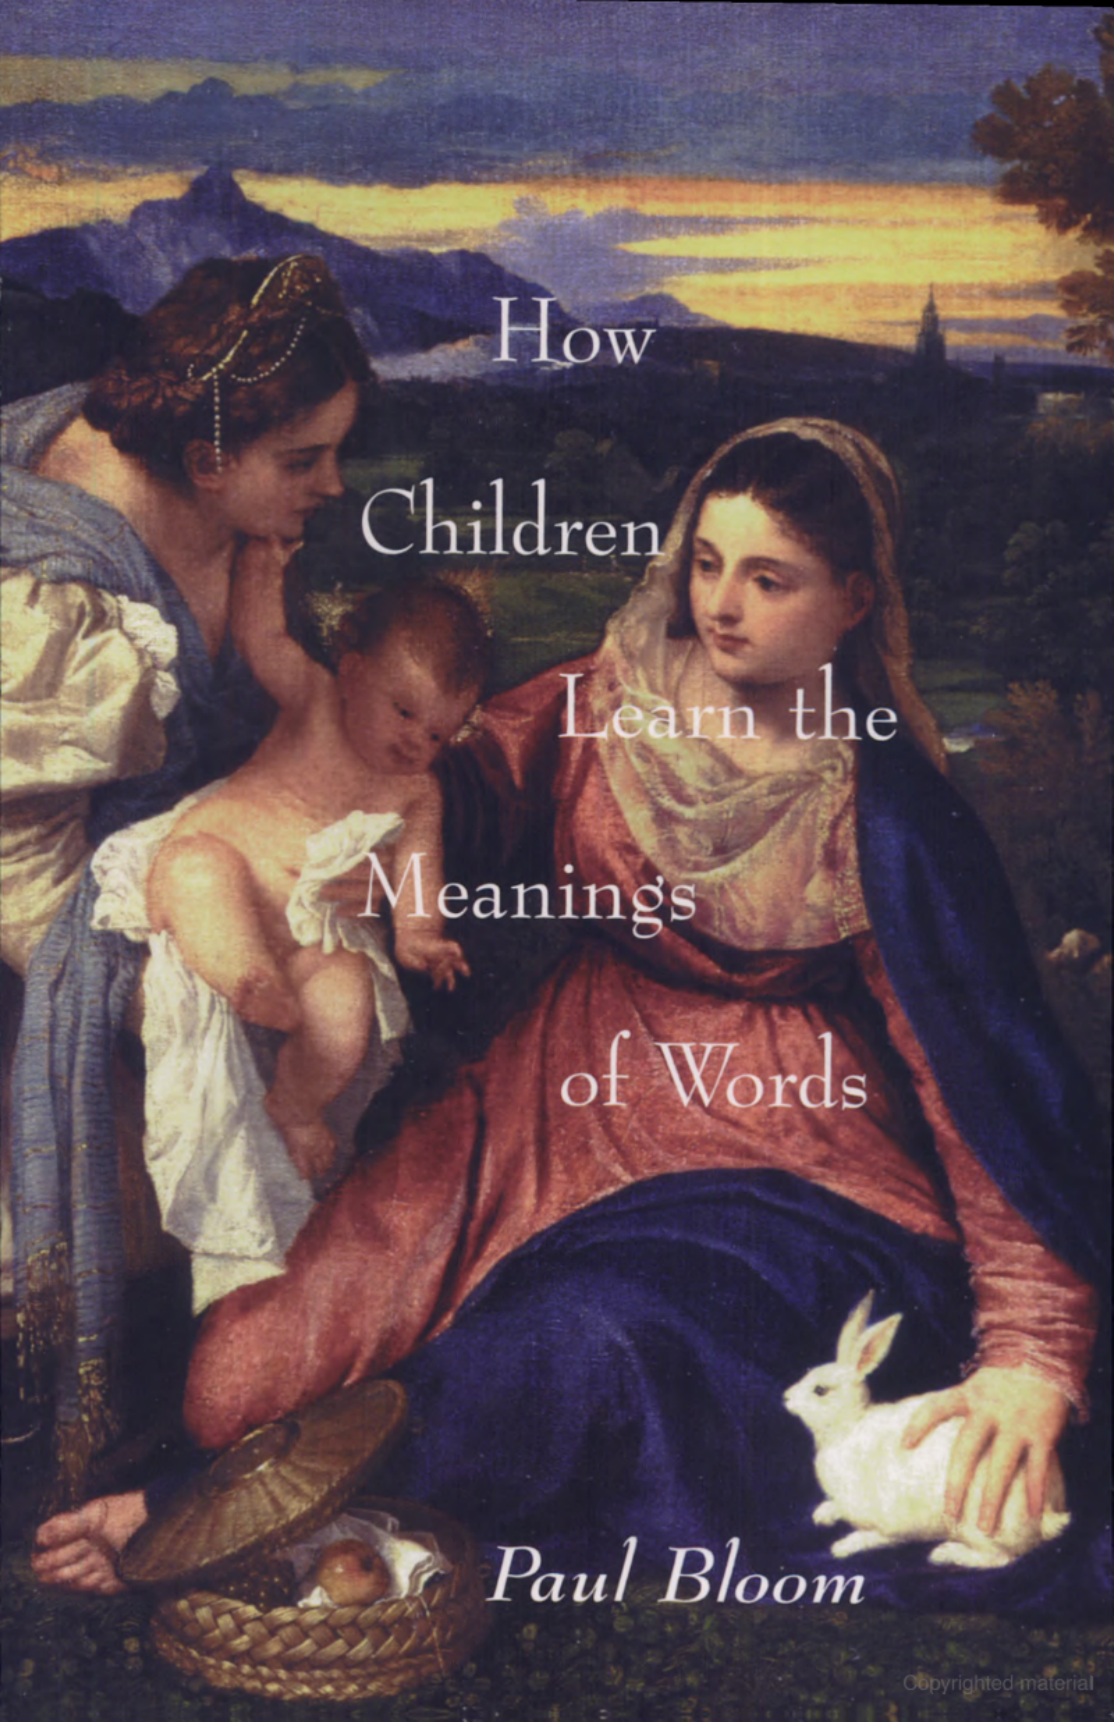

How Children Learn the Meanings of Words

Paul Bloom

How Children Learn the Meanings of Words

Paul Bloom

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

First MIT Press paperback edition, 2002

This book was set in Palatino by Achorn Graphic Services, Inc. and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Bloom, Paul, 1963–
How children learn the meanings of words / Paul Bloom.
p. cm.
Includes bibliographical references and index.
ISBN 0-262-02469-1 (hc : alk. paper), 0-262-52329-9 (pb)
1. Language acquisition. 2. Semantics. I. Title.
P118.B623 2000
401'.93—dc21 99-23901
CIP

10 9 8 7 6

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 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 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

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.

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 meaning of each of the words you are now reading. And although the vast majority of words we know are not frequently encountered, once heard or read—*putrid, centrifuge, apostle, fawning*—they are immediately understood. There is no denying the impressiveness of word learning.

This chapter addresses three general issues. The first concerns the nature of the word-learning process. How much input do children need to learn a new word? Do children learn words better than adults? And to what extent does word learning differ from other types of learning? The second issue is the time course of word learning—when it starts, when it stops, and what happens in between. And the third issue is individual differences: What differences exist, what causes them, and what do they tell us about the process of word learning?

