For Amy and Beth

Patterns in the mind

Language and human nature

RAY JACKENDOFF

BasicBooks

1 Finding our way into the problem: The nature/nurture issue

Why are we the way we are? Are we born that way, or are we products of our environment? Or some mixture? These basic questions lie at the root of any inquiry into human nature.

These questions can be interpreted in various ways. Most often, I find, people tend to think of "the way we are" in terms of differences among individuals: one's "nature" is seen as an issue of metabolism or intelligence or personality. What makes one person fat and another skinny, one sociable and one shy, one good at math and another good at art? Could they have been different if they had been brought up differently? Which things about ourselves can we change, and which are we fated to live with?

Another frequent interpretation of "the way we are" is in terms of differences among groups. Could people differ in intelligence, social behavior, or moral qualities along lines of race or gender or culture? If such differences exist, are they products of heredity or the environment? Far too often, alleged hereditary differences among groups have been used to justify repression, then "supported" with pseudo-scientific evidence. For the moment, let me only observe that even if such differences should exist, they provide no grounds, scientific or moral, for wholesale repression.

The main issues of human nature I want to think about in this book, though, are at the level of the species: What makes human beings the way they are? How are we different from animals? How are we like other animals and different from computers?

In order to find out what *makes* us the way we are, it stands to reason that we have to look closely at the way we are. If we want to know the balance of responsibility between nature and nurture—and how much about ourselves we *can* change—it helps to have a better idea of what the combination of nature and nurture is responsible *for*. In this book, I want to use human language as a vehicle for examining "the way we are."

I have two reasons for choosing language as a focus. First, the possession of language has always been regarded as one of the major differences between us and the beasts, so it's important to find out just exactly what we've got and they haven't. (We'll see in the next chapter how human language differs from other animal communication.) Second, and to me more important, the modern study of language has uncovered complexities of the mind far beyond what anyone would have imagined thirty years ago—complexities that draw on evidence from, and have implications for, fields as disparate as neuroscience, child development, philosophy, and literary criticism. Consequently, understanding language offers the prospect of integrating biological and humanistic views of "the way we are."

How might we bring language to bear on questions of human nature? One natural way is to ask: How is human experience affected by the fact that we can all speak and understand a language? A number of answers come to mind pretty easily. Most obviously, by virtue of having language, we have access to history: our ancestors have conveyed to us, through either written documents or oral tradition, a record of what happened before we were born. Along with history, we get our culture's accumulation of technology, world views, and rituals—not to mention legal systems, propaganda, gossip, and jokes. Little of this, if any, could be transmitted without language.

Another thing that language does for us is make it possible to coordinate the actions of large numbers of people. A bird's alarm call can make a whole flock flee at once. But people can communicate more differentiated things such as: "When I give the signal, you people over there pull on your ropes, and you people here let go of your ropes, and you other guys over there push like crazy." This kind of directed and coordinated action is hard to imagine without language, and it's necessary in order to do things like erect large structures, a hallmark of advanced civilizations.

The advantage that language is perhaps most often said to confer on us is that it enables us to think. While there is a great deal of truth to this idea—language certainly is invaluable in helping us sharpen certain kinds of thoughts—we should be a little cautious about endorsing it entirely. For one thing, we probably don't want to deny the capability of thought to at least some animals. For another, not all *human* thought requires language. Did it take thought for Beethoven and Picasso to produce their masterpieces? (I think so.) Did it take *language*? (I don't think so.)

Whatever the precise relation of language and thought, though, it is undeniable that human existence is deeply affected by the ability to speak and understand language.

In this book, however, I want to ask a different question about the relation of language and human nature: *What does human nature have to be like to account for the fact that we can all speak and understand a language?* That is, I want to discuss not the *consequences* of having language but rather the *prerequisites* for language: What do we need in order to be able to talk?

It's hard to think up plausible answers to this question. Or rather, the answers that spring immediately to mind turn out to be less than persuasive. For instance, one possible answer is that we have language because we have bigger brains than (other) animals. Let's be a bit more careful, though. After all, there are other animals with big brains—elephants and whales have brains bigger than ours, and the brains of bottlenose dolphins are larger in proportion to body size than ours—but they don't have language (or if they do, it's nothing like human language).

It's natural to think that a big brain makes us more intelligent, and because we're more intelligent we've figured out how to talk. But in what ways does a big brain make us smarter? As we'll see, it's not so obvious how being smart in and of itself makes talking possible.

In fact, there is a basic difficulty with an explanation that relies just on brain size. For now, it can be stated like this: you can't always get an entirely new function out of a device just by adding more of the same parts. To take a crude example, you can't get your car to fly by adding more cylinders to the engine, or more speeds to the transmission, or more wheels or bigger windows. Its existing function of carrying you along the road comfortably may be improved in some way, but the damn thing still only travels on the ground. To get it to fly, you need some sort of structural innovation like wings or a helicopter rotor. A major theme of this book is that the same is true of the brain and language: expanding a monkey's brain to the size of ours would still not enable it to talk. Beyond size, there has to be some difference in the way our brains are put together.

For the moment, the main thing is to appreciate how hard a problem this is. The fact that we can talk (and cats can't) seems so obvious that it hardly bears mention. But just because it's obvious doesn't mean it's easy to explain. Think of another perfectly obvious, well-known phenomenon: the fact that metals turn red when you heat them enough. Why does this happen? It could be otherwise—

they might just as well turn green or not change color at all. It's a simple phenomenon, easily observable, but the explanation isn't simple at all. It turns out to involve at the very least the theories of electromagnetic radiation and quantum mechanics, two of the more amazing intellectual advances of the past century. So it is, I want to suggest, with the human ability to use language.

The basic parameters underlying a theory of language ability were first laid out in the late 1950s and early 1960s by Noam Chomsky, who can justifiably be called the creator of modern linguistic theory (and who is, at the time of this writing, still doing pathbreaking research). I am going to lay these parameters out in a form that I'll call the two Fundamental Arguments. Just to give you an idea of where we are going, let me state them in very abbreviated form:

> **The Argument for Mental Grammar:**
> **The expressive variety of language use implies that a language user's brain contains a set of unconscious grammatical principles.**
>
> **The Argument for Innate Knowledge:**
> **The way children learn to talk implies that the human brain contains a genetically determined specialization for language.**

These two arguments lead us to the conclusion that the ability to speak and understand a human language (say English) is a complex combination of nature and nurture. Moreover, the part coming from nature involves more than a big brain: it is a specific human adaptation for language learning and use. The next two chapters will work through the Fundamental Arguments; Parts II and III will be spent explaining, qualifying, and elaborating them.

Part IV places the Fundamental Arguments in a larger context. It asks: If the human brain contains unconscious grammatical principles and a genetically determined specialization for language, what are the implications for other aspects of human behavior and experience? We will see that language is a revealing microcosm of the mind as a whole—that similar characteristics emerge in activities as disparate as seeing, thinking, listening to music, and taking part in a social environment. Along the way, we will present a third Fundamental Argument, whose consequences are perhaps even more radical than those of the first two:

> **The Argument for the Construction of Experience:**
> **Our experience of the world is actively constructed by the unconscious principles that operate in the brain.**

I am shaping this book around the Fundamental Arguments because, of all the starting points I know for investigating language, they motivate the deepest possible scientific inquiry. On one hand, as we will see, they force us to integrate all sorts of issues from the broadest range of sources. And on the other hand, they yield the greatest insight into human nature in general: they allow us to see language as part of an integrated whole.

