

particular, indeed unique, abilities in many animals. These differences are not a matter of philosophy, theology, or misplaced humanist sympathies; they are empirical features of nature. We may not be able to take flight by flapping our upper extremities, but we are the only species known that can rationally discuss our inability to do so. As Bertrand Russell famously put it, "A dog cannot relate his autobiography; however eloquently he may bark, he cannot tell you that his parents were honest though poor."

2

Language and Communication



At tea-time, when the dog, Jip, came in, the parrot said to the Doctor,
"See, HE's talking to you."

"Looks to me as though he were scratching his ear," said the Doctor.
"But animals don't always speak with their mouths," said the parrot in a high voice, raising her eyebrows. "They talk with their ears, with their feet, with their tails—with everything. Sometimes they don't WANT to make a noise. Do you see now the way he's twitching up one side of his nose?"

"What's that mean?" asked the Doctor.

"That means, 'Can't you see that it has stopped raining?'" Polynes-

Language and Communication

sia answered. "He is asking you a question. Dogs nearly always use their noses for asking questions."

—*The Story of Doctor Dolittle*

Communication is virtually universal among living things. Even bacteria communicate. Some classes of bacteria secrete distinctive organic molecules, for which they have specialized receptors. This apparatus allows the bacteria to detect the presence of others of the same species, a system known in the literature as quorum sensing. "Bacteria, it turns out, are like bullies who will not fight unless they are backed up by their gang. An attack by a small number of bacteria would only alert the host's immune system to knock them out. So bacteria try to stay under the radar until their numbers are enough to fight the immune system." The molecules secreted by one bacterium serve to communicate its presence to the others. Yet surely not all communication is of a piece with all other communication: the use of the word *talk* in the title of the *New York Times* story about quorum sensing is simply the journalist's effort to be clever.


To determine the true issue here, consider an example. One evening I returned home to find my wife correcting papers for her French class. When I asked her what we were doing for dinner, she said, "I want to go out." That is, she produced a certain sequence of sounds, and as a result I knew that she wanted us to get in the car and drive to a restaurant, where we would have dinner.

When I came home the following night, I found my cat in the kitchen. She looked at me, walked over to an oriental rug in the next room, and began to sharpen her claws on it. She knows I hate that . . . and as I came after her, she ran to the sliding glass door that leads outside. I yelled at her, but my wife said, "Don't get mad; she's just saying, 'I want to go out.'"

We conclude that both my wife and my cat can say "I want to go out." Do we want to assert that they both have language? Surely that is at best an oversimplification, although it is clear that both can communicate. Each can behave in such a way as to convey (somewhat similar) information to me.

Here is a sketch of how "real" communication takes place: One organism has a message in mind that he or she wants to communicate to another organism. He or she emits some behavior (makes a noise, scratches the carpet) that encodes that message. The other organism (me, for example) per-

Language and Communication



ceives the behavior, identifies it in terms of the meaning encoded, and treats the result of that decoding as the meaning of the message.

Sometimes called the Message Model of Communication, this description may seem fairly obvious, but is it a valid general definition of communication? Communication can take place even when there is no evident basis for saying the communicator "intends" to communicate anything. Think of our bacteria above, or a blush, or the visible signs in many species when a female is in estrus and receptive to mating: there is no intention on the part of the signaler, but a message is communicated all the same.

On the other side, it may be that the recipient interprets the message only in part on the basis of its literal content and relies also on various non-overt contextual or social factors. Consider "Can you pass me the salt?" Here the literal content is an inquiry about the listener's physical capacity to perform an action, but the message usually conveyed is a request that the salt indeed be passed. Or perhaps I ask my colleague what she thinks of the candidate we have just interviewed for a job, and she says "He seems very diligent." In an academic context, this implies a very negative recommendation. If a candidate's best quality is diligence, it is *not* creativity, imagination, or inspirational teaching. In both examples, clearly the linguistic content of what we say may be quite different from what we communicate.

The little story about my wife and my cat illustrates the characteristics of any communication system. First, what is the nature of the behavior or other signal? The cat scratches the carpet and runs to the door to convey a message we might interpret as similar to one my wife conveys by moving her vocal organs to produce sound. Second, what is the range of messages the system can convey? Evidently, my cat can say fewer things than my wife: what is the basis of this difference in expressivity? Third, what relation, if any, is there between the message expressed and the communicator's intentions? The cat certainly intends *something*, but her behavior

Language and Communication

actually reflects her internal state: my wife can say what she does even if she doesn't really want to go out. Finally, what is involved on the receiving end? Obviously, you have to know the code in order to get the message, but what else? My wife and I understand the cat's scratching behavior as attention seeking in the context of my evident and constant displeasure at it, but is there some kind of underlying code that all three of us share?

Another important aspect of communication systems (not significant in this case) is how the communication system came into being. Did it evolve gradually out of something else, or did it spring into operation fully formed? My cat scratches the carpet basically to sharpen her claws; whatever additional meaning may accrue to that action has grown up ad hoc between us. Most systematic means of communication have more interesting and far longer histories.

