

CHAPTER FOUR



What Children Learn About Language

The problem of Language, like the External World problem, is largely invisible in everyday life. While we know abstractly that we are hearing a sequence of arbitrary sounds, it feels as if thoughts are simply pouring into our minds. Suppose your spouse were to come into the room now and say, "You know, that woodwork could use some sanding." It would take, at most, a second or two for you to understand that sentence. Yet during that second you would have to do a remarkably complex set of calculations.

First, you have to break up the continuous stream of sounds into separate pieces and identify each sound accurately. Very small differences in sounds can make a big difference in meaning; *you know* means something quite different from *you go* or *who knows*. Then you have to string the sounds together into words. Because the average English speaker knows more than seventy-five thousand words, there are a lot of possibilities. Then, assuming you know the words, you have to combine them to make a sentence. Just as small differences in sound

can make a big difference in meaning, so small differences in the arrangement of words can make a big difference in meaning. The sentence "You know, could this sand use some wood-working?" would mean something quite different from your partner's remark. (It is one of life's tragedies that "John loves Mary" does not mean the same thing as "Mary loves John.") Then you need to understand all the nuances of meaning each word can have. You need to know that the word *sand* in that sentence refers to an action and not the stuff on the beach, and that *you know* doesn't actually refer to your knowledge at all. And, finally, you have to figure out something about the larger intent of the sentence. Is your partner reproaching you for reading instead of doing household chores? Or announcing an intention to spend the next hour sanding the woodwork, so you might want to move to a less noisy location to read? You figure all this out instantly and without any conscious effort.

Just as a magic show makes us realize how much we take for granted about things, visiting a foreign country makes us realize how much we take for granted about words. You greet your spouse's comment with effortless understanding, but if he were speaking a foreign language, you would feel baffled incomprehension instead. One of the brilliant aspects of the film *The Third Man* is that the actors who play the inhabitants of postwar Vienna actually speak German, with no translations or subtitles. When you watch the film, you find yourself experiencing just the same vertiginous incomprehension as Joseph Cotten, the innocent American hero. Simply hearing what are clearly important sentences spoken in a strange language rocks our usual calm assurance that we have some idea about what's going on (even without Orson Welles looming out from behind bullet-riddled baroque statues in the background). It isn't just that you don't know what the words

mean; you don't even know what the words, or even the sounds, are, or where one sound ends and the next begins. Everyone seems to talk so fast. This, of course, is where babies start out. In fact, the babies are worse off than Joseph Cotten in some ways, because they have no other language in which they can express their bafflement.

The Sound Code

Learning to understand a language is like cracking a deeply encrypted code. We all crack this code effortlessly, at an age we can't even remember, and we use it effortlessly as adults. But it turns out that the code is far more baffling than any spymaster's cryptogram. No computer has been able to figure it out.

When people comment on the scientific impossibilities of *Star Trek*—light speed and warp drive and even holodecks and replicators—they rarely mention what seems like a small technological detail. On *Star Trek* people talk to the ship's computers and the computers understand them (in fact, they even talk to the ship's doors). That technology may not be quite as distant as warp drive, but it's not close. Today, computer software companies all over the world are trying to create a machine that can understand spoken language. Companies and governments have spent billions of dollars on speech technology over the last fifty years, but no computer in the world has solved the Language problem yet. Our bathroom scales and elevators can produce a bit of understandable if annoyingly unnatural speech, but there is no computer that can do what every three-year-old can do: understand a conversation.

To nonscientists—even to the guys investing the billions of dollars—it isn't obvious why the problem is so hard. Pat was returning home from a conference in 1991, and, conquered by

the usual airplane exhaustion, she slumped into her seat next to a young guy with a backpack. The young guy turned out to be her fellow Seattleite Bill Gates, CEO of Microsoft. Pat spent the next four hours answering Bill's questions about why it was so difficult for his computers to understand speech. At the time, Microsoft and other computer companies were struggling to liberate computer users from keyboards. What the science could tell them was largely why the problem was so hard, rather than how to solve it.

The core of the problem, as in the problem of the External World and that of Other Minds, is the mysterious gap between the sound waves that actually reach our ears and the sounds and words we create in our minds. We can make a sort of photograph of a sound called a spectrogram. The spectrogram shows the actual physical properties of the sound waves: how loud they are, what pitch they are, and how they change. Just as we must translate the two-dimensional pattern of light on our retinas into the three-dimensional solid objects we perceive, so we must translate this pattern of sounds into language. The distance from there to here is just as great.

There are some glaring problems that become obvious as soon as you compare the spectrogram with the words we perceive. First, the sounds of human speech aren't like beads stacked next to one another on a string: there are no gaps or pauses between the sounds in the spectrogram. Instead, the sounds are continuous, and we have to divide them into units. Second, all voices are different because our mouths are all different sizes and shapes, so even simple sounds (like *ah*) sound different depending on who says them. And when we speak more quickly or more slowly, which we do all the time, the sound waves change again. Moreover, each time a consonant sound, such as *b* or *d*, is placed in front of a different

vowel, the sound changes. The *d*'s in front of the words *dude* and *deed* are so different that a spectrogram of the *d* in *dude* actually looks more like a *g* than like the *d* in *deed*.

Finally, and most complicated of all, people speaking different languages hear sounds totally differently. A sound with exactly the same spectrogram will be heard differently by someone speaking Japanese and someone speaking English. Two physically different sounds (like *r* and *l*) may sound identical to a Japanese speaker but completely different to an English speaker. It isn't only that you must figure out how to get from the sensations to the representations, as you do when you translate the two-dimensional retinal images to a three-dimensional world. You also must do that translation differently for each different language.

Three-year-olds have solved all these problems. They can recognize a *d* sound whether it's spoken by Mom or Dad, whether it's in *deed* or *dude*, whether it's quickly whispered or slowly sung, and they make just the right discriminations for English. Computer systems that can do some speech recognition can't match the three-year-old. As we mentioned, most English speakers know more than seventy-five thousand words. If you limit your speech to the ten digits of telephone numbers, or even to the cities in America, computer speech recognition functions well. But real conversations can't happen using ten words, or even a thousand words.

One of the biggest problems for computers is segmenting speech into separate units for analysis. Early computer programs solved this by having speakers separate each word. People using this voice-recognition technology had to speak very slowly, SEPARATING (one-second pause) EACH (one-second pause) WORD (one-second pause) WITH (one-second pause) A (one-second pause) VERY (one-second pause) ANNOYING (one-second pause) ONE- (one-second pause) SECOND (one-

second pause) PAUSE. To solve the problem of different voices, computers are programmed to recognize only one particular person's voice and then have to be reprogrammed for each separate user. In the same way, computers are programmed to treat the *d*'s in *deed* and *dude* separately, as if each were a completely different sound. In 1998 the first continuous speech recognition software for dictation became available, but it still requires training for each separate user and limits your vocabulary.

So much for the notion of a *Star Trek* computer that answers your every question accurately in a soothing if somewhat chilly voice. By the time any current computer could understand "Implement evasive maneuver Alpha Theta and fire photon torpedoes at Romulan vessel on my mark," the *Enterprise* would be a noncorporeal energy particle pattern.

