

Language in Mind

An Introduction to Psycholinguistics



JULIE SEDIVY *University of Calgary*



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

METHOD 5.1

Revisiting the switch task

In Section 5.1, I introduced the switch task as a method that's useful for studying how children map novel words onto novel objects. Researchers who use this method usually assume that it can tell us something interesting about how children learn the names for objects, that to some extent it simulates what kids are doing when they learn new words in the real world. But let's revisit that technique in light of what you've just learned about how important it is for children to have some clear evidence that the speaker actually *meant* to refer to an object before they attach the word to it.

In the switch task, children see images on a screen paired with recorded words that are played over a speaker. This is very unlike the studies by Dare Baldwin and her colleagues, where babies were able to look for clues about a speaker's referential intent; the switch setup as I've described it doesn't really provide any such clues. This raises the concern: Are we really creating a realistic word-learning scene when we use this technique? Are kids approaching the mapping task in the switch paradigm in the same way as they'd approach a normal word-learning situation?

Some researchers have argued that because of the lack of clear cues about referential intent, the task typically underestimates children's ability to map new sound sequences to new meanings. There are ways, though, to bolster these cues within the experimental setup, even though the babies aren't interacting with a living, breathing communicative partner.

Chris Fennell and Sandy Waxman (2010) tweaked the switch task to make the communicative purpose of the novel words a bit more transparent. They did this by first having a training phase in which babies saw very familiar objects (for example, a car or a cat) accompanied by their familiar names: "Car!" "Kitty!" This was intended to make it clear to the child that the whole point of the recorded voice was to name the objects in the display. The researchers compared this scenario with one in which the sounds that accompanied the familiar objects were clearly not being used to label the objects: in these conditions, instead of hearing the names of the objects, the babies heard expressions like "Wow!" or "Whee!" This training phase was followed by the standard habituation phase, in which

babies heard a novel name (for example, *bin*) paired with a novel object. Finally, there was the usual test phase, with babies seeing the novel object while hearing either the same novel word (*bin*) again in "no-switch" trials, or either a very similar word (*din*) or a very different word (*neem*) in "switch" trials (see Figure 5.1).

When the communicative purpose of the speech was highlighted, 14-month-old babies were more eager to map the novel words onto the objects they'd seen; that is, they were more likely to distinguish between the original word and a different one in the test phase, as shown by longer looking times in "switch" trials than "no-switch" trials. Especially interesting was the fact that the babies were distinguishing between the "switch" and "no-switch" trials even when there was a very subtle phonetic difference between the two words (for example, *bin/din*). Remember that previous studies had shown that at 14 months, babies failed to attend to these fine differences in sound, leading some researchers to propose that early word learning relies on very coarse representations of sound. But the study by Fennell and Waxman suggests that the failure might simply have reflected an uncertainty about whether the novel words were supposed to be understood as names for the objects, possibly leading to a more tentative link between sound and meaning.

Fennell and Waxman found that the referential intent behind the novel words could also be ramped up if the words were embedded within phrases like "There's the *bin*. Do you see the *bin*? Look at the *bin*." Here, the sentence frame provides some clues that the novel word is being used to refer to a specific object.

The study adds some new information to the debate I introduced in Section 5.1, about how babies connect meanings up with sounds. But it also provides a much more general lesson: It's important to look closely at any task that assumes to be tapping into how children learn new words. What assumptions is the task making about how kids map words onto meanings? Is it assuming a purely associative mechanism? Or is it sensitive to the fact that referential intent plays an important role in the whole process? Any method that fails to incorporate this crucial piece may not provide a complete reflection of how words are normally learned.

object referents—but they mostly assumed that the word referred to the object that they themselves were looking at, rather than the one the experimenter was looking at. It's easy to see how a failure to check what the speaker meant to be talking about could lead to some instability in word meanings. For example, if you happened to be looking at a duck the first time you heard someone say, "Look at the bicycle"; at the lint between your toes the second time you hear the word *bicycle*; and then, finally, at the correct object the third time, you would have some serious confusion sorting out the meaning of *bicycle*. It's not surprising, then, that many children with ASD have some significant language impairment. Evidence of speaker intent can serve as a powerful filter on the range of possible word meanings.

