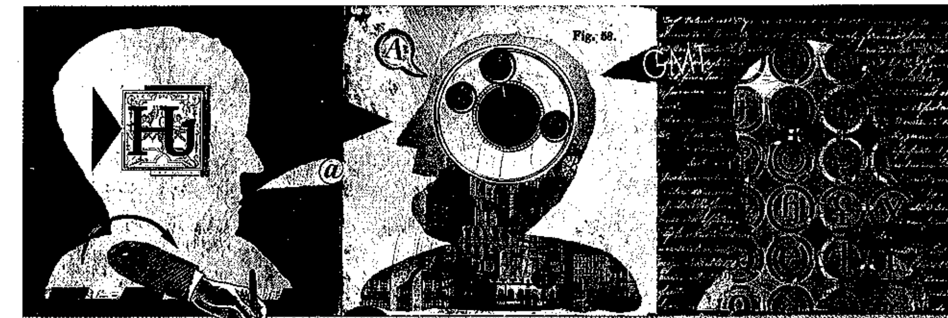

Language in Mind

An Introduction to Psycholinguistics

JULIE SEDIVY *University of Calgary*

SINAUER ASSOCIATES, INC. • Publishers • Sunderland, Massachusetts • USA

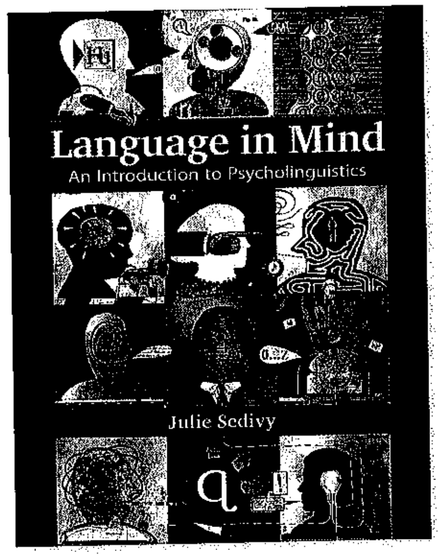

About the Cover Artist

Bruno Mallart is a talented European artist whose work has appeared in many of the world's premier publications, including *The New York Times*, *The Wall Street Journal*, and the *New Scientist*, to name a few. A freelance illustrator since 1986, Mallart first worked for several children's book publishers and advertising agencies using a classic, realistic watercolor and ink style. Later he began working in a more imaginative way, inventing a unique mix of drawing, painting, and collage. His work speaks of a surrealistic and absurd world that engages the viewer's imagination and sense of fun.

Despite the recurring use of the brain in his art, Mallart's background is not scientific (though both his parents were neurobiologists). He uses the brain as a symbol for abstract concepts such as intelligence, thinking, feeling, ideas, and, as seen here, human language and communication. His abstract representations beautifully illustrate the topics of *Language in Mind*; his work is seen not only on the book's cover, but in the distinctive chapter-opening images and as "icons" marking the book's various features. To see more of Bruno Mallart's art, please visit his website: www.brunomallart.com.

Language in Mind: An Introduction to Psycholinguistics

 Address inquiries and orders to:

Sinauer Associates, Inc., 23 Plumtree Road PO Box 407, Sunderland, MA 01375 U.S.A.

Inquiries: publish@sinauer.com

Orders: orders@sinauer.com

FAX: 413-549-1118

www.sinauer.com

Library of Congress Cataloging-in-Publication Data

Sedivy, Julie.
Language in mind : an introduction to psycholinguistics / Julie Sedivy -- First Edition.
pages cm
Includes index.
ISBN 978-0-87893-598-7
1. Psycholinguistics. 2. Cognition. I. Title.
BF455.S3134 2014
401'.9--dc23 2014002893

Printed in China

5 4 3 2 1

For My Students

1 Science, Language, and the Science of Language

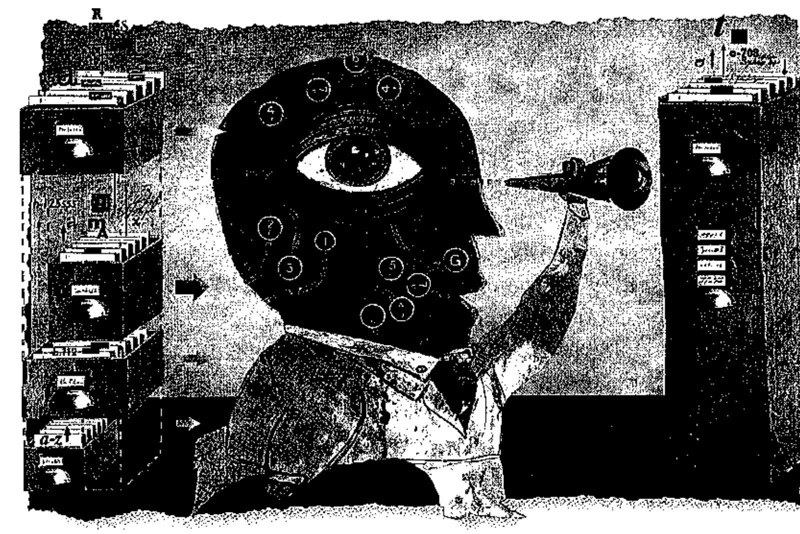

Before you read any further, stand up, hold this book at about waist height, and drop it. Just do it. (Well, if you're reading this on an electronic device, maybe you should reach for the nearest unbreakable object and drop it instead.)

Now that you've retrieved your book and found your place in it once more, your first assignment is to explain why it fell down when you dropped it. Sure, sure, it's gravity—Isaac Newton and falling apples, etc.

Your next assignment is to explain: How do you know it's gravity that makes things fall down? What's the evidence that makes you confident that gravity is better than other possible explanations—for example, you might think of the Earth as a kind of magnet that attracts objects of a wide range of materials. Chances are you find it much easier to produce the right answer than to explain why it's the right answer. It's possible, too, that throughout your scientific education you were more often evaluated on your ability to remember the right answer than on being able to recreate the scientific process that led people there. And I have to admit that there's a certain efficiency to this approach: there's a wheel, learn what it is, use it, don't reinvent it.

The trouble with the "learn it, use it" approach is that science hardly ever has "the" right answers. Science is full of ideas, some of which stand an extremely good chance of being right, and some of which are long shots but the best we've got at the moment. The status of these ideas shifts around a fair bit (which partly explains why textbooks have to be revised every couple of years). If you have a good sense of the body of evidence that backs up an idea (or can identify the gaps in the evidence), it becomes much easier to tell where a certain idea falls on the spectrum of likelihood that it's right.

This was a point made by scientist and author Isaac Asimov in his well-known essay "The Relativity of Wrong" (**Box 1.1**). In this 1988 essay, Asimov challenged an English student who wrote to accuse him of scientific arrogance. The letter-writer pointed out that, throughout history, scientists have believed that they understood the universe, only to be proven wrong later. Hence, concluded Asimov's correspondent, the only reliable thing one could say about scientific knowledge is that it's bound to be wrong.