Making Meanings

All this complex code breaking is necessary just to figure out the words of a language from the sounds you hear. This part of understanding language happens so effortlessly and so quickly that it's hard even to recognize the problem at first, and scientists have only started tackling it quite recently. The next part of the process, getting from words to meanings, is more obviously difficult. Philosophers have pondered for millennia the question of how words can mean things.

It almost seems as if there is a magic link between the words we use and the outside world. Say a word and suddenly you are in touch with the thing the word refers to, no matter how distant or strange. Many cultures and religions explicitly believe in word magic; knowing the true name of a thing gives you power over it. But when you think about it, our everyday language is just as mysteriously powerful. Consider what is happening as you read this book. By casting your eye over a bunch

of arbitrary printed shapes, you are suddenly in contact with the minds of three people thousands of miles away from you. You visit a laboratory you've never seen and meet children who grew up long ago. And as every novel reader or letter writer or Internet cruiser knows, words can not only take you to other worlds, they can create worlds of their own. You don't need *abracadabra*; *once upon a time* will do. How can words defy all the limits of space and time and possibility in this way? How could anyone learn to wield this kind of power?

Almost two millennia ago Saint Augustine proposed one solution to the problem, perhaps the most obvious one: as children, we heard our parents say words and point to things, and we associated the words with the things. But the more you think about it, the less adequate this solution seems. Over the centuries other philosophers have demonstrated the difficulties. Bertrand Russell showed that meanings go beyond just the things we point to. How do we learn words like *unicorn*, words that refer to things that don't exist? How do we learn all the words—verbs and adjectives and prepositions—that don't refer to things at all? Ludwig Wittgenstein raised another set of questions. How do we learn not just what the words refer to but what the speaker wants us to do about them? After all, even to understand pointing you need to know something about the intention of the person who points. You have to know that the gesture of extending your index finger is a way of picking out an object to be named rather than, say, casting a curse or conferring a blessing. The philosopher Willard Quine raised yet another set of questions. How do we know that a name refers to the thing someone points at rather than to the thing plus a bit of the empty space near it or some part of the thing? Computers still can't solve the problem of decoding sounds, but they are even further away from solving these sorts of problems.

The Grammar We Don't Learn in School

In addition to the challenge of decoding sounds and meanings, there is still another set of problems. In the 1960s Noam Chomsky raised a whole new set of questions which people hadn't paid much attention to earlier. How do we combine words to make new sentences? Almost all of the sentences we hear are brand-new, and yet we have no difficulty figuring them out. Even when we know individual words, arranging those words in different ways can lead to different meanings. Chomsky's answer to that question created a new field, modern linguistics, and a new way of thinking of the old idea of grammar.

Traditional grammar, the sort of thing we used to learn in school, vacillated between telling you what speakers of a language actually did and saying what speakers ought to do. Chomsky argued that knowing a language involved knowing a set of unconscious rules, but they weren't like the rules of traditional grammar. These rules weren't socially imposed like the rules of traffic codes or Monopoly or the old elementary school grammar books. Rather, they were natural, unconscious rules. They were like the rules we use when we turn visual information into representations of objects. They were more like the law of gravity than like the law of the land (or the law of the English teacher).

Chomsky's solution to the problem of Language is much like the modern solution to the External World and Other Minds problems. We are designed to take in sequences of sounds and translate them into representations of meanings just as we are designed to take in sensory information and translate it into representations of objects and to take in facial expressions and translate them into representations of feelings. We have an implicit set of rules that allows us to transform the sequences of sounds we hear into sequences of ideas.

We actually know quite a lot about how some parts of this system work. For instance, we know a great deal about how we translate the sounds we hear into meaningful units like words, though still not enough to let us program a computer to do it. We know something, though less than Chomsky originally hoped, about how we combine words to make sentences. Other parts of the problem remain deeply mysterious. In particular, the problem of how words come to refer to the external world, of how meanings are made, is almost as mystifying to contemporary linguists as it was to Saint Augustine.

We assume that this system was designed by evolution, and it is certainly distinctively human. Perhaps the most obvious advantage of language is that it lets us communicate and coordinate our actions with other people in our group. But language also has other less obvious but equally distinctive advantages. The fact that we speak different languages also lets us differentiate between ourselves and others; it's as good a way as any of knowing who is part of your group and who is an outsider (keeping information away from your enemies may be almost as important as sharing it with your friends). And the development of language is probably linked to the development of our equally distinctive ability to learn about people and things. It allows us to take advantage of all the things that people before us have discovered about the world. We can see so much further than other species because we stand on the shoulders of our mothers and fathers (who at least look like giants to babies).

Chomsky's solution raises a deeper developmental problem: Where does this linguistic system come from? By the time they are in kindergarten, children have mastered almost all of the complexities of their particular language, with no conscious effort or instruction. How do they do it? The broad lines of the developmental answer to this question should be familiar

by now. Babies are born knowing a great deal about language. They also have powerful learning procedures that allow them to add to that knowledge and, in particular, to learn all the details and peculiarities of the language of their own community. Finally, adults play an especially crucial role in language learning.

The analogy to science works very well when it comes to explaining how babies solve the Other Minds problem and the External World problem. There really is a world of objects and minds out there. Babies make up theories about that world, but those theories can always be revised if new evidence comes along. In the case of language, however, the problem is rather different. It is not about discovering an independent reality but about coordinating what you do with what other people do. There isn't any abstract "language" out there that is independent of what people say. We could find out (in fact, we do find out) that we are all wrong about some important aspect of the world or other people. But we couldn't find out that we are all speaking English the wrong way; English just is the language we speak. So the babies' Language problem is not so much the scientist's problem—find out what the world is really like—as it is a kind of sociological or even anthropological problem: find out what the folks around here do and learn how to do it yourself. The other folks are crucial.

The problem is difficult because different communities speak different languages, sometimes quite radically different languages. Babies don't know beforehand which language they are going to be exposed to. Potentially, they have to be able to master any one of thousands of different languages. And yet, by the time they are four or five, children have figured out precisely which language is spoken in their community.

Grown-ups are both the teachers and the subject matter.

What they say is the only source of evidence about what the language is like. And for the children the aim of the enterprise is not just to find out about the grown-ups' language but also to make that language their own.

What Newborns Know

Ask anyone when children start to learn language. Almost everyone will say that language begins when babies say their first words. But the new techniques for uncovering what babies know have led to a surprising discovery. Babies know important things about language literally from the time they are born, and they learn a great deal about language before they ever say a word. Most of what they learn at that early age involves the sound system of language. We decode the sound cryptogram, and solve many of the problems that still baffle computers, before we can actually talk at all.