Fast Mapping

Young children can grasp aspects of the meaning of a new word on the basis of a few incidental exposures, without any explicit training or feedback—in fact, even without any explicit act of naming. In the first experiment to systematically explore this phenomenon, Carey and Bartlett (1978) casually introduced a new color word to three- and four-year-old children who were involved in another, unrelated activity. The children were asked by the experimenter to walk over to two trays, a blue one and an olive one, and to "Bring me the chromium tray, not the blue one, the chromium one." All of the children retrieved the olive tray, correctly inferring that the experimenter intended "chromium" to refer to this new color. When tested a week later on their comprehension of the word, over half of the children remembered something about its meaning, either that it named olive or that it named a color that resembled olive. In another study, children were taught the word, tested after a week, and then given a production test five weeks later. Even six weeks after hearing the new word, children typically retained some understanding of its meaning, if only that it was a color term (Carey, 1978). Susan Carey (1978) has dubbed this process of quick initial learning *fast mapping*.

Heibeck and Markman (1987) expanded on this research in certain ways, obtaining some interesting results. First, even two-year-olds can fast map new words. Second, fast mapping is not limited to color words; children can fast map shape and texture terms as well. Shape is easiest of the three to learn, texture hardest, and color falls in between. Finally, explicit linguistic contrast is not necessary. Children don't need to hear something like "Bring me the chromium tray, not the blue one," in which the novel word is explicitly contrasted with an existing color term; they do just as well if they simply hear "Bring me the chromium tray, not the other one."

There is one way in which the Heibeck and Markman (1987) study was quite different from the original Carey and Bartlett (1978) study. In the original study, children were tested after a gap of one week and six weeks. For Heibeck and Markman, the gap between teaching and testing was 10 minutes. This short gap is common in word learning research, even in studies that are specifically designed to study fast mapping (e.g., Dollaghan, 1985; Rice, 1990). In fact, given the prominence of the Carey and Bartlett study, and even though there have been countless subsequent studies in which children have been taught new words and then tested for what they think these words mean, there are almost no experiments that replicate one of the most interesting features of the original study—the long delay between exposure to the new word and testing for its retention. Teaching children new words and immediately testing them on what they think the words mean is a fine way of exploring their conceptual and linguistic biases. But it tells us little about the retention capacity so manifest in normal word learning, and so certain basic questions about word learning that were raised by the Carey and Bartlett study remain unanswered.

First, color names constitute a relatively narrow domain, with universal constraints on the terms that exist and the order in which they are acquired (e.g., Berlin & Kay, 1969). In contrast, there are thousands of object names, and these show vastly more variability across cultures, as their referents include foods (*apple, bagel*), animals (*dog, snake*), artifacts (*clock, car*), and so on. It is a natural question whether fast mapping applies for these words as well.

Second, is there any difference between children's and adults' abilities to fast map? One might expect adults to be better, since they have had more experience acquiring words and are usually more adept at learning and memory tasks. Nevertheless, young children are notably superior to adults at successful acquisition in the linguistic domains of phonology, morphology, and syntax (Newport, 1990), and it is conceivable that a similar critical period might exist in lexical acquisition.

Finally, is the ability to fast map limited to word learning? Children's capacity to quickly learn new words on the basis of limited experience is frequently cited as showing how good they are at language learning. But it is possible that this capacity emerges through more general capacities of learning and memory. This issue can be addressed by seeing whether fast mapping of arbitrary information occurs with equal force in contexts other than word learning.

Lori Markson and I have addressed these questions in a series of studies designed to be similar to the original Carey and Bartlett study.

to the word task. Four of the objects that they see in the test phase have familiar names. Since people tend to infer that a novel word refers to something that does not already have a name (Markman & Wachtel, 1988; see chapter 3), they might assume that the word must refer to one of the six unfamiliar objects, even if they have no memory of having heard it before. Here, then, chance for the word task should be calculated as a more conservative one-sixth, or 17 percent. The results are shown in figure 2.1.