To this, Asimov countered that what matters isn't knowing whether an idea is right or wrong, but having a sense of which ideas might be more wrong than others. He used the flat-Earth theory as an example of how scientific theories develop. In ancient times, the notion that the Earth was flat wasn't a stupid or illogical one—it was the idea that happened to be most consistent with the available body of knowledge. Eventually, people like Aristotle and others observed things that didn't quite mesh with the flat-Earth theory. They noticed that certain stars disappear from view when you travel north, and certain others disappear if you travel south. They saw that Earth's shadow during a lunar eclipse is always round, and that the sun casts shadows of different lengths at different latitudes. In short, the available body of evidence had expanded. The flat-Earth theory was no longer the best fit to the observations, causing it to be abandoned in favor of the notion that the Earth is a sphere. As it turned out, when even more evidence was considered, this theory too had to be abandoned: the Earth is *not* exactly a sphere, but an oblate spheroid, a sphere that's been squished toward the center at the North and South Poles.

As Asimov put it, "when people thought the Earth was flat, they were wrong. When people thought the Earth was spherical, they were wrong. But if *you* think that thinking that the Earth is spherical is *just as wrong* as thinking the Earth

BOX 1.1

Wrong or insightful?

Isaac Asimov on testing students' knowledge

"Young children learn spelling and arithmetic, for instance, and here we tumble into apparent absolutes.

How do you spell "sugar?" Answer: s-u-g-a-r. That is right. Anything else is wrong.

How much is 2 + 2? The answer is 4. That is right. Anything else is wrong.

Having exact answers, and having absolute rights and wrongs, minimizes the necessity of thinking, and that pleases both students and teachers. For that reason, students and teachers alike prefer short-answer tests to essay tests; multiple-choice over blank short-answer tests; and true-false tests over multiple-choice.

But short-answer tests are, to my way of thinking, useless as a measure of the student's understanding of a subject. They are merely a test of the efficiency of his ability to memorize.

You can see what I mean as soon as you admit that right and wrong are relative.

How do you spell "sugar?" Suppose Alice spells it p-q-z-z-f and Genevieve spells it s-h-u-g-e-r. Both are wrong, but is there any doubt that Alice is wronger than Genevieve? For that matter, I think it is possible to argue that Genevieve's spelling is superior to the "right" one.

Or suppose you spell "sugar": s-u-c-r-o-s-e, or $C_{12}H_{22}O_{11}$. Strictly speaking, you are wrong each time, but you're displaying a certain knowledge of the subject beyond conventional spelling.

Suppose then the test question was: how many different ways can you spell "sugar?" Justify each.

Naturally, the student would have to do a lot of thinking and, in the end, exhibit how much or how little he knows. The teacher would also have to do a lot of thinking in the attempt to evaluate how much or how little the student knows. Both, I imagine, would be outraged."

From Isaac Asimov (1988). The relativity of wrong. In *The relativity of wrong: Essays on science*. New York: Doubleday. Used with permission.

is flat, then your view is wronger than both of them put together." Without the distinction that one is more wrong than the other, for example, you could be left with the belief that, for all we know, in 50 years, scientists will "discover" that the oblate spheroid theory was wrong after all, and the Earth is cubical, or in the shape of a doughnut. (In actual fact, the oblate spheroid theory *is* wrong: the Earth is very, very slightly pear-shaped, with the South Pole being squished toward the center just a bit more than the North Pole. Still, not a cube.)

Asimov's point about scientific progression and the graded "rightness" of ideas seems fairly obvious in the context of a well-known example like the flat-Earth theory. But unfortunately, the way in which people often talk about science can blot out the subtleties inherent in the scientific process. In many important discussions, people *do* behave as if they think of scientific ideas as being right or wrong in an absolute sense. For example, you've probably heard people express frustration upon reading a study that contradicts earlier health tips they've heard; a common reaction to this frustration is to vow to ignore *any* advice based on scientific studies, on the grounds that scientists are constantly "changing their minds." And when people talk about evolution as "just a theory" (and hence not something we need to "believe"), or object that the science of climate change "isn't settled," they're also failing to think about the *degree* to which these scientific ideas approach "rightness." Naturally, being able to identify whether an idea is very likely to be wrong or very likely to be right calls for a much more sophisticated body of scientific knowledge than simply having memorized what the supposedly right answer is. But ultimately, the ability to evaluate the rightness of an idea leaves you with a great deal more power than does merely accepting an idea's rightness.

I have a 3-year-old niece who is definitely on to something. Like many preschoolers, her usual response to things you tell her is to ask a question. But in her case, the question is almost always the same: whether you've told her that eating her carrots will make her healthy or that the sun is many, many miles away, she demands, "How do you know?" She makes a great scientific companion—in her presence, you can't help but realize where it is that your understanding of the world is at its shallowest. (Conversations with her have a way of sending me off on an extended Google search.) One can only hope that by the time she hits college or university, she hasn't abandoned that question in favor of another one, commonly heard from students: Which theory is the right one?

1.1 What Do Scientists Know about Language?

In studying the language sciences, it's especially useful to approach the field with the "how do you know?" mindset rather than one that asks which theory is right. The field is an exceptionally young one, and the truth is that its collection of facts and conclusions that can be taken to be nearly unshakable is really very small. (The same is also true of most of the sciences of the mind and brain in general.) In fact, scientific disagreements can run so deep that language researchers are often at odds about fundamentals—not only might they disagree on which theory best fits the evidence, they may argue about what kind of cloth a theory should be cut from. Or on very basic aspects of how evidence should be gathered in the first place. Or even the *range* of evidence that a particular theory should be responsible for covering. It's a little bit as if we were still in an age when no one really knew what made books or rocks fall to the ground—when gravity was a new and exciting idea, but was only one among many. It needed to be tested against other theories, and we were still trying to figure out what the best techniques might be to gather data that would decide among the competing explanations. New experimental methods and new theoretical approaches crop up every year.

All this means that language science is at a fairly unstable point in its brief history, and that seismic shifts in ideas regularly reshape its intellectual landscape. But this is what makes the field so alluring to many of the researchers in it—the potential to play a key role in reshaping how people think scientifically about language is very, very real. A sizable amount of what we "know" about language stands a very good chance of being wrong. Many of the findings and conclusions in this book may well be overturned within a few years (so you might make a habit of visiting the book's companion website to check for some important updates). Don't be surprised if at some point, your instructor bursts out in vehement disagreement with some of the material presented here, or with the way in which I've framed an idea. In an intellectual climate like this, it becomes all the more important to take a "how do you know?" stance. Getting in the habit of asking this question will give you a much better sense of which ideas are likely to endure, as well as how to think about new ideas that pop up in the landscape.

The question also brings you into the heart of some of the most fascinating aspects of the scientific process. Scientific truths don't lie around in the open, waiting for researchers to stub their toes on them. Often the path from evidence to explanation is excruciatingly indirect, requiring a circuitous chain of assumptions. Sometimes it calls for precise and technologically sophisticated methods of measurement. This is why wrong ideas often persist for long periods of time, and it's also why scientists can expend endless amounts of energy in arguing about whether a certain method is valid or appropriate, or what exactly can and can't be concluded from that method.