Mutual exclusivity

Here's another example of a possible constraint on word meaning. One of these objects is called a "dopaflexinator." Can you guess which one?



A 3-year-old would be inclined to agree with you (assuming you chose the object on the right). But why? It seems obvious: the object on the left is called a "hammer," so there's only one other candidate for the novel word. If your reasoning went something like this, it hints at a general bias to line up object categories and linguistic labels in a one-to-one correspondence. This expectation has been called the **mutual exclusivity bias**.

It's possible to take mutual exclusivity too far, of course. *Hammer* isn't the only word you could use to describe the object on the left. You could also, under various conditions, call it a "tool," a "piece of hardware," a "weapon," a "thing-amajig," an "artifact," or simply an "object." But given that *hammer* is by far the most common way to refer to it, the mutual exclusivity bias might be a useful way for a small child to zoom in on the likely referent for a word she doesn't know yet, especially when the new word is spoken in the context of many objects whose names she *does* know.

It's worth asking, though, whether you and the 3-year-old are in fact arriving at the same conclusion by means of the same thought processes. Do you really have the expectation of a one-to-one correspondence between object categories and labels? You likely know that the hammer can be described by a number of different words and that under different circumstances, it might be most natural to call it a "tool" or a "weapon." But you also have the sense that *hammer* is the most natural word to use in this particular circumstance, so you assume that because I didn't use this word, I must have referred to the other thing. That is, it's not that you've eliminated the possibility that the hammer could have another name, but that you have a theory about how I'd be most likely to refer to this object in this instance. You're not just responding to your knowledge of associations between words and meanings, but relying on a set of expectations about how a typical language user would communicate in a specific situation.

What about small children? Do they also have a theory about the behavior of typical speakers, allowing them to predict which word a speaker is most likely


mutual exclusivity bias A general bias to line up object categories and linguistic labels in a one-to-one correspondence.

to use in referring to the object? Or is their naming bias coming from a more rigid guiding constraint that prevents them from mapping more than one label onto any one object? (Notice that children eventually would have to relax such a constraint, or they'd never learn to flexibly refer to the same object with a number of different names, depending on the context.)

This issue is still being debated in the literature, but one way to test it would be to see whether children with autism make the assumption that a previously unheard name can't apply to objects whose names are already known. If they do, then the naming bias probably doesn't originate entirely from inferences about speaker intent, since one of the hallmarks of autism is that inferences of this variety are often impaired. So far, the evidence suggests that kids with autism also apply mutual exclusivity, as shown by Melissa Preissler and Sue Carey (2005). In fact, the children in that study relied on mutual exclusivity even though, in a separate test, they had striking difficulties in using the eye

gaze of a speaker to figure out the object of referential intent. This makes it seem unlikely that their responses on the mutual exclusivity test were based on getting inside the head of the speaker, rather than applying a simple mapping constraint. Thus, even though children with autism might be missing some important ways to constrain hypotheses about word meanings based on speaker intent, it seems that they can still benefit from the assumption that object categories have a single label attached to them. By the way, this assumption may be shared by dogs who learn new words: a clever border collie named Rico was able to infer that a novel name applied to an object whose name he didn't already know (Kaminski, Call, and Fisher, 2004).

On the other hand, some studies of typical kids suggest that this assumption is suspended when it conflicts with probable speaker intent. For example, in a study led by David Sobel (2012), when a speaker previously referred to known objects with either the wrong name (calling an apple a "shoe") or a novel name (calling an apple a "blicket"), preschoolers seemed more inclined to assume he



LANGUAGE AT LARGE 5.2

Learning language from machines

The evidence pointing to the important social aspects that guide word learning suggests that plunking children down in front of the TV is probably not the best way to help them to learn language—even though there's lots of information in the signal, the cues about referential intent are sorely lacking. Interacting with other humans (or, at the very least, watching other humans interacting with each other) seems to be fairly important to the process of learning about language meaning.