In the new word task, children and adults in all three delay conditions performed significantly above chance. Adults were better than children when tested immediately, but after a week and a month, three-year-olds and four-year-olds were doing as well as adults, and all age groups remembered which object was the koba over half of the time, far better than chance. In the linguistically presented fact task, the results are basically identical. All age groups again performed significantly better than chance in all delay conditions. There were no age differences and no significant decline across delay conditions. In contrast, in the visually presented fact task, adult performance was considerably diminished after a week and a month, and the three-year-olds and four-year-olds, taken together, showed a significant decline over time and did significantly worse than in the koba and uncle conditions. Only the adults performed significantly better than chance after a one-month delay; the children's performance was indistinguishable from guessing.

We can now return to the three questions raised earlier. First, fast mapping does apply to object names. Even after a one-month interval, most children and adults retained the meaning of an object name that was presented to them in an incidental context. Second, there does not appear to be a critical period for fast mapping: adults and children tend to perform equally well. Finally, fast mapping is not limited to word learning. Children and adults were as good at remembering an arbitrary linguistically presented fact about an entity as they were at remembering what the entity was called. Nevertheless, fast mapping does not apply to any arbitrary memorization task, as illustrated by the children who, after a one-month delay, were unable to recall which object a sticker was applied to.

One possible objection to this experiment is as follows. Suppose children really are better at fast mapping words than facts. This difference could be obscured in the experiment above because learning the new word condition involves storing novel phonological information (the sound *koba*) while the uncle condition does not. If the uncle condition also involved memorizing a novel sound, perhaps learning the word would be easier.

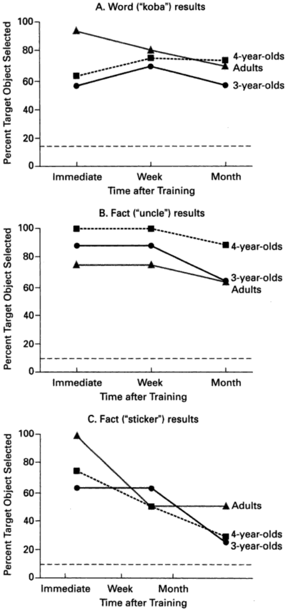

Figure 2.1
Three-year-olds', four-year-olds', and adults' recall of novel words (A) and novel facts (B, C) (from Markson & Bloom, 1997)

Table 2.2
Children's and adults' choice of the original object when asked about words and facts different from those that they were originally taught

	Percent Choice of the Original Object	
	Children	Adults
Taught "koba," asked for "koba" (from original study)	69%	65%
Taught "koba," asked for "modi"	19	19
Taught "from uncle," asked "from uncle" (from original study)	75	93
Taught "from uncle," asked "from sister"	13	13

Table 2.3
Two-year-olds' recall of novel words and novel facts

	Percent Correct	
	Immediate	One-Week Delay
Word	81%	81%
Fact	88	75

When tested on a word or fact that was not originally taught, subjects tended not to choose the original object. These results show that, in general, when tested a month after hearing a specific word or fact applied to an object, subjects tend to remember what that word or fact is (Markson, 1999).

What about fast mapping by younger children? A group of two-year-olds was tested on a simpler version of the original study, using only four objects, all of which were novel. They were tested either immediately or after a week delay, and the results are shown in table 2.3. In all conditions, children did significantly better than chance, and there was again no difference between the new word and the new fact (Markson, 1999).

What do these findings tell us about the mechanisms of word learning? It is conceivable that two distinct capacities or mechanisms explain our results—one explaining performance on the new word task and the other explaining performance on the nonlexical uncle task. But this is an unparsimonious way to explain the data. Given two patterns of learning that are virtually identical, it is simpler to see them as emerging from one mechanism and not two. More generally, since children and adults seem to have a general ability to fast map, there is no motivation to posit a separate mechanism that applies only in the domain of word learning.

One might object that similarities in how words and facts are remembered do not directly bear on the question of how words and facts are learned. Perhaps words are learned in an entirely different way from facts; it is just that in both cases the information is stored in the same sort of memory.