This is an area of inquiry where the questions that can be raised are potentially more interesting than the answers currently available. Historical evidence for the sounds of language is minimal; even the soft tissue of tongues, ears, and brains leaves no trace in the fossil record.

The original nineteenth-century constitution of the Société linguistique de Paris is famous for explicitly prohibiting the discussion of matters concerning the origin of language at the society's meetings. This was no mere quirk of the founders: they introduced this limitation for precisely the reason that there could apparently be no real science that bore on the topic. Since the late 1990s, interest among linguists and others has reawakened, and conferences are now regularly devoted to the subject. To my mind, this revival is not based on additional data, but rather on the mistaken impression that if we can pose an important question, we ought in principle to be able to find an answer. Fortunately, we need not resolve this vexatious problem before studying communication systems and communicative abilities *comparatively* across animal species. We will return to these matters in Chapter 11.

Notions of Language and Communication

How might we distinguish between "language" and "communication"? One way of approaching the distinction is to note that communication is something we *do*, whereas language is a *tool* we can use. We can, of course, communicate without language, though the range of material we can transmit is limited in significant ways. Most of the amusement value of the game of charades, for instance, lies in trying to circumvent these limitations. In fact,

Language and Communication

a desirable skill in this game consists in referring to words without actually using them (using gestures interpreted as "short word," "sounds like," and so on).

For comparison, the activity of building houses is also something we do, and we use particular tools to do it. Without hammers, nails, saws, and levels, we could not practice the construction trade as we know it. Yet that does not mean we could not construct shelters. We can do a certain amount of building without tools, or using different tools, as other societies do. Still, the structure of the tools makes certain sorts of construction easy and natural. We can study the structure of the hammers and saws and ask where they come from. We see, of course, that there is a close connection between the structure of the tools and what we can do with them, but we should not confuse the activity of carpentry or construction with the tools we use in pursuing it.

Suppose we want to open a nut. We do it by exerting force on the shell through a hard object—either with leverage, using a nutcracker, or by hitting it, for instance with a hammer. Chimpanzees in the wild open nuts by putting them on one rock, then hitting them with another rock—a technique similar to one used by humans. The tools are not identical, but they have the same structure in the relevant respects. There is an activity, and similar means are used in carrying it out. As far as communication is concerned, we do a lot with facial expressions, grunts, and the like. Again, considerable similarity among human and nonhuman primates exists in the activity and in the means for executing it.

Orangutans in nature do not use tools equivalent to those of human carpentry. But if we give an orangutan a claw hammer, and he knows that something good to eat is inside a wooden box that is nailed shut, he can use the claw hammer to remove the nails and open the box, much as a human would. Provide him with the tool, and his cognitive abilities are certainly adequate for using it in some of the ways humans do—ways that depend on the essential structure of the tool.

I imagine that chimpanzees can learn fairly quickly to open nuts with a nutcracker by utilizing the structure of the tool, which is novel to them but suited in form to the task. Yet if we give a chimpanzee a small tape recorder, I seriously doubt that the ape could use it to record grunts and send them to be replayed for another chimpanzee in order to communicate a message. The principal use of a tape recorder might be to serve as the base on which to put a nut in order to smash it.

Language and Communication

These distinctions are important when asking whether another species (say, monkeys or apes) can use language. Provided with the proper tools, an ape can use them to engage in at least some "carpentry." What about language and communication? When we ask whether animals other than humans can engage in communication, the answer is, obviously. What is the structure of the means they use to that end, and how closely does their communicative activity resemble human natural language? If we supplied an ape with a human natural language, how much communication could he or she achieve? We need to know a certain amount about the structure of human natural language if we are to make these questions precise; the more we know, the more precise we can be.

In nature, the range of ways in which animals (especially other primates) communicate with one another is certainly not limited to vocalization. Smell (particular substances, such as those secreted by specialized scent glands in both lemurs and rhesus monkeys, as well as normal smells), sight (facial expression, posture), and touching (grooming behavior), among other modalities, also supply information, sometimes intentionally on the part of the communicator and sometimes not.

In terms of the structure of the tools involved, none of these systems seem to fall within a range that might usefully be compared to language. Signals in media such as smell and touch typically are individually simple (that is, they lack a relevant internal structure such that parts of the signal correspond to distinct parts of the message), and in some cases (especially olfactory communication) they are not very flexible in their temporal pattern. There are exceptions: the chemical signals produced and perceived by lemurs may include substances from multiple individuals, deposited at different times; the animals are apparently sensitive to this complexity. Chemical signals in the insect world can be even more complicated. But even where some internal organization is present in the signal, these systems appear to be rather different from human languages.

Characteristics of Language

Now when I speak of "talk" between animals and myself, you who read this must understand that I do not always mean the usual kind of talk between persons. Animal "talk" is very different. For instance, you don't only use the mouth for speaking. Dogs use the tail, twitches of the nose, movements of the ears, heavy breathing—all sorts of

Language and Communication

things—to make one another understand what they want. Of course, the Doctor and I had no tails of our own to swing around. So we used the tails of our coats instead. Dogs are very clever; they quickly caught on to what a man meant to say when he wagged his coat-tail.