2 The argument for mental grammar

The communicative situation

Let's start with a fairly crude picture of the communicative situation—what goes on when one person says something to another.

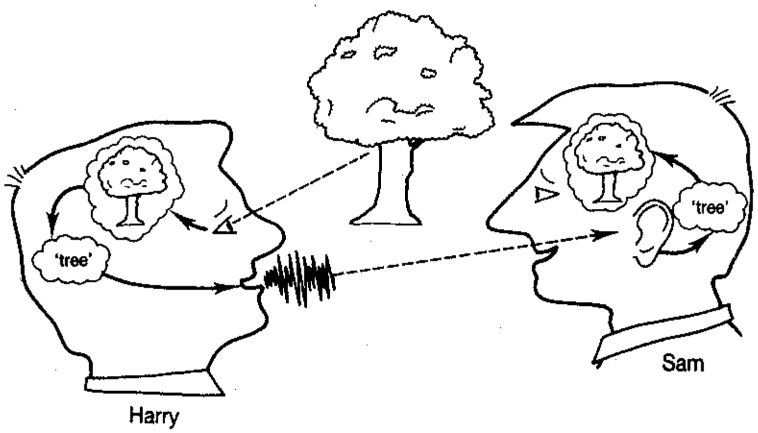

Figure 2.1 *The communicative situation: Harry tells Sam about a tree*

In this picture (Figure 2.1), a pattern of light reflected off of a tree strikes the eyes of the person on the left (let's call him Harry). As a result of activity in Harry's nervous system, he comes to see the tree out there in the world. This is indicated in the picture by a tree inside a little cloud in Harry's head. Of course we know there are no clouds or trees in people's heads, and eventually (Chapter 13) we'll ask what's really there, but let this stand for the moment.

Once Harry has perceived a tree, it may occur to him that the

word "tree" describes what he has seen—that is, the word "tree" is evoked from his memory. (If Harry spoke French rather than English, of course, the word "arbre" would be evoked instead.) This is indicated in the picture by the word "tree" in another little cloud in Harry's head. Again, we know there are really no little clouds, but this will have to do for now (we'll come back to it in Chapter 4).

Perhaps Harry decides to say something about the tree to the person on the right (let's call him Sam). Then Harry's nervous system causes his lungs to expel air, his vocal cords to tighten, and his tongue and jaw and lips to go through some gyrations. As a result, he produces some sound waves which travel through the air, striking Sam's ears, Sam's eyes, the furniture, and everything else.

But unlike Sam's eyes and the furniture, Sam's ears react to these sound waves by activating Sam's nervous system, so that he comes to perceive Harry uttering the word "tree." Assuming Sam also speaks English, his nervous system very likely goes on to produce a visual image of a tree—Sam is able to imagine what Harry sees, though probably not in many of its particulars.

Even this little dissection of the obvious has revealed quite a lot of complexity. There are a lot of parts to this simple communicative act, and each one of them involves tough puzzles. (For example: What is really in Harry's and Sam's brains instead of the little clouds? Exactly what gyrations of Harry's tongue, jaw, and lips take place? What happens in Sam's ears?) But we still haven't seen the full difficulty of the problem.

Suppose Harry wants to say something a little more interesting about what he sees than the single word "tree." Here are some things he might say (I'll put numbers and letters in front of example sentences so we can refer back to them later):

(1) *a* There's a bird in the tree.
b A bird was in the tree yesterday.
c Are there any birds in that tree?
d A bird might be in the tree.
e Birds like that tree.
f That tree looks like a bird.

This time it isn't so easy to draw pictures in little clouds that depict what Harry has in mind. What difference can we make in the pictures in order to distinguish sentences (1a), (1b), (1c), and (1d)? (If we start putting question marks and writing in a picture, that's cheating—it's not just a picture anymore!) For sentence (1e), how do we show that the birds *like* the tree rather than, say, merely swarm around it?

In all the cases so far, the picture at least has both a bird and a tree in it, whatever its other failings. But what about sentence (1f)? What seems to come to mind is something like a bird-shaped tree. But such a picture has only one object corresponding to both words—yet another complication.

These examples illustrate some of the *expressive variety* of language—the number of different things we can say by combining words in different ways. Moreover, this expressive variety in many respects can't be conveyed by pictures, whether on a piece of paper or in the head. That is, significant parts of the messages that language conveys are abstract, or nonsensory, in nature.

Here we see a significant difference between human language and any of the forms of animal communication. To be sure, many kinds of animals convey information to each other. But in none of the known systems—birds, bees, whales, nonhuman primates, or whatever—is there an inventory of elements like words that can be combined and recombined in limitless new ways to express new messages. There are no elements that indicate points in time ("yesterday"), a desire for information ("are there ... "), or possibility ("might"). Animals may have a way to indicate their own desires or feelings, but they can't convey someone else's, as in "Birds like that tree." Nor can animal communication systems explicitly draw resemblances among different objects, as in "That tree looks like a bird."

So, although people often speak loosely of animal communication as a kind of language, in fact the way animals communicate is orders of magnitude different from the way humans do. To make this distinction clear, I will adopt the policy of using the word "language" to mean only "human language (Spanish, Chinese, Navajo, etc.)," and I will use the more general word "communication" for any means by which information is conveyed, including both language and animal systems. (I will mention some attempts to teach human languages to apes in Chapter 10.)

The argument for mental grammar: The expressive variety of language use implies that a language user's brain contains unconscious grammatical principles

The expressive variety of language is the springboard for the first of the Fundamental Arguments. Any normal human being can understand and create an indefinitely large number of sentences in his or her native language. Aside from stereotyped utterances like "Hi, how are you?" and "Please pass the salt," most of the sentences we speak in the course of a day are sentences we have never heard or spoken in their entirety before. The same is true of most of the sentences we hear. For example, I doubt that you have ever heard or spoken any of the sentences on this page before. Yet you have no difficulty understanding them.

Let's think about what must be going on in your head that makes this possible. In the previous section we assumed that Harry and Sam could simply pull the word "tree" out of their memories when needed. Could this be true of whole sentences as well?

No. The number of sentences we are capable of using is just too large to store them individually. Let me run up the number in some rather stupid ways, just as a sample. Consider this series of sentences, all of which are perfectly comprehensible.

(2) Amy ate two peanuts.
Amy ate three peanuts.
Amy ate four peanuts.
...
Amy ate forty-three million, five hundred nine peanuts.
...
...

There are as many sentences in this series as there are nameable integers. The biggest number name listed in my Webster's Collegiate is a vigintillion (10^{63} in US/French usage; 10^{120} in British/German usage). With all the numbers up to this at our disposal, we can create more sentences in this series than there are elementary particles in the universe.