In language research, many of the *Eureka!* moments represent not discoveries, but useful insights into how to begin answering a certain question. Language is a peculiar subject matter. The study of chemistry or physics, for example, is about phenomena that have an independent existence outside of ourselves. But language is an object that springs from our very own minds. We can have conscious thoughts about how we use or learn language, and this can give us the illusion that the best way to understand language is through these deliberate observations. But how do you intuit your way to answering questions like these:

- When we understand language, are we using the same kind of thinking as we do when we listen to music or solve mathematical equations?
- Is your understanding of the word *blue* exactly the same as my understanding of it?
- What does a baby know about language before it can speak?
- Why is it that sometimes, in the process of retrieving a word from memory, you can draw a complete blank, only to have the word pop into your mind half an hour later?
- What does it mean if you accidentally call your current partner by the name of your former one? (You and your partner might disagree on what this means.)
- What exactly makes some sentences in this book confusing while others are easy to understand?

To get at the answers to any of these questions, you have to be able to probe beneath conscious intuition. This requires acrobatic feats of imagination, not only in imagining possible alternative explanations, but also in devising ways to go about testing them. In this book, I've tried to put the spotlight not just on

the conclusions that language researches have drawn, but also on the methods they've used to get there. As in all sciences, methods range from crude to clever to stunningly elegant, and to pass by them with just a cursory glance would be to miss some of the greatest challenges and pleasures of the study of language.

psycholinguistics The psychology of language; the study of the psychological and neurobiological factors involved in the perception, production, and acquisition of language.

1.2 Why Bother?

At this point, you might be thinking, "Fine, if so little is truly known about how language works in the mind, sign me up for some other course, and I'll check back when language researchers have things worked out a bit better." But before you go, let me suggest a couple of reasons why it might be worth your while to study **psycholinguistics**, such as it is.

Here's one reason: Despite the fact that much of the current scientific knowledge of language is riddled with degrees of uncertainty and could well turn out to be wrong, it's not nearly as likely to be wrong as the many pronouncements that people often make about language without *really* knowing much, if anything, about it (see **Table 1.1**). The very fact that we can have intuitions about language—never mind that many of these are easily contradicted by closer, more systematic observation—appears to mislead people into believing that these intuitions are scientific truths. Aside from those who *have* formally studied the language sciences, or have spent a great deal of time thinking analytically about language, almost no one knows the basics of how language works, or has the slightest idea what might be involved in learning and using it. It's simply not something that is part of our collective common knowledge at this point in time.

TABLE 1.1 Some things people say about language (that are almost certainly wrong)

You can learn language by watching television.
People whose language has no word for a concept have trouble thinking about that concept.
English is the hardest language to learn.
Texting is making kids illiterate.
Some languages are more logical/expressive/romantic than others.
People speak in foreign accents because their mouth muscles aren't used to making the right sounds.
Some languages are spoken more quickly than others.
Saying *um* or *er* is a sign of an inarticulate speaker.
Failure to enunciate all your speech sounds is due to laziness.
Sentences written in the passive voice are a sign of poor writing.
Swearing profusely is a sign of a poor vocabulary.
Deaf people should learn to speak and lip-read in spoken language before they learn sign language, or it will interfere with learning a real language.
Speech errors reveal your innermost thoughts.
You can't learn language by watching television.

Try this: ask your mother, or your brother, or your boyfriend, or your girlfriend, "How can you understand what I'm saying right now?" Many people happily go through their entire lives without ever asking or answering a question like this, but if pressed, they might answer something like, "I recognize the words you're using." Fine, but how do they even know *which* bunches of the sounds you're emitting go together to form words, since there are no silences between words? And once they've figured that out, *how* do they recognize the words? What do "word memories" look like, and is there a certain order in which people sort through their mental dictionaries to find a match to the sounds you're emitting? Moreover, understanding language involves more than just recognizing the words, or people would have no trouble with the phrase *words I you're the using recognize*. Obviously, they're responding to the right order of words as well. So, what *is* the right order of words in a sentence—not just for this one, but more generally? How do people know what the right order is, and how did they learn to tell whether a sentence they've never heard before in their lives has its words strung together in the "proper" order?

Most people have a decent sense of how they digest their food, but the common knowledge of many educated people today does not contain the right equipment to begin answering a question as basic as, "How can you understand what I'm saying?" Imagine how it must have been for people, before awareness of gravity became part of common knowledge, for people to be asked, "Why does a rock fall to the ground?" A typical answer might have been, "It just does." Most people probably never thought to ask themselves why. Many might have stared and stammered—much as they do now when asked about how language works. So, by studying the psychology of language, you're entering a world of new questions and of new ways of thinking that isn't visible to most people. You'll be privy to discussions of ideas before they've become the officially received "right answers" that everyone knows. You might find this all so stimulating that you eventually wind up being a language researcher yourself. But the vast majority of readers of this textbook won't. Which brings me to the second reason to study psycholinguistics.

There are few subjects you can study that will have such a broad and deep impact on your daily life as the study of language. While you're unlikely to ever become a professional language researcher, you're *extremely* likely to use language in your daily life. Inevitably, you'll find yourself asking questions like these:

- How can I write this report so that it's easier to understand?
- What kind of language should I use in order to be persuasive?
- If I sit my kid in front of the TV for an hour a day, will this help her to learn language?
- Why do my students seem incapable of using apostrophes correctly?
- How can I make my poem more interesting?
- Should I bother trying to learn a second language in my thirties?
- Why is this automated voice system so infuriating?

Even a basic understanding of how language works in the mind will provide you with the tools to approach these and many, many other questions in an intelligent way.

Unfortunately, those of us who are deeply immersed in studying language don't always take the time to talk about how our accumulated body of knowledge might provide a payoff for daily users of language. This would be a poor textbook indeed if it didn't help you answer the questions about language that will crop up throughout your life. Ultimately, whether or not you become a professional psycholinguist, you should feel well equipped to be an amateur language scientist. And to do that, you need much more than "answers" to questions that researchers have thought to ask. One of the goals of this book is to give you the conceptual framework to address the questions that *you* will think to ask.

Throughout this book, you'll find many activities and exercises that are designed to immerse you in the scientific process of understanding language. The more deeply you engage in these, the more you'll internalize a way of thinking about language that will be very useful to you when faced with new questions and evidence. And throughout the book, you'll find discussions that link what can sometimes be very abstract ideas about language to real linguistic phenomena out in the world. Many more such connections can be made, and the more you learn about how language works, the more you'll be able to generate new insights and questions about the language you see and hear all around you.

