1991).

Wanna contraction is an example of leftward cliticization, but the same story holds for rightward cliticization, in the use of reduced *is*, *will*, and so on. Thornton and Crain conducted a similar experiment to elicit another kind of long-distance question. The target productions were evoked by the following protocols for rightward cliticization:

(10) *Experimenter*: Ask Ratty if he knows what that is doing up there.
Child: Do you know what that's doing up there?
Rat: It seems to be sleeping.

(11) *Experimenter*: Ask Ratty if he knows what that is up there.
Child: Do you know what that is up there?
Rat: A monkey.

In (10) the child is invited to produce a sentence where *what* is understood as the object of *doing*: *Do you know what$_x$ that is doing x up there?* Therefore, *is* may be cliticized to the immediately right-adjacent element, *doing*. However, in (11) the child produces a sentence where *what* is understood as the complement of *is*, i.e. between *is* and the following item: *Do you know what$_x$ that is x up there?* (cf. *That is a bottle up there*). The intervening *x* marker prevents the *is* from cliticizing on to *up* in adult speech; no adult would say **Do you know what that's up there*, with the reduced form (cf. *That's a bottle up there*). Thornton and Crain found that young children behaved just like adults, manifesting the hypothetical genetic constraint. The children tested ranged in age from 2 years, 11 months, to 4 years, 5 months, with an average age of 3 years, 8 months. In the questions elicited there was not a single instance of the reduced form where it is impossible in adult speech. Children produced elaborate forms like those of (12), but never with the reduced *is*.

12 (a) Do you know what that black thing on the flower is? (4 years, 3 months)
(b) Squeaky, what do think that is? (3 years, 11 months)
(c) Do you know what that is on the flower? (4 years, 5 months)
(d) Do you know what that is, Squeaky? (3 years, 2 months)

The experiments just described deal with elicited production, but comprehension studies also show that hypothetical genetic constraints are in effect in very young children, at the earliest stage at which they can be tested. Thornton (1994) reported children's comprehension of yes/no questions containing negation, such as (13). The difference between the two forms lies in the structural position of the negative: in (13a) the negative is inside the IP (partial structure given in (13a′)), whereas in (13b) it has cliticized to *did* and moved out of the IP to C (13b′).

13 (a) Did any of the turtles not buy an apple?
(a′) $_{CP}$[did $_{IP}$[any of the turtles **not** buy an apple]]
(b) Didn't any of the turtles buy an apple?
(b′) $_{CP}$[**didn't** $_{IP}$[any of the turtles buy an apple]]

The position of the negative corresponds to two distinct interpretations. That correspondence between meaning and structural position follows from principles of UG, which we need not go into here; essentially, a negative "scopes over" any element within its complement. The phenomenon is clear. Suppose that turtles A and B bought an apple but turtle C did not. Then if somebody asked question (13a), an appropriate answer would be that turtle C did not. If somebody asked (13b), then the appropriate answer would be very different: turtles A and B did. So children's responses to questions like (13a) and (13b) reveal how they interpret negatives. In particular, responses to (13b) show whether children interpret the negative in the higher structural position. This is worth testing, because Thornton found that all her children *produced* non-adult negative questions. Most doubled the auxiliary verb (*What do you don't like?*), and one failed to move the auxiliary to the position of C (see structure (4)): *What you don't like?*

In testing comprehension, Thornton found that the children had no difficulty interpreting negative questions in adult fashion; significantly, all children were able to access interpretations like (13b), where the negative needs to be interpreted in the position of C. She tested children between the ages of three-and-a-half and four-and-a-half. The comprehension test used a modified form of the Truth Value Judgment Task (Crain 1991). A story was acted out by one experimenter and was watched by the child and a second experimenter, who played the role of a puppet, in this case "Snail." At the end of each story, the experimenter asked Snail a targeted question. Snail had difficulty with the question ("That's a hard one . . ."), and requested help from the child. If the child was cooperative, she answered the question for Snail.[7] The scenarios used to test children's comprehension of questions like (13a) and (13b) were designed so that either (13a) or (13b) could be asked appropriately; children's answers, however, indicated their analysis of the structural position of the negative. Thornton found that, while these children made *production* errors with expressions like adult *What don't you like?*, their *comprehension* was adult-like and manifested the UG principles which determine the scope of negatives.

So there is a clear production/comprehension asymmetry, which, of course, is no surprise on the modular view of mind that I have articulated. Whatever it is that causes the delay in producing the adult forms, the fact that children interpret the negative questions in adult fashion shows that they have access to whatever principles of UG assign scope relations. The difficulty evidently lies in the behavior of the clitic *n't*: children produce non-adult questions which retain the *n't* in the inflection phrase IP until they figure out that *n't* may cliticize to *did* and move with it outside the IP to the higher complementizer position C.