Here's another way to make lots of sentences. There are at least some tens of thousands of nouns in English. Let's be conservative and say we know ten thousand (10^4). Now let's construct all the sentences we can by putting in different nouns for X and Y in "An X is not a Y." Here are some of them.

(3) A numeral is not a numbskull.
A numeral is not a nun.
A numeral is not a nunnery.
...
A numbskull is not a numeral.
A numbskull is not a nun.
A numbskull is not a nunnery.
...

A nun is not a nursery.

. . .

An oboe is not an octopus.

. . .

These are all completely absurd, but they *are* sentences of English nevertheless. There will be something like $10^4 \times 10^4$ of them = 10^8. Now let's put pairs of these sentences together with "since," like this:

(4) Since a numeral is not a numbskull, a numbskull is not a nun.
Since a numeral is not a numbskull, a numbskull is not a nunnery.
Since a numeral is not a numbskull, a numbskull is not a nuptial.
. . .
Since a numeral is not a nursery, a numbskull is not a nun.
. . .
Since an oboe is not an octopus, a numeral is not a numbskull.
. . .

And so on it goes, giving us $10^8 \times 10^8 = 10^{16}$ absolutely ridiculous sentences. Given that there are on the order of ten billion (10^{10}) neurons in the entire human brain, this divides out to 10^6, or one million sentences per neuron. Thus it would be impossible for us to store them all in our brains, in the unlikely event that we should ever want to use or understand any of them. But still, you did just understand a sampling of them. And these lists are only a minute proportion of the sentences you can understand. What lists include the sentences of this paragraph, for instance?

In short, we can't possibly keep in memory all the sentences we are likely to encounter or want to use—not to mention all the unlikely ones such as the sentences in (2)–(4). On the other hand, we are apparently ready to encounter them—we seem to know what the possibilities are.

The way the brain seems to achieve expressive variety is to store not whole sentences, but rather words and their meanings, plus *patterns* into which words can be placed. For example, it is only by using patterns that we can reasonably store the sets of sentences of which (2), (3), and (4) form a tiny sample: the pattern for the sentences in (2) is "Amy ate N peanuts"; that for the sentences in (3) is "An X is not a Y"; and that for the sentences in (4) is "Since an X is not a Y, a Z is not a W." With such patterns, plus a list of words to insert into them, we can specify a large number of possibilities at minimal cost in storage. Moreover, such a system is prepared for *novelty*: it can recognize or create examples of the pattern on the spur of the moment, whether or not they have been encountered before.

But even using these kinds of fixed patterns isn't quite good enough. Consider the list of sentences in (5).

(5) *a* Bill thinks that Beth is a genius.
b Sue suspects that Bill thinks that Beth is a genius.
c Charlie said that Sue suspects that Bill thinks that Beth is a genius.
d Jean knows that Charlie said that Sue suspects that Bill thinks that Beth is a genius.
. . .

This sequence can be extended as long as we have the patience—that is, it is effectively infinite. (To be more precise, there is no longest sentence in this sequence, because we can always add one more.) As a result, we can't specify a single pattern for this list the way we could for the lists sampled in (2)–(4). Rather, each sentence has to come from a different pattern, and the patterns get longer and longer. (6) shows the first three of these patterns; the term "Verbs" stands for one of the words "thinks," "suspects," "knows," and so forth.

(6) X Verbs that Y is a Z.
W Verbs that X Verbs that Y is a Z.
T Verbs that W Verbs that X Verbs that Y is a Z.
. . .

Can we store all these patterns in our heads? Again, no, because no matter how many we store, there is always a longer one. On the other hand, there is clearly a more basic pattern involved: given any declarative sentence, we can make another declarative sentence by placing "X Verbs that . . . " in front of it. For instance, we can apply this pattern to any of the sentences in (2)–(4) above to get whole new classes of sentences. Here are some of them (marking in italics the sentence we started with): "Bill knows that *Amy ate two peanuts*," "Wolfgang realizes that *an oboe is not an octopus*," "Ludwig suspects that *since a numbskull is not a nunnery, a nun is not a nuptial*," and so on. This pattern can be summarized as the formula given in (7).

(7) X Verbs that S. (where *S* is any declarative sentence)

Going back to the sequence of sentences in (5), we can apply formula (7) to the sentence "Beth is a genius" to get "Bill thinks that *Beth is a genius*," sentence (5a). And then comes the fun: we can use our new sentence as the sentence *S* in formula (7), giving us "Sue suspects that *Bill thinks that Beth is a genius*," sentence (5b); then we can use *this* sentence as *S* in (7), giving us "Charlie said that *Sue suspects that Bill thinks Beth is a genius*," and so on as long as we want. That is, we get longer and longer sentences by applying formula (7) over and over to its own output, or *recursively*. What makes (7) different from the earlier patterns is that it contains another pattern within it: instead of just putting words into the slots in the pattern, we insert another pattern—in this case a whole declarative sentence.

This is a typical case of what we find in the course of investigating the expressive variety of language. The sequences in (8) and (9) show two more patterns with patterns inside them; as in (5), we can go on applying them recursively till our patience runs out.

(8) *a* Ben's father is a linguist.
b Ben's father's older brother is a linguist.
c Ben's father's older brother's best friend is a linguist.
d Ben's father's older brother's best friend's former lover is a linguist.
. . .

(9) *a* This is the house that Jack built.
b This is the refrigerator that sits in the house that Jack built.
c This is the cheese that fell out of the refrigerator that sits in the house that Jack built.
d This is the mold that grew on the cheese that fell out of the refrigerator that sits in the house that Jack built.
. . .

In short, in order for us to be able to speak and understand novel sentences, we have to store in our heads not just the words of our language but also the patterns of sentences possible in our language. These patterns, in turn, describe not just patterns of *words* but also patterns of *patterns*. Linguists refer to these patterns as the *rules* of language stored in memory; they refer to the complete collection of rules as the *mental grammar* of the language, or *grammar* for short.

This demonstration of the expressive variety of English, complete with recursive patterns, can be reproduced in any of the human languages of the world. The particular patterns of mental grammar may not be the same from one language to the next, but patterns of comparable complexity can always be found. In this respect, there is no difference between the languages of contemporary Western societies, those of present-day "primitive" cultures, and those of the distant past that can be recovered from written records. (An important exception arises in "pidgin" languages, to be discussed in Chapter 10.)

Clarifying the notion of mental grammar

The notion of a mental grammar stored in the brain of a language user is *the* central theoretical construct of modern linguistics. So it's important to make it as clear as possible before going on. Let me engage in a dialogue with an imaginary skeptic who raises some of the most common questions and objections.

Why should I believe that I store a grammar in my head? I just understand sentences because they make sense.

In reply I ask you: Why do some combinations of words "make sense" and others not? For instance, if we interchange adjacent words in the sentences in (2)–(5), to form chains of words like (10), we find that the sentences don't "make sense" anymore.

(10) Amy two ate peanuts.
A is numeral not a numbskull.
Bill that thinks Beth is a genius.
etc.

Why don't they make sense?

Well, these are sentences I've never heard before.