GO TO
sites.sinauer.com/languageinmind
for **web activities, further readings, research updates, new essays,** and other features

2 Origins of Human Language

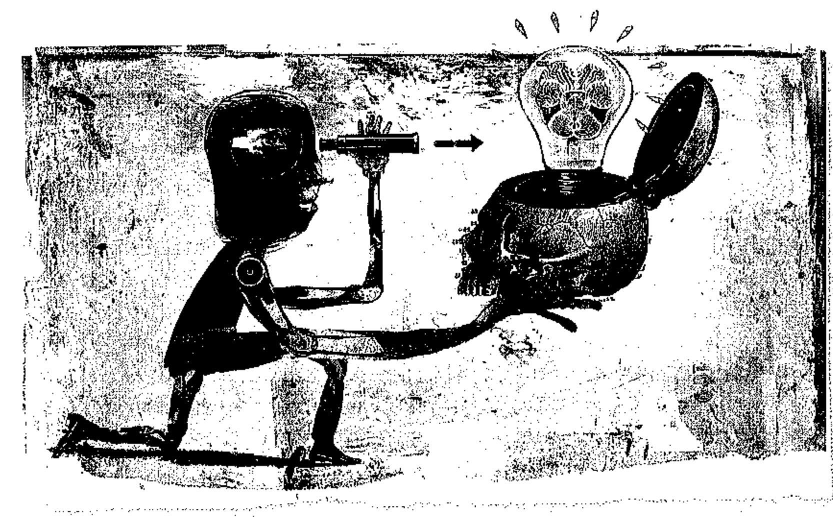

As far as we know, no other species on Earth has language; only humans talk. Sure, many animals communicate with each other in subtle and intricate ways. But we're the only ones who gossip, take seminars, interview celebrities, convene board meetings, recite poems, negotiate treaties, conduct marriage ceremonies, hold criminal trials—all activities where just about the only thing going on is talking.

Fine, we also do many other things that our fellow Earth-creatures don't. We play chess and soccer, sing the blues, go paragliding, design bridges, paint portraits, drive cars, and plant gardens, to name just a few. What makes language so special? Here's the thing: language is deeply distinct from these other activities for the simple reason that *all humans do it*. There is no known society of *Homo sapiens*, past or present, in which people don't talk to each other, though there are many societies where no one plays chess or designs bridges. And all individuals *within* any given human society talk, though again, many people don't play chess or design bridges, for reasons of choice or aptitude.

So, language is one of the few things about us that appears to be a true defining trait of what it means to be human—so much so that it seems it *must* be part of our very DNA. In fact, language has often been described as an innate instinct, something that we are inherently programmed to do. This **nativist view** is presented in Steven Pinker's book *The Language Instinct* (1994). In its strongest version, the nativist position says that not only do our genes program us to have a capacity for language, we're genetically programmed for the thing itself—its general structures, the building blocks that go into it, the mental process of acquiring it, and so on. One way to express this way of thinking is to say that as children, we don't so much *learn* language (the way we learn chess or

nativist view The view that not only are humans genetically programmed to have a general capacity for language, particular aspects of language ability are also genetically specified.

anti-nativist view The view that the ability of humans to learn language is not the result of a genetically programmed "language template," but is an aspect (or by-product) of our extensive cognitive abilities, including general abilities of learning and memory.

piano-playing) as *grow* language based on a genetic blueprint, much as birds grow wings, elephants grow trunks, and female humans grow breasts. This view of language as a genetically specified mental "organ," or as a preprogrammed instinct, captures why it is that language is not only common to all humans but also is unique to humans—no "language genes," no talking.

But many language researchers see it differently. The **anti-nativist view** is that language is not a specialized "organ," but a magnificent by-product of our impressive cognitive abilities. Humans alone learn language, not because we inherit a preprogrammed language template, but because we are the superlearners of the animal kingdom. What separates us from other animals is that our brains have evolved to become the equivalent of swift, powerful supercomputers in comparison with our fellow creatures, who are stuck with more rudimentary technology. Current computers can do qualitatively different things that older models could never aspire to accomplish. This supercomputer theory is one explanation for why we have language while squirrels and chimpanzees don't.

But what about the fact that language is universal among humans, unlike chess or trombone-playing (accomplishments which, though uniquely human, are hardly universal)? Daniel Everett, a linguist who takes a firm anti-nativist position, puts it this way in his book *Language: The Cultural Tool* (2012): Maybe language is more like a tool invented by human beings than an organ that's genetically programmed to grow. What makes language universal is that it's an incredibly *useful* tool for solving certain problems that all humans have—foremost among them being how to efficiently transmit information to each other. Everett compares language to arrows. Arrows are nearly universal among hunter-gatherer societies, but few people would say that humans are genetically programmed to make arrows *specifically*. More likely, making arrows is just part of our general tool-making, problem-solving competence. Bows and arrows can be found in so many different societies because at some point, people who didn't grow their own protein had to figure out a way to catch protein that ran faster than they did. Since it was well within the bounds of human intelligence to solve this problem, humans inevitably did—just as, Everett argues, humans inevitably came to speak with each other as a way of achieving certain pressing goals.

The question of how we came to have language is a huge and fascinating one. If you're hoping that the mystery will be solved by the end of this chapter, you'll be sorely disappointed. It's a question that has no agreed-upon answer among language scientists and, as you'll see, there's a range of subtle and complex views among scientists beyond the two extreme positions I've just presented.

In truth, the various fields that make up the language sciences are not yet even in a position to be able to resolve the debate. To get there, we first need to answer questions like: What *is* language? What do all human languages have in common? What's involved in learning it? What physical and mental machinery is needed to successfully speak, be understood, and understand someone else who's speaking? What's the role of genes in shaping any of the above behaviors? Without doing a lot of detailed legwork to get a handle on all of these smaller pieces of the puzzle, any attempts to answer the larger question about the origins of language can only amount to something like a happy hour discussion—heated and entertaining, but ultimately not that convincing one way or the other. In fact, in 1866, the Linguistic Society of Paris decreed that no papers about the origins of language were allowed to be presented at its conferences. It might seem ludicrous that an academic society would banish an entire topic from discussion. But the decision was essentially a way of saying, "We'll get nowhere talking about language origins until we learn more about language itself, so go learn something about language."

A hundred and fifty years later, we now know quite a bit more about language, and by the end of this book, you'll have a sense of the broad outlines of this body of knowledge. For now, we're in a position to lay out at least a bit of what might be involved in answering the question of why people speak.

~~2.1 Why Us?~~

The "language" of bees

Let's start by asking what it is about our language use that's different from what animals do when they communicate. Is it different down to its fundamental core, or is it just a more sophisticated version of what animals are capable of? An interesting starting point might be the "dance language" of honeybees, as identified by Karl von Frisch (1967).