We mentioned that part of what makes learning language difficult is that languages carve up sounds and different languages carve them up differently. A wide variety of different sounds, with very different spectrograms, will all seem like the same sound to us, and, in turn, that sound will seem sharply different from other sounds that are actually quite similar to it physically. Suppose you use a speech synthesizer to gradually and continuously change one particular feature of a sound, such as the consonant sound *r*, and play that gradually changing sound for people. You very gradually and continuously change the *r* sound to *l*. What is actually coming into the listeners' ears is a sequence of sounds, each of which is just slightly different from the last. But what they perceive is someone saying the same sound, *r*, over and over, and then suddenly switching to a new sound, *l*, over and over. The listeners have divided up the continuous signal into two sharply defined categories: either it's an *r* or an *l*, not anything in

between. They can't distinguish between all the different *r*'s, even though the sounds themselves are quite different. Scientists call this categorical perception, because a continuously changing set of sounds is perceived categorically as being either black or white, *r* or *l*, with nothing in between.

The way we categorically perceive speech is unique to each language. In English we make a sharp categorical distinction between *r* and *l* sounds. Japanese speakers don't. In fact, Japanese speakers can't hear the distinction between American *r* and *l*, even when they are listening very hard. (Hence all the dubious jokes about Japanese speakers ordering what sounds like "fled lice" instead of "fried rice.") Pat was in Japan to test Japanese adults and their babies on the *r-l* distinction. She had carefully carried the computer disk with the *r* and *l* sounds to Japan, and when she arrived in the laboratory in Tokyo, she played them on an expensive Yamaha loudspeaker. She thought that such clearly produced sounds would surely be distinguished by her Japanese colleagues, who were quite good English speakers as well as being professional speech scientists. As the words *rake*, *rake*, *rake* began to play out of the loudspeaker, Pat was relieved to know that the disk worked and the sound was perfect. Then the train of words changed to an equally clear *lake*, *lake*, *lake*, and Pat and her American assistant smiled, looking expectantly at her Japanese colleagues. They were still anxiously straining to hear when the sound would change. The shift from *rake* to *lake* had completely passed them by. Pat tried it over and over again, to no avail.

This happens to all of us, of course, as we try to hear distinctions that are used in languages that are not our own. Say we use the speech synthesizer again to change the sound *b* gradually and continuously to *p*, and we test speakers of many languages on these sounds. Americans will hear two sharp cat-

egories, *b* and *p*. Spanish and French speakers listening to these same sounds hear two categories also but divide the continuous stream of sounds in a different place than the Americans do. What sounds like *b* to a Spanish speaker will sound like *p* to an English speaker. Thai speakers hear three categories. In each case listeners hear sharp changes—quantum leaps—between categories, with no in-between. But the speakers of Spanish, French, and Thai hear those quantum leaps at different places. And we English speakers don't even notice the categorical shifts Spanish or French or Thai speakers hear, just as the Japanese speakers didn't notice the change from *rake* to *lake*. It's as if the speakers of each language have a very different way of transforming the actual sound waves that come into their ears into a set of language sounds.

Why do the speakers of different languages hear and produce sounds so differently? Ears and mouths are the same the world over. What differs is our brains. Exposure to a particular language has altered our brains and shaped our minds, so that we perceive sounds differently. This in turn leads speakers of different languages to produce sounds differently. When and how do babies start to do this? Do they start out listening like a computer, with no categorical distinctions? Or do they start out with the categorical distinctions of one particular language, say English or Japanese or Russian?

We can't ask babies directly whether they think two sounds are the same or different, but we can still find out. Very young babies can tell us what they hear by sucking on a special nipple connected to a computer. Instead of producing milk, sucking on this special device produces sounds from a loudspeaker, one sound for each hard suck. Babies love the sounds almost as much as they love milk: they may suck up to eighty times a minute to keep the sounds turned on. Eventually, though,

they slow down; they get bored hearing the same thing over and over again. When the sound is changed, however, infants immediately perk up and suck very fast again to hear the new sound. That change in their sucking shows that they can hear a difference between the new sound and the sound they heard before. Using this technique we can do the same *r* and *l* experiment we just described with adults. We can use a speech synthesizer to present the babies with a slowly and continuously changing consonant sound. Then we can test the babies to see which sounds they think are the same and which sounds they think are different.

Scientists anticipated that these tests would show that very young babies initially can't hear the subtle differences between speech sounds and only slowly learn to distinguish those that are important in their particular language, such as *r* and *l* in English. In fact, the results were just the reverse. In the very first tests of American infants listening to English, babies one month old discriminated every English sound contrast we threw at them. Moreover, the babies demonstrated the categorical perception phenomenon. They thought all the *r*'s were the same and different from all the *l*'s, just as adult English speakers do.

But then shortly afterward speech scientists discovered something even more remarkable. Kikuyu babies in Africa and Spanish babies in Mexico were also excellent at discriminating American English sounds as well as the sounds of Kikuyu and Spanish, and American babies were just as good at discriminating Spanish sounds—much better than American adults. The sophisticated Japanese scientists who strained to hear the difference between *rake* and *lake* would not have had any trouble doing so when they were forty or fifty years younger. Very young babies discriminated the sounds not only of their own language but of every language, including languages they had

never heard. Infants were as good at listening to American English as they were at listening to African Kikuyu, Russian, French, or Chinese regardless of the country they were raised in. Pat also discovered that babies, unlike computers, make these distinctions no matter who is talking—a man or a woman, a person with a high squeaky voice or one with a deep resonant voice.

So babies start out knowing much more about language than we would ever have thought. Newborn babies already go well beyond the actual physical sounds they hear, dividing them into more abstract categories. And they can make all the distinctions that are used in all the world's languages. Babies are "citizens of the world." Perhaps we grown-up scientists failed to predict this because our skills are so much more limited. Our citizen-of-the-world babies clearly outperform their culture-bound parents.

Taking Care of the Sounds: Becoming a Language-Specific Listener

Providing the answer to one puzzle creates another. If babies are born listening like universal linguists, how do they grow up to be culture-bound language specialists? Japanese babies learn English if they're raised in America and Japanese if they're raised in Japan. When does a Japanese baby learn that in her language it doesn't matter whether she produces *r* or *l* because the adults in her culture can't hear the difference anyway and it won't change the meaning of the word in Japanese? When does the American baby learn that the difference between the two Spanish *b*'s doesn't matter in English? Most scientists thought at first that babies would appreciate these language-specific distinctions only after they had already learned quite a lot of meaningful language.

To answer this question we needed a way to test babies once

they had listened to their particular language for a while. After four months or so, many babies are much less likely to suck to turn on computer sounds, so that technique doesn't work as well. But another test works very well with six- to twelve-month-olds. The babies sit on a parent's lap. On their right a person keeps their interest by slowly manipulating toys—dangling a plastic spider, turning a toy horse upside down, and doing other visually interesting things. On the babies' left there is a loudspeaker with a black box on top of it. A sound repeats out of the loudspeaker, something like *oo, oo, oo*. Every once in a while the sound is changed to *ee, ee, ee*. When babies hear the sound change, they tend to get distracted from the interesting person and look over toward the loudspeaker. When they do, the black box on top of the loudspeaker lights up. Inside the box a bear dances or a monkey pounds a drum, delighting the babies. Then the display goes off and the babies turn back to the interesting person on their right. Soon the babies figure out that if they turn their heads toward the loudspeaker when the sound changes, they'll see something interesting. Whether they turn their heads or not tells us whether they heard the sound change or not.

