Chapter 1
First Words

It looks simple. A 14-month-old toddles after the family dog, smacking it whenever she gets close. The dog wearily moves under the table. "Dog," the child's mother tells her. "You're chasing the dog. That's the dog." The child stops, points a pudgy hand at the dog, and shrieks, "Daw!" The mother smiles: "Yes, dog."

Many parents—and many philosophers and psychologists—would say that word learning is as simple as it looks. It can be explained in part by the processes of association and imitation and in part by the efforts of parents who want their children to learn how to speak. A child starts by listening to her parents use words and comes to associate the words with what they refer to. When she starts to use words herself, her successful acts of naming are rewarded, and her mistakes are gently corrected.

From this perspective, word learning is the easiest part of language development. The rest of language emerges without the support of "negative evidence"; children do not receive consistent feedback on the grammaticality of what they say (Brown & Hanlon, 1970; Marcus, 1993). But word learning may be a different story. While parents tend to be unconcerned if their child says "goed" instead of "went," they are likely to notice, and react, if their child was to use *dog* to refer to a chair. Another difference is that much of language is productive. An understanding of syntax, for instance, allows us to produce and understand a potential infinity of new sentences. But word learning is merely the memorization of a series of paired associates: *dog* refers to dogs; *water* refers to water, *Mommy* refers to Mommy, and so on.

This is one picture of word learning. This book presents another. I will argue that a careful consideration of what children know and how they come to know it reveals that word learning is actually far from simple. Children's learning of words, even the simplest names for things, requires rich mental capacities—conceptual, social, and linguistic—that interact in complicated ways.

John Macnamara defends this alternative in the first paragraph of his 1982 book *Names for Things.* He remarks that the learning of simple names

> is a surprisingly complicated matter. And much of the complexity has eluded the abundant literature on language learning. Complexity is as much a nuisance as gout, but sometimes just as real and inevitable. Like gout one avoids introducing it to the system, but confronted with it one has no reasonable alternative but to deal with it. So far psychologists have failed to deal with what strikes me as the very real complexity of name learning.

I think this is basically true, with two qualifications. First, the situation in psychology has changed over the last several years (largely as a result of Macnamara's own work), and there has been renewed interest in the topics he lists as being unfairly neglected—reference, meaning, intentionality, hierarchies, and the role of grammar. And second, a better analogy for complexity might be cholesterol; gout is always a nuisance, but there is bad cholesterol and there is good cholesterol. Some psychological problems are complex in bad ways: they cut across domains in a chaotic and messy fashion; they have no clean answers; and their solutions, to the extent they have any, impart no illumination about the mind in general. But a reader of Macnamara's book is drawn toward the conclusion that the learning of names is complex in a different, more positive, sense. Word learning is complex because it involves different cognitive capacities working together in an elegant fashion. Hence the study of word learning might provide insight into these capacities and how they interact in the course of development.

This brings us back to the question of how word learning relates to other aspects of language acquisition. In the sections that follow, I suggest that deep similarities exist between word learning and other aspects of language development. But there is one major difference. Under many analyses, systems such as syntax and morphology have a highly modular flavor; they are self-contained, with their own rules and representations, and interact in a highly circumscribed fashion with perceptual and motoric systems, as well as with other aspects of language. In contrast, it is impossible to explain how children learn the meaning of a word without an understanding of certain nonlinguistic mental capacities, including how children think about the minds of others and how they make sense of the external world. To the extent that Leibniz was right in saying that language is "a mirror of the mind," he was talking about words.

The Problems of Word Learning

Word learning, and especially the learning of names for things, certainly *seems* like a simple process, at least to scholars who are not directly engaged in its study. To take a typical example, in the midst of an otherwise fine discussion of primate drawing abilities, Maureen Cox (1992, pp. 17–18) makes the following remark: "Now, chimps cannot speak because they lack the necessary vocal apparatus, but they can be taught to use sign language. They may not be able to use it in quite the same creative way as humans, but at least they can use it to name things."

I am not concerned (not here, at least) with the empirical claims about what chimps can or cannot do. And while I believe that naming really is a creative act, it is reasonable to say that it is not creative in the same sense as other parts of language. The part of this passage that grates is the phrase "at least." To me this is like saying that chimps can't play checkers—but at least they can play chess! If chimps could use signs for the act of naming, it would show that they have remarkable mental powers.

What is so impressive about word learning? In a classic discussion, Willard V.O. Quine (1960, p. 29) asks us to imagine a linguist visiting a culture with a language that bears no resemblance to our own and trying to learn some words: "A rabbit scurries by, the native says 'Gavagai,' and the linguist notes down the sentence 'Rabbit' (or 'Lo, a rabbit'), as tentative translation, subject to testing in further cases."

Quine goes on to argue that it is impossible for the linguist to ever be certain such a translation is right. There is an infinity of logically possible meanings for *gavagai.* It could refer to rabbits, but it could also refer to the specific rabbit named by the native, or any mammal, or any animal, or any object. It could refer to the top half of the rabbit, or its outer surface, or rabbits but only those that are scurrying; it could refer to scurrying itself, or to white, or to furriness.

The linguist could exclude some of these interpretations through further questioning (assuming some means of figuring out when the native was saying yes or no). For instance, if the native denies that a rat is *gavagai,* the linguist could be confident that the word does not refer to all animals; if he agrees that *gavagai* could be used for a gray rabbit, then it could not mean white, and so on. Other interpretations are harder to exclude. How could the linguist know that the native isn't using the word *gavagai* to refer not to rabbits but to time slices of rabbits—to entities that exist only for the instant that the word is used? Or that the native isn't talking about, as Quine puts it, "all and sundry undetached parts of rabbits"?

There are actually several different problems here. The first is how the linguist knows that *gavagai* is a name at all, as opposed to the native clearing his throat, or making a noise to warn the animal away, or talking to himself, or saying the equivalent of "Look!" or "I'm bored." How does the linguist know that it is one word and not two—*gava* and *guy?* This segmentation problem is a real one when one considers less idealized examples of translation. People do not typically use words in isolation; most words are used in the context of sentences. Even if the linguist can be certain that an act of naming is going on, he or she has to somehow parse the utterance so as to extract the name (which might itself be more than one word, as in *chinchilla rabbit*).