This is a valid point, but one needs to keep in mind that fast mapping, as I have been using the term here, includes not only the ability to store information for a long period of time but also the ability to learn the information in the first place. So there are two parallels between words and facts: one has to do with memory; the other has to do with learning. In the studies above, the fact conditions were virtually identical to the word conditions. Most notably, the facts and the words were taught in the same linguistic and pragmatic context ("Let's use the ______ to measure which is longer"). The finding that both types of information could be learned under these quite minimal circumstances suggests that the cognitive mechanisms that underlie word learning also apply more generally.

This isn't to deny that words have special properties. They do, and some of the next chapter discusses what they are and how children learn about them. But the findings reviewed above do suggest that the act of learning and storing words in memory—the act of fast mapping—involves cognitive mechanisms that apply to facts as well.

What is the scope of fast mapping? What sorts of information can be fast mapped? Given how little we know about the process, and since we only have one reliable case in which fast mapping does not fully apply (the sticker condition, in which the fact is learned but not recalled after a long delay), there is no shortage of possibilities. Perhaps it applies only to information conveyed through language. Or perhaps salience is the key factor, and our subjects simply found the name of the object and its origin to be more *interesting* than the placement of the sticker and hence easier to remember. A third possibility (independently suggested by Dare Baldwin and Susan Carey) is that fast mapping applies only in circumstances in which the information is seen as relevant to the category that an object belongs to. This would be consistent with the notion that fast mapping is intimately tied to our understanding of categories or kinds, and it raises the intriguing possibility that fast mapping should not apply to proper names, as these refer to individuals and not to kinds.

A different possibility is that fast mapping applies only to information that cannot be accessed through observation. There is, after

Table 2.4
Three-year-olds' and adults' recall of a novel word, a hidden fact, and an observable fact

	Percent Correct		
	Word	Inside Color	External Color
Three-year-olds	87%	81%	50%
Adults	63	75	44

all, a functional motivation for a fast-mapping mechanism to apply in exactly these cases. The location of a sticker is an observable property; this information is "stored in the world," and there is no need to store it in the head. If you forget where the sticker is, just have another look. And in fact, adults are often oblivious to dramatic changes across a visual scene, even for objects that are the direct focus of attention, a phenomenon known as *change blindness*. In one striking demonstration of this, an experimenter started a conversation with a pedestrian and then, during a distraction, was surreptitiously replaced by a different experimenter. Only about half of the pedestrians noticed the change (Simons & Levin, 1998). Unless there is conscious encoding, most observable information does not enter memory.

An object's name and where it comes from, however, are social and historical facts, accessible only through attending to what others say. If you forget this information, you might never encounter it again. For these sorts of facts, fast mapping could be crucial.

We have just begun to test this hypothesis, but one finding is promising. In one study, we showed three-year-olds and adults a novel object, commented on its external color ("You can start with the blue one"), its name ("This is a koba"), and its internal color, which could not be observed ("which is white on the inside"). A week later, subjects were shown a black-and-white picture of the object and asked about its name and its internal and external color. In this situation, more information was actually given during training about the external color: subjects were told about it *and* they could see it. But our prediction was that, since the external color was available to perception and the internal color was not, it would be the internal color that would be fast mapped. The results are as shown in table 2.4.

Subjects tended to recall the object's name and its internal color but did significantly worse on its external, observable, color. These findings are consistent with the view that fast mapping emerges from a general capacity to learn socially transmitted information—including, but not of the contexts in which they are used. For instance, if a child points and says "daw" when a dog enters the room, this would be counted, by a parent or a psychologist, as showing that he or she has learned the English word *dog*. This method of attribution is what Lois Bloom (1970) calls *rich interpretation*.

Considerable debate has been engaged in over the nature of children's earliest words. Many investigators have suggested that the word spurt said to occur later in development is caused by a dramatic change in the ability to learn and understand words; this change makes adult-like word learning possible. Prior to the word spurt, children's first words differ radically from those of adults (e.g., Dromi, 1978; Nelson, 1988). They are not semantically constrained in the right way and are learned in a slow associative fashion, without the support of lexical constraints.