When Pat went to Japan to test adults on the American *r* and *l* sounds, she also tested babies. Japanese and American seven-month-olds discriminated *r* from *l* equally well. But just three months later, the two groups of infants were as different as night and day. At ten months, Japanese infants could no longer hear the change from *r* to *l*. American infants not only could do so but had actually gotten much better at making this distinction. A previous study of babies being raised in English-speaking homes had similar results. That study showed that at six months Canadian babies could discriminate Hindi speech sounds that Canadian adults can't distinguish. But by twelve months the Canadian babies could no longer do so.

As they hear us talk, babies are busily grouping the sounds they hear into the right categories, the categories their particular language uses. By one year of age, babies' speech categories begin to resemble those of the adults in their culture. Pat conducted some even more complicated experiments with Swedish babies using simple vowels to see how early they start organizing the sounds of their language in an adultlike way. She showed that at six months the process has already begun. The six- to twelve-month time span appears to be the critical time for sound organization.

What might be happening to the babies between six and twelve months? One way of thinking about it is in terms of what Pat calls prototypical sounds. After listening to many *r* sounds in English, for example, babies develop an abstract representation of *r*—a prototypical *r*—that is stored in memory. When we want to identify a new sound, we seem to do it by unconsciously comparing the new sound to all of the prototypes stored for our language and picking the one that's the best overall fit. Once we've unconsciously done this, we distort the way we hear a sound to make it more like the prototype stored in memory than like the sound that actually hits our ears.

It's similar to what happens when you show people a drawing of something they've seen very often, a house, for example, and then ask them to copy it from memory. If the house you show them doesn't have a chimney, many people will add one to their drawing anyway, even though it wasn't in the original drawing they saw. Once they coded the picture as a house, they distorted their memory of it to make it more like what they think of as the prototypical house. We can do complicated analyses to show just what the prototypes of our speech sounds are and just how we distort what we hear to suit them. Our language prototypes "filter" sound uniquely

for our language, making us unable to hear some of the distinctions of other languages. Pat's tests suggest that babies' language prototypes begin to be formed between six and twelve months of age.

It isn't just that younger babies have a skill they lose later on. Rather, the whole structure of the way babies organize sounds changes in the first few months of life. Before they are a year old, babies have begun to organize the chaotic world of sound into a complicated but coherent structure that is unique to their particular language. We used to think that babies learned words first and that words helped them sort out which sounds were critical to their language. But this research turned the argument around. Babies master the sounds of their language first, and that makes the words easier to learn.

When babies are around a year old, they move from sounds to words. Words are embedded in the constant stream of sounds we hear, and it is actually difficult to find them. One problem computers haven't yet solved is how to identify the items that are words without knowing ahead of time what they are. Try to find the words in a string of letters like *theredona-teakettleoftenchips*. The string contains many different words: *The red on a teakettle often chips* or *There, Don ate a kettle of ten chips* and so on. Of course, in written language there are normally spaces between words. But in spoken language there aren't actually any pauses between words. That's why foreign speech sounds so fast and continuous, and that makes the Language problem very hard for computers to solve.

Babies seem to learn some general rules about the words in their particular language before they learn the words themselves. By nine months, for example, they've learned that English contains words that have a certain emphasis pattern: words with a first-syllable stress pattern, like *BASEball* or

POPCorn, are more common than the reverse (a word like *surPRISE*). In some other languages it's the other way around: first syllables are stressed less often than last syllables. By nine months babies have all this sorted out. American babies prefer to listen to words with the American pattern, while babies from other countries prefer to listen to words typical of their own languages, even though the babies at this age don't understand the meanings of these words.

After babies learn which sounds are possible in their language, they learn which sound combinations are possible. In English, for instance, the sound combination *zb* is not possible. No English word contains this combination. In Polish, however, this combination is possible. (Zbigniew Brzezinski, President Carter's national security adviser, could never have become president because, for one thing, hardly any Americans could pronounce his name.) By nine months babies show a preference for listening to sound combinations that are possible in their language, even if the sound combinations don't form real words in the language but are only potential words. American nine-month-olds already have trouble with words like *Zbigniew*, while Polish nine-month-olds would think it's no problem. Knowing which words are possible in your language helps you begin to divide the continuous stream of speech into words, even if you don't know what those words mean. If you are American, you can already eliminate the strings that have *zb* in them or that have the wrong stress pattern.

The Tower of Babble

So babies are learning about speech a long time before they begin to talk. But, of course, what parents notice most is what babies actually say. Whether babies are born in Paris, Zimbabwe, Berlin, or Moscow, they start to coo when they're about three months old. They make delightful little *oooh* and *aahh*

sounds when a parent is face-to-face with them, talking and smiling. Babies seem to grasp intuitively that humans take turns in this kind of exchange. They coo, we goo, and thus we have our first conversations with our children. Babies already know something about how dialogue works.

A short time later, at about seven or eight months, babies begin to babble. They start producing strings of consonant-vowel syllables, *dadadada* or *babababa*. Babies across cultures babble at first in an identical way, producing consonant-vowel combinations using sounds like *b*, *d*, *m*, and *g* with the vowel *ah*. Pat and Andy vividly remember when their daughter, Katherine, started to babble. As a speech scientist Pat had waited for months for the moment babbling began. She'd even set up a recorder hoping to catch it on tape. One morning Pat took Kate out for a walk in the stroller and stopped at the local Starbucks for a latte (it's Seattle, after all). Kate was happily cooing and gurgling at the line of customers as Pat ordered, when suddenly a "Babababa" rang out. Pat froze and waited to hear it again, asking the cashier and other Starbucks habitués to listen for confirmation. With the customers poised on the edge of their seats, Kate blissfully went on, "Babababa." She babbled right on time, like clockwork, like all babies around the world.

Once babies reach the babbling milestone, the universal phase of language production ends. Babies from different cultures, learning different languages, start to make the distinctive noises of their own community sometime between a year and a year and a half. The Chinese baby starts to babble in a way that sounds Chinese. She uses very rapid pitch changes just like adult Chinese speakers. Swedish babies babble in a way that sounds distinctly Swedish, using the rising intonation patterns typical of adult speakers of Swedish. (They sound a bit like the Swedish chef on *The Muppet Show*.)

The First Words

So far we've talked about how children come to understand the system of sounds of the language they hear. This apparently simple problem turns out to be extremely complicated. Babies are hard at work on it throughout infancy. But we haven't even started in on what most of us think of as the central problem of language: learning what words mean.

Remember that Augustine thought there was a simple answer to this problem: children saw things and heard their parents name the things and then associated the things with the names. That idea has a strong intuitive appeal. If you ask the average parent, or for that matter the average psychologist, when babies begin to talk, they will tell you about when they use their first names. Usually, parental egocentricity being what it is, the parents report "Mama" and "Dada." It turns out, in fact, that across a variety of very different languages the "baby words" for mother and father are very similar; "Mama" and "Dada" are joined by "Mati" and "Tata." They are also, of course, precisely the sounds that babies are very likely to produce spontaneously when they babble. So it's not entirely clear whether babies say "Mama" and "Dada" because that's what their beloved parents call themselves, or whether parents call themselves Mama and Dada because that's what the babies say anyway.