But look: You never heard the sentences in (2)–(5) before either, and even so, they "make sense" (albeit of a stupid sort).

What's the difference?

The difference is that the sentences in (2)–(5) are examples of patterns of English that we know, and the strings of words in (10) are not. That is, "making sense" involves, among other things, conformity to known patterns. In other words, the mental grammar plays some sort of role after all.

This is not to say that conformity to the patterns of English is the only factor involved in "making sense." Lots of sentences conform to the grammatical patterns of English but still don't "make sense."

(11) Colorless green ideas sleep furiously.
Bill elapsed three times this month.

I'm memorizing the score of the sonata I hope to compose someday.
The harvest was clever to agree with you.

These examples (drawn from early writings of Chomsky's) are certainly nonsense. But they do conform to the grammatical patterns of English, as we can see by substituting one or two more "sensible" words in each one:

(12) Large green lizards sleep soundly.
Bill sneezed three times this month.
I'm memorizing the score of the sonata I hope to perform someday.
The lawyer was clever to agree with you.

On the other hand, notice that if we exchange adjacent words in the sentences in (11), so that they violate the grammatical patterns of English, they sound far worse: "Colorless green sleep ideas furiously," "Bill three elapsed times this month," etc. In this case, it's not that they have strange meanings; rather, it's hard to say what they mean at all. So the mental grammar seems to be involved even in sentences like (11) that don't make sense.

In fact, we can recognize patterns of English even if not all the words are real English words. This is the basis of Lewis Carroll's famous poem *Jabberwocky*:

'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe . . .

These lines are clearly an example of the same pattern as the following, which contains all real words:

'Twas evening, and the slimy toads
Did squirm and wiggle in the cage . . .

This shows that the patterns themselves have a degree of life independent of the words that make them up. Indeed, if you start exchanging words at random in *Jabberwocky*, again the patterns fall apart altogether.

Why do you want to call it a **grammar** *that I store in my head? Why couldn't I just have a bunch of* **habits** *that I follow in speaking and understanding English?*

My return question is: What is a habit anyway? It's something stored in memory that guides behavior on appropriate occasions. If the "habitual" behavior varies from occasion to occasion, as it does in the case of language, what is stored in memory has got to be a pattern. Why? Because the brain can't store all the individual examples—and even if it could, there would be no reason to call this random collection of behaviors a unified "habit."

Once we realize that habits must themselves be stored patterns, we shouldn't have a problem acknowledging that the "habits" of speaking English involve storing the patterns of English. That is, claiming that our knowledge of English is a kind of habit doesn't eliminate the need for us to have grammars in our heads.

What about people who speak ungrammatically, who say things like "We ain't got no bananas"? They *don't have grammars in their heads.*

This question points up an important difference between the ordinary use of the term "grammar" and the linguists' theoretical construct "mental grammar." In ordinary usage, "grammar" refers to a set of rules taught in school that tell us how we should speak in order to conform to the norms of polite (roughly, educated middle-class) society. "Proper grammar" frowns on the use of "ain't," the use of "got" for "have," and the use of double negatives; the "proper" way to say this sentence is "We don't have any bananas" or "We have no bananas." In the sense of "school grammar," then, speaking ungrammatically is a violation of a social norm, sort of like spitting in public.

The concept of "mental grammar" provides a different perspective on this issue. The mental grammar in our heads is what enables us to put words together into sentences. So it has to specify not just which patterns are socially acceptable and which are not, but *all* the patterns of the language. This includes some patterns that are much more basic than they ever had to teach us in school, for instance that the subject precedes the verb in English, or that adjectives precede the nouns they apply to ("ripe banana," not "banana ripe"). But it also includes some patterns that are much more complex than those taught in school, as we will see in the next few chapters.

What about the people who don't speak "correct English"? A moment's reflection suggests that their speech does in fact fall into consistent patterns. Someone who says "We ain't got no bananas" still doesn't produce monstrosities like "ain't no we got bananas" or "no got ain't bananas we": the words come in a well-defined order. More subtly, such a speaker won't substitute the so-called correct term "have" for "got," saying "We ain't have no bananas." In other words, there are principles that govern the use of "incorrect" English too, even if it violates the canons of school grammar.

This means that such speakers don't *lack* a mental grammar; they just have a mental grammar that is slightly different from that of speakers of "correct" English. Setting aside the issue of social approbation, the situation is exactly parallel to the difference between "proper" British and American English. Speakers of these two dialects have slightly different mental grammars, so the patterns they produce don't match up exactly. Consequently, each sounds somewhat exotic (or sloppy) to the other.

In short, although my imaginary critic may wish to deplore certain people's language from the point of view of school grammar, it is hard to deny that they have a mental grammar in their heads that governs their patterns of speech.

> *When I talk, the talk just comes out—I'm not consulting any "grammar in my head." If I look into my mind, I may find some scraps of school grammar, but you're trying to tell me that's not what mental grammar is supposed to be. So what is it supposed to be?*

The answer to this question is potentially the most troubling. Here's the situation. We have just seen that an explanation of language ability demands that the patterns of language be stored in our memory somehow. We're now faced with the apparently conflicting fact that our memory reveals no such patterns to us. So something has got to give.

Can we give up the idea of a mental grammar? No: I've tried to convince you that just about any other way of thinking about the expressive variety of language amounts to the same thing. So let's grasp the other horn of the dilemma, and explore the hypothesis that the rules of language are not conscious, and are not available to introspection.

What could such a hypothesis mean? In this post-Freudian age, we are certainly accustomed to speaking of unconsciously (or subconsciously) guided behavior: "Willy has low self-esteem because he unconsciously identifies with his father." The premise of Freudian analysis, as well as most subsequent forms of psychotherapy, is that unconscious beliefs of this sort can be made conscious through suitable therapeutic procedures, and that in becoming conscious they cease to exert the same pernicious influence on one's experience and behavior.

Freud's notion that parts of the mind are not accessible to consciousness challenges the standard Cartesian identification of the

mind with *consciousness*: there is a lot more going on in our minds than we are ever aware of. This is upsetting not only because it goes against intuition ("I *know* what I think!") but also because it tells us we are not altogether in conscious control of our behavior. What's more, the Freudian unconscious is full of dark and uncomfortable motives. (Freud stressed the sexual underpinnings of those motives, perhaps because of his social milieu; modern psychodynamic theory recognizes many other themes as well.)

In a way, the unconsciousness of mental grammar is still more radical than Freud's notion of the unconscious: mental grammar isn't available to consciousness under *any* conditions, therapeutic or otherwise. On the other hand, an unconscious mental grammar that guides our behavior is a good deal less personally threatening than an Oedipus Complex or a Death Instinct. Unlike these Freudian constructs, mental grammar doesn't have pernicious effects. On the contrary, we couldn't speak without it, except in terms of stereotyped fixed expressions. It is mental grammar that makes possible the expressive variety of our language.