When a worker bee finds a good source of flower nectar at some distance from her hive, she returns home to communicate its whereabouts to her fellow workers by performing a patterned waggle dance (see **Figure 2.1**). During this dance, she repetitively traces a specific path while shaking her body. The elements of this dance communicate at least three things:

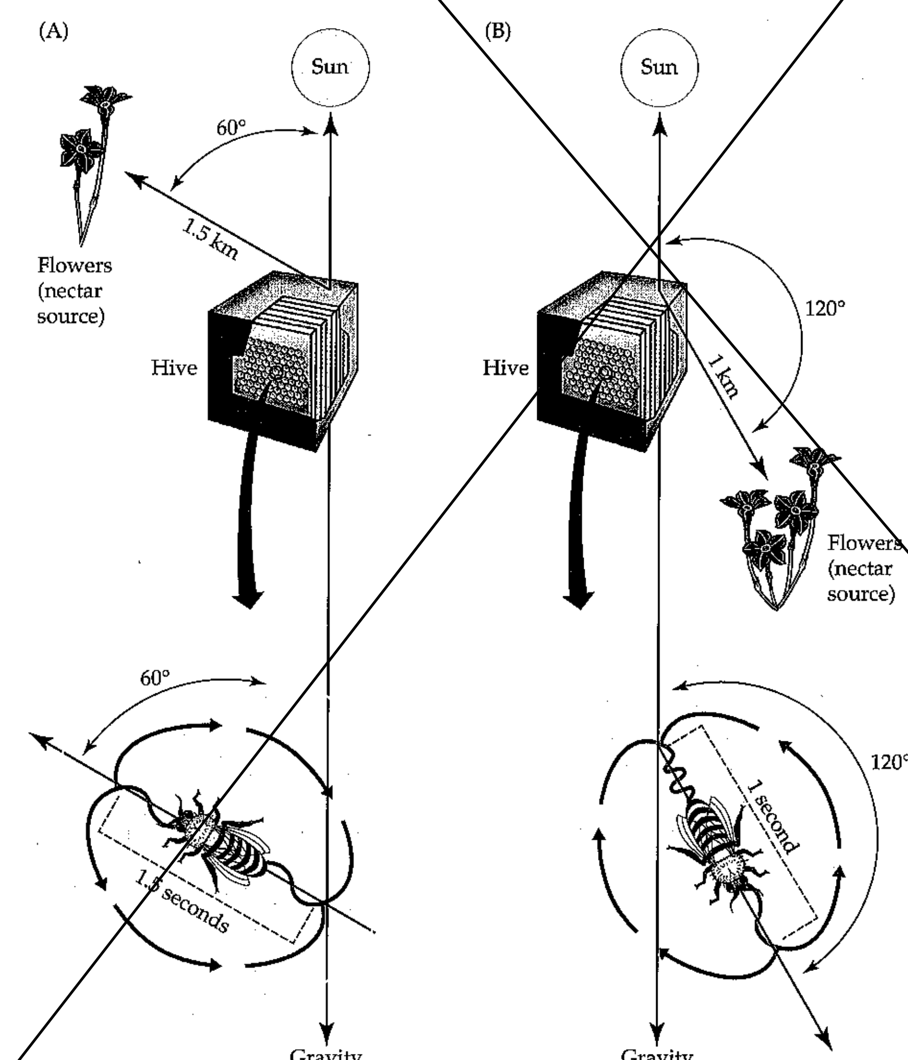

Figure 2.1 The waggle dance of honeybees is used by a returning worker bee to communicate the location and quality of a food source. The worker dances on the surface of the comb to convey information about the direction and distance of the food source, as shown in the examples here. (A) The nectar source is approximately 1.5 km from the hive flying at the indicated angle to the sun. (B) The nectar source is closer and the dance is shorter; in this case the flowers will be found by flying away from the sun. The energy in the bee's waggles (orange curves along the line of the dance) is in proportion to the perceived quality of the find.

METHOD 2.2

Exploring what primates can't (or won't) do

Discussions about the evolution of language rely heavily on comparative studies, in which the abilities of different species are evaluated against each other in targeted experiments. These comparisons usually involve designing a test that seems appropriate for all of the species being studied, then applying the same test to members of all the species. The performance of the test subjects is then presumed to reveal something about their cognitive abilities.

But, as any schoolteacher knows, test results don't always tell the whole story when it comes to assessing a student's abilities or knowledge. Some students are highly intelligent but simply don't care, and so their test results are mediocre. Other students become anxious in test situations and freeze up, unable to fully access their knowledge. Test results for all of these students will underestimate what they know or are capable of doing. It's worth keeping this in mind when it comes to using performance on experimental tests as a way to gauge the abilities of various species.

Suppose we're trying to compare a human 2-year-old to an adult chimpanzee. Can we assume that the same test given to both will be an equally good measure of the targeted skills of each? Not necessarily. For instance, the very fact that a *human* experimenter is in charge might have very different effects on the two subjects. The child might be more motivated than the chimp to perform well for the human. The chimp might be more anxious than the child in the presence of a human.

Eliminating these potential problems isn't always possible—after all, we can't readily train a chimp to administer an experiment. But some potential problems can be avoided by paying close attention to the design of the experiments. Specifically, in testing to see whether primates have a particular skill, it makes sense to ask in what situations that skill would be useful to them in the wild. These are the situations in which we'd expect that they'd be motivated to use the relevant skill, and be less disoriented by the experimental setup. This strategy has led researchers to conclude that some past experiments have underestimated the social cognition of chimpanzees.

For example, in one study (Povinelli & Eddy, 1996), chimpanzees appeared to be unable to imagine the visual experience of another individual. The test involved putting a chimpanzee subject in the presence of two humans with food placed between them. One of the humans had a bucket over his head, and hence couldn't see the chimp. The experimenters were interested in learning whether the chimps would choose to beg for food from the human who could see them. Interestingly, their study subjects begged randomly from either human, as if oblivious to the fact that the person wearing the bucket was temporarily blind.

But a number of researchers have argued that, in the wild, a chimpanzee's social understanding is most useful in competitive situations than cooperative ones. A later study (Bräuer et al., 2007) built its experimental design around this premise. In the competitive setup, a lower-ranking chimpanzee competed with a higher-ranking chimpanzee for food. Two pieces of food were present, one of which could be seen by both chimps but the other of which was hidden from the higher-ranking chimp by a barrier. More often than not, the subordinate chimp decided to go for the food that the dominant chimp couldn't see.

The moral of the story seems to be, if you want to find out what a chimp is thinking, you have to think like a chimp as you set up your study.

2.3 The Structure of Language

Combining units

Being able to settle on arbitrary symbols as stand-ins for meaning is just one part of the language puzzle. There's another very important aspect to language, and it's all about the art of combination.

Combining smaller elements to make larger linguistic units takes place at two levels. The first level deals with making words from sounds. In principle, we *could* choose to communicate with each other by creating completely different sounds as symbols for our intended meanings—a high-pitched yowl might mean "arm," a guttural purr might mean "broccoli," a yodel might mean "smile," and so on. In fact, this is very much how vervets use sound for their alarm calls. But at some point, we might find we'd exhausted our ability to invent new sounds but still had meanings we wanted to express. To get around this limitation, we can take a different approach to making words: simply recruit a relatively small number of sounds, and combine them in new and interesting ways. For example, if we take just ten different sounds to create words made up of five sounds each without repeating any of the sounds within a word, we can end up with a collection of more than 30,000 words. This nifty trick illustrates Hockett's notion of *duality of patterning*, in which a small number of units that don't convey meanings on their own can be used to create a very large number of meaningful symbols. In spoken language, we can take a few meaningless units like sounds (notice, for example, that there are no specific meanings associated with the sounds of the letters *p*, *a*, and *t*) and combine them into a number of meaningful words. Needless to say, it's a sensible approach if you're trying to build a beefy vocabulary.