While philosophers, psychologists, and parents were so sure they knew how babies started to speak, no one consulted the babies themselves until around the 1970s. With the advent of videotape you could actually watch what children said and when they said it. The results were surprising. Babies did say "Mama" and "Dada" (parents weren't utterly self-deluded), and also "juice" and "ball" and "doggie." But they said many other things that grown-ups didn't notice. Perhaps parents didn't notice because the words children used were so pecu-

liar. Babies consistently said things like "gone," "there," "uh-oh," "more," and "what's that?" among their very first words. Why these rather odd words? And what did they mean to the babies?

Alison set out to find out. She had come to Oxford to study philosophy. But she was interested in how language begins, so she spent many hours each week watching babies in the big, drafty rooms of North Oxford villas. The grand houses have been divided into apartments, but they still have some of the same atmosphere they had when Lewis Carroll told stories about Wonderland to Dean Liddell's daughters. Oxford is still gray and wet and dim and full of elegant, chilly buildings and faces. There is still nowhere more elegant and more chilly than Logic Lane, where the philosophy classes are held. And yet Oxford also still has lamplit rooms full of luminously red-headed toddlers gathered around tables of cream cakes and milky tea, and little courtyard gardens with iron gates where pink-cheeked three-year-olds play ball. The children had much the same effect on Alison, a young American woman philosopher, that they had had on Lewis Carroll, the elderly bachelor logician, a hundred years earlier. They were a bright glimpse of clarity and warmth compared with the vaguely threatening and deeply eccentric creatures of Logic Lane. And yet, at the same time, the private children's world had its own mystery and strangeness, even if the strangeness often seemed more sensible than the accepted craziness of the world outside.

For while the babies' language was superficially simpler than the convoluted paragraphs of Logic Lane, it was, in its own way, just as peculiar. Almost always your first guess at what a word meant would turn out to be wrong. For example, *gone* was one of the most common words these babies used. When parents notice this at all, they assume that it has something to

do with finishing up food. But, in fact, the babies rarely used the word this way. In one taped hour Henry, a particularly cherubic eighteen-month-old, said “gone” turning over a small piece of wax paper with a bit of brown sugar attached to it so that the sugar became invisible (seven times in a row), turning the page of a picture book so that each baby animal was no longer in sight, hiding a ring under the edge of Alison’s skirt (twelve times in a row), putting a block inside a toy mailbox (seven times in a row), and plaintively searching for Mum, who had gone to a neighbor’s (“Mummy gone!”—think Jackie Coogan in *The Kid*).

It turned out that *gone* didn’t have much to do with food at all, or with the way grown-ups use *gone*. Instead, Henry, and the other babies Alison studied, used *gone* to describe the many and varied ways that objects disappear from view. They commented on the fact that they couldn’t see something they knew was still out there somewhere.

Alison discovered that there were a number of other words that seemed to work the same way as *gone*. For instance, babies often use a word to indicate whether they succeed or fail in doing something. American babies use *there!* to note their successes and *uh-oh* to describe failures, while the babies in the Oxford villas used the more genteel *oh dear* (although one British baby did briefly but memorably say *oh bugger*). Most parents don’t really think of *uh-oh* as a word at all, let alone an important word. But it is consistently one of the very first words that American and English babies use. As we found out later, Korean and French babies also have their own special words to talk about when they succeed and fail (French one-year-olds who manage to build a rather shaky tower of blocks produce a splendidly Gallic *voilà!*)

Even the children’s early names turn out to be more complicated than Augustine, or parents, might think. A father’s

delight at hearing the baby say his name may fade a bit as the baby hails the dad’s best friend with the same jubilant “Dada!” And the mailman. And the TV repairman. It’s a slight comfort perhaps to see that the family pet suffers from the same fate: any animal from an anteater to a zebra is a “doggie.” One linguist reported that her daughter used *moon* to talk about the actual moon, but also lamps, oranges, and crescent-shaped fingernail parings. Just as the babies extend *uh-oh* or *gone* to new situations, they also extend their early names. They are trying to make sense of the language they hear by applying it to concepts that seem important to them. They use words in a way that makes sense to them, even if grown-ups don’t use the words that way.

Initially children use just a few names, mostly for familiar things and people. But when they are still just beginning to talk, many babies will suddenly start naming everything and asking for the names of everything they see. In fact, *what’sat?* is itself often one of the earliest words. An eighteen-month-old baby will go into a triumphant frenzy of pointing and naming: “What’sat! Dog! What’sat! Clock! What’sat juice, spoon, orange, high chair, clock! Clock! Clock!” Often this is the point at which even fondly attentive parents lose track of how many new words the baby has learned. It’s as if the baby discovers that everything has a name, and this discovery triggers a kind of naming explosion.

It turns out you can show experimentally that babies at this stage have a new approach to learning words. You can give a baby just one example of a new nonsense word naming a new kind of thing (“Look, a dax!” you say, pointing to an automatic apple corer), and it will become a permanent part of the baby’s vocabulary. Weeks or even months later, he’ll correctly identify the “dax.” Just one salient instance and babies will internalize a word forever (sometimes, of course, with

rather embarrassing consequences). The process is called fast mapping. The babies seem to assume at once that the new name they hear names the new object they've just seen. Babies start to fast-map at about the time they have their naming explosion.

Language is as much invented as learned. Babies don't simply soak up associations between names and things or mimic adults' use of words. Instead, they actively restructure language to suit their own purposes. If they need a word for disappearance or failure, they'll happily press *all gone* or *uh-oh* or even *oh bugger* into service. If they need a word for all animals, they'll make *doggie* fit the bill.

If you can make some assumptions about what people are trying to say, that also gives you a substantial leg up in decoding their language. Experiments show that children know something about other people's intentions and use that knowledge to help figure out what words mean. These experiments also show that Augustine was wrong in another way.

If Augustine were right, what would happen if children just happened to be looking at an apple when Dad said, "Where are the pears?" They ought to be stymied. They should mistakenly think that *pears* means *apple*. However, even toddlers don't make this kind of mistake. Suppose you get an eighteen-month-old to look at one new object, for example, a potato masher, while his mother is looking at another object, say a bulb baster. The mother says, "Oh, look, a dax!" Then you put both objects in front of the baby and ask, "Show me the dax." The baby, it turns out, assumes that *dax* means the bulb baster, the thing his mother was looking at, rather than the potato masher he was looking at when he heard the word.

You can set up even more complicated situations like this one. Suppose the toddler and the experimenter sit at a table full of toys and the experimenter picks up each toy and looks

at it. Then the experimenter leaves the room, and while he's away another person brings a new toy into the room and leaves it on the table. Now the experimenter returns and cries out, "Oh, look, a dax!" The child assumes that the new object is the dax. Of course, that's what we would assume, too. But when you think about it, that assumption requires a lot of sophisticated knowledge about the other person and about communication. The baby seems to know that people talk about things that are new to them, rather than things that are familiar. And once again Augustine's theory that children learn language by associating a name and a thing turns out to be wrong. In this case the child is looking at many different things when she hears the name, but she connects the name to the thing that is new to the other person.