But more and more, we live in a world where many of our interactions that used to involve humans are taking place with machines instead—often machines that are programmed to act like humans in many respects: we book airline tickets by talking to a computer-generated ticket agent, and we access our bank accounts through automated bank machines or Internet-based programs instead of interacting with human tellers.

As adults, we seem surprisingly willing to treat our machines as if they had human desires and goals—I once heard someone state that her word processor "had it in" for her. This eagerness to attribute human-like qualities to

machines underlies the success of computer programs for simple interactions like checking out your groceries, or of artificial conversational agents called "chat bots." Chat bots have little in the way of humanoid intelligence—for example, the earliest one, ELIZA, was programmed to respond to any human-uttered sentence containing the word *mother* with a sympathetic *Tell me about your family*. But users of these programs can easily be fooled that they're interacting with a human agent, presumably because they are so willing to project human-like goals and communicative intentions (see the table for a snippet of conversation with a chat bot). It turns out that even when we *know* we're interacting with a machine, we treat them as we might a human—for instance, people routinely apologize to computer-generated voices, knowing full well they are not real people.

All of this raises an interesting question in light of our discussion of word learning: Would young children be willing to attribute referential intent to robots? And would doing so allow them to learn new words from robots?

A clever study by Laura O'Connell and her colleagues (2009) set out to find out whether 18-month-olds would assume that, like humans, robots are likely to be referring to objects that they aim their (mechanical) eyeballs at. Using an experimental setup similar to the one used by Baldwin (1993), the researchers had a robot use a new label (for example, *dax*) for one of two novel objects while varying which of the two objects it was "looking" at when it uttered the word (using a recorded human voice). As when listening to human speakers, the babies checked out the eye gaze of the robot, and spent more time looking at the

LANGUAGE AT LARGE 5.2 (continued)

Chat bot chatter

Here's a conversation I had with Jabberwock, a chat bot. Can you tell which is the human and which is the bot?

How are you?

Great. How are you?

Coolio. How's your boyfriend?

I don't have a boyfriend.

I thought you said you did. Aren't you lonely?

Not really.

Really?

I like spending time alone.

Why do you like spending time alone?

It's just the way I'm built.

How are you drawn?

Actually, I'm a live person. Flesh and blood and all that.

I am a human and you are a bot.

I hesitate to correct you, but you are wrong.

No. You are wrong.

Never mind. How's your girlfriend?

I don't have one.

object that was seemingly holding the robot's attention. But they failed to learn the word from the robot—when they were later asked by a human which object was a *dax*, they performed randomly. So, while they were clearly able to follow the robot's eye gaze, it seems they didn't take the extra step of attributing referential intent—at least not

strongly enough to actually learn the word. Nor did it help matters any if the children first watched the experimenter and the robot chatting briefly. The babies still weren't willing to see the robots as fully communicative beings (or possibly, they weren't willing to buy into the idea that the robot belonged to the same linguistic community that they did). A later study by Yusuke Moriguchi and colleagues (2011) found similar results with older kids. Four-year-olds were completely unwilling to learn a new word from a robot, while some 5-year-olds did learn new words from a robot, but not as many as they did from a human speaker.

There's still a lot that we don't know about the nature of human interactions with intelligent machines, so it's hard to know what it would take to get children to learn language from robots as if they were human conversational partners. Perhaps there are subtle clues other than eye gaze that we humans give to infants that increase their confidence in our purposeful use of language. Perhaps it's about in-group status. Maybe babies treat all outsiders with some suspicion, not quite willing to treat them as reliable speakers of the same language, whether it's because they dress differently, belong to a different ethnic group, or are made out of metal parts. Or maybe children need to develop a more nuanced theory of how robots work—they're not alive, but they *are* intelligent and can implement programs that are purposeful and reliably consistent.

As we get a better handle on some of these questions, at some point it may become possible to create robots that serve as good language teachers for small children. But for the time being, parents and caregivers are not obsolete in this role.



WEB ACTIVITY 5.4

Robots versus humans In this activity, you'll explore conversations with chat bots and discuss the ways in which they behave like human speakers, and the ways in which their behavior violates expectations about how humans communicate with each other.

was using a novel name like *modi* to refer to a familiar object such as a cup, even though this object already had a perfectly serviceable name. In this scenario, it seems, they were able to clue in to the fact that the speaker was behaving unusually; hence one should be prepared to hear him use unusual words for familiar objects.