> *You're telling me that a mental grammar is present in my mind but that I'll never find it by looking there? Aren't you trying to pull a fast one?*

Well, consider: there are lots of other things going on in our brains of which we aren't conscious either. Think about getting from an intention such as "I think I'll wiggle my fingers now" into commands to be sent to the muscles, so that our fingers wiggle. Just how do we do it? From the point of view of introspection, the experience is entirely immediate: we decide to wiggle the finger, and the finger wiggles, unless there is some obstruction or paralysis. How the mind actually accomplishes this is entirely opaque to awareness. In fact, without studying anatomy, we can't even tell which muscles we've activated. So it is, I want to suggest, with the use of mental grammar.

If mental grammar can't be studied by introspection, then we have to find some other, less direct way to study it. I will take up this problem in Part II, showing how the investigation of mental grammar is an experimental science, and describing some of the organization that has been revealed by linguistic research. For now, the point is that if at least some other processes in the mind are not open to consciousness, it shouldn't be too distasteful to say that parts of language ability are unconscious too.

This, then, is our first inference about human nature on the basis of the nature of language. In order to account for the human ability to speak and understand novel sentences, we must ascribe to the speaker's mind a mental grammar that specifies possible sentence patterns. But in order to account for the fact that we have no direct access to this mental grammar, we must admit the possibility that some essential and highly structured parts of our abilities are completely unconscious.

3 The argument for innate knowledge

The character of language acquisition

We now turn to the preliminaries to the second Fundamental Argument. Suppose, following the discussion of the previous chapter, that we have mental grammars in our heads. The next question is: How did they get there?

Observation: All normal human children end up being able to speak whatever language is spoken in the community where they grow up. (If more than one language is spoken regularly, they usually end up speaking them all—but let's stick to the monolingual case for now.) And the language they speak has nothing to do with where their parents came from: a child of American parents growing up in Israel as part of a Hebrew-speaking community will become a native speaker of Hebrew; a Vietnamese baby adopted in Holland will become a native speaker of Dutch. So it's pretty obvious that children learn their language from the other speakers around them.

How do children do it? Many people immediately assume that the parents taught it. To be sure, parents often engage in teaching *words* to their kids: "What's this, Amy? It's a *BIRDIE*! Say 'birdie,' Amy!" But language learning can't be entirely the result of teaching words. For one thing, there are lots of words that it is hard to imagine parents teaching, notably those one can't point to: "Say 'from,' Amy!" "This is *ANY*, Amy!"

Think also about children of immigrants, say the Americans who move to Israel. The adults often never feel comfortable with the language of the adopted country. They speak with an accent, they express themselves with hesitation, they admit to not quite following the news on television, and so forth. Yet their children become fully fluent native speakers of the new language. Evidently the children have learned something their parents don't know. So the parents couldn't have taught them. Nor is the children's knowledge necessarily a result of teaching in school—and of course in nonliterate societies it *can't* be the result of teaching in school. More often, the children just "pick up" the language from being with other children. (This example also touches on another phenomenon, the fact that adults usually have much more difficulty learning a new language than children do. I'll return to this in Chapter 9.)

Although children often learn *words* as a result of parental instruction, it is less clear that they learn *grammatical patterns* this way. Anyone who has attempted to correct a two-year-old's grammar will know that it can't be done. The following dialogue, recorded by the linguist David McNeill, is a famous illustration.

CHILD: Nobody don't like me.
MOTHER: No, say "nobody likes me."
CHILD: Nobody don't like me.
...
(eight repetitions of this dialogue)
...
MOTHER: No, now listen carefully; say "*nobody likes me.*"
CHILD: Oh! Nobody don't likes me.

(Of course, we can be sure that this child eventually got it right. But it may well have been at a time when the mother wasn't even paying attention.)

It is true that certain grammatical patterns *are* taught as part of school grammar, for example the rule that a preposition is something you must never end a sentence with. However, English speakers violate this rule all the time, and have for hundreds of years. I just did, two sentences ago. The idea that a preposition shouldn't occur at the end of a sentence seems to have arisen during the eighteenth century, when for the first time "authorities on English usage" sought to determine the "correct" way to speak, on the basis of the models of the classical languages Latin and Greek.

Now Latin and Greek genuinely do not allow sentences that end with prepositions. Neither do most modern European languages (for instance French, Italian, Spanish, and, with some caveats, German; Swedish, however, is more like English). If we translate "Who did she arrive with?" word for word into those languages—say, "Qui est-elle arrivée avec?" in French—it sounds as barbarous as "Harry ate peanuts a hundred" does in English.

By analogy, the "authorities" ruled that prepositions shouldn't end sentences in English either. Since that time, generations of children have been drilled on this rule, with little effect except in their

formal writing. And ending sentences with prepositions is still very much alive in English.

Such proscriptive teaching of grammar, which evidently doesn't work very well, contrasts strikingly with aspects of English sentence patterns that probably nobody has ever thought to teach. Here's an example. Look at the four sentences in (1).

(1) *a* Joan appeared to Moira to like herself.
b Joan appeared to Moira to like her.
c Joan appealed to Moira to like herself.
d Joan appealed to Moira to like her.

Without thinking about it consciously, you have automatically inferred that each of these sentences has a different combination of who is to like whom. In (1a), Joan likes Joan; in (1b), Joan likes Moira or some unspecified third party; in (1c), Moira is to like Moira; in (1d), Moira is to like Joan or a third party.

How do we come to understand these sentences this way? It obviously depends somehow on the difference between ordinary pronouns such as "her" and reflexive pronouns such as "herself," and also on the difference between the verbs "appear" and "appeal." But how? Whatever reasons there may be, I'm sure no one is ever taught about contrasts like this by their parents or teachers or anyone else. Yet this aspect of English grammatical patterns is deeply ingrained, much more so than the taught prohibition against ending a sentence with a preposition.

I can't resist another example, because it's so striking. There is an alteration called "expletive infixation" that many speakers perform on words of English under conditions of extreme exasperation, as in (2).

(2) How many times do I have to tell you? I'm not talking about the *Allegheny* River! Can't you get it into your stupid head that I'm talking about the *Susque-goddam-hanna*?

Even if you're too refined ever to use an expression like this, I'm sure you recognize it. Now the interesting thing is that we have pretty clear intuitions about how to use this infix. It sounds natural in the examples in (3), but decidedly odd in those in (4).

(3) uni-goddam-versity
manu-fuckin-facturer
(4) Jacken-bloody-doff
ele-goddam-phant

In addition, for those words that allow us to use the infix, we are very particular about where it has to go. If we try moving the infix to different places in the words in (3) ("un-goddam-iversity," "manufac-fuckin-turer," etc.) we can see that only the versions given in (3) sound at all acceptable.

I'm fairly certain none of us was ever *taught* the principle (or pattern) that says where it is possible to insert an expletive infix into English words. Yet we readily use this principle to make intuitive judgments about new cases. At the same time, the principle is not so obvious to conscious introspection.

(In case you're wondering, the infix sounds right only when it immediately precedes the syllable of the word with main stress—"Susque*han*na," "uni*ver*sity," and "manu*fac*turer." Since "*Jack*endoff" and "*el*ephant" have main stress on the first syllable, there is no place to put the infix. But this is only a first approximation; there are further complexities that we can't go into here.)