This sort of discontinuity between early and later words is sometimes proposed from the standpoint of a constructivist theory of development, in which children gradually learn, though experience with language, just how words work. But a discontinuity is also consistent with the idea that there is a distinct language capacity that underlies word learning, one that emerges through neural maturation. Such a capacity might not have matured by the time children start to speak. Noam Chomsky (1975, p. 53), for instance, suggests that "It is possible that at an early stage there is use of language-like expressions, but outside the framework imposed . . . by the faculty of language—much as a dog can be trained to respond to certain commands, though we would not conclude, from this, that it is using language."

But does such a discontinuity really exist? Consider the claim that early words have bizarre meanings. Lois Bloom (1973) reports that Allison used the word *car* only when watching cars move on the street below the living room window. Melissa Bowerman (1978) notes that Eva used *moon* to talk about, among other things, a half grapefruit, the dial of a dishwasher, and a hangnail. Eve Clark (1973) gives the example of a child who called a doorknob *apple*. My son, Max, at 20 months, put a slice of yellow pepper on his head during dinner one evening, and said "hat." A month later, when we were perusing a book called *Trucks*, I pointed to the ice cream cone affixed to the top of an ice cream truck, and asked him what it is. I expected him to say "ice cream," a word he knows well, but instead he said "pee-pee," which is his word for penis. And indeed, the cone *was* shaped like a penis, though I hope that wasn't the intent of the illustrator.

First words are also said to blur the semantic distinctions between objects, properties, and actions. Children have been observed to use *hot* to talk about both the property of being hot and also to refer to certain objects, such as ovens and radiators. Esther Dromi (1978), in her careful diary study of the acquisition of Hebrew, found that her daughter's early words frequently referred to both objects and actions, such as the use of the verb meaning "to fly" to refer to birds.

These funny utterances may be signs of an incomplete understanding of word meanings, either because children have not yet figured out how words work or because they possess an immature faculty of language. But there are other interpretations.

First, all these examples are based on observations of productive speech, not comprehension, and some might be speech errors, slips of the tongue. There is independent evidence, after all, that children have greater problems with lexical retrieval than adults and are more error prone in their speech (Marcus et al., 1992).

Second, some of the examples might not be errors at all. When one takes into account that these children don't know many words and have no productive syntax, many of these utterances might be perfectly reasonable. For instance, children have been observed to point at a cookie jar and say "cookie." One might say that this reflects a profound confusion: the child thinks that the word refers not only to cookies but to all cookie-related entities. But perhaps the child thinks that there is a cookie in the cookie jar, doesn't know the words *jar* and *in*, and is expressing her thought in a sensible way (Huttenlocher & Smiley, 1987). Similarly, a child who says "cookie" only while in a highchair might have a bizarre contextualized meaning for the word, in which it can be used only in a certain situation—but might also know that the highchair is the place for eating, and requesting, cookies. And there is nothing at all wrong about calling a stove "hot" or saying "flying" about a bird, so long as one is expressing the opinions that the stove is hot and the bird is flying. After all, if I see a nice car go by, I might say "Nice!" This doesn't mean that I have my object names and property names mixed up.

Related to this, it has often been pointed out that when children call a doorknob "apple," it could mean that they are observing that the doorknob is *like* an apple. Overextensions are especially likely if children don't know the right word for what they wish to talk about. As Lois Bloom (1973, p. 79) puts it, using another example: "It is almost as if the child were reasoning: 'I know about dogs, that thing is not a dog, I don't know what to call it, but it is like a dog!'"

Sometimes it is clear that children are using words in a nonliteral fashion, as when a two-year-old says that wheeling searchlights on top of a building are "like a helicopter" (Macnamara, 1982). And I suspect that when Max put the pepper on his head, he didn't think it was *really* a hat. He found the situation extremely amusing, after all, much more the same question by pointing and saying something like "Wha?," "Tha," or "Eh?" Interestingly, this sometimes occurs early in lexical development. Most of the children Nelson studied had a deictic pronoun that was used in just this way by the time they learned 50 words—and six of the children had one such pronoun among their first 10 words. Going back to Chomsky's speculation about the nature of early speech, it's worth noting that asking about the names of things isn't something that dogs do.