A more serious problem was noted above: How can the linguist know what the word is describing? It could be the whole rabbit, the rabbit and the ground it is on, a part of the rabbit, its color, shape, size, and so on. And this raises the final problem—figuring out how to extend this word in the appropriate way in new circumstances. Suppose the linguist can be sure that *gavagai* is a name and that it refers to the whole rabbit. How should the word be used in the future?

This problem of generalization is a specific instance of a more general dilemma. Nelson Goodman (1983) has pointed out that for any act of induction, there is an infinite number of equally logical generalizations that one can make, each equally consistent with the experience one has had so far. If you burn your hand on a large white stove, for instance, one has to decide which objects to be more careful around. The right answer is *stoves* and not *large things* or *white things.* Goodman points out that there is no logical reason to favor this conclusion over any of these alternatives, as well as over some truly bizarre hypotheses, such as *stoves—but only to the year 2000, then carrots.*

Word learning is a paradigm case of inductive learning. Something has to explain why the linguist, as well as any child, should favor the hypothesis that *gavagai* should be extended to other rabbits as opposed to the hypothesis that it should be extended to other white things, or other rabbits plus the Eiffel Tower, and so on.

These problems of reference and generalization are solved so easily by children and adults that it takes philosophers like Quine and Goodman to even notice that they exist. If we see someone point to a rabbit and say "gavagai," it is entirely natural to assume that this is an act of naming and that the word refers to the rabbit and should be extended to other rabbits. It would be mad to think that the word refers to undetached rabbit parts or rabbits plus the Eiffel Tower. But the naturalness of the rabbit hypothesis and the madness of the alternatives is not logical necessity; it is instead the result of how the human mind works.

Some Facts about Words and How They Are Learned

Since humans learn words, we somehow solve the problems sketched out above. But how often and under what circumstances? If children knew few words, for instance, and had to learn each word under extensive tutelage, this would motivate a different psychological theory of their abilities than if they knew many words and could learn them under very impoverished circumstances.

How many words do people learn? This is a hard question to answer. It requires a robust notion of what a word is, some understanding of what it is to know a word, and a good method with which to test whether such understanding exists. If you simply ask educated people how many words they know, you will get very low estimates (Seashore & Eckerson, 1940), and they are even stingier when estimating the vocabularies of others. Jean Aitchinson (1994) remarks that one respected intellectual in the nineteenth century claimed that peasants have a vocabulary that does not exceed 100 words; they make do with such a small lexicon because "the same word was made to serve a multitude of purposes, and the same coarse expletives recurred with a horrible frequency in the place of every single part of speech." The linguist Max Müller proposed that highly educated people use 3,000 to 4,000 words, other adults know about 300 words, and "the child up to the eighth year probably confines himself to not more than 150 words" (cited by Nice, 1926). More recently, the writer Georges Simenon explained that he makes his books so simple because most Frenchmen know fewer than 600 words. (Simenon also claims to have slept with 10,000 women in his life, leading Aitchinson to suggest that he suffers from a general problem with numerical cognition.)

More sensible estimates emerge from studies that use the following methodology (Miller, 1996). Words are taken from a large unabridged dictionary, including only those words whose meanings cannot be guessed using principles of morphology or analogy. (Even if you never learned *restart,* for instance, you can guess what it means, and so it shouldn't be included in a test of how many words you learned.) Since it would take too long to test people on hundreds of thousands of words, a random sample is taken. The proportion of the sample that people know is used to generate an estimate of their overall vocabulary size, under the assumption that the size of the dictionary is a reasonable estimate of the size of the language as a whole. For example, if you use a dictionary with 500,000 words, and test people on a 500-word sample, you would determine the number of English words they know by taking the number that they got correct from this sample and multiplying by 1,000. The typical test is a multiple-choice question with four

or five alternatives (which introduces a chance factor that must be controlled for); studies with young children use other, more sensitive, methods as well (see Anglin, 1993).

This procedure yields estimates of about 45,000 words for American high school graduates (Nagy & Herman, 1987). This is roughly the same number as those found by scholars in the late 1800s and early 1900s, though they used somewhat cruder methods (Nice, 1926). It is inevitably a low estimate, as it excludes proper names for people and places, idiomatic expressions, and undecomposible compounds. Taking these into account, the estimate jumps to 60,000 or 80,000, and people who do a lot of reading might know twice this many (Aitchinson, 1994; Miller, 1996).

Children start to produce words at about the age of 12 months, which, if we stick to the more conservative estimate of 60,000, equates to about 10 new words a day up until the end of high school. Steven Pinker (1994b, p. 151) remarks that this sort of learning of arbitrary pairings is unprecedented: "Think about having to memorize a new batting average or treaty date or phone number every ninety minutes of your waking life since you took your first steps." And while the recovery of most arbitrary facts is slow and hard, access to words and their meanings is fast and effortless. In normal speech, we produce about three words a second and can recognize a word about one-fifth of a second after its onset (Marslen-Wilson & Tyler, 1980).

What about the circumstances under which words are learned? Consider the example that began this chapter, in which a child interacts with a dog, hears it called "dog" as she is looking at it, and is rewarded (and corrected) in her efforts to name the animal. How much of this is necessary? That is, how much can we take away from this situation and still have children learn the meanings of words?

First, words can be learned without a strict spatial and temporal cooccurrence between the word and the meaning. It is true that in many Western cultures parents often speak to their children in contexts in which the referents of words are easily recoverable. In particular, they often use words to refer to what the child is attending to at the moment the word is spoken, and for one-year-olds this may be the case as much as 70 percent of the time (Collins, 1977; Harris, Jones & Grant, 1983). But if so, this still leaves 30 percent of cases in which no such cooccurrence occurs. Some of the time that children hear "Time for your milk" they will be looking at their milk, and it would be reasonable to map the word onto the substance, but some of the time they will be looking at a fork or a person's face, and a mapping based on a sensitivity to spatiotemporal association would lead them into grief. It is revealing, then, that children are capable of learning words on the

basis of a single trial and that serious mistakes—such as a child thinking that *milk* means "fork"—virtually never occur.