—*Doctor Dolittle and the Secret Lake*

What are the essential properties of human language, and how does it differ from other communications? On the face of it, it ought to be possible to generate criteria that would make the difference clear. It turns out, though, that this task is harder than it appears.

One well-known attempt to specify just what properties define a "language" in the human sense was made by the linguist Charles Hockett in the 1950s. Some of Hockett's Design Features of Language may also be found in nonhuman communication systems, but he argued that the whole set is found together only in human languages. Ultimately, the effort to define language in this way poses more problems than it solves, but at least it provides a basis for discussion.

Vocal-Auditory Channel

Language is expressed and perceived in sound. This property is not unique to human language: many species use auditory signals to communicate. Hockett suggests, though, that the particular acoustic spectral features that differentiate messages in human languages really are unique. Vowel color (the property that distinguishes the vowels of *bet*, *bat*, *boat* from one another), for instance, is a characteristic of acoustic signals that does not appear to be exploited by the system of any other animal. That is not just coincidental: the development of a vocal tract capable of making distinctions of vowel color is one of the physical specializations for speech that appeared in the course of human evolution. This point is developed at some length in various works by Philip Lieberman and his colleagues.

Just as humans are not the exclusive users of sound for communicative purposes, the vocal-auditory channel is not the only one in which communication occurs. Many others are used as well: signals can be visual, tactile, olfactory, chemical, electrical, . . . In fact, humans communicate with one another (intentionally or not) in most of these ways, though we do not usually confuse that communication with language.

In part because of this diversity, it is vital to distinguish one system from another, so that we speak of a single species as employing multiple

communication systems rather than just as "communicating." Each system has its own internal coherence, which can be studied independently of the other systems. Sometimes a single behavior may involve multiple systems in complementary ways. Think of the role played by facial expression in understanding the messages conveyed by accompanying language. The messages conveyed by facial expressions can be explored systematically in one way; language in another. The totality of the communication results from both taken together, but nothing is gained (and coherence is lost) if we attempt to study both at the same time by the same methods.

Indeed, the same channel can convey information from more than one system at the same time. The pitch of the voice on individual syllables serves in many languages (such as those of China and most of the languages of Africa, known as *tone languages*) to distinguish words from one another, in the same way that the difference between one vowel or consonant and another serves this function. Voice pitch is also an aspect of the expressive system of paralinguage (about which I will say more in Chapter 4), which conveys a variety of information about our emotional state, attitude, and so on. The fact that a speaker of Cantonese distinguishes words by tone does not alter the fact that other aspects of the overall contour of that speaker's voice quality and pitch serve simultaneously as paralinguistic cues to excitement or boredom, contempt or admiration.

Separating language from paralinguage is critical to achieving a coherent understanding of the way both systems work. Every time we say something, we communicate much more (or sometimes less) than the literal content. Paralinguistic features (pitch range, loudness, breathiness) have very different properties from sentence structure and meaning. The way pitch is used paralinguistically is inseparable from the way it is used linguistically, even though the two are quite distinct logically.

Hockett took it as self-evident that true language is executed in the vocal-auditory medium. By this he meant to distinguish spoken language especially from writing, which he regarded (correctly) as secondary and parasitic on the spoken language. In fact, not all language (even disregarding writing) *does* involve the vocal-auditory channel. The signed languages used and acquired natively among hearing-impaired individuals have the structural properties of a language such as Chinese or Kiswahili, although they involve the visual channel. Understanding of the richness of the structure of signed languages, and their basic similarities to spoken languages,

did not really develop (among linguists, at least) until after Hockett's paper appeared. Chapter 9 is devoted to the properties of these languages.

Broadcast Transmission, Directional Reception

Signals travel generally to any potential receiver, and their properties can help to determine the location of the originating source. This characteristic seems at first glance to apply to just about any communication system, if we disregard the fact that communication can take place even when we cannot locate the source. Think of a disembodied voice backstage in a play, for instance. Of course, in some cases the nature of the medium makes broadcast transmission rather narrow. One example is tactile signals, such as tapping a dancer on the shoulder to indicate a desire to cut in. But not all communication shares this property even in a limited way. Consider the marking of an animal's territory by olfactory signals. At the time the communication takes place—when another animal perceives the scent—the source is not necessarily present, but the signal nevertheless plays its role perfectly well.

Rapid Fading

Many kinds of communication are transitory, in the sense that the signal is not available for inspection for very long after it is produced. Even if you are in a cave with a remarkably persistent echo, the sound fades away within a few seconds. The same is true of signed language or any system of visually perceived gestures; there is not even any obvious analogue of an echo. For logical purposes we can disregard the modern possibility of recording sound or images for later playback, as well as that of writing down what we have heard. Under those circumstances we could consult the transcribed record at leisure, but these special cases are in no way intrinsic to the way speech (or sign) communication works. In this way the modalities of natural languages (both spoken and signed) differ from those of chemical or olfactory signals, or from outwardly visible physical changes that communicate one animal's internal state to another.