This last experiment illustrates the obvious fact that young children are not entirely like adults in their syntactic behavior, even if they seem to manifest at the earliest stage possible the hypothesized genetic constraints that we have been discussing. In another study, Thornton analyzed an intriguing type of non-adult speech. She observed wh- questions like (14) in some three- and four-year-old children, where a copy of the wh- word shows up at the front of the clause where it is understood; such sentences do not occur in adult English, but analogues to them occur in dialects of German and in other languages (Thornton 1995).

14 (a) What do you think what pigs eat?
   (b) Who do you think who eats trash?

Furthermore, even after children stop using medial wh- words in object extraction questions like (14a), they persist in using them in subject extractions like (14b). That developmental sequence illustrates principles of UG: the extra wh- word in subject extraction questions is an overt manifestation of the "agreement" properties needed to license an empty subject. That won't make much sense until we get to chapter 9 (section 9.3) and consider subject extraction in the context of the evolution of UG in the species,

but my point here is more primitive: some children systematically produce things like (14), even though they never hear such sentences uttered by adult speakers of English. Children do not simply imitate models; they develop a system, a grammar, which at certain stages of development yields things which no child hears from English-speaking adults. When we study the development of these grammars, we can often understand the properties they have, even when they do not reflect what children have heard.

There is, of course, much more to be said about grammars and their acquisition, and there is an enormous technical literature. Here I have tried to sketch the essence of the biological view of grammars, because it is central to my account of language change. Now we shall proceed to investigate language change as change in individual grammars, and we'll see where that gets us.

Meanwhile, we have an approach to the riot of differences that we find in the languages of the world and even within individual languages: there's a universal language, and it is this UG which makes us accessible to one another. As children, our linguistic experience varies tremendously; no two children experience the same set of sentences, let alone the same pronunciations. Nonetheless, the approach I have sketched enables us to understand the universality of our development, why we categorize the linguistic world so similarly and can talk to each other, despite the enormous variation in our childhood experience.

For centuries people have held that organisms grow according to the prescriptions of an internal regulatory program. As a result of twentieth-century work in molecular biology, we now understand much about how that regulatory mechanism works, how it is transmitted, how it can be amended, and so on. But the idea that there must exist an internal regulatory program long antedates recent work in molecular biology. For centuries scientists theorized that the sperm contained a perfect miniature creature, a "Russian doll," or "homunculus," which simply grew bigger as time went on. This was preformationism, which was quite a reasonable theory to hold in the eighteenth century. Earlier, Plato held that the basic ideas and elements of thought were innate, and that at birth we drank of the River Lethe, the river of forgetfulness, which rendered this equipment subconscious. We no longer hold such ideas, but they do suggest that the question of how to account for the way in which living things develop has been a basis for theorizing. People have long thought that the development is internally directed in some way. The reason for this belief has always been that environmental factors alone cannot determine certain features of the mature makeup of an organism.

Recently, theoretical developments have brought an explosive growth in what we know about human languages. Linguists can now formulate interesting hypotheses, account for broad ranges of facts in many languages, with elegant abstract principles. Work on human grammars has paralleled work on the visual system and has reached similar conclusions, particularly with regard to the existence of highly specific computational mechanisms. In fact, language and vision are the areas of cognition that we know most about; we know much more in these domains than we know about memory, the emotions, or, for heaven's sake, consciousness. Much remains to be done, but we can show how children attain certain elements of their language organs by exposure to only an unorganized, haphazard set of simple utterances; for these elements we have a theory which meets basic requirements. Eventually, the growth of language in a child will be viewed as similar to the growth of hair: just as hair emerges with a certain level of light, air, and protein, so, too, a biologically regulated language organ emerges under exposure to a random speech community.

Our focus here is grammars, not the properties of a particular language, or even general properties of many or all languages. A language on the view sketched here is an epiphenomenon, a derivative concept, the output of certain people's grammars (perhaps modified by other mental processes). Relegating the notion of a language in this way avoids various problems of classification: we no longer need to decide whether dialects of Swedish and Norwegian belong to one of two distinct languages, or whether two dialects of Chinese would be better classified as different languages. A grammar has a clearer status: it is the finite system that characterizes an individual's linguistic capacity and that is represented in the individual's mind/brain. No doubt the grammars of two individuals whom we regard as speakers of the same language will have much in common, but there is no reason to worry about defining "much in common," or about specifying when the outputs of two grammars constitute one language. Just as it is unimportant for most work in molecular biology whether two creatures are members of the same species (as emphasized, for example, by Monod (1972: ch. 2) and Dawkins (1976)), so too the notion of a language is not likely to have much importance if our biological perspective is taken.