Moreover, while parents might name objects ("This is a cookie"), they do not name actions. Most of the time that adults use verbs, the actions that the verbs refer to are not taking place (Gleitman, 1990; Tomasello, 1992). As Lila Gleitman (1990, p. 19) puts it,

> When, every evening, Mother opens the door upon returning from work, what does [the child] hear? I would venture that he rarely hears her say *Hello, Alfred, I am opening the door!*, but very often hears *Hello, Alfred, whatcha been doing all day?* . . . In short, any scheme for learning from observation must have some machinery for dealing with the fact that caretaker speech is not a running commentary on scenes and events in view.

This point is worth stressing, since the standard scenario that finds its way into discussions of word learning is that the child is observing a scene and hearing words that describe it. But as Gleitman points out, this isn't typically what happens: opening often occurs without anyone using the word *opening*, and the word *opening* often occurs without anything being opened. As we will see later, there is also experimental evidence that children are able to learn words for objects and actions that are not observable to them at the time the words are being used.

Second, children do not need a full complement of sensory abilities to learn words. Deaf children learning a signed language such as ASL do so at exactly the same pace as hearing children learning spoken languages; the age of milestones of word learning, such as the first word and the first 50 words, is identical (Petitto, 1992). More surprising, given the tradition of viewing visual experience as a driving force in language learning, is how well blind children learn language. Such children cannot identify objects that are not within reach, cannot follow the direction of a parent's gaze, and cannot use pointing as a cue. Yet one extensive longitudinal study of the language development of three blind children found that two of them showed only a small initial delay in the onset of word use—which may be in part due to a more general lag in motor development found in blind children—and the third was actually linguistically precocious (Landau & Gleitman, 1985). By each of the children's third birthday, their language was indistinguishable from that of sighted children.

Landau and Gleitman also explored in detail certain aspects of the lexical knowledge of one blind child named Kelli. In one set of studies, they found that her knowledge of color words was similar to that of sighted children of the same age. That is, she knew that color words belong to a single domain, that they apply only to concrete objects, and

that they map onto a property that she could not herself identify. At the very least, these observations suggest that visual perception might not play as large a role in language development as many have suggested.

Third, children do not need feedback to learn word meanings. Although some Western parents correct their young children if they use words incorrectly, this is not universal; there are cultures in which adults do not even speak to children until the children are using at least some words in a meaningful manner (see Lieven, 1994, for review). Yet such children nevertheless come to learn language. Consider also studies of children who for various reasons cannot speak but can hear and are otherwise neurologically intact. In one study of a four-year-old who could produce only a few sounds, it was found that he could understand complex syntactic structures, could make appropriate grammaticality judgments, and had a normal vocabulary (Stromswold, 1994). Needless to say, if someone cannot talk, they cannot get feedback on their speech, and so the fact that this child developed a normal language proves that parental reactions such as correction cannot be necessary for vocabulary development.

Fourth, children do not need ostensive naming for word learning. The paradigm case for the study of word learning—both in philosophy and psychology—is the sort of example that began this chapter: the child is looking at a dog, someone says "dog," and she somehow connects the word with the object. But, just as with feedback, this pattern of naming is not a human universal; children can learn language without it. For example, Bambi Schieffelin (1985, pp. 531–532) describes the cultural context of children acquiring Kaluli. These children grow up in a rich linguistic environment, surrounded by adults and older children who are talking to one another, including making observations about the infant himself, as in: "Look at Seligiwo! He's walking by himself." Furthermore, Kaluli adults explicitly teach children *assertive language,* such as teasing, shaming, and requesting, by modeling the appropriate sentence to the child and adding the word *elema*—an imperative meaning "Say like that." (Appealing or begging for something is never part of an *elema* sequence; according to Kaluli ideology, assertiveness has to be taught, but begging is innate.) But object labeling is never part of an *elema* sequence; there is no naming of objects and no labeling interactions: when a child names an object for an adult, the adult's response is disinterest. This lack of object labeling has been observed in other cultures as well (Lieven, 1994).

All of these considerations show how robust the word learning process is. Nobody doubts that for children to learn words they have to be exposed to them in contexts in which they can infer their meanings:

this is a truism. But the words do not need to be presented in a labeling context (they can be learned from overheard speech), nor do children need to be able to see what the words refer to or have their efforts at using the words encouraged and corrected. And children learn words over and over again, coming to build a vocabulary in the tens of thousands, each word available in an instant for production and understanding.

Finally, consider the nature of what is learned. Noam Chomsky (1993, p. 24) has often maintained that vocabulary acquisition poses learning problems akin to those posed by the acquisition of other aspects of language: "The pervasive problem of 'poverty of stimulus' is striking even in the case of simple lexical items. Their semantic properties are highly articulate and intricate and known in detail that vastly transcends any relevant experience."

To take a simple example (see Keil, 1979; Pustejovsky, 1995), consider the word *book.* This can refer to a material entity, as in the sentence *There are five books on the floor.* But if you say that "John wrote a book," *book* refers to an abstract entity, one that need not correspond to any material object. (All of the five books on the floor might be copies of the book that John wrote). If you say that you are "beginning a book," it will normally be taken as meaning that you are beginning to read a book and will have a different interpretation than the phrase *beginning a sandwich.* Similarly the adjective in *a hard book* or *a long book* has a different meaning than it does in *a hard cookie* or *a long flagpole.* An adequate theory of language acquisition must explain how we come to know all this without any explicit tutelage. Similar puzzles arise when we consider the subtle ways in which verbs and prepositions are used to denote both concrete and abstract relations and events (e.g., Jackendoff, 1990; Lakoff, 1987; Pinker, 1989).

One particular case of interest that is discussed later on is the use of words to name representations of what the words depict. We use *dog* to refer not only to dogs but to statues of dogs, photographs of dogs, and drawings of dogs—including those that bear no resemblance at all to actual dogs. As we will see, this common use of words poses some surprisingly complicated problems from the standpoint of learning and development.