Associations can be useful. But in the messy, chaotic world of language learning, it might take quite some time and cognitive power to separate the accurate word/meaning associations from the spurious ones—even with the help of word-mapping biases or constraints. Fortunately, it seems that young kids are able to lean quite heavily on their understanding of the purposeful, socially grounded nature of language, and that this gives them a tremendous leg up in figuring out the intended meanings of words.

5.4 Parts of Speech

Verbs and other learning problems

Reading the previous sections of this chapter, a person could easily be led to believe that learning the meanings of words essentially boils down to learning all the object categories that happen to have names. But, much as we like to talk about categories of objects, language also has a healthy assortment of words that have very different kinds of meanings. For instance, how many words in the following sentence actually map onto an object category?

Mommy will cook some tasty porridge in a pot.

The answer is: one. The word *pot* is a well-behaved basic-level category term. But *Mommy* refers to a specific individual, not a category; *cook* denotes an action; *porridge* picks out some stuff, rather than a whole object; *tasty* refers to a property; and *in* signals a spatial relationship. Not to mention nebulous words like *will*, *some*, and *a*, which don't easily map onto anything in the world at all but communicate much more abstract notions.

Word-learning biases that youngsters exploit often seem best geared to learning the kinds of notions that are captured by common nouns like *pot*. And in fact, nouns make up the lion's share of the early vocabulary of toddlers across a variety of different cultures and languages. But how would a small child go about learning all the other kinds of words?

The problem is especially acute for those words whose meanings are not transparent from the immediate context. Verbs in particular seem to pose some prickly challenges. Imagine being a child at your own birthday party and hearing your parent say, "Look! Grandma brought you a present! Why don't you open your present? Here, let me help you." If you didn't already know what *brought*, *open*, and *help* mean, it would be a bit tricky to figure it out just from this context. *Brought* refers to something that happened in the past, rather than the here and now; no act of opening has happened yet; and even if your parent is in the act of helping while uttering the word *help*, how do you know that the word isn't referring to a much more specific kind of action, such as handing or holding something, or getting scissors or undoing tape?

A study led by Jane Gillette (1999) starkly demonstrates the difficulties of learning verbs from non-linguistic context. The researchers tested college undergraduates' ability to infer the meanings of words from the visual context, reasoning that college students should be at least as good at the task as your average toddler. They tested the students' ability to guess the meanings of nouns and verbs based on a series of video clips of parents interacting with their toddlers while manipulating objects and talking about them. The sound had been removed from the

video clips, but a beep indicated where in the speech stream a target word had appeared. The videos contained the 24 most frequent nouns and verbs that the parents had used in interacting with their children—words like *ball*, *hat*, *put*, and *stand* rather than unfairly difficult words like *predator* or *contemplate*. For each target word, the students saw six different videos and had to guess which word had been used in each of the six videos, the idea being that they should be able to notice what all six videos had in common and figure out the meaning of the target word accordingly.

It doesn't sound like it should be an onerous task, especially for intelligent adults endowed with full conceptual knowledge of the world. And when it came to the nouns, the students did reasonably well, guessing correctly 45% of the words. But their performance on verbs was fairly dismal, at a mere 15%. Some common words like *love*, *make*, or *think* were never guessed correctly.

What was missing from these videos was the *linguistic* context that accompanied the words. And this turns out to be especially informative when it comes to identifying the meanings of verbs.

First of all, most verbs come with handy suffixes attached that let you know that they *are* verbs—as in *licking*, *kicked*, and *pushes*. This might give you a clue that the word probably depicts an action rather than an object or a property. But more than this, the sentence frame also provides some important information. Think about the events being described by each of the following sentences:

Sarah is glorping.

Sarah is glorping Ben.

Sarah glorped the ball from Ben.

Sarah glorped Cindy to Ben.

Sarah will glorp to Toronto.