So when we think about change over the course of time, diachronic change, we shall now think not in terms of sound change or language change, but in terms of changes in these grammars, which are represented in the mind/brains of individuals and emerge and grow in the way that we have discussed in this chapter. Under certain conditions, if the linguistic environment is a little different, a child's brain may grow a grammar somewhat different from that of her mother. We'll investigate what those conditions are, and how that growth takes place.

Notes

1 Other examples of specifiers are *American* in *American attempts to dispose of Saddam Hussein*, *have* in *They have read the article*, *straight* in *Go straight to Chicago*.

We see that *will visit Chicago* is represented as a constituent, but it is labeled a I′ in 3. The I′ is an intermediate inflection projection; it designates an inflectional projection which is larger than $_I$[will] but which itself projects to a still larger inflectional structure, the IP. The category IP does not project to a larger type of inflectional expression and is, in this sense, maximal. Similarly, *Students from Denver* would be a NP, hence maximal, when part of a sentence *I met students from Denver*. However, these nodes are recursive, so one also finds expressions like *Students from Denver with green trumpets*. When embedded in this way *Students from Denver* is not maximal and is a N′, part of a larger NP. We reserve the XP notation for a node which is maximal. An element might be both a N and a NP, as *Denver* and *Chicago* in (3).

2 Government was an important notion in work conducted within the Government-Binding (GB) model of the 1980s, but it is not a primitive in more recent work conducted under the Minimalist Program (Chomsky 1995; Radford 1997; Uriagereka 1998). Nonetheless, the distinctions captured by the idea of government also need to be made in all models of grammar. Reformulating government in ways congenial to minimalist goals involves technical issues that I don't want to get into here, so I maintain the anachronism, although I define it in a more strictly local fashion than was used in much GB work.

3 Binding Theory, in fact, divides nouns into three types: anaphors (*herself*, *each other*, etc.), pronouns, and names. Each noun type has different indexical properties, defined by a principle of the Binding Theory. So Principle A applies to anaphors and says that they must be locally co-indexed; Principle B applies to pronouns, as we have seen; and Principle C applies to names. On this view, the child classifies each noun she encounters as belonging to one of the three classes; that is the learning involved. The indexing conventions apply only to the largest NP/DP available; that restriction is characterized in the technical literature as a function of "c-command" relations.

Until 1976 grammarians made the traditional assumption that structural conditions determined when pronouns might be co-indexed with other phrases. Then in a seminal paper, Howard Lasnik (1976) turned things around and argued that more elegant analyses would postulate structural conditions under which pronouns might *not* be co-indexed with other phrases.

4 In Old English, pronouns like *her*, *me*, etc. did double duty and behaved sometimes as pronouns subject to Principle B and sometimes as anaphors (subject to Principle A of the Binding Theory – see n. 3). So we find expressions like *Ic behyde me*, "I hid myself," which we discussed in chapter 1. Here *me* acts as an anaphor and needs to be glossed "myself" in modern English.

5 For good discussion of the viability of these neural networks as a basis for explaining language acquisition, see Lachter and Bever 1988 and Pinker and

Prince 1988. Jerry Fodor also has an excellent, accessible article on this topic in *The Times Literary Supplement* (1997). Where these connectionist models are worked out, contrary to their advertisements, they make strong, language-specific innateness assumptions throughout (Marcus 1997).

"Learning" is a slippery term. Languages are *learned* in the sense that children derive information from their linguistic environment and then converge on some grammar in accordance with specific principles. That, however, is a specialized sense of "learning," and not the way the term is used in informal discourse. If we understand learning in a nontechnical sense, then it is not clear that learning plays any role in the acquisition of language. In that case, we may be better off not to talk about children "learning" a language, but to talk instead in terms of the growth of language. This was Chomsky's point in the passage that Bates and Elman misrepresent.

6 One idea about limiting the trigger to simple structures is the "degree-0 learnability" of Lightfoot 1989, 1991, 1994, which posits that triggers are limited to elements from unembedded binding domains.

7 The child was not asked the question directly, in order to alleviate any feeling she might have of being tested; in this setup, Snail is being quizzed, not the child. Here I am giving only the briefest description of the experiments, but the experimental techniques used for this kind of work require great ingenuity and are of enormous interest in themselves. For excellent discussion, see Crain and Thornton 1998.