But the combinatorial tricks don't end there. There may be times when we'd like to communicate something more complex than just the concepts of leopard or broccoli. We may want to convey, for instance, that Johnny kicked Freddy in the shin really hard, or that Simon promised to bake Jennifer a cake for her birthday. What are the options open to us? Well, we could invent a different word for every complex idea, so a sequence of sounds like *beflo* would communicate the first of these complex ideas, and another—say, *gromi*—would communicate the second. But that means that we'd need a truly enormous vocabulary, essentially containing a separate word for each idea we might ever want to communicate. Learning such a vocabulary would be difficult or impossible—at some point in the learning process, everyone using the language would need to have the opportunity to figure out what *beflo* and *gromi* meant. This means that as a language learner, you'd have to find yourself in situations (probably more than once) in which it was clear that the speaker wanted to communicate the specific complex idea that *gromi* was supposed to encode. If such a situation happened to *never* arise, you'd be out of luck as far as learning that particular word.

A more efficient solution would be to combine *meaningful* elements (such as separate words) to make other, larger meaningful elements. Even better would be to combine them in such a way that the meaning of the complex idea could be easily understood from the way in which the words are combined—so that rather than simply tossing together words like *Jennifer, birthday, promised, bake, Simon,* and *cake,* and leaving the hearer to figure out how they relate to each other, it would be good to have some structured way of assembling sentences out of their component parts that would make their meanings clear from their structure. This added element of predictability of meaning requires a **syntax**—or a set of "rules"—about how to combine meaningful units together in systematic ways so that their meanings can be transparent. (For example, once we have a syntax in place, we can easily differentiate between the meanings of *Simon promised Jennifer to bake a birthday cake* and *Jennifer promised Simon to bake a birthday cake.*) Once we've added this second level into our communication system, not only have we removed the need to learn and memorize separate words for complex ideas, but we've also introduced the possibility of combining the existing elements of our language to talk about ideas that have never before been expressed by anyone.

Structured patterns

It should be obvious that the possibility of combining elements in these two ways gives language an enormous amount of expressive power. But it also has

syntax In a given language, the set of "rules" that specify how linguistic elements are put together so that their meaning can be clearly understood.

some interesting consequences. Now, anyone learning a language has to be able to learn its underlying structural patterns. And since, for the most part, human children don't seem to learn their native language by having their parents or teachers explicitly teach them the rules of language (the way, for example, they learn the rules of arithmetic), they have to somehow intuit the structures on their own, simply by hearing many examples of different sentences. You might think of language learning as being a lot like the process of reverse engineering a computer program: Suppose you wanted to replicate some software, but you didn't have access to the code. You could try to deduce what the underlying code looked like by analyzing how the program behaved under different conditions. Needless to say, the more complicated the program, the more time you'd need to spend testing what it did.

In case you're tempted to think that language is a fairly simple program, I invite you to spend a few hours trying to characterize the structure of your own native tongue (and see **Box 2.3**). The syntactic structures of human languages involve a lot more than just basic word orders. Once you start looking up close, the rules of language require some extremely subtle and detailed knowledge. For example, how come you can say:

Who did the hired assassin kill the mayor for?

meaning "Who wanted the mayor dead?" But you can't say:

Who did the hired assassin kill the mayor and?

intending to mean "Aside from the mayor, who else did the hired assassin kill?" Or, consider these two sentences:

Naheed is eager to please.
Naheed is easy to please.

These two sentences look almost identical, so why does the first involve Naheed pleasing someone else, while the second involves someone else pleasing Naheed? Or, how is it that sometimes the same sentence can have two very different meanings? As in:

Smoking is more dangerous for women than men.

meaning either that smoking is more dangerous for women than it is for men, or that smoking is even more hazardous than men are for women.

It's not only in the area of syntax that kids have to acquire specific knowledge about how units can be put together. This is the case for sound combinations as well. Languages don't allow sounds to be combined in just any sequence whatsoever. There are constraints. For instance, take the sounds that we normally associate with the letters *r, p, m, s, t,* and *o*. If there were no restrictions on sequencing, these could be combined in ways such as *mprots, stromp, spormt, tromps, rpmsto, tormps, torpsm, ospmtr,* and many others. But not all of these "sound" equally good as words, and if you and your classmates were to rank them from best- to worst-sounding, the list would be far from random.

Here's another bit of knowledge that English speakers have somehow picked up: even though they're represented by the same letter, the last sound in *fats* is different from the last sound in *fads* (the latter actually sounds like the way we usually say the letter *z*). This is part of a general pattern in English, a pattern that's clearly been internalized by its learners: if I were to ask any adult native speaker of English exactly how to pronounce the newly invented word *gebs*, it's almost certain that I'd get the "z" sound rather than the "s" sound.

WEB ACTIVITY 2.2

Finding structure in language In this activity, you'll put yourself in the shoes of a language learner, and you'll try to find some of the patterns inherent in the structure of English and other languages.

BOX 2.3
The recursive power of syntax

Among the most impressive linguistic accomplishments of human-reared apes is their ability to combine signs in new ways (often producing them in some regular order) and their ability to distinguish the meanings of sentences based solely on how the words are combined. We might conclude from this that apes "get" the idea of syntactic structure, and that the only difference between apes and humans when it comes to combinatorial structure is that we can produce and understand longer, more complicated sentences than they can. But we'd be leaping to unwarranted conclusions. In Method 2.1, you were cautioned against the temptation to overinterpret the data, and to consider the possibility that two species behaving in similar ways at times doesn't mean that the behaviors are driven by the same underlying cognitive apparatus. This applies to evaluating apes' syntactic abilities just as well as it does to their symbolic powers.

Let's suppose that a chimp regularly produces neatly ordered sequences of signs that combine action signs with person/chimp signs. For example:

Washoe tickle. Lana eat. Sue give. Dog growl.

These look a lot like perfectly well-behaved (if minimalist) sentences of English that even a linguistics professor might sometimes utter:

John stinks. Bianca jogs. Students complain.

But this doesn't mean that the system of rules that has generated these sentences looks the same for the chimp and the linguistics professor. All of these sentences *could* be generated by a rule that says:

Symbol for person/chimp/animal doing the action + action

But let's suppose that the professor, in addition to making simple sentences like the ones above, is also capable of producing sentences like these:

John's brother stinks.

Bianca and Leila jog.

Students who are too lazy to do the readings complain.

In fact, our professor can probably even say things like:

John's brother's dog stinks.

Bianca and the girl who lives next door to John's brother jog.

Students who are too lazy to do the readings but want to get an A anyway complain.