We see, then, that much that we know about the grammatical patterns of English has not been taught. But this leads to a further problem about how children acquire language. Chapter 2 showed not only that we have a mental grammar, but that most of it isn't available to conscious introspection. Since adults aren't consciously aware of the principles of mental grammar (and the examples just presented provide further illustration), they certainly can't explain these principles to children—if children could understand the explanations in any event!

In fact, the most an adult can do is supply the child with *examples* of the patterns, in the form of grammatical sentences, or corrections to the child's sentences. For instance, notice that in the dialogue I quoted above, the mother isn't saying " 'Nobody' and 'not' are both negative words, and you shouldn't use two negatives in a sentence." She is just supplying the child with a correct *form*. This means that the child has to *figure out* the patterns of the language—that is, *the child has to construct his or her own mental grammar*. How?

Children are probably no more conscious of the patterns than adults. For instance, it doesn't make much sense to think that a child would confront sentences like "Joan appeared to Moira to like herself" by thinking "Hmm. I wonder who 'herself' is supposed to be. Well, 'herself' is a reflexive pronoun, so that probably makes a difference . . ." To be sure, children eventually learn the words "noun" and "verb," and maybe even "reflexive pronoun," but usually not until the age of ten or so, long after gaining command of the grammatical distinctions these words refer to.

Even simpler phenomena show the disparity between children's command of language and their *conscious* command of it. For instance, by the age of three or four, children can be taught to count syllables in a word, but they are certainly making use of syllables long before that. Similarly, learning to read depends in part on being conscious of sequences of speech sounds, in order to sound out words. For many children, this is difficult at age six and even later; that's why *Sesame Street* spends so much time on it. But at the same time, children couldn't discriminate and understand thousands of words by this age—not to mention appreciate rhymes—if they didn't have a sensitive ability to discriminate and sequence speech sounds. (We will see in Chapter 5 how this ability is organized.) So we're evidently faced with the same problem for children as for adults: their learning is backed by unconscious principles that are unavailable for conscious introspection. And if anything, we're tempted to suspect that children's abilities at introspection are less well developed than adults'.

Where does that leave the learning of language? On the basis of what the child hears in the environment, and in the (near-) absence of teaching and of conscious awareness of what is being learned, the child manages to acquire a command of the grammatical patterns of the language—that is, manages to construct a mental grammar. This isn't the way we're accustomed to thinking of language learning. We usually think of it in terms of something like French class in school, a highly structured situation in which teacher and learner bring a lot of conscious attention to bear on rules and regulations. The child's learning of grammatical structure just doesn't seem to be like that. The child learns just by speaking and being spoken to.

As a result, we can draw another conclusion about human nature: *We can acquire unconscious patterns unconsciously, with little or no deliberate training*. Perhaps we shouldn't even call such a process "learning," but for lack of a better word, let's leave the terminology alone.

A suggestive parallel to the unconscious learning of language might be the process of learning to skip, which requires complicated patterns of muscle coordination. It's impossible to describe to a child how to do it; the best we can do is demonstrate. And when the child figures out how to skip, it will be impossible to get him or her to explain it. Rather, the process of constructing the patterns takes place outside of consciousness; the major part of the learning is experienced as "just intuitive."

The argument for innate knowledge: The way children learn to talk implies that the human brain contains a genetically determined specialization for language

Here is what makes the child's acquisition of language even more remarkable. Thousands of linguists throughout the world have been trying for decades to figure out the principles behind the grammatical patterns of various languages, the very same grammatical principles that children acquire unconsciously. But any linguist will tell you that we are nowhere near a complete account of the mental grammar for any language. In other words, an entire community of highly trained professionals, bringing to bear years of conscious attention and sharing of information, has been unable to duplicate the feat that every normal child accomplishes by the age of ten or so, unconsciously and unaided. This contrast is so striking and so fundamental that it deserves a name. I like to call it the Paradox of Language Acquisition.

! ! !

What are we to make of this? How could linguists apparently be so inept compared to children, including the children they once were? Unfortunately, one commonly held attitude is that in fact linguists are just misguided, and that the complications they are struggling with simply don't exist. "Language just has to be simple: even a child can do it."

But if language is so simple, why hasn't anybody else, maybe someone without linguists' methodological blinders, figured it out either? As a case in point, one of the early predictions of the computer revolution was that we would have computers that talked to us and understood us within five years or so, as soon as we could build a machine big and fast enough.* But at the time of writing, forty years later, state-of-the-art computer understanding of spoken

* In fact, Chomsky's first book, *Syntactic Structures*, which appeared in 1957, acknowledges support from the US armed forces, who were at that time funding research on computer analysis of language. Why were the armed forces interested? Among other things, a "voice-writer," a computer that could take dictation, would be pretty handy for tapping phones.

and written language is pretty rudimentary, and one still often hears the prediction that a full solution is only five years off! So the computer people are evidently no better than linguists at figuring out the organization of language—they are just more optimistic.

A more romantic approach to the Paradox might go something like this: "Children are just so wonderfully open and unselfconscious about the world around them! Look! They can pick up language without thinking about it, while we poor adults are hobbled by our self-conscious hangups." Now while there may be a grain of truth in this, it has to be an oversimplification. Why can we think more clearly than children about simple things like income taxes and going to the dentist, but not about the organization of language? Along with our "self-conscious hangups" does come some sophistication, after all. So the Paradox remains: there is something special about language learning that isn't available to adults, and it still remains to be explained what mechanisms permit children to pull off the feat. Saying that it's wonderful and unconscious doesn't *explain* it, it just restates the problem. We still want to know how it works.

There are three steps involved in escaping the Paradox. The first two have already been touched on. First, as shown in Chapter 2, what the child ends up with is a mental grammar that is completely inaccessible to consciousness. Hence adult linguists can't figure out the principles of mental grammar just by looking into their minds. Second, as shown in the last section, a substantial part of the language-learning process is also unconscious, so linguists can neither directly observe it nor ask children about it.

But to escape the Paradox, a third step is needed. Remember: children can't just "absorb" mental grammar from the surroundings. All they can hear in the surroundings are sentences; they must (unconsciously) discover for themselves the patterns that permit them both to understand these sentences and to construct new sentences for other people to respond to. Whether this process of discovery goes on unconsciously in the child or consciously in the linguist, the very same problems have to be solved. That is, doing it unconsciously still gives the child no advantage over the linguist.

About the only way anyone has devised to overcome this difficulty is to suppose that *children have a head start on linguists*: children's unconscious strategies for language learning include some substantial hints about how a mental grammar ought to be constructed. These hints make it relatively easy for them to figure out principles that fit the examples of language they are hearing around them. (Only *relatively* easy: it still takes them eight or ten years!) By contrast, though, linguists have no such hints at their conscious disposal, so the problem is much harder for them. To invoke a crude metaphor, linguists are in the position of searching for a needle in a haystack, but children have a powerful magnet that pulls the needle right out.