Interchangeability

Competent language users both produce and comprehend the same range of signals, at least under normal conditions and barring pathology (such as deafness or blindness). In this respect human language is different from some other systems. Birdsong, for instance, is typically produced by the

male and not by the female (although there are species in which both sexes sing) and comprehended by male and female (but in different ways). The honeybee dance is interchangeable among the workers in that these bees both dance and understand the dances of others. The situation is different for queens and drones, who do not dance but do understand the dances—at least to the extent that these indicate possible sites for a new hive and not the location of a food source.

This connection between production and perception of the signals may seem adventurous, but it turns out to be rather important, at least according to some theories. The motor theory of speech perception claims that the way we perceive the speech of others involves a direct reference to what we might have done ourselves. If our perceptual system is truly organized in this way, the fact that hearers are also talkers (and vice versa) is no accident. Humans and songbirds appear to provide strong evidence for this claim.

Total Feedback

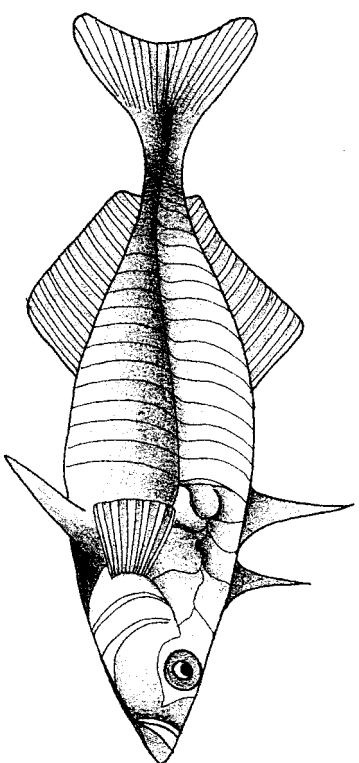
Senders can monitor their own signals. Some nonlinguistic signals do not have this property: consider physical changes in an organ outside the range of the individual's vision but visible to others, such as a blush. Blushing certainly conveys information to the viewer, but without a mirror not to the one who blushes.

Some famous instances of communication do not allow for feedback. The stickleback (a fish that was an object of great interest to ethologists in the 1950s) communicates about mating through a characteristic change of color in the belly and eyes of the male and a characteristic distension of the female's belly. But neither can see his or her own belly or eyes.

With respect to human languages, we might feel that feedback, while usually present, is not necessary. Even in the case of deaf speakers or blind signers, though, there is feedback from other modalities (kinesthetic, proprioceptive). Feedback of this sort seems to be significant in learning; both birds and humans (the only well-studied cases of learned communicative behavior) get seriously off the track when it is not available. In normal speech, laboratory conditions that disrupt, distort, or prevent proper feedback can make fluent speech virtually impossible to produce.

Specialization

Hockett's definition of this property refers to the fact that the communicative activity involved in human language does not serve any other func-



tion: "the direct energetic consequences of an act of communication serve no other biological purpose." Compare this concept, for instance, with the fact that we can derive information from events that *do* have other functions, perhaps more basic than communication. When a dog pants with his tongue hanging out, he is cooling off through evaporation, but he may also be supplying information (especially to other dogs) about the location, state, and identity of the panter. The female stickleback's distended belly, which communicates to the male her readiness to breed, is the result of the development of roe, not of any intent to communicate.

Specialization for communication is tied up with the range of uses for the organs involved. It is sometimes claimed that there are no true speech organs: the organs we use for speech production all have other functions (vegetative, respiratory), and the ears that perceive much besides speech are what we use to hear. The notion that this multifunctionality excludes a specialization for speech overlooks a great deal. In fact, the vocal organs have changed over the course of evolution in the direction of greater functionality in speaking, even at the cost of being less suited to their other tasks.

For example, the natural position of the human larynx is considerably lower in the throat than in other primates, or in mankind's earlier ancestors, or indeed in newborn babies. Among other differences, the tongue is large and rounded, as opposed to the short, flat tongues of other primates; and the vocal tract makes a 90-degree turn, as opposed to the nearly straight vocal tracts of others. Because of its construction, the human vocal tract has a portion that is necessarily involved in the transfer of food and drink, on the one hand, and of air on the other. In our primate relatives (as well as in babies) it is possible to isolate the digestive channel from the respiratory, making it possible to eat or drink and breathe at the same time—but

we all know what happens when *we* try that. Evolution, in other words, has modified a nice, serviceable system so as to make it possible for us to choke on our food.

Nonetheless, the resulting system is much more flexible than that of other primates in terms of the range of sound types we can produce. Many basic varieties of vowels and consonants are beyond the articulatory capacity of nonhumans, or of earlier hominids. Evolution has specialized us as speakers. As eaters and drinkers, we simply have to make the best of it.

The perceptual system, as well, seems to have a specialized mode of operation that applies to auditory inputs that have the overall properties of speech. This mode is quite distinct from the one that comes into play in perceiving nonspeech. Under unusual conditions—thoroughly unnatural, but neither impossible nor painful—it is possible to engage both systems with respect to the same stimulus. In the laboratory phenomenon of “duplex perception,” we seem to hear both a speech signal (a syllable, such as *ka*) and a nonspeech signal (a sort of falling or rising pitch whistle) in response to the same sound input, when parts of that input are provided to one ear and parts to the other ear. This phenomenon confirms the notion that the human auditory system has indeed evolved a distinctive specialization for dealing with speech, even if we use the same physical ears to hear both the announcer and the crack of the bat when we listen to a baseball game. The idea that there is no speech apparatus *per se* turns out to be a misconception.