Clearly, these sentences can't be produced by the simple rule that is perfectly sufficient to capture the chimp's combinations. It's not just a matter of stringing more words together, or even of having a larger number of rules that allow words to be strung together in a greater variety of ways. It's that the rules have to be different *in kind* from the one I proposed as sufficient for describing the chimp's combinations. For the professor, the linguistic units that get combined with the action words can't just be symbols corresponding to people/chimps/animals, etc. They have to be much more abstract objects—like, whatever it is that you might call the grouping *Bianca and the girl who lives next door to John's brother* or *students who are too lazy to do the readings but want to get an A anyway.* These abstract objects have to in turn be composed of other abstract objects such as *the girl who lives next door to John's brother*—which, by the way, also has to be an abstract object, composed of combinations such as *John's brother.* And so on. Rules that operate at this level of abstraction are called **recursive** because they allow linguistic objects of the same kind (for example, *John's brother*) to be nested within other linguistic objects of the same kind (*the girl who lives next door to John's brother*)—notice, for example, that either of these phrases by itself can be combined with an action word, which is a hint that they are the same kind of phrase.

Rules that work to combine not just words, but higher-order groupings of words, take the language to a whole new level. Whereas word-combining rules allow language users to put together existing words in ways they haven't seen before, rules that allow recursion literally allow users to produce an *infinite* array of meanings because small, abstract units can always be combined with others to make larger abstract units, which can always be combined with others to make larger abstract units, which can always be combined... you get the picture.

Although apes can manage a small number of word-combining rules, recursive rules appear to be beyond their cognitive reach, leading a number of researchers (e.g., Hauser, Chomsky & Fitch, 2002) to suggest that the property of recursion is part of a uniquely human set of skills for language.

recursion "Nesting" of related clauses or other linguistic units within each other.

So there's structure inherent at the level of sound as well as syntax, and all of this has to somehow be learned by new speakers of a language. In many cases, it's hard to imagine exactly *how* a child might learn it without being taught—and with such efficiency to boot. For many scholars of the nativist persuasion, a child's almost miraculous acquisition of language is one of the reasons to suspect that the whole learning process must be guided by some innate knowledge.

One of the leading nativists, Noam Chomsky (1986), has suggested that the problem of learning language structure is similar to a demonstration found in one of Plato's classic dialogues. (The ancient Greek philosopher Plato used "dialogues" written as scenarios in which his teacher Socrates verbally dueled with others as a way of expounding on various philosophical ideas.) In the relevant dialogue, Socrates is arguing that knowledge can't be taught, but only *recollected*, reflecting his belief that a person's soul has existed prior to the current lifetime and arrives in the current life with all of its preexisting knowledge. To demonstrate this, Socrates asks an uneducated slave boy a series of questions that reveals the boy's knowledge of the Pythagorean theorem, despite the fact that the boy could not possibly have been taught it.

Chomsky applied the term "Plato's problem" to any situation in which there's an apparent gap between experience and knowledge, and suggested that language was such a case. Children seem to know many things about language that they've never been taught—for instance, while a parent might utter a sentence like "What did you hit Billy for?" she's unlikely to continue by pointing out, "Oh, by the way, notice that you can't say *Who did you hit Billy and?*" Yet have you ever heard a child make a mistake like this?

Moreover, Chomsky argued that children have an uncanny ability to home in on exactly the right generalizations and patterns about their language, correctly chosen from among the vast array of logical possibilities. In fact, he's argued that children arrive at the right structures even in situations where it's extremely unlikely that they've even *heard* enough linguistic input to be able to choose from among the various possible ways to structure that input. In reverse-engineering terms, they seem to know something about the underlying language program that they couldn't have had an opportunity to test. It's as if they are ruling out some types of structures as impossible right from the beginning of the learning process. Therefore, they must have some prior innate knowledge of linguistic structure.

week 4

Are we programmed for language structure?

If we do have some innate knowledge of linguistic structure, what does this knowledge look like? It's obvious that, unlike vervets with their alarm calls, or bees with their dances, humans aren't born programmed for specific languages, since all humans easily learn the language that's spoken around them, regardless of their genetic heritage. Instead of being born with a preconception of a specific human language, Chomsky has argued, humans are prepackaged with knowledge of the kinds of structures that make up human languages. As it turns out, when you look at all the ways in which languages could possibly combine elements, there are some kinds of combinations that don't ever seem to occur. Some patterns are more inherently "natural" than others. The argument is that, though different languages vary quite a bit, the shape of any given human language is constrained by certain universal principles or tendencies. So, what the child is born with is not a specific grammar that corresponds to any one particular language, but rather a **universal grammar**, which specifies the bounds of human language in general. This universal grammar manifests itself as a predisposition to learn certain kinds of structure and not others.

universal grammar An innately understood system of combining linguistic units that constrains the structural patterns of all human languages.

If the idea that we could be genetically programmed to learn certain kinds of language patterns more easily than others strikes you as weird, it might help to consider some analogous examples from the animal kingdom. James Gould and Peter Marler (1987) have pointed out that there's plenty of evidence from a broad variety of species where animals show interesting learning biases. For example, rats apparently have trouble associating visual and auditory cues with foods that make them sick, even though they easily link smell-related cues with bad food. They also have no trouble learning to press a bar to get food but don't easily learn to press a bar to avoid an electric shock; they can learn to jump to avoid a shock but can't seem to get it through their heads to jump to get food. Pigeons also show evidence of learning biases: they easily learn to associate sounds but not color with danger, whereas the reverse is true for food, in which case they'll ignore sound but pay attention to color. So, among other animals, there seems to be evidence that not all information is equal for all purposes, and that creatures sometimes harbor useful prejudices, favoring certain kinds of information over others (for instance, for rats, which are nocturnal animals, smell is more useful as a cue about food than color, so it would be adaptive to favor scent cues over color).

Proposing a universal grammar as a general, overarching idea is one thing, but making systematic progress in listing what's in it is quite another. There's no general agreement among language scientists about what an innate set of biases for structure might look like. Perhaps this is not really surprising, because to make a convincing case that a specific piece of knowledge is innate, that it's unique to humans, and furthermore, that it evolved as an adaptation *for language*, quite a few empirical hurdles would need to be jumped. Over the last few decades, claims about an innate universal grammar have met with resistance on several fronts.

First of all, many researchers have argued that nativists have underestimated the amount and quality of the linguistic input that kids are exposed to and, especially, that they've lowballed children's ability to learn about structure on the basis of that input. As you'll see in upcoming chapters, extremely young children are able to grasp quite a lot of information about structure by relying on very robust learning machinery. This reduces the need to propose some preexisting knowledge or learning biases.

Second, some of the knowledge that at first seemed to be very language specific has been found to have a basis in more general perception or cognition, applying to non-linguistic information as well.

Third, some of the knowledge that was thought to be language-specific has been found to be available to other animals, not just humans. This and the more general applicability of the knowledge make it less likely that the knowledge has become hardwired specifically because it is needed for language.