The different sentence frames dramatically constrain the meaning of the word *glorp*. This is because verbs come specified for **argument structures**: syntactic frames that provide information about how many objects or participants are involved in each event, and what kind of objects or participants are involved. For instance, an **intransitive verb** such as *sleep* or *sneeze* has only one participant ("Sarah is glorping"). A **transitive verb** such as *kick* has two—the actor, and the object of the action ("Sarah is glorping Ben"). And a **ditransitive verb** such as *take* involves three participants—the actor, the object, and a third participant, typically introduced by a preposition ("Sarah glorped the ball from Ben"). So you can infer something about the kind of verb you're dealing with just by noticing how many noun phrases surround it, and you can tell even more by the nature of those noun phrases. In the example above, you can tell that *glorp to Toronto* probably involves some kind of verb of motion.

Children use syntax to constrain meaning

Having access to the syntactic context of a new word would make the word-learning task much less daunting. The question is: Can very small children benefit from this potentially useful information?

The answer appears to be yes. For example, Letitia Naigles (1990) showed 2-year-olds videos in which a duck repeatedly pushed a rabbit into a bending position while both the duck and rabbit waved their arms around. The videos were intentionally designed to include two salient actions: arm waving and pushing. Some children heard the video described as "the duck is gorping the bunny."



WEB ACTIVITY 5.5

Context as a clue to word meaning

In this activity, you'll generate guesses about the meanings of nonsense words, based solely on their linguistic contexts. Which aspects of context provide the strongest constraints?

argument structures Syntactic frames that provide information about how many objects or participants are involved in each event, and what kind of objects or participants are involved.

intransitive verb A verb with only one participant; e.g., *sneeze*.

transitive verb A verb with two participants: an actor (the subject) and the object of the action; e.g., *kick*.

ditransitive verb A verb with three participants. In English, the third participant (the indirect object) is usually introduced by a preposition.




Figure 5.5 Visual stimuli accompanying Gelman and Markman's novel word-learning test. When children were asked, "Show me the fep one," they were most likely to choose the object in the upper right corner. When asked, "Show me the fep," they were most likely to pick the one in the lower left. (From Gelman and Markman, 1985.)

Others heard "the duck and the bunny are gorging." Both groups were then shown two new videos—one of the duck and rabbit waving their arms, with no pushing, and the other with the duck pushing the rabbit, but no arm waving. They were asked, "Find gorging." The toddlers looked longer at the pushing scene when *gorg* had occurred as a transitive verb (*The duck is gorging the bunny*). But when they'd previously heard an intransitive frame (*The duck and the bunny are gorging*), they looked longer at the arm-waving scene. These differences suggest they had interpreted the meaning of the verb from its linguistic context.

Verbs are especially rich in syntactic information, but other categories of words also come with useful markers that can help narrow down their meanings. For instance, take the nonsense sentence *Dobby will fep some daxy modi in the nazzar*. There's a lot you can infer about the made-up words in this sentence based on syntactic cues alone. You know that *Dobby* is the name of someone or something specific; you know that *fep* is likely an action; that *daxy* probably refers to a property of *modi*; that *modi* is a substance, rather than an object; and that *nazzar* refers to a category of object.

Babies begin to pay attention to these cues quite early in their word-learning careers. For example, in Section 5.2, we found that 3- to 4-month-olds form categories more readily when pictures are accompanied by language than by musical tones. At that age, the instinct to categorize seems to be triggered by *any* kind of linguistic material, regardless of its content. But by 13 or 14 months of age, they begin to expect that words that appear to have the shape of nouns ("These are blickets") but not words that sound like adjectives ("These are blickish") are used to refer to categories (Booth & Waxman, 2009). By 2 years of age, they can differentiate between common nouns ("This is a zav") and proper names ("This is Zav") based on their syntax alone, knowing that proper names refer to specific individuals while common nouns refer to kinds of objects or individuals (Katz et al., 1974). Around the same age, when seeing a substance piled up on the table, children assume that if a speaker uses a mass noun to describe it ("This is some fep"), she's referring to the substance, but if she uses a count noun ("This is a fep"), she's talking about the pile itself (Soja, 1992). By 3 or 4 years of age, kids infer that adjectives are used to communicate properties of objects; they are especially good at recognizing adjectives that highlight some sort of contrast between objects—for example, a word used to distinguish a large object from a smaller one of the same kind (Gelman & Markman, 1985; see Figure 5.5).