Semanticity

Linguistic forms have denotations: That is, they are associated with features of the world, as opposed to many nonlinguistic signals that refer only to themselves (think once more of the stickleback’s belly). I deal further with this issue in Chapter 7, in connection with the meaning of alarm calling behavior. In understanding the workings of language, we want to distinguish semantic signals (which refer to events and objects in the world outside of the signaler) from expressive ones (which simply reflect to the outside world some aspect of the internal state of the subject).

Arbitrariness

It is conventional to observe that the linguistic signal has no necessary relation to what it denotes. Speech signals, that is, are not in general iconic. *Cat* refers in my speech to instances of *Felis domesticus* not because of some

perceived resemblance between the sound of the word and some aspect of a real cat, but merely because that is the English word for it. The arbitrariness is reinforced when we observe that other languages have quite different words for the same thing. In Navajo, for example, a cat is a *ma’i*, but the cats themselves are just the same.

Arbitrariness is often thought to be falsified in the case of onomatopoeia: thus, a cat says “meow” because . . . well, because that is the noise a cat makes. In fact, though, different languages have at least partially conventionalized onomatopoeic words for animal noises. Cats say “ngeong” in Indonesian, for example. A rooster says “cock-a-doodle-doo” in English, but “cocorico” in French or “kikiriki” in German. A turkey says “gobble, gobble” in English but “glu, glu” in Turkish. A pig says “oink” in English but “groin groin” in French, “röh röh” in Finnish, “chrum chrum” in Polish, “nöff” in Swedish, or “soch, soch” in Welsh. Although these words generally are inspired by sounds made by the animals in question, they are nonetheless words of particular languages, and with very few exceptions they conform to the principles of words in those languages. A pig could not say “groin groin” or “röh röh” in English, because English does not have the nasal vowel [ɛ̃] of the French word or the front rounded [ö] of Finnish. English cats could not mimic their Indonesian counterparts because English words cannot begin with [ŋ] (*ng*), and so on.

As opposed to the words of spoken languages, paralinguistic vocal features are less arbitrary, in that their dimensions tend to be related iconically to those of the internal states they express. Thus, when we are angry, our voice may get loud. When we are angrier, it gets LOUDER—and when we are extremely angry, EXTREMELY LOUD. The dimension of loudness can vary in a continuous way, showing (in principle, at least) as many degrees as does our potential anger or other internal state to which the loudness corresponds. This continuous and iconic character is one of the basic ways in which paralinguage differs from language.

Even apparently transparent iconic communication may have some arbitrariness, though, in the sense that it may have to be acquired in order to be understood. Thus, we take the gesture of pointing for granted as a way to call another’s attention to something, but not all cultures use similar gestures in this way.

A story (probably apocryphal) that I heard in an undergraduate class illustrates this point. A missionary is dropped into the jungle and tries to learn the language of the surrounding community. Eventually she learns

how to express "What's that?" and sets out to expand her vocabulary. She points to a house and asks "What's that?" and hears "Boogoo-boogoo," so she writes in her notebook: *house*: [bugubugnu]. Then she points to a tree and asks "What's that?" and again hears "Boogoo-boogoo." She decides she must have been wrong the first time, and that [bugubugnu] means *wood*, not *house*. But she points to a passing dog and asks "What's that?" to which the response is, once again, "Boogoo-boogoo." Eventually she learns that [bugubugnu] actually means *right index finger*: In the local culture pointing is done with the chin, and every time she asked "What's that?" her position had been such that her chin was directed toward the pointing finger.

Although various nonhuman primates assuredly have a sense of drawing attention to an object, most do not understand finger-pointing gestures as the way to do so, at least not without extensive training.

Discreteness

The linguistic signal is subdividable into separate units (sounds, syllables, words, phrases . . .), and relatively small inventories of these basic elements can be combined in various ways to generate a much greater variety of messages. As Steven Pinker puts it in *The Language Instinct*, and as I discuss in greater detail in Chapter 8, a language is a *discrete combinatorial system*. Other communication systems have a kind of discreteness—birdsongs are made up of parts, for example—but it is not the same. Some birds can do a certain amount of recombining of the basic elements (analogous to syllables) of their song, but the result is always a variant expression of the same message. The key here is that birds cannot combine syllables in different ways to produce substantively different messages.

The signs used by baseball coaches and managers constitute a different system that also displays discreteness. This system can express a broader array of messages than can birdsongs, but it still lacks the meaningful re-combinability we find in language. The coach touches the letters on his uniform twice, then spits (baseball players and coaches do a great deal of spitting), then tickles his right ear, then pulls the lobe . . . and this means *don't swing*. Other combinations mean *hit and run* and the like. The system is based on a set of discrete elements that can be combined in different ways, but each combination is a single unit: that is, there is not some part of the message *hit and run* that is associated, say, with the spitting. In language, on the other hand, the word "hit" in (spoken) "hit and run" is associated with a specific subpart of the total meaning.