What about the children's ability to learn words? As we will see, young children learn new words at a much slower rate than older children. But it is not clear that they learn words in a different way. There is evidence that when the task is made easy enough, fast mapping occurs even with very young children. In a study by Oviatt (1980), one-year-olds were introduced to a live animal, such as a rabbit, and told its name ("rabbit"). When later asked a question such as "Where's the rabbit?," half of the 12- to 14-month-olds looked and gestured toward the rabbit—but they did not do so when asked about a nonsense word: "Where's the kawlow?"

In another set of studies, 13- and 18-month-olds were told the name of a novel object nine times in a five-minute session ("That's a tukey. See, it's a tukey. Look, it's a tukey."). Another novel object was present and commented on ("Oooo, look at that. Yeah, see it? Wow, look at that.") but was not named. When the children were later asked to point out "the tukey," even 13-month-olds could do so better than chance. Remarkably, they even succeeded with a 24-hour delay between being taught the new word and being tested (Woodward, Markman & Fitzsimmons, 1994).

The Word Spurt

After the first words are acquired, there is often said to be a dramatic change in how words are learned. As Dorothea McCarthy (1954, p. 526) wrote, "After the appearance of the first few words used consistently with meaning in appropriate situations, there occurs a rapid increase in vocabulary." This is sometimes said to occur just with names (in which case it is called a *naming explosion*) or with vocabulary as a whole (and called a *word burst, word spurt*, or *vocabulary spurt*). It is said to occur once children have learned about 50 words, at the age of about 16 to 19 months (Benedict, 1979; Goldfield & Reznick, 1990; Nelson, 1973), though some investigators have found it occurs later, sometimes just prior to the emergence of productive syntax (Dromi, 1987; Mervis & Bertrand, 1995). This spurt might be caused by—among other things—the child's insight that language is symbolic (Dore, 1978; McShane, 1979), the emerging ability to categorize in a mature fashion

(Gopnik & Meltzoff, 1986), the onset of word-learning constraints (Behrend, 1990), or the nonlinear dynamics of a connectionist learning procedure (Plunkett, Sinha, Møller & Strandsby, 1992).

The existence of a word spurt is often presented as an undisputed fact. Several studies look at longitudinal data, find when the word spurt occurs, and then explore its relationship to other facets of language and cognition, such as children's object sorting or their understanding of syntax. But these studies have an odd definition of what counts as a spurt. They count children as going through a spurt when they learn words at a certain *rate,* such as 10 or more new object names in a three-week period (Gopnik & Meltzoff, 1986), 12 or more new words in a three-week period (Lifter & Bloom, 1989), 10 or more words in a two-and-a-half week period (Goldfield & Reznick, 1990), or 10 or more new words in a two-week period, at least five of which are object names (Mervis & Bertrand, 1995). Reznick and Goldfield (1992) have a similar criterion for a word spurt in comprehension: two new words in 2.5 weeks, taken from a list of words that are expected to be difficult.

What's wrong with these criteria? Note first that, by adult standards, children who have had a word spurt according to these investigators are still learning words at a relatively slow rate—less than a word a day. If they stuck to this rate, they would end up with an adult vocabulary of about 4,500 words—a far cry from the 60,000 estimate. So when one defines *word spurt* in this way, it is no longer an empirical question whether it exists. Its existence is a mathematical necessity, simply because 17-year-olds know far more than 4,500 words.

A more serious problem is that the criteria above have nothing to do with a spurt (or burst or explosion) in any normal sense of the term. To see why, consider figure 2.2, which represents the rate of vocabulary development in two imaginary children.