A further aspect of the poverty-of-stimulus problem is our grasp of word meanings that correspond to things that do not exist. As an example, Chomsky (1995, p. 25) cites John Milton: "The mind is its own place, and in itself can make a Heaven of Hell, a Hell of Heaven." One can find this perfectly intelligible, even true, without being committed to the idea that any of these names actually refer either to things in the natural world or to entities in some abstract mental world.

How Children Learn the Meanings of Words

An argument often made in the cognitive sciences starts by describing how hard a task is (such as object recognition, for instance) and then uses this consideration to argue that there is a dedicated part of the mind that does this task. This is not the argument I am making here. To the extent that this book has an overarching theme, it is this. Word learning really is a hard problem, but children do not solve it through a dedicated mental mechanism. Instead, words are learned through abilities that exist for other purposes. These include an ability to infer the intentions of others, an ability to acquire concepts, an appreciation of syntactic structure, and certain general learning and memory abilities. These are both necessary and sufficient for word learning: children need them to learn the meanings of words, and they need nothing else.

This proposal is not original. Many scholars who look at word learning from the standpoint of social cognition argue that word learning is the product of children's ability to figure out what other people are thinking when they use words. And scholars interested in syntactic cues have made a similar claim for the role of syntax, just as those discussing the cognitive prerequisites of word learning have been concerned with the conceptual and logical underpinnings of the process. I argue that a complete explanation for how children learn the meanings of words requires all of these capacities.

There are two ways in which such a proposal could be wrong. It might be attributing too much to young children. It could be argued, for instance, that children do not need an elaborate theory of mind to determine which objects words refer to because they can use statistical information instead. Perhaps a theory that posits fewer resources on the part of the child can explain the developmental facts just as well.

Alternatively, the capacities I have proposed might not be enough. Perhaps lexical constraints (or principles, assumptions, or biases) specifically earmarked for word learning are needed to explain how children learn the meanings of words. There has been a proliferation of these constraints over the last decade or so. They include the whole-object bias, the taxonomic bias, and the mutual-exclusivity bias (Markman, 1989), the noun-category linkage (Waxman, 1994), the shape bias (Landau, Smith & Jones, 1988), the principles of contrast and conventionality (Clark, 1993), and the principles of reference, extendibility, object scope, categorical scope, and novel name-nameless category (Golinkoff, Mervis & Hirsh-Pasek, 1994).

When some of these constraints were first proposed, critics, such as Nelson (1988), argued that they attributed too much preexisting mental structure to young children. These criticisms have been taken to heart

by developmentalists. Few proponents of the constraints view are rash enough to propose that they are innate. Instead, they are said to be learned (or better, to develop or emerge), although—with the important exception of the shape bias (Landau, Smith & Jones, 1988)—nobody has much to say about how this learning, development, or emergence supposedly takes place. In fact, even the mild suggestion that constraints on word learning exist at all is seen as an extreme view, and researchers are careful to insist that they mean *constraint* in a weak sense, not at all like the sorts of principles that linguists talk about (e.g., Golinkoff, Mervis & Hirsh-Pasek, 1994).

All of this caution reflects the empiricist prejudices of the field, and it seems to me to be unjustified. There is nothing biologically implausible about innate constraints on language learning, and we would be unsurprised to find innate constraints underlying the development of analogous systems in other species, such as bee dance, monkey cries, or birdsong. My objection to these special constraints isn't that they are nonbiological or not developmental enough; it isn't that there is some *a priori* reason to believe that they cannot exist. It is that the evidence suggests that, in fact, they don't exist.

By rejecting the idea of special constraints, I am not denying that young children know a lot about words—about their phonology, morphology, syntax, and meaning—and that this knowledge can facilitate the learning of language. For instance, two-year-olds have a tacit appreciation that words referring to objects are typically count nouns. This is part of their understanding of the relationship between meaning and form, and it can help them learn new words. I am not denying that such knowledge of language exists or even that some of it might be innate. The proposal I am arguing against is that there exist *additional* constraints of the sort proposed by Markman and others, constraints whose sole role is to facilitate the process of word learning.

Note also that a rejection of the special-constraint proposal does not entail rejecting the view that children must be constrained as to the inferences they make. This point is often misunderstood. For instance, Roberta Golinkoff and her colleagues (1995, p. 192) discuss Lois Bloom's position that lexical constraints are the inventions of researchers, not actually mental entities on the part of the child, and they suggest that her view ''begs the question of how children determine the meanings of words without considering a myriad of hypotheses.'' But Bloom doesn't beg the question; she just denies that its answer lies in special constraints on word learning.

Elsewhere, Golinkoff, Mervis, and Hirsh-Pasek (1994) suggest that lexical principles ''enable the child to avoid the Quinean (1960) conundrum of generating limitless, equally logical possibilities, for a word's

meaning." But the problem of "limitless, equally logical possibilities" arises for any act of induction. If a dog jumps onto a stove and gets burned, it is likely to infer that stoves are hot—not that undetached stove parts are hot or that stoves until the year 2000 are hot, and so on, even though these alternatives are logically consistent with its experience. So at least *some* constraints on induction are independent of language learning. The issue, then, isn't whether children's inferences about word meaning are somehow constrained (they must be, since word learning is a form of inductive learning); it is whether these constraints are special to the learning of words.

I suggest in the chapters that follow that the phenomena that such constraints have been posited to explain (such as children's tendencies to treat words as object names, to avoid words with overlapping references, and to generalize object names on the basis of shape) are better explained in terms of other facts about how children think and learn.

Preliminaries

The question "How do children learn the meanings of words?" needs clarification. I briefly discuss each of its four content words—children, learn, meanings, and words—to make clear some foundational assumptions and to raise some of the issues that are discussed in subsequent chapters. Meaning is the thorniest issue of all and so is saved for last.

Children

Most research on word learning focuses on two-year-olds to five-year-olds. Why? Why study young children at all, instead of older children and adults?

This would be a silly question to ask about other aspects of language. By the time children are about four, they have mastered just about all of the phonology, syntax, and morphology they are ever going to know, at least for their first language. If you want to study these aspects of language learning, there is no alternative to studying children. But words are different. A six-year-old knows about 10,000 words (Anglin, 1993)—which is less than one-sixth of the number she will know when she graduates from high school.