This difference is sometimes referred to as that between "phonological" syntax and "semantic" syntax. In a system with phonological syntax, the individual signals have an internal structure and are made up of component parts that can be combined in various ways for different signals or variations on the same signal. The parts themselves do not make discrete contributions to the signal's meaning, however. To the extent that we can find discrete components within internally complex signals, a number of animal and other communication systems can be said to display phonological syntax, but only a system like that of English has semantic syntax.

Animal systems are either discrete or continuous. If they are discrete, they are made up of a small number of possible signals (on the order of five to fifty) that are not semantically recombinable. If they are continuous, different messages correspond to different values on some dimension. The notion of "continuous" here comes from the mathematical sense of the word. It refers to a physical scale (such as direction or distance) with the property that for any two values, there is always (at least in principle) another possible value intermediate between them. The bee dances described in Chapter 4 are examples of a continuous communication system.

Displacement

With language, we can refer to objects and events that are distant in space and time from the location of the speaker or the hearer. Other signaling systems do not in general have this property. To the extent that it makes sense to describe animals' signals as "referring" to something, it is always to the here and now—the attitude or the internal state of the animal doing the signaling.

Even rather rich systems devised and used by humans share this limitation, to the degree that they are not basically parasitic on language itself. A baseball coach may have a signal for *hit and run*, but there is none for *if we're still ahead in the seventh inning, I'll be able to take you out for a pinch hitter*. Bee dances perhaps are an exception, if we think of the bee as "describing" the properties of a distant food source to her fellow workers. Still, it may also make sense to think of this system as reflecting the bee's own internal state, a state that results (here and now) from the foraging flight she has just undertaken. If we think of the hive as both the location of the dance and the origin of the flight vector it indicates, the putative spatial displacement is less evident. In any event, there is no question of a temporally displaced referent: Bees' dances relate to food sources available within a very short

temporal horizon from the present, not to where they found a nice patch of hollyhocks last season.

Openness and Productivity

An open or productive system is one that is capable of expressing an unbounded range of possible meanings. Most animal communication systems serve to convey at most a few dozen different possible messages. Once again, bee dances are a possible exception. Since the parameters of the dance can potentially distinguish a continuous range of possible food source locations, it follows that the number of distinct "messages" is unlimited, at least in principle.

If we ignore the point that the bees themselves may not be infinitely precise in producing and interpreting these dances, there can thus be an infinite number of dances. This kind of productivity is completely different from what we find in human language. Even on the most charitable interpretation, the bees are always "talking about" the same thing (however many subtly distinct variants there may be), whereas the variety of things humans can refer to when talking is not limited in that way. This difference in the productivity of communication systems requires us to distinguish *continuous* openness—as illustrated by bee dances—and *discrete* openness of the sort we find in natural language.

Duality of Patterning

In every language, units can combine to make larger structures at two quite distinct levels of analysis. Sounds combine to make words, the smallest meaningful units, and words combine to make sentences or whole propositions.

This duality makes for a certain efficiency of language, in that there are not very many different units that have to be kept apart in production and perception. If we needed to learn completely different signals for every word of our language, we would quickly reach the limits of our acoustic and auditory abilities. Compare alphabetic with ideographic writing systems, for example. A child learning to read and write a Chinese language, or the kanji characters necessary to read a Japanese newspaper, can testify to the immense burden of having to memorize all the individual signs separately, as opposed to a child learning the twenty-six signs of the roman alphabet. We avoid that problem by having only a small number of basic sound types (each with no meaning in itself) and combining them in an unlimited num-

ber of ways to make larger structures—to which we can assign individual meanings, and then combine *those* into more complex expressions.

Traditional Transmission

A person's language is *learned* rather than completely built in. For example, a child of any genetic background will learn whatever language is spoken in the surrounding community. Certain birdsongs (in three out of twenty-seven orders of birds) are also learned, and possibly also some cetacean vocalizations. In contrast, the (nonsong) calls of birds, primates, and other animals, as well as the broad scope of their applicability, are not learned. As far as we know, all other communicative behavior on the part of nonhumans is innate.

Although most communicative behavior thus is *not* learned, it is not necessarily independent of experience. For instance, some tuning of the appropriate conditions for use of vervet monkey alarm calls apparently takes place during growth and development, although the signals themselves need not be learned. This pattern is characteristic of a great many systems. It is responsible, for instance, for the development of numerous local "dialects": a given vocalization may have a range of possible realizations rather than just one, or a range of possible uses. Individuals may attune their choice from within such a range to the usage of those around them, even though the basic system develops in them without reference to the behavior of others. The claim of innateness in many communicative systems is thus subject to the qualification that the ways in which such behavior is used may be modified somewhat over the life span of the individual.

Prevarication

Some theorists of language origins, fond of paradoxes, have said that language "must have been invented for the purpose of lying." Charitably, we can interpret this statement as emphasizing that language can describe things that are not literally realized or true. We can talk about unicorns and squared circles, even if we cannot ever point them out. We can also use language to lie more literally, referring to states of affairs that are contrary to what we know about the world—doing so not just to exercise our theoretical imaginations, but to actively mislead our listener. Yet insofar as a communicative signal is simply an external manifestation of an animal's internal state, it is not really possible for the animal to "lie."