Fourth, earlier claims about universal patterns have been tested against more data from a wider set of human languages, and some researchers now argue that human languages are not as similar to each other as may have been believed. In many cases, apparent universals still show up as very strong tendencies, but the existence of even one or a couple of outliers—languages that seem to be learned just fine by children who are confronted with them—raises questions about how hardwired such "universals" can be. In light of these exceptions, it becomes harder to make the case that language similarities arise from a genetically constrained universal grammar. Maybe they come instead from very strong constraints on how human cognition works, constraints that tend to mold language in particular ways, but that can be overridden.

Finally, researchers have become more and more sophisticated at explaining how certain common patterns across languages might arise from the fact

LANGUAGE AT LARGE 2.1

Engineering the perfect language

In all likelihood, the earliest users of languages never sat down to deliberately invent a system of communication in the way that humans invented Morse code or even the system of writing that we use to visually express our spoken languages. More likely, people intuitively fumbled around for the most natural-seeming way to express themselves, and language as we now know it was the eventual result. And the result *is* impressive. The combinatorial properties of human languages make them enormously powerful as communicative systems. By combining sounds into words and then words into sentences, we can create tens of thousands of meaningful symbols that can be combined in potentially infinite ways, all within the bounds of human ability to learn, use, and understand. An elegant solution indeed. But can it be improved upon?

It's interesting that when modern humans have turned their deliberate attention to language, they've often concluded that naturally occurring language is messy and poorly constructed. Many have pointed out its maddening irregularities and lapses of logic. For example, why in the world do we have the nicely behaved singular/plural forms *dog/dogs, book/books, lamp/lamps*, and *toe/toes* on the one hand—but then have *foot/feet, child/children, bacterium/bacteria, fungus/fungi*, and *sheep/sheep*? Why is language shot through with ambiguity and unpredictability, allowing us to talk about noses running and feet smelling? And why do we use the same word (*head*) for such different concepts in phrases like *head of hair, head of the class, head of the nail, head table, head him off at the pass*?

In a fascinating survey of invented languages throughout history, Arika Okrent (2010) described the various ways in which humans have sought to improve on the unruly languages they were made to learn. Many of these languages, such as Esperanto, were designed with the intention of creating tidier, more predictable, and less ambiguous systems. But unlike Esperanto, which was based heavily on European languages, some invented languages reject even the most basic properties of natural languages in their quest for linguistic perfection.

For example, in the 1600s, John Wilkins, an English philosopher and ambitious scholar, famously proposed a universal language because he was displeased with the fact that words arbitrarily stand in for concepts. In a more enlightened language, he felt, the words themselves should *illuminate* the meanings of the concepts. He set about creating an elaborate categorization of thousands of concepts, taking large categories such as "beasts" and subdividing them down into smaller categories so that each concept fit into an enormous hierarchically organized system. He assigned specific sounds to the various categories and subcategories, which were then systematically combined to form words. The end result is that the sounds of the words themselves don't just arbitrarily pick out a concept; instead, they provide very specific information about exactly where in the hierarchical structure the concept happens to fall. For example, the sounds for the concept of dog were transcribed by Wilkins as *Zita*, where *Zi* corresponds to the category of "beasts," *t* corresponds to the "oblong-headed" subcategory, and *a* corresponds to a sub-subcategory meaning "bigger kind."

Wilkins's project was a sincere effort to create a new universal language that would communicate meaning with admirable transparency, and he held high hopes that it might eventually be used for the international dissemination of scientific results. And many very educated people praised his system as a gorgeous piece of linguistic engineering. But notice that the Wilkins creation dispenses with Hockett's feature of *duality of patterning*, which requires that meaningful units (words) are formed by combining together a number of inherently meaningless units (sounds). Wilkins used intrinsically meaning*ful* sounds as the building blocks for his words, seeing this as an enormously desirable improvement.

But the fact that languages around the world *don't* look like this raises some interesting questions: Why not? And do languages that are based on duality of patterning somehow fit together better with human brains than languages that don't, no matter how logically the latter might be constructed? As far as I know, nobody ever tried to teach Wilkins's language to children as their native tongue, so we have no way of knowing whether it was learnable by young human minds. But surely, to anyone tempted to build the ideal linguistic system, learnability would have to be a serious design consideration.

that all languages are trying to solve certain communicative problems. We can come back to our much simpler analogy of the seeming universality of arrows. Arrows, presumably invented independently by a great many human groups, tend to have certain striking similarities—they have a sharp point at the front end and something to stabilize the back end; they tend to be similar lengths, etc. But these properties simply reflect the optimal solutions for the problem at hand. Language is far more complex than arrows, and it's hard to see *intuitively* how the specific shape of languages might have arisen as a result of the nature of the communicative problems that they solve—namely, how to express a great many ideas in ways that don't overly tax the human cognitive system. But an increasing amount of thinking and hypothesis testing is being done to develop ideas on this front.

week 2

2.4 ~~The Evolution of Speech~~

The ability to speak: Humans versus the other primates

In the previous sections, we explored two separate skills that contribute to human language: (1) the ability to use and understand intentional symbols to communicate meanings, perhaps made possible by complex social coordination skills; and (2) the ability to combine linguistic units to express a great variety of complex meanings. In this section, we consider a third attribute: a finely tuned delivery system through which the linguistic signal is transmitted.

To many, it seems intuitively obvious that speech is central to human language. Hockett believed human language to be inherently dependent on the vocal-auditory tract, and listed this as the very first of his universal design features. And, just as humans seem to differ markedly from the great apes when it comes to symbols and structure, we also seem to be unique among primates in controlling the capacity for speech—or, more generally, for making and controlling a large variety of subtly distinct vocal noises. In an early and revealing experiment, Keith and Cathy Hayes (1951) raised a young female chimpanzee named Viki in their home, socializing her as they would a young child. Despite heroic efforts to get her to speak, Viki was eventually able to utter only four indistinct words: *mama, papa, up*, and *cup*. To understand why humans can easily make a range of speechlike sounds while great apes can't, it makes sense to start with an overview of how these sounds are made.

Most human speech sounds are produced by pushing air out of our lungs and through the vocal folds in our larynx. The vocal folds are commonly called the "vocal *cords*," but this is a misnomer. Vocal sounds are definitely not made by "plucking" cord-like tissue to make it vibrate, but by passing air through the vocal folds, which act like flaps and vibrate as the air is pushed up. (The concept is a bit like that of making vibrating noises through the mouth of a balloon where air is let out of it.) The vibrations of the vocal folds create vocal sound—you can do this even without opening your mouth, when you making a humming sound. But to make different speech sounds, you need to control the shape of your mouth, lips, and tongue as the air passes through the vocal tract. To see this, try resting a lollipop on your tongue while uttering the vowels in the words *bad, bed*, and *bead*—the lollipop stick moves progressively higher with each vowel, reflecting how high in your mouth the tongue is. In addition to tongue height, you can also change the shape of a vowel by varying how much you round your lips (for instance, try saying *bead*, but round your lips like you do when you make the sound "w"), or by varying whether the tongue is extended forward in the mouth or pulled back. To make the full range of consonants and