Nevertheless, most studies on this topic, including my own, follow the practice of developmental research in general and focus on two- to five-year-olds. Such children are the right blend of the exotic and the accessible: they are different in their mental habits from adults, with funny beliefs and immature patterns of thought, and yet they are easy to find, relatively good company, and can be studied without

expensive and time-consuming procedures. And by looking at their capacity to learn words, one can gain insight into the different components of the word learning process, particularly if these emerge at different points in development.

In the end, every age is relevant. Prelinguistic infants are interesting because they lack whatever capacities are necessary to start talking, one-year-olds are interesting because they are word-learning novices, older children and adults are interesting because they are word-learning experts, and preschoolers are interesting because they represent an illustrative midpoint between novice and expert. And by any account, children at these different age groups really do differ in their patterns of word learning. If you look at how many words children learn per day, the difference between a one-year-old and a two-year-old is striking, as is the difference between a two-year-old and a four-year-old, and a four-year-old and 10-year-old. And when you get to adults, the rate of word learning often drops to close to zero, perking up only for the learning of proper names and names of cultural and terminological innovations, such as *Internet, karaoke,* and *Tickle Me Elmo.*

Why do these age differences exist? This is an issue that is discussed in detail in the next chapter, but one obvious consideration is that as children get older, they have increasing access to information concerning what words mean. One-year-olds start off with little or no syntactic knowledge to guide their interpretation of a new word. And before they have learned their first word, they are obviously not going to be able to learn new words from linguistic context. Slightly older children have more syntactic understanding and know a few more words; once their vocabulary and syntax takes off, they can learn words by hearing them used in sentences. They hear words in more diverse contexts; in some cultures, this includes exposure to television and videos and, most important, through literacy.

Learn

Talk about learning has an old-fashioned flavor, and many scholars prefer expressions such as development, emergence, growth, and acquisition. Some argue that there is no such thing as learning—that the notion reflects an out-of-date way of looking at mental processes. The argument often goes like this. We know from the biological sciences that the brain, and all that the brain can do, emerges from the interaction of genes and environment. Since there is no sense in which the environment can have an effect on the brain that is not strongly constrained by our genetic endowment, the whole idea of something being *learned* (as opposed to the ethological notion of being *triggered,* for instance) is an archaic idea that should be expunged from the cognitive sciences.

This is the nativist attack on learning; ironically, the empiricist attack is similar. We know from the biological sciences that the brain, and all that the brain can do, emerges from the interaction of genes and environment. Since there is no sense in which genes can have an effect on the brain that is not strongly constrained by the environment, the whole notion of something being *innate* (as opposed to its *developing* or *emerging*) is an archaic notion that should be expunged from the cognitive sciences. And since nothing can be unlearned (which is the usual meaning of *innate*), then the notion of learning either is incoherent or applies to every aspect of human knowledge. Either way, we should get rid of it too.

If we were to accept these arguments, it would have dramatic consequences and, I think, unfortunate ones. It would be a poor psychology that insists that the same developmental story be told about the emergence of Down syndrome and how people come to learn the word *rabbit*. In fact, even though both arise from the interaction of genes and environment, it is nonetheless entirely reasonable to conclude that Down syndrome is innate and the meaning of the word *rabbit* is learned. There is a sensible dichotomy that should be maintained.

The problem with the arguments does not lie in their premises. It is true that any effect of the environment on how one thinks can occur only if the right innate abilities are in place (some "instinct to learn"), and it is also true that the action of genes on brain and behavior comes about through considerable interaction with all sorts of environments, from the cell to the society. Nobody has ever doubted this. The arguments go wrong in concluding that these facts show that nothing can be explained as caused by the environment or caused by the genes. This doesn't follow; the notion of causal responsibility, both in science and in normal usage, is more sophisticated than that. If Fred throws a cup to the floor and it breaks, this breaking is a profoundly interactionist affair; the cup would not have broken if gravity didn't exist, if the floor hadn't been made of a hard surface, if the cup wasn't made out of a fragile material, if Fred's parents hadn't decided to have sexual intercourse at a certain time, and so on. But none of this takes away from the banal fact that it was Fred who caused the glass to break, not his parents and not the person who built the floor.

By the same token, the interaction between genes and environment does not make it any less reasonable to say things like "Down syndrome is caused by certain genetic factors" or "Joe knows the word *rabbit* because he heard his father use it to refer to Flopsy." These are reasonable things to say; in fact, they are true. Some things are caused by the genes, and others by the environment. There are even genuine cases of interaction, in which both sources play a substantial causal

role; cases like alcoholism and syntax come to mind. This division between genetics and environment is more than common sense; it is good science. If you want to see why a child has Down syndrome, you would look for a genetic cause, but if you want to see why he or she thinks rabbits are called *rabbits,* you would look toward the environment.

Distinguishing between genes and environment is not enough to save the notion of learning, however. After all, bullet wounds and tenure are caused by the environment, but there is no sense in which they are learned. The notion of learning picks out a subset of environmentally caused events, those in which the organism comes to store and represent information through a rational process (Fodor, 1981) of interaction with the environment. The caveat of ''rational'' is present to capture the intuition that not any interaction counts: if you get smacked in the head and miraculously come to know the rules of baseball, this wouldn't count as learning. But if you come to know baseball by observing other people play the game or by having someone explain the rules to you, then this does count as learning—even though, of course, this process would be impossible without the innate ability to learn.

This is a crude definition, but it captures the sense in which word learning counts as learning. In fact, word learning is the clearest case of learning one can imagine. Nobody was born knowing the meaning of the English word *rabbit.* Everyone who knows the word has heard *rabbit* used in a context in which its meaning could be recoverable from the environment using a rational process; that is, everyone who knows the meaning of *rabbit* has learned it. If you can stomach the terminology, I suspect this might be the least controversial claim in the study of language development.

Words

There are different notions of what a word is, not all of them appropriate for the study of word learning. One notion is that of a syntactic atom, something that can be a member of a category such as a noun or a verb and that can be the product of morphological rules (Pinker, 1994b). This notion is what morphologists have in mind, and it corresponds roughly to our intuitive notion of a word: a sound or sign that, if written down, corresponds to a string of letters that has spaces or punctuation marks on either side (Miller, 1996). Under this definition, the sentence *John stayed in the poker game until he got cleaned out* has 11 words, and this assessment is confirmed by the word count tool of my word processor, which uses this algorithm.