Putting It Together

Putting words together to make new sentences and more complex meanings is another central part of language. Before they are three, children are working out this part of the language problem as well. Many English-speaking children go through a stage where they start putting words into two-word combinations, like poor Henry's plaintive "Mummy gone!" Two interesting observations suggest that even these very young children already have some idea of grammar. First, they seem to recognize that only some word orders are possible in their language. They say "Mommy gone" but not "gone Mommy"; "more cookie" but not "cookie more." Second, they already use different word orders to express different meanings. "Kiss Teddy" means Mommy should kiss the teddy bear, while "Teddy kiss" means the teddy bear is going to kiss Mommy (undoubtedly assisted by the speaker). These very simple two-word sentences already follow certain rules, even though a two-year-old would never have heard these sentences from any-

one else. Just as babies invent meanings, they also invent grammatical rules.

As most English-speaking babies grow older, they start to produce longer and more complicated sentences, but those sentences still sound very different from the sentences of the adult language. We all recognize that two- and three-year-olds have a kind of distinctive “cookie monster” talk (that, of course, is precisely why *Sesame Street*’s Cookie Monster talks that way). What may be less obvious is that “cookie monsterese” is very systematic. Young children systematically leave out word endings, such as the plural *s* or the past tense *ed*, and they omit “grammatical” words such as *the* or *of*. Even if you try to get a toddler to repeat a grown-up sentence word for word, what comes out will be very different. “I don’t want the broccoli, I want the cookies” becomes “No want bwocwi, want cookie!” The children have largely made up their own language with its own rules and grammar, just as they decide themselves what the words they use will mean. But the important thing is that they are rules and it is grammar.

Some children, though, especially younger siblings, take quite a different route toward grammar. Rather than starting out with a bunch of individual words and gradually combining them into more complex sentences, these babies seem to take the opposite approach. They seem to get hold of whole sentences and then take them apart into separate words. They start out by grasping the intonation patterns of whole adult sentences, and they babble in a way that mimics those intonation patterns. Often it sounds as if they’re quite fluent in a language their parents just don’t happen to know, like Klingon or Vulcan. Sometimes startled parents will suddenly hear a whole sentence of English emerge—“I-want-some-cookies-please”—or will hear a few recognizable English words em-

bedded in the Klingon sentence. Eventually and gradually the odd sentences turn into English sentences.

In addition to learning verb and noun endings, babies also have to learn the details of how those verb and noun endings are used. It isn’t as straightforward as just adding *s* to a word to indicate a plural. How about *boxes* (with an *iz* sound) and *rods* (with a *z* sound)? And let’s not even think about *women* and *children* and *sheep*.

Children learn, and create, systematic rules for dealing with these variations. One of the very first experimental studies of language development demonstrated this. You can show a toddler a picture of an imaginary creature and say, “This is a wug.” Now you show him a picture of two of the imaginary creatures and say, “There are two of them, what are they?” By four or five, though not earlier, children will happily say, “They’re wugs.”

Looking at the children’s mistakes, paradoxically, also shows they are learning in an intelligent way. Preschoolers often use invented words like *womans* and *childs* (Alison’s sister referred to her large family by saying, “All of we’s is childs.”) In fact, children often begin by getting some of these words right—a two-year-old may say “children,” much to the pleasure and pride of grown-ups, and only later produce the invented form *childs*. The “mistake” actually shows, though, that the toddler has learned the more general rule for making plurals.

Children learning different languages vary even more radically in the ways they approach grammar. We’ve seen that very young infants are already sensitive to the particular sounds of their particular language. In the same way, even very young toddlers seem to be sensitive to the particular grammar of their particular language.

We just mentioned that English “cookie monsterese” leaves

out noun and verb endings like plurals and past tenses. In fact, English doesn't use those endings very much compared with other languages. Any English speaker who's tried to learn French, Russian, Spanish, Latin, or practically any other language for that matter will remember the countless conjugations and declensions with fear and loathing. (Of course, speakers of other languages find English prepositions and articles equally baffling and repulsive.) Children who are learning other languages pick up and use noun and verb endings much earlier than English-speaking children. Korean-speaking children, for instance, not only use many more early verbs than English speakers, they also use verb endings correctly even when they are using only single words. French-speaking children seem to have almost no trouble picking up the system of grammatical gender, a feat that will seem absolutely astonishing to anyone who has tried to learn French.

How Do They Do It?

How do children manage to do it? There is clearly some genetic foundation that enables human beings to acquire language. That fact has been the focus of much attention in linguistics. But children must also have powerful learning mechanisms, particularly in order to learn the specific properties of their own language. Moreover, grown-ups seem designed to help babies learn.

Word-Blindness: Dyslexia and Dysphasia

Just as there are genetic disorders that make it difficult to understand the mind and the world, there are genetic disorders that make language difficult. And once again, these tragic examples demonstrate that we have an innate endowment that lets us understand and speak.

While all normally developing children effortlessly perform

minor miracles with sound, not all children can do this. There are children who hear perfectly well and who are perfectly intelligent but who still have a hard time with language. We don't know exactly what causes these language disorders, but we do know that the disorders run in families, which points to genetic factors.

Often, language problems become evident only when children start to learn to read and write. Children who are having trouble with the sound system of the language may be able to compensate enough to understand everyday speech. After all, in everyday communication there are lots of cues about what someone is trying to say, including tone, inflection, facial expression, and context. But in order to read and write, you have to translate the system of language sounds directly into a system of written letters. If you don't already have a secure mastery of the sound system, this can be a very tough job. Many dyslexic children, children who have trouble reading and writing, also turn out to have trouble with sounds as well as with written letters. They can't hear the simple distinctions between *r* and *l* or *b* and *p* that most of us could hear at birth. It can be helpful to such children to artificially make the basic sounds of language more distinct. For example, you can use computer programs that alter speech to exaggerate the sound distinctions. Some recent studies suggest that listening to this kind of altered speech improves dyslexic children's reading and writing.

Some children seem to have genetically determined difficulties with other aspects of language. We've seen how normally developing children must master the system of word endings such as *ing* and *ed*. Some people never master this system at all. They may eventually learn word endings but only painfully and one at a time. If you ask them the wugs question ("This is a wug, what are two of them called?"), they fail hope-

lessly. Their reaction to such linguistic problems is much like the autistic child's reaction to questions about emotion. If you ask them to name a picture of two cats, they may say something like, "I know that . . . s . . . s . . . more than one is s. Cat [pause] s." The rest of us would say "cats" without even thinking about it. In some of these cases we not only know there is a family history of these disorders, we can actually trace the site of the defective genes.

Learning Sounds

No matter how rich the genetic basis for language is, babies still have to disentangle all the particularities of Japanese or English or whatever language they are learning. Children clearly must have some powerful abilities to abstract patterns and discover regularities in the language around them. We know less about these learning mechanisms than about what children learn when. But we have some ideas.