Some reports suggest that animals other than humans use supposedly

communicative behavior to deceive. One celebrated case is that of the (mother) piping plover (*Charadrius melodus*), who apparently pretends to be hurt to distract predators from the nesting site where her relatively helpless offspring would be endangered. But is this undeniably effective strategy really an instance of intentional misleading? The plover's behavior is not just a reflex: the bird clearly tries to lead the intruder away from the nest. I discuss the interpretation of this case in Chapter 5; to anticipate the conclusion, there is no reason to believe that the bird is lying so much as engaging in behavior that she knows will attract a predator away from her nest.

Verret monkeys sometimes behave in ways that could be seen as an attempt to deceive their fellows about the presence of predators, and Dorothy Cheney and Robert Seyfarth have explored ways to disentangle intended deception from other interpretations. There is also a substantial (and controversial) literature on apparently deceptive behavior in primates, under the heading "Machiavellian intelligence." A substantial corpus of incidents has accumulated, but the evidence remains at the level of intriguing anecdotes rather than systematic patterns of behavior.

To say that some communication is genuinely deceptive, we would want to establish that the communicator has some sense that the recipient of the message has a view of the world, and that the communicator is attempting to manipulate that view (rather than directly manipulating the behavior itself). It is a thorny issue, and one that has been much discussed. Do any animals other than humans have a *theory of mind*, in the sense that they see other individuals not merely as acting but as holding opinions that underlie their actions? There seems to be no valid evidence for this claim in any species, and some evidence in higher primates that argues against it. For current purposes it suffices to mention that this philosophical question is relevant to the notion of deceptive communication.

The fact that we can use language to talk about things that are not true, or not possible, or simply imaginary, is qualitatively quite different from this point. We can also use language to lie and deceive, but it is difficult to see that as its principal role in our lives.

Reflexiveness

We can use language to talk about language itself. This is a property of human language, and indeed of no other communication system. Birdsong cannot be used to make a comment on another birdsong, for instance, but

only to put out a message. The significance of this ability is far from clear, however. Yes, it is possible to write a book such as the present one about the nature of language, but even its author has to admit that this is a fairly marginal activity. To consider it one of the defining properties of human language puts the cart before the horse.

Learnability

Any normal human can learn any one of a variety of languages, depending on the data available in infancy when the person's first language is being formed (or when the language organ is growing). This point seems essentially the same as the one made above about the "traditional transmission" of language, but a subtle difference exists.

When we say that something is learned, that statement can be interpreted in a number of ways. At a minimum, it means that the requisite knowledge or skill developed on the basis of some interaction with the environment. But that is not enough, because the environment alone cannot suffice to explain what is learned. Information in the environment can only lead to learning if the learner is constructed so as to be able to make use of it. And an organism's inherent structure can interact in several different ways with environmental information in order to result in the development of some ability we would call learned.

What we usually have in mind when we talk about learning is what goes on in school: information about facts or methods is presented by a teacher, and as a result the learners develop a more or less accurate representation of the same knowledge or skill. Whatever is learned is, obviously, a direct function of what was taught. Much can be learned in that way, but it is necessary to emphasize that the learning is based on a general intelligence that is quite broadly applicable. And that is not the only kind of learning that we need to recognize.

Paradoxical as it may seem, learning in some instances must be said to take place even though only one possible thing will be learned; different experiences all lead to the same result. All that really matters is that there be *some* relevant experience to trigger learning.

This type of learning is what appears to take place in the lower-level visual system, for example. The visual system is not completely formed at birth. The neural organization that allows for the recognition of specific features develops as neurons form specific synaptic connections with

one another over time, as the organism matures. The particular features that emerge in the visual system seem to be essentially uniform across the species, yet the apparatus for recognition can develop only in the presence of visual data. When crucial features of the visual environment are absent from early experience, the system does not develop as it should—even though there is really only one course it can follow. This fact was demonstrated some years ago in classic experiments by David Hubel and Torsten Wiesel on kittens, and much the same point emerges from more recent work on the maturation of the human visual system.

It makes sense to say that we "learn" how to see, then, even though no traditional transmission is involved. In contrast, we apparently do not need to "learn" how to taste. The neural apparatus for detection of the basic taste sensations is completely in place at birth. My colleague Linda Bartoshuk, who works in this area, points out that most of what we call taste is actually a matter of *smell*, or at least of detection and analysis by the olfactory system. Smell, as opposed to taste in the strict sense, is a faculty that involves a lot of learning—and "traditional transmission." We do not need experience to detect sweet, sour, and bitter sensations, but we certainly need to be taught to tell a Bordeaux from a Burgundy.

In the case of language, environmental input is needed in order for the language faculty to develop, although it turns out that, in extremis, infants will seize on remarkably little information as the basis for linguistic maturation. The nature of the system that emerges is strikingly uniform in remarkable ways, and in that sense it is similar to vision; but variations are also present. Different languages have different words, but there is much more than that to say about the range of potential variability. The specific properties of the individual language learned must be fixed in ways contingent on the available evidence.