Using the Mervis and Bertrand (1995) criteria, both children go through a word spurt in the period right before they reach 18 months; they each learn 10 new words. But only Zack undergoes a spurt in any meaningful sense of the term. Zoe shows a gradual increase in the rate of word learning. The criteria used by these investigators tell us when children start to learn words at a certain rate (something that may be of interest for other purposes); they tell us nothing about the nature of the change in vocabulary growth.

This is more than a quibble about terminology; it really matters. The theories of the word spurt mentioned above all assume that *something interesting* happens at the point of the word spurt. This seems true for Zack. But it is not at all true for Zoe; nothing interesting happens to her prior to 18 months, nothing that distinguishes this point in time from the period before and the period after.

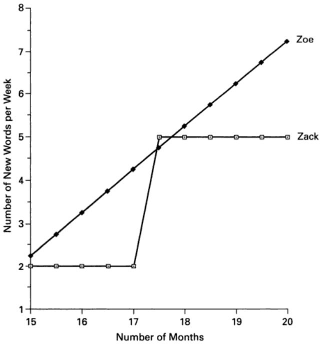

Figure 2.2
Changes in the vocabulary growth of two imaginary children

Consider an analogy. If Joe has eaten 40 french fries in 10 minutes, there has to be some period of time in which he eats french fries at the rate of more than 2.5 per minute. But this does not entail a "starch spurt" or "french fry explosion." One *might* occur: Joe might nibble at his fries at a leisurely rate for the first nine minutes and 30 seconds and then gobble down the rest in the remaining half minute. But it is also possible that Joe could eat his fries at a constant rate. Or he could eat one fry in the first minute and then slowly speed up his rate of fry eating. There would be no spurt, just a gradual increase. To point to the moment he starts eating fries at the rate of 2.5 per minute and say "Aha! That's an eating explosion" is worse than bad terminology. It leads to bad theorizing, since it gives the false impression that something special is happening at this point, something that has to be explained.

In a perceptive discussion of these issues, Jeffrey Elman and his colleagues (1996) suggest that vocabulary development might in fact

Awarded the 2002 Eleanor Maccoby Book Award,
American Psychological Association

How Children Learn the Meanings of Words
Paul Bloom

How do children learn that the word "dog" refers not to all four-legged animals, and not just to Ralph, but to all members of a particular species? How do they learn the meanings of verbs like "think," adjectives like "good," and words for abstract entities such as "mortgage" and "story"? The acquisition of word meaning is one of the fundamental concerns of the study of mind.

According to Paul Bloom, children learn words through sophisticated cognitive abilities that exist for other purposes. These include the ability to infer others' intentions, the ability to acquire concepts, an appreciation of syntactic structure, and certain general learning and memory abilities. Although other researchers have associated word learning with some of these capacities, Bloom is the first to show how a complete explanation requires all of them. The acquisition of even simple nouns requires rich conceptual, social, and linguistic capacities interacting in complex ways.

This book requires no background in psychology or linguistics and is written in a clear, engaging style. Topics include the effects of language on spatial reasoning, the origin of essentialist beliefs, and the young child's understanding of representational art. The book should appeal to general readers interested in language and cognition as well as to researchers in the field.

Paul Bloom is Professor of Psychology at Yale University.

"*How Children Learn the Meanings of Words* is a tour de force, by far, the best book written on the subject."
—Steve Pinker, Peter de Florez Professor of Psychology, Massachusetts Institute of Technology

"A state-of-the-art description and discussion of the complexity of first-language vocabulary acquisition, this book is as clear, thorough, and unassuming as its title suggests."
—*Choice*

"A tremendous book that brings together a huge body of findings in a way that is both readable and scholarly."
—*Perception*

Learning, Development, and Conceptual Change series

A Bradford Book

The MIT Press
Massachusetts Institute of Technology
Cambridge, Massachusetts 02142
http://mitpress.mit.edu

0-262-52329-9

90000

9 780262 523295