Babies growing up in different language environments will hear very different sounds in the speech around them. Pat estimates that by six months of age, the average American baby has heard hundreds of thousands of instances of the vowel *ee* (as in the words *baby*, *daddy*, *mommy*, *cookie*). On the other hand, that same child will have heard hardly any examples of the nasal sound at the end of the French word *non*. We think that babies abstract the prototypical *ee* sound from all these examples. They unconsciously figure out what the ideal *ee* should be like.

Abstracting these mental prototypes has an enormous impact both on how babies hear speech and on the way they coo and babble. The babies unconsciously compare other incoming sounds to the prototypes. If the sound they hear is anything close to the prototype, they simply ignore the differences and assume it was the prototypical sound. So when

different people, some of whom may have sore throats or colds, speak to a baby, the baby doesn't actually attend to the distorted sounds they make but "smooths them over" and hears what they meant to say. Babies act as though they heard the prototype. Babies get the speech prototypes from the adults they hear around them, but then they turn around and use the prototypes to decode what the adults are saying, even when they don't speak clearly.

So forming prototypes has great benefits. The downside is that these same prototypes prevent the babies from perceiving what a foreign-language speaker is saying. The babies now hear sounds through the filter of the native-language prototypes. And now the babies' own noises start to sound like the sounds of their particular language. By six to twelve months of age, the baby is no longer a citizen of the world but a culture-bound language specialist, like you and me.

To use our earlier example of the picture of the house, many Americans think the prototypical house has a chimney because that's been true of most of the houses they've seen. Once they form this prototype, it influences the very way they remember chimneyless houses. Presumably people growing up in a place where very few houses have chimneys, such as equatorial Africa, would not make the chimney part of their prototype. And Americans might also have a harder time than a native African at discriminating and remembering African houses.

There are other mechanisms at work in learning the sound system of a language. We mentioned that by the time they're a year and a half, babies raised in Chinese homes sound Chinese and those raised in Swedish homes sound Swedish. We saw that that depends on their understanding how the sounds of their language work. But it also depends on being able to imitate and produce those sounds.

Imitating a sound is a lot more complicated than it seems, however. If you just hear a sound, you don't know what to do with your mouth to produce it yourself. When we hear a Swedish vowel sound like *eu*, we don't really know what to do to make it. Should I raise my tongue or lower it? Should I pucker my lips or not? To make it, you raise your tongue as if to produce English *ee* but pucker your lips as if to produce English *oo* and there you have Swedish *eu*. But if we didn't tell you how to do it, and if you don't speak Swedish or French or one of the other languages that use that sound, you would be clueless. How do babies link the sounds they hear others make to the movements they must make to produce those same sounds?

One idea is that when babies are cooing and babbling, they aren't just exercising their vocal cords and moving their mouths randomly. We believe they are creating a kind of mouth-to-sound map, relating the movements of their speech articulators (their lips, tongue, mouth, and jaw) to the sounds they produce. We know babies play with their arms and legs, moving them to and fro and watching in fascination. In much the same way, they also seem to play with their mouths and listen to the sounds they can produce. Babies will lie in their cribs all by themselves and play with sounds, squealing with delight and producing *ee*'s and *aa*'s and *ba*'s and *ga*'s and even just raspberries for long stretches. By playing in this way, they learn how to make the sounds they hear us produce. They learn that to create a sound like *ee*, they have to raise their tongue, whereas for *ah* they have to lower it.

Moreover, babies aren't just able to imitate us, they are driven to do so. Babies love to copy adult sounds. Pat and Andy found that when five-month-olds listen to a simple vowel like *ee* for fifteen minutes in the laboratory, they will coo back with a vowel of their own that resembles the one they heard.

They can't make perfect *ee*'s, but they already have an idea of what to do with their mouths to make a sound that resembles *ee*. They have learned that *ee* is produced when people raise their tongue and retract their lips. Just hearing a grown-up produce the sound motivates babies to try to produce it themselves.

Remember also that in the last chapter we saw that babies at the same age do something akin to lipreading. They prefer to look at the face of a person mouthing a vowel that matches one they are listening to, rather than at a face mouthing a different vowel that doesn't match. This is another sign that they are linking up the sounds they hear with the mouth movements that make those sounds. This combination of abstracting prototypes, playing with sounds, and imitating sounds seems to help children break the speech code.

Learning How to Mean

Why do babies use odd words like *gone* and *uh-oh*, and why do they start to fast-map names? We saw in the last chapter that when babies are about eighteen months old, just when they're learning to talk, they are also learning a great deal about the way objects can appear and disappear, about how they can use tools, and about how objects fit into categories, and they're fascinated by all these problems. Alison suspected that these changes in the way children solve problems might actually be connected to their early words. A fascination with disappearance, and not the rituals of cereal eating, could account for the otherwise mysterious prevalence of *gone*.

How could you test the idea that these weird early words were the result of the problems children were trying to solve? You can visit a baby every few weeks and give him different kinds of problems to solve. At the same time, you keep track of his new words. Alison and Andy found out that babies start

to use *gone* within a week or two of the time they first solve the hardest keys-under-the-washcloth hiding problem, sometimes a little before and sometimes a little after. The word for disappearance and the concept of disappearance seem to emerge together.

Just as *gone* is related to object disappearances, *uh-oh* is related to children's ability to use tools. Remember in the last chapter we talked about how babies learned to use a rake to get a faraway toy. Babies worked out how to use tools like that within a few weeks of the time they started to use words such as *uh-oh*, just as they solved disappearance problems at about the time that they first used *gone*.

Naming turns out to be connected to understanding a rather different aspect of the world. We saw in the last chapter that children learn about how objects fit into categories at about this age. They start to divide mixed-up groups of objects into several different piles, with a different kind of thing in each pile. Children start to do this at about the same time they begin to fast-map new words and to use lots of new names. When we visited babies once every few weeks, we found that just when they suddenly used a lot of new names, they also started to sort mixed-up objects in a new way. Babies "get" the idea that everything has a name and that everything belongs in a category at the same time. So early words often appear at the same time children are solving relevant new problems.

What's going on here? We think it's a bit like college, really. Think about the first time you learned about some new concept in a class like introductory physics. If you were really interested in the course, and not just in getting into medical school, you went to lectures and read the textbook, trying to understand how physics worked. As you studied, you came

across peculiar new words such as *entropy* that at first you understood only vaguely. On the other hand, you could see that they were relevant somehow to the physics problems you were trying to work out. Entropy had something to do with heat loss and something to do with disorder, though you couldn't quite tell what being cold and disorderly had in common (aside from being characteristic of your dorm room, not to mention your boyfriend). Then one day there was a magic moment when everything clicked; you "got" entropy. Part of getting it was really understanding the word and really being able to use it convincingly on the exam, but part of it also was really getting the idea, really getting the concept that brings heat and chaos together and being able to solve problems that require it. ("Briefly define . . ." and "Solve, show your work" both are likely to appear on the exam.) You probably still didn't entirely get it, though, and you undoubtedly used the word in peculiar ways that revealed your continued ignorance to the godlike and omniscient TA's.

Baby Henry's *gone* seems about like the freshman's *entropy*, but without test anxiety interfering. Henry was working on these baffling problems of appearance and disappearance and kept hearing the people around him say *gone* just as some peculiar disappearance took place. One fine day he "got" both the word and the concept of disappearance.