But this notion of word is unsuitable for certain psychological purposes, particularly if you are interested in what children have to learn. For instance, children do not have to learn the word *stayed.* What they

do need to have learned is the verb *stay* and the morphological rule that adds *-ed* to transform verbs into the past tense. In general, it is clear that we can use and understand far more words (in the morphological sense) than we have learned. As soon as one learns the verb *stay,* then *stayed, staying,* and *stays* all come for free.

Idioms pose another problem. To understand the above sentence, it is not enough to have learned the verb *clean* and the preposition *out;* you also have to learn something else; the meaning of the idiom *clean out,* which means, roughly, to be totally deprived of something, usually money. As with many idioms, the meaning of the whole bears some relationship to the meaning of the parts (Gibbs & Nayak, 1989), but to fully understand the idiom you have to learn its meaning in much the same way as you would learn the meaning of the syntactic atoms *clean* and *out.*

Finally, consider *poker.* From a learning perspective, the string of letters is at least two words—a card game and a fireplace tool—and each meaning has to be learned separately. The individuation of words, then, must make some reference to meaning. This point is sometimes missed. As Miller and Wakefield (1993) point out, when studies ask how many words children and adults know, they often mean by *word* what lexicographers call a *lemma*—a listing in a dictionary. This has the advantage that *stay, stays, stayed,* and *staying* count as a single word, as do *zeugma* and *zeugmas.* But it has the disadvantage that *poker* is also counted as one word, despite its ambiguity.

The relevant sense of *word* from the standpoint of language acquisition should include all and only those forms whose meanings must be learned. This sense corresponds to *listemes,* units of a memorized list (Di Sciullo & Williams, 1987), or *minimal free forms* (Miller, 1996) or "the smallest semantic units that can move around in an utterance" (Clark, 1993) (though note that the second and third definitions exclude idioms). All these definitions have as their basis the notion of a Saussurian sign (Saussure, 1916/1959)—an arbitrary entity that has on one side a concept and on another, a form.

This is the sense I adopt here; when I talk about children learning words, I mean Saussurian signs. *Dog* is a word, then, but so is *clean out, hat trick, capital gains, kit and kiboodle,* and *Citizen Kane. Poker* has to be learned twice; it is two separate words. On the other hand, certain units that are words from the standpoint of other theories do not count as words for the purposes here, as they are not Saussurian signs. So while I will have a lot to say about how children come to know *dog,* I have nothing to say about how they come to know *dogs* or *dogcatcher.*

What makes this complicated is that words do not come with tags that they are Saussurian signs. A child who hears *poker* used to refer

to a game and then, days later, *poker* used to refer to a tool has to figure out that these are two words and not one. A child who knows the meanings of *clean* and *out* and who hears that someone was "cleaned out" in a poker game has to figure out that this expression is an idiom and hence a sign that has to be learned. And a child who hears *stayed* has to realize that this is not itself a word that has to be learned, though it includes one. (Phonology is a good cue here, but the child does have to be wary; the adjective *staid* sounds the same as the verb, but it really is a Saussurian sign.)

In the end, then, both senses of *word* are relevant. The morphological sense—the sense that people use when they count the number of words in a manuscript—describes the input to the child. Long before learning the meanings of words, children have partially solved the problem of segmenting the sound stream into words in the morphological sense (Jusczyk, 1997). But what they have to *learn* are words in the Saussurian sense, arbitrary signs. This makes the task of word learning even more complex.

Meaning

What is it to know the meaning of a word? Some philosophers say there is no such thing. Quine did not use his Gavagai example to encourage developmental psychologists to search for cognitive constraints on children's inferences (though when Macnamara, 1972, introduced this example to the developmental community, this was ironically its effect); he used it to argue against the very idea of meanings in the head. What the problem of radical translation shows, Quine argued, is that the only robust notion of meaning is the behaviorist one of *stimulus meaning*—a person's disposition to respond to certain sensory stimulation. Other philosophers share this skepticism about meaning, and still others propose that while sentences have meanings, such as their truth conditions or their methods of verification, words do not.

Since I am talking about how children learn word meanings, this commits me to the view that such things as word meanings exist and can be learned and known. In particular, to know the meaning of a word is to have

1. a certain mental representation or concept
2. that is associated with a certain form.

Under this view, two things are involved in knowing the meaning of a word—having the concept and mapping the concept onto the right form. This is the sense of "knowing the meaning of a word" implicit in most discussions of language development, both scientific and informal. Saying, for instance, that a two-year-old has mixed up the

meanings of *cat* and *dog* implies that the child has the right concepts but has mapped them onto the wrong forms. On the other hand, saying that the two-year-old does not know what *mortgage* means implies that the child lacks the relevant concept. People can also possess concepts that are not associated with forms. A child might have the concept of cat but not yet know the word, and even proficient adult users of a language can have concepts, such as of a dead plant or a broken computer, that they don't have words for.

If a concept is to constitute a word's meaning, it has to include some aspects of knowledge but not others. Consider what it is to know the meaning of *dog.* I once owned a dog named *Bingo,* but this knowledge cannot be part of the meaning of *dog* (at least as we normally talk about meaning) since someone could know *dog* even if they've never heard of Bingo. More generally, if the meaning of *dog* were determined by all thoughts related to dogs, then there would be no sense in which two people, or even a single person over time, could ever have the same meaning of a word. This is an undesirable consequence of an extremely holistic theory of meaning (Fodor & LePore, 1992).

Intuitively, then, only some aspects of knowledge are relevant to meaning. What are they? The traditional view, emerging first in Aristotle, is that the meaning of a word is what determines its reference. A word like *dog* has an extension (which entities the word refers to—dogs) and an intension (what the entities share—what all dogs have in common). Meaning is identified with the intension. While the intension of a word is not itself a psychological entity (Frege, 1892), it can be learned and understood. Hence the meaning of *dog* determines which things are and are not dogs, and knowing the meaning of *dog* entails knowing what things are and are not dogs.