The learning of language (for a human, at least) is thus intermediate between the kinds of learning discussed above. It is certainly not like the development of the visual system, in that only one possible kind of knowledge can emerge. The specific language a child learns depends on the language spoken in the environment during childhood, not on that spoken by the parents. It is not like learning calculus at school either, because what is learned is not solely a product of what is heard. Human language learning involves a major contribution from the learner. It is more a matter of selecting from a large but limited range of possibilities than of developing an ability from scratch.

Problems with Lists

If we take a set of characteristics such as those just reviewed (perhaps massaged a bit, to allow for possibilities such as signed languages, which Hockett's discussion did not accommodate), we can probably come up with a list of tests with the desired property. All human languages "pass" these tests, and nothing that is not a human language can pass them all. Such lists have tended to define much of the content of current animal communication research. That is, researchers want to establish that such and such an organism either "has language" or at least "has the cognitive ability to acquire language." Their approach is to take a list of characteristics and show that their animal does indeed exhibit all of them, or at least that they can teach it enough to "pass all the tests." Gary Bradshaw calls this the Signature Characteristic Strategy.

Such a strategy raises a number of conceptual problems, as Bradshaw points out. For one thing, any specific set of signature characteristics is liable to need constant revision. For instance, now that linguists understand that manual languages such as ASL have all the structure of other natural languages, we have to rethink the role of sensory modality in the basic nature of language. In general, the tests suggested by Hockett reflect the understanding linguists had of these matters in the 1950s. Much has changed since then—presumably for the better, at least in terms of our understanding of what a language is. Although the cognitive abilities and communication systems of some animals might well "pass" at a given stage in the development of science, the same animal might be said to "fail" one or more tests as the actual battery evolves in accordance with changes in what we know about language.

Researchers exploring the cognitive capacities of other species, especially those involved in ape language, have often complained about what seems to them to be a double standard. Sue Savage-Rumbaugh, a prominent figure in this field, objects that "they keep raising the bar. First the linguists said we had to get our animals to use signs in a symbolic way if we wanted to say they learned language. OK, we did that, and then they said 'No, that's not language, because you don't have syntax.' So we proved our apes could produce some combinations of signs, but the linguists said that wasn't enough syntax, or the right syntax. They'll never agree that we've done enough."

It is certainly true that scientific notions of what constitutes the essence

of language have evolved over the years. Contrary to Savage-Rumbaugh's impression, these changes have not been arbitrary and capricious, designed to ensure that nonhuman animals will never be able to succeed. In particular, the centrality of syntactic organization to the expressive power of human language has become much more obvious than it once was.

It is fair to say that in the 1950s, linguists had little understanding of the syntax of human natural languages. Research focused almost exclusively on the systematic role of sound in language and on the structure of words. Only with the development of generative theories of grammar in the 1960s and thereafter have we come to a better appreciation of syntactic structure. That structure is amazingly intricate, and remarkably uniform across languages despite superficial differences. It is undoubtedly based in the biologically determined maturation of human beings, and it contributes in essential ways to linguistic expressivity. If we place more emphasis today on the role of syntax in the evaluation of ape language experiments, it is because we understand language better, not because we have found that this is a way to deny it to chimpanzees, bonobos, and others.

What is essential is recognizing that the evaluation of cognitive research on other species is not some sort of computer game, in which passing a certain number of tests means that you "win." We are trying to understand what animals do and what they do not do, what they can do and what they cannot do. Is it true that they really can learn all of what seem to be the essential structural properties of a human language? Most of these? Only a few? None at all? Whatever the answer to these questions, it constitutes a result, and one in which linguists are interested. If it turns out that either naturally occurring communicative behavior or abilities induced in laboratory experiments indeed have the characteristic properties of a human language, linguists in general will be fascinated, not repelled, by that fact.

Even if it were possible to agree on a more or less definitive set of signature characteristics of language, it might still be difficult to measure the extent to which an organism does or does not exhibit those characteristics. The more what we are looking at diverges from language in its essence, the harder it is to develop meaningful comparisons in matters of detail.

For example, the utterances of young children become progressively longer as they mature. The measure of Mean Length of Utterance (MLU) is presumed to reflect the emerging complexity of the child's linguistic behavior, and it is common to assume that MLU actually measures linguistic maturity. In the ape language experiments, however, the number of signs

produced in an "utterance" seems to have quite a different basis; their multiplicity reflects repetition, rather than elaboration of a basic idea. If we assess the linguistic maturity of chimpanzees by measuring MLU, we get the impression that they reach a very high level quite quickly—then get stuck. In actuality, though, MLU measures rather different things in baby humans and baby apes.

If our goal is to understand what a natural language actually is at its core, any checklist of the sort we have been exploring is almost certain to be superficial. A set of criteria that picks out a given class within the world as we know it may not accurately express its essence. Consider the definition of humans as "featherless bipeds." Who would be content with that as an expression of the nature of humanness (or of primateness, if you think of chimpanzees and gorillas as bipeds)?

To understand how human languages are (and are not) related to other forms of communication among animals, we need to do more than devise a set of tests. Chapter 3 explores some of the pitfalls that arise in understanding and appreciating what we find when we study the cognitive capacities of various animals.