For a more biological analogy, perhaps language acquisition is something like reproduction. Just about everybody figures out how to reproduce—it seems pretty easy! But it has nevertheless taken centuries of research to understand the actual mechanisms of reproduction, and we don't have a complete account even yet. Does that mean that biologists are inept? Of course not. We don't expect them to have conscious access to the biological mechanisms of reproduction. Somehow, though, because language is in the *mind*, we are more prone to thinking we should be able to understand it readily. Again I want to stress that we have to give up this preconception. We must realize that some parts of our minds are as distant from awareness as our chromosomes are.

Somewhat more technically, the claim is that all of us as children come to the task of language learning equipped with a body of innate knowledge pertaining to language. Using this knowledge, children can find patterns in the stream of language being beamed at them from the environment, and can use these patterns as a mental grammar. Because this innate knowledge must be sufficient to construct a mental grammar for any of the languages of the world, linguists call it *Universal Grammar* or *UG*.

Well, fine, but how is that different from the "romantic" solution? The difference is that it doesn't just revel in the mystery of the child's ability. Rather, it leads to three important questions for research:

1. What do children know (unconsciously) about language in advance of language learning? That is, what is Universal Grammar?
2. How do they use Universal Grammar to construct a mental grammar?
3. How do they acquire Universal Grammar?

I'll set questions 1 and 2 aside until subsequent chapters, when we have a better idea of what a mental grammar is. At the moment I want to think about question 3, the issue of how there could be such a thing as "innate knowledge"—knowledge that is not learned.

First I have to deal with a couple of simple matters. For one thing, we have to keep remembering that Universal Grammar is every

bit as unconscious and inaccessible to introspection as the final mental grammar the child achieves and we adults use. So we have to tolerate a certain degree of strangeness in the use of the term "knowledge."*

"Innate" is also used a little loosely, in that I am not necessarily committed to its presence immediately at birth. Like the teeth or body hair or walking, Universal Grammar could just as well develop at some considerable time after birth; what is important is that its development is conditioned by a biological timetable. In fact, children usually begin acquiring grammatical patterns sometime toward their second birthday (although, as will be mentioned in Chapter 8, there are earlier precursors).

However we describe it, though, the point is that Universal Grammar is not learned. Rather, it is the machinery that makes learning possible. So question 3 amounts to this: How can knowledge or cognitive organization be available to the child *before learning*?

Fortunately, the outlines of a mechanism behind innate knowledge are available. Two components are involved: the determination of brain structure by genetic information, and the determination of mental functioning by brain structure. Let me take these up briefly in turn.

First component: Until relatively recently, it was a major mystery how organisms reproduce their own kind—how it is that people give birth to little people and pigs to little pigs, but not the other way around. One of the major achievements of twentieth-century science is some understanding of the mechanisms that determine inheritance of the physical structure of organisms: genetic material, coded in the cells' DNA and passed on from generation to generation, determines the physical arrangement and functioning of the body. Although the precise steps by which the genetic material guides the development of the body are as yet mostly unknown, we have for the first time a way of describing the physical basis of reproduction, inherited characteristics, mutation, and evolution.

* The philosopher Gilbert Ryle has made a distinction between "knowing that" (for instance, knowing that Grant's wife is buried in Grant's Tomb) and "knowing how" (for instance, knowing how to swim). The latter might be called "operational knowledge" or "skill"—it is not necessarily verbalizable. And perhaps we can best think of the child's knowledge of *how* to learn language as like this. (There are some tricky caveats in this, though, because Ryle himself intended the term purely behaviorally: it's not clear he would have been willing to say that someone who has been paralyzed still "knows how to swim." I would.) In any event, for lack of a better term, I'll continue to use the term "knowledge," though it should be clear that I intend the term in the very special sense we have been working out here.

Among the parts of the body determined by the DNA is, of course, the brain. Its anatomical structure is highly complex—at least as complex as that of, say, the little finger. So, although there is some plasticity in the brain's physical organization, there is good reason to believe that substantial aspects of this organization are genetic. As Chomsky often puts it, we don't *learn* to have arms rather than wings. Why, then, should we suppose that our brains acquire their fundamental structure through learning rather than genetic inheritance?

Second component: The way we think is partly constrained by the way our brains are built. Hardly anyone disputes this: for instance, you are relying on this assumption when you claim that we are smarter than animals because we have bigger brains. Now the idea of innate knowledge of Universal Grammar can be rephrased, if you like, as saying that children have a certain "way of thinking" that enables them unconsciously to construct a mental grammar, given appropriate inputs in the surroundings. The hypothesis, then, is that this "way of thinking" is a consequence of the physical organization of some part of the brain—which is in turn determined by genetic structure. In short, *the mechanism for acquiring innate knowledge is genetic transmission, through the medium of brain structure.*

This hypothesis—let me call it the *Genetic Hypothesis*—leads us into a rich range of issues. For it says that the ability to learn language is rooted in our biology, a genetic characteristic of the human species, just like an opposable thumb and a pelvis adapted for upright stance. This means that we can draw freely on biological precedents in trying to explain language.

For example, think of all the surprising structural specializations in the organisms of the world—the elephant's trunk, the bat's sonar, or the little bones of our middle ear. Given such biological precedents, it hardly seems outlandish that there might be a structural specialization in the brain for language (and language learning).

Next consider the fact that the "innate knowledge of language" doesn't seem to be present at birth, but begins to manifest itself at around the age of two. According to the Genetic Hypothesis, this knowledge is determined by brain structure, so it is present only when the supporting brain structures are present. Now development of the physical structure of the body, including the neural structure of the brain, is by no means complete at birth. Among other things, in the newborn's brain the myelin sheaths that serve to electrically insulate the neurons from each other are not yet fully developed. Further, as suggested earlier, the physical growth of various parts of

the body follows a fairly predictable timetable: think of the developmental sequence of baby teeth followed by adult teeth, or the body changes of puberty, not to mention more dramatic developments in other species such as tadpoles turning into frogs. There is no reason why the development of the particular brain structures that support innate knowledge about language couldn't be like that. In other words, gradual development of innate knowledge over several years of life is very much in line with other developmental phenomena.

This is not to say that we should go out and start looking for the "gene for language." The connections between DNA sequences in the chromosomes and the body's structure are exceedingly indirect. Biologists are just beginning to figure out how the genes guide the differentiation of the embryo into head, middle, and tail, about the crudest of all body structures. When it comes to the exquisite differentiation of the brain (or heart, or ear, or wrist), we are very much in the dark. In addition, we know almost nothing about how brain structure governs the nature of thought. So both components of the Genetic Hypothesis leave a lot of questions at the moment. Still, the Genetic Hypothesis seems to be a plausible way—and maybe the *only* plausible way—of providing the child with innate knowledge. It can hardly be said to be *proven*, but all the pieces of it look reasonable.

Let's try to put this all together, tracing our argument back to our initial questions. (1) We arrived at the Genetic Hypothesis as a potential answer to the question of how there could be such a thing as innate knowledge. (2) Why did we need innate knowledge? We needed it in order to solve the Paradox of Language Acquisition —how it is that all children can unconsciously "pick up" a mental grammar on their own, while linguists as a community can't figure out completely how the mental grammar works. That is, the Paradox shows how difficult the task is that children accomplish. Innate knowledge of some aspects of language would give children a head start on learning the language spoken in the environment.