As Murphy (1991) points out, this conception is implicit in almost all psychological discussions of the learning and representation of word meaning. One traditional view, for instance, is that meanings are pictures. The meaning of *dog* is a picture of a dog, and you know the meaning of *dog* if you have a mental representation of that picture that lets you tell the dogs from the nondogs. Another view is that meanings are identified with sets of necessary and sufficient conditions. The meaning of *bachelor* is said to be "unmarried man" because all and only bachelors are unmarried men, and hence you know *bachelor* when you map the form onto a concept that includes this definition. A currently popular notion is that meanings are sets of weighted feature and hence knowing the meaning of *dog* is to have a mental representation of the appropriate feature set, which will allow you to judge the extent to which different objects in the world are dogs. Other views include the idea that meanings are mental models, nodes in a semantic

network, or sets of specific exemplars. These accounts all share the assumption that knowing the meaning of *x* involves being able to tell the differences between those things that are x and those things that are not.

There is, for instance, debate over how well prototype theory can explain our knowledge of words such as *chair* and *mother* (e.g., Armstrong, Gleitman & Gleitman, 1983; P. Bloom, 1996a; Malt & Johnson, 1992; Rosch et al., 1976), but the one thing that is agreed by all sides is that for prototype theory to work it must adequately capture patterns of categorization: it has to explain why we think that some things, but not others, are chairs and mothers. And when developmentalists talk about constraints on word meaning (Markman, 1989), inductions about word meaning (Soja, Carey & Spelke, 1991), or cues to word meaning (Gleitman, 1990), they are talking about constraints, inductions, and cues that pertain to the sorts of things children think words refer to.

This fits our commonsense idea of what it is to know the meaning of a word. If someone consistently uses *dog* to talk about tables, then this person does not know the meaning of this word. Conversely, if someone uses *dog* to talk about dogs and only dogs, then they do know the word, even if they have a lot of otherwise bizarre beliefs about dogs. Anyone who believes that dogs are expert chess players has a serious psychological problem, but we would not usually say that their problem is a *lexical* one. We are comfortable translating a word from an ancient Greek text into the English word *star,* even though the ancient Greeks believed that stars were holes in the sky. It is enough that we all use the word to refer to the same things; further cognitive overlap is not necessary.

For these reasons, relating word meaning to categorization seems like a reasonable strategy, and it is the one that I adopt throughout the discussions that follow. But serious objections to this view have been raised, and these have important ramifications for any psychological theory of concepts and meanings.

The main problem is this. It is true that when we talk about knowing the meaning of a word, we are usually thinking in terms of sameness of reference: we and the Greeks both mean the same thing by *star* because we are referring to the same things. But there are many cases in which our mental representations do not determine reference, and so if reference is central to meaning, then meaning is not determined by mental representation.

Consider some examples from Hilary Putnam (1975, 1988), which I slightly modify for my purposes here. Imagine a normal eight-year-old girl who uses the word *water* to refer to the stuff that she drinks, washes with, and swims in. She has clearly learned the meaning of the

English word and uses it to refer to the stuff that happens to be made up of H_2O (though she doesn't know this). Now imagine that there is another world, Twin Earth, that is exactly the same as ours, except that instead of being composed of H_2O the stuff that they call *water* is made up of different chemicals: XYZ. The eight-year-old will have an identical twin on Twin Earth, who uses *water* to refer to the substance that she drinks, washes with, and swims in. But her word does not refer to H_2O; it refers to XYZ (though she doesn't know this). If reference determines meaning, then the two girls use the words with different meanings, and, as Putnam (1975, p. 227) famously put it: "cut the pie any way you like, 'meanings' just ain't in the *head!*"

One can find the same situation without resorting to science fiction. Consider two boys, one raised in Boston, the other raised in London, and assume that they each have the same concept associated with the word *robin;* they each believe that robins are red-breasted birds, and this is all they know about them. The twist here is that *robin* in American English and *robin* in British English refer to different species of red-breasted bird. So which meaning do the boys have associated with *robin*—the American one or the British one? The most natural solution is that the Boston boy knows the American English meaning of the word and the British boy knows the British English meaning of the word. But this would again imply that there is more to meaning than human psychology.

One might be tempted to argue that neither boy knows the meaning of *robin* because this requires the actual ability to determine its precise reference; incomplete knowledge is not enough. But by these standards, most people do not know the meanings of many words: Can *you* tell robins from nonrobins? Or consider the word *gold.* Many people who use the word do not know what distinguishes gold from other metals and have no ability to make the distinction in practice. Similarly, I know the name *Moses* and so do you, but we might have entirely different beliefs about who the person is, and it is perfectly possible that we could discover that all of our beliefs about Moses are false (Kripke, 1980). Still, although we might lack the knowledge to pick out the unique substance that is gold or the unique person that is Moses, we nonetheless know the meanings of both words and use them to talk about the substance gold and the person Moses.

Matters get worse when we try to make sense of what is meant by *the* meaning of a word. What is the meaning of *disinterested?* One possibility is that it means "unbiased," as most dictionaries say; another is that it means "uninterested," as many English-speaking adults believe. Whatever the answer is—assuming that there is an answer—is not going to be solved by neurological and cognitive research. It is more of a

sociological issue, related in complicated ways to notions of authority, expertise, community standards, and so on (Chomsky, 1995).

What moral should we draw from these cases? Putnam (1988) takes them as showing that a psychological theory cannot capture certain basic facts about meaning and reference. Hence such approaches are inherently limited; better to study word meanings in a social and embodied context. Chomsky (1995) draws the opposite conclusion. He argues that the above examples show that trying to build up a theory of semantics from notions such as reference is a waste of time. The only scientific theory will be an "internalist" one, the same sort that holds for aspects of language such as phonology and syntax. This point has been elaborated and extended by Ray Jackendoff, who distinguishes I-concepts (internal concepts) from E-concepts (external concepts) and who sees only the former as "a fruitful framework in which to conduct scientific inquiry" (1989, p. 100).

In either case, the conclusion is that there are two ways to think about meaning—the psychologist's way (as something identified with mental representations) and the philosopher's way (as something identified with reference). And, as David Lewis (1972, p. 170) remarks: "Only confusion comes from mixing these two topics."