So babies' guesses about the meaning of their first words are informed by the other kinds of cognitive progress they've already made in infancy. Their ability to solve the Language problem is closely tied to the particular ideas they've already developed in solving the External World problem. The mechanisms that drive children to make coherent sense of the world also lead them to pay attention to the words they hear and to learn how to use those words themselves.

"Motherese"

Grown-ups are the third component in the solution to the Language problem. We've mentioned that we sound positively silly when we talk to babies. If you listen to mothers talking to other grown-ups and then to their babies, you hear a strange shift in their voices. A mother says to a friend, "The traffic, it was awful, and I had to park and there was a delay and I didn't have change for the meter . . .," droning on about the events of the day. Then, with hardly a pause, she turns to the baby in her arms and coos, "Hiiiiii, sweeeetie. How's my baaaaby?" She swoops in with her voice and face. "Ooooh [tickling the baby's cheek], open up thoose eeyes. Ooooh, you're sooooo cute. Can Mommy have a big, big smile? Mmmmm, give me big blue eeyes, toooo!"

Anyone listening to a parent talking to a baby knows that this is definitely not a job-interview voice. It's the voice of a playful, animated, warm, and practically giddy person totally absorbed in the little bundle in front of him or her. Out of context it sounds ridiculous. But put us in front of a baby and we all do it, mothers, fathers, grandparents, friends—even four-year-olds speak motherese to their baby brothers and sisters. (We have occasionally heard a macho, deep-voiced dad say, "I never talk baby talk, don't believe in it," and in the next breath turn to the baby and say, in a voice a full octave higher, "Dooo I, sweetie? No, I doon't talk baby talk to YOOOOU.")

And babies love it. When you give babies a choice of what to listen to, a kind of baby Nielsen rating, they choose to listen to mothers talking to infants over mothers talking to other adults. In these tests the babies sit in an infant seat, and slight turns of the head give them eight-second sound bites of either a mother talking to a baby or that same mother talking to

another adult. Babies get to choose which tape to listen to simply by turning their heads in one direction or the other.

The tests show that babies' preferences have nothing to do with the actual words mothers use. Babies choose motherese (or "parentese" or "caretakerese") even when the speaker is talking in a foreign language so infants can't understand the words, or when the words have been filtered out using computer techniques and only the pitch of the voice remains. Apparently they choose motherese not just because it's how their mother talks but because they like the way it sounds. Motherese is a sort of comfort language; it's like aural macaroni and cheese. Even grown-ups like it. Pat's graduate students discovered that listening to the lab tapes of motherese in a foreign language was a wonderful therapy for end-of-term stress. The mother's voice is an acoustic hook for the babies. It captures babies' attention and focuses it on the person who is talking to them.

The elaborate techniques of computer voice analysis reveal exactly what it is we do when we talk to an infant. The pitch of our voice rises dramatically, sometimes by more than an octave; our intonation becomes very melodic and singsongy; and our speech slows down and has exaggerated, lengthened vowels.

Motherese is a universal language. People across all cultures do it when they talk to their infants, even though they usually aren't aware of doing it at all. When mothers listen to recordings of themselves producing motherese, the reaction is: That can't be me. I sound really stupid. Should I be doing that? But they do it intuitively, without conscious awareness.

Why do we do it? Do we produce motherese simply to get the babies' attention? (It certainly does that.) Do we do it

just to convey affection and comfort? Or does motherese have a more focused purpose? It turns out that motherese is more than just a sweet siren song we use to draw our babies to us. Motherese seems to actually help babies solve the Language problem.

Motherese sentences are shorter and simpler than sentences directed at adults. Moreover, grown-ups speaking to babies often repeat the same thing over and over with slight variations. ("You are a pretty girl, aren't you? Aren't you a pretty girl? Pretty, pretty girl.") These characteristics of motherese may help children to figure out the words and grammar of their language.

But the clearest evidence that motherese helps babies learn comes from studies of the sounds of motherese. Recent studies show that the well-formed, elongated consonants and vowels of motherese are particularly clear examples of speech sounds. Mothers and other caregivers are teachers as well as lovers. Completely unconsciously they produce sounds more clearly and pronounce them more accurately when they talk to babies than when they talk to other adults. When mothers say the word *bead* to an adult, it's produced in a fraction of a second and it's a bit sloppy. But when mothers say that same word to their infants, it becomes *beeeeed*, a well-produced, clearly articulated word. This makes it easier for infants to map the sounds we use in language.

In fact, adults may even adapt motherese to suit the characteristics of their particular language. Pat recently discovered that Swedish, Russian, and English mothers each make subtle, unconscious variations in the way they talk that are tailored to the particular language they use. Swedish motherese makes the vowels of Swedish sound much clearer than does ordinary Swedish speech to adults. Similarly, English motherese seems particularly well designed to make the vowels of English sound

clear. The Swedish and English mothers provide the babies with just the range and variety of sounds they need to abstract the right prototypes for each language. This is a particularly important result because it makes it especially likely that babies are taking advantage of motherese to learn the sounds of their language. If motherese were no more than a universally attractive and comforting set of sounds, it might not play much of a role in the details of what babies learn. But in addition it is exquisitely adapted to help babies solve the particular problems of their particular language. That makes us think it must be having a real effect.

Studying babies leads us to realize that, however effortless and instinctive our adult ability to speak may seem, it is actually the outcome of a great deal of learning. There is nothing contradictory about saying this and saying that language also has an important innate component. In fact, the point is not that language is the product of both nature and nurture, innate knowledge and learning. Rather, nature and nurture are inseparably intertwined. The innate endowments enable babies to use their powerful learning mechanisms to take advantage of the information they receive from grown-ups. The fact that babies can already make the right distinctions between sounds at birth enables them to reorganize and reshape those distinctions in light of what they hear their parents say. The fact that babies already organize their world, and are motivated to make sense of it in new ways, also motivates them to learn new words and shapes the meanings they give those words.

Linguists sometimes use the term *bootstrapping* to describe this process. Babies take what they already know and use this as a basis to learn more: they pull themselves up by their own bootstraps. Although language learning is different from scientific-theory formation in many ways, both kinds of learn-

ing involve this sort of bootstrapping: scientists also use their current theory as a basis to formulate new theories. As anyone who has actually helped a young child pull up real boots knows, a few tugs and nudges from grown-ups come in handy, too.

CHAPTER FIVE



What Scientists Have Learned About Children's Minds

So far we've been talking in detail about what children know about a vast array of different topics, from broccoli preferences to toy-car trajectories to the difference between *p* and *b*. We've seen that in three brief years there are enormous changes. What newborns know is very different from what one-year-olds know, which is different in turn from what three-year-olds know. We've also seen that children tackle profound and significant problems. They learn that other people have minds, that the world exists independent of their subjective experience, and that words have meaning. These are *hard* problems. In each case there is a gap—in fact, a yawning chasm—between the data that enters the children's eyes and ears, the light and sound waves, and the conclusions children reach about people and the world and language.

We know more about what children learn than about how they learn it. The mechanisms of learning may be quite different for different problems. We've already seen that understanding how words sound is very different from under-