(3) Why does the Paradox of Language Acquisition arise? Because, as we showed, children mostly don't learn language by being taught. Rather, they must be unconsciously figuring out a mental grammar that gives them the patterns for forming sentences. (4) How do we know that children must be figuring out an unconscious mental grammar? Because that's what they have to end up with as adults, in order to account for their ability to speak and understand an unlimited range of sentences they have never heard before.

Questions about innate knowledge

For some reason, the hypothesis of a genetically determined Universal Grammar has provoked various degrees of astonishment, disbelief and outrage since the time it was proposed by Chomsky. Let me try to defuse some of the more common reactions, once again through a conversation with my imaginary skeptic.

The child's acquisition of language clearly depends on exposure to language in the environment. So why should we believe that it is genetically determined?

The answer is that one's language ability is a complex combination of nature and nurture. A biological comparison may be helpful here. Our bone structure is obviously genetically determined, but it can't develop properly without nourishment and exercise. In this case it's clear that environmental interaction complements genetic endowment: both are necessary. Why shouldn't the same be true of the brain structure that supports language, where "nourishment" includes a sufficient quantity and variety of incoming information, and "exercise" includes the opportunity to converse with people?

I don't mind the idea of a genetic component to learning, so long as none of it is specifically linguistic, that is, if it consists only of general-purpose learning strategies such as stimulus–response learning or principles of association or analogy. But why do you insist that there is a genetic component of learning that has to do specifically with language?

The problem is that general-purpose learning strategies alone can't solve the Paradox of Language Acquisition. Adults, including adult linguists, have access to plenty of general purpose strategies, but they can't figure out the organization of mental grammar. We have to suppose that children know something more, something specifically about *language*. Still, this is not to say that language acquisition doesn't make use of more general learning strategies. It is just that this cannot be all there is.

But how could a brain structure for Universal Grammar have originally come to be coded into the genes?

The only possible answer is evolution. Unfortunately, there isn't any record of the evolution of language: we can't dig up fossil vowels or verbs, and the earliest written documents already display the full expressive variety and grammatical complexity of modern languages. So the route by which language evolved is pretty mysterious. It is easy to see how having language would confer a selective advantage in the course of evolution, but presumably it didn't spring into existence full-blown. What are the steps on the way? There are not just one but many missing links.

On the other hand, evolution gives us an interesting angle on the Paradox of Language Acquisition. For it says that language acquisition doesn't just take ten years of the child's life. Those ten years are backed up by a couple of million years that evolution has spent developing in the brain the Universal Grammar that children start with—more time than linguists will ever have!

Isn't the Genetic Hypothesis just a "null hypothesis," a desperate move to explain away all this embarrassing complexity?

Remember: we're stuck between a rock and a hard place. On one hand, the expressive variety of language demands a complex mental grammar that linguists can't entirely figure out. But on the other hand, children manage to acquire this grammar. Thus, in a sense the Genetic Hypothesis *is* a move of desperation. As I said earlier, it's the only answer anybody has been able to think of; different schools of thought disagree mainly in exactly what and how much they think is innate.

Still, I don't think the Genetic Hypothesis is an attempt to explain the complexity *away*. One can imagine a similar criticism of the theory of gravitation: "The theory postulates an occult, invisible force; it just restates the facts of the interaction of physical bodies without explaining them." In fact, the Genetic Hypothesis plays much the same role in linguistics as the hypothesis of gravitation does in physics. It is a construct which, as we will see, serves to unify a large body of diverse facts from language structure, language universals, and language acquisition.

In turn, like the theory of gravitation, the Genetic Hypothesis calls for eventual deeper explanation. But remember, it is over three centuries since Newton postulated a gravitational force, and we don't yet have an entirely satisfactory theory of how gravity works. So I'm inclined to counsel patience.

Conclusions

Putting together all the considerations of the past sections, our proposed account of mental grammar takes the following overall form:

Mental Grammar = Innate Part (Universal Grammar)
+ Learned Part

In trying to understand the mental grammar of English (or Chinese, or whatever), linguists try always to find the simplest possible account, consistent with the complexity of the facts of the language. At the same time, rather than insisting that language is all learned (or all innate), we leave it as an empirical question to determine how the mental grammar is parceled out between innate and learned parts. Three basic criteria are involved.

1. If the language in question is different from other languages in some respect, the child must be able to acquire this difference, so it must fall into the learned part.
2. If certain aspects of all languages we have examined are alike, these aspects are *likely* to fall into the innate part. Of course, there is always the possibility that they are alike purely by accident. In practice, this can be checked out by examining more languages, preferably unrelated ones.
3. Suppose there is some aspect of language that children couldn't possibly figure out from the evidence in the speech they hear around them. Then this aspect can't be learned; it has to fall under the innate part of the language.

The last of these criteria has been called the "poverty of the stimulus" argument. Its use requires a certain amount of care, and in fact there is a running debate on what sorts of evidence children are capable of using. We have already encountered this debate in discussing the character of language learning; there is more to come in succeeding chapters.

We can go a step further and decompose the innate part of language like this:

Innate part of language = Part due to special purpose endowment for language
+ Part due to general properties of the mind

Again, rather than insisting that language is based entirely on general-purpose principles, or entirely on principles peculiar to language, we leave it to be decided by research how the work is divided up.

I sympathize with those who are suspicious of a specific language capacity: we *should* try to minimize the first factor. A special-purpose endowment for language, after all, demands an evolutionary jump during the time since we diverged from the apes, and we would like to think that the jump was not too extraordinary. But that doesn't mean we can eliminate it altogether: *something* has to account for the Paradox of Language Acquisition.

To close this chapter, let's recall our initial question: What does human nature have to be like in order for us to be able to use language? Two more answers, having to do with the nature of learning, have emerged from the Argument for Innate Knowledge.

First, the learning of language isn't just a passive "soaking up" of information from the environment. Rather, language learners actively construct unconscious principles that permit them to make sense of the information coming from the environment. These principles make it possible not just to reproduce the input parrotlike, but to use language in novel ways. What is learned comes as much from inside the learner as from the environment.

Second, we have spent considerable time chewing over the idea that certain aspects of our knowledge of language must be derived genetically, rather than through learning *per se*. We have concluded that the child's language ability comes from a combination of environmental influence, which is obvious, and heredity, which is far less so. The fact that language learning is supported by a genetic component is what makes the task possible for every normal child, despite the complexity of the resulting knowledge.

Is the learning of language just a curious exception in the story of human learning, or are other kinds of learning like this too? If they are, there are strong implications for one's approach to education: one should see the learner as an active agent of learning, not just a vessel to be filled with facts. Education should stress the learner's engagement and creativity, for ultimately the learner must construct the knowledge in his or her own mind.

Similarly, we can ask if other kinds of learning are, like language acquisition, supported by some sort of special-purpose innate endowment. If this is the way language is, what about all the other things we do? We will return to this question in Part IV, after working out our ideas about language itself more clearly. But in the meantime, it should remain lurking in the background, for this is ultimately the issue that makes the study of language absolutely crucial to understanding ourselves.