But this seems rather extreme. After all, any adequate philosophical theory of reference and meaning must explain what it is about people that enables us to use words that have semantic properties. The fact that I can use the word *gold* to refer to gold, while not entirely explainable in terms of my cognitive structure, plainly has *something* to do with it. Conversely, one of the motivations for a psychological theory of concepts and meaning has always been to explain how we can think about, talk about, and categorize entities in the external world. It may well be that Chomsky and Jackendoff are correct that this is a wrongheaded approach. But there is not, at present, any worked-out alternative.

So what is meant by *the meaning of a word?* Following the lead of *two-factor theories* of semantics (e.g., Block, 1986), we can assume that there are two aspects (or determinants) of the meaning of a word—an internal psychological aspect, sometimes called *narrow content,* and an external social and contextual aspect, sometimes called *broad content.* These work together to determine what words refer to. In what follows, I use the expression *meaning of a word* to correspond to narrow content—to the psychological aspect of meaning. In the end, nothing rests on the terminology, and philosophical purists should feel free to replace my expression *meaning of a word* with the unwieldy but more modest expression *knowledge associated with a word that is relevant to explaining people's intuitions about reference and categorization.*

Still, even if one does adopt a two-factor theory, the phenomena pointed out by Putnam and others still seriously constrain any account, however internalist, of the meanings of words. They show that while possession of a concept might be intimately related to categorization ability, it does not reduce to it. One cannot say that children have learned the meaning of *gold* only when they can tell gold from nongold or have learned a proper name like *Moses* when only they can correctly pick out the person referred to by the name. By these standards, nobody knows the meanings of such words. The psychological part of knowing the meaning of a word has to be more subtle than this.

One final point. The program of relating word learning to issues of reference and categorization works best for common nouns like *dog* and *gold* and proper names like *Moses* and *Fido.* And it can be readily extended to some adjectives and verbs. But it works very poorly for words such as determiners, prepositions, and modals. The semantics of these terms is substantively different; these words get their meanings not by reference but by the roles that they play in modulating the meanings of other, referential, terms. It does complicate matters to say that there are (at least) two types of word meanings, one for *dog* and one for *the.* But the alternative—that all words are understood and learned in the same way—is not very promising.

Outline

Each of the following chapters is self-contained enough to be read on its own, but they do have a logical progression, and each rests to some extent on evidence and arguments introduced earlier.

The next chapter explores fast mapping—the rapid acquisition of the meanings of new words. It presents some data about the nature and scope of fast mapping and then turns to questions about the time course of word learning and individual differences in how words are learned. Why does word learning start when it does? Why does it speed up in the years to follow and slow down in adolescence? How do people differ in their word learning, and why?

Chapter 3 discusses children's appreciation of the mental states of others. Evidence is presented that this understanding underlies several aspects of the learning process, including how children know which entities in the world certain words refer to. When an understanding of intentions of others is partially absent, as with autistic children, there are devastating results.

Once children know what the word refers to, they have the further problem of figuring out whether the word is a common noun, referring to the kind (as in *rabbit*), a proper name (*Flopsy*), or a pronoun (*her*).

Common nouns are the topic of chapter 4, and pronouns and proper names are dealt with in chapter 5. While these chapters focus mostly on object names, they also discuss more abstract expressions, such as *family* and *London.*

Chapter 6 concerns the conceptual foundation of word learning—the nature of the concepts that constitute certain word meanings. Chapter 7 focuses on an important case study for any theory of concepts and naming—visual representations.

The idea that there are linguistic cues to word meaning is introduced early in the book, as such cues help explain how children learn names for objects, substances, and specific individuals. But they are far more important when it comes to learning the meaning of more abstract nouns, such as *mortgage* and *idea,* as well as for learning other parts of speech, such as verbs like *think* and adjectives like *blue.* This is the topic of chapter 8.

Chapter 9 addresses the learning of number words. These words show an interesting pattern of development and illustrate both the importance of linguistic cues and their limitations. Toward the end of this chapter, I explore the idea that the learning of number words might affect how we think about numbers. This raises the more general question, addressed in chapter 10, of how the words we learn affect our mental life. I suggest that language can affect thought, but only in certain circumscribed ways. The chapter concludes by arguing that the rich mental life of humans is the foundation of word learning; it is not the product of it.

Chapter 11 contains a brief summary and some general remarks.

BLANK PAGE

Chapter 2

Fast Mapping and the Course of Word Learning

The average American or British high school graduate has learned about 60,000 words (Aitchinson, 1994; Miller, 1996; Pinker, 1994b). This is a rough estimate, and there are considerable individual differences. Some people learn many more words, others somewhat fewer, and those who know two or three languages might know two or three times as many. But 60,000 is a good conservative number. Since word learning starts at about 12 months of age, this averages to learning 3,750 new words a year, or 10 words a day—a word every waking 90 minutes.

This statistic is impressive, but it is misleading in a number of ways. Learning the precise meaning of certain words, especially verbs, might be a long process requiring many trials, as shown by the fact that even some relatively frequent verbs, such as *pour* and *fill,* are not fully understood until middle childhood (Gropen, Pinker, Hollander & Goldberg, 1991). On any given day, then, it might not be that children are learning 10 words; they might instead be learning one-hundredth of each of a thousand different words.

Also, word learning does not proceed at an even pace. It does take some of the drama away to realize that, despite what is often said in language-acquisition textbooks, three-year-olds are not learning even close to 10 words a day; it is more like 10 words a week. But in another sense this somewhat slow start makes the word-learning task all the more impressive—because it means that older children have to learn words at an even faster rate, such as 12 or 15 words per day, a word every waking hour.

Sixty thousand words are a lot to learn and remember. Learning a word requires memorizing an arbitrary relationship between a form and a meaning, and the rote learning of paired associates is notoriously slow and difficult. Consider how hard it is to learn the capitals of different countries or the birthdays of particular people. The recall of such arbitrary facts is also relatively slow. What is the capital of Spain? When is your mother's birthday? If it took you a half second to answer these questions, this is much slower than it took you to access the