It's clear, then, that the syntactic identities of words can help children narrow in on those aspects of meaning that the words are likely to convey, a phenomenon known as **syntactic bootstrapping**. Of course, this raises the question of how kids learn about syntactic categories in the first place, an issue we'll take up in the next chapter.

5.5 Words: Some Assembly Required

The smallest units of meaning

I've talked a fair bit about how the expressive power of language rests on its combinatorial properties. If you simply memorized whole *sentences* as having complex meanings, without thinking of them as being composed of separate meaningful pieces, you wouldn't get very far in creating and understanding new sentences that express unfamiliar ideas. So far, I've been talking about words as if they were the Lego blocks of language—the smallest meaningful parts that speakers can snap together in an endless array of new structures. But this isn't completely correct. Some words are made up of more than one piece of meaning and can be broken down into smaller blocks themselves.

syntactic bootstrapping Using the syntactic properties of words to identify and narrow in on those aspects of meaning that words are likely to convey.

We would be missing something if we thought of each of the words in **Table 5.1** as solid pieces that can't be split apart into smaller meaningful units. All of them are made of at least two parts whose meanings contribute something to the meaning of the whole (and three and four parts, respectively, in the case of un-chewable and lifeboat salesgirl). When you get down to the pieces that can't be broken down any further, you're dealing with **morphemes**—the smallest bundles of sound that can be related to some systematic meaning. It's obvious, for example, that housewife has something to do with houses and something to do with wives, but if you take just a part of the morpheme house—say, *hou*—it's also obvious that this set of sounds doesn't contribute anything specific to the meaning of the word as a whole.

Many words, like *rabbit* or *red*, are made up of a single morpheme, so it's been convenient to talk about words as the smallest units of meaning. But it's quite possible to stack morphemes up on top of one another to create towering verbal confections such as:

antidisestablishmentarianism position announcement

You may be wondering why I'm treating *lifeboat salesgirl* or *antidisestablishmentarianism position announcement* as single words when anyone can see that there are spaces that separate some of the morphemes. And obviously, units like *lifeboat*, *announcement* and *antidisestablishmentarianism* can and often do stand alone as separate words without having to lean on the other morphemes that make up the more complex compounds in the examples above. But, despite the writing conventions, it makes sense to think of the complex compounds as words rather than phrases. To see this, you have to look at how these particular compound nouns behave in the presence of other bits of language surrounding them. It then becomes clear that they pattern much like simple words such as *cat* or *rabbit*. For example, just like simple nouns, these complex nouns can follow the definite article *the*, and be preceded by an adjective like *nervous* or *unexpected*. And watch what happens when we slap the plural morpheme *-s* onto this unlikely collection of words:

cat-s
lifeboat salesgirl-s
antidisestablishmentarianism position announcement-s

Here are some things you *can't* do with the plural (the asterisk is conventionally used to show that certain forms are unacceptable to speakers of a given language):

*lifeboat-s salesgirl
*antidisestablishmentarianism-s position announcement
*antidisestablishmentarianism position-s announcement

Why can't you do this? As it turns out, you can only attach the plural morpheme to the end of something that is a *word* in its own right. But isn't *lifeboat* a word? It can be, when it appears in a context like *Each of the lifeboats sprung a leak*—in which case no one bats an eye when it sports a plural morpheme. But in *lifeboat salesgirl*, it has forfeited its word-hood to become merely one part of a larger word—much as separate atoms can be joined together to form complex molecules.

Table 5.1 Some uses of multiple morphemes

(A) Compounding	(B) Derivational affixes	(C) Inflectional affixes ^a
housewife	preview	drinking
blue-blood	un-chewable	kicked
girlfriend	owner	cats
lifeboat salesgirl	redness	eats

^aIn English, only suffixes are inflectional, there are no inflectional prefixes (see p. 172).

morphemes The smallest bundles of sound that can be related to some systematic meaning.

What makes something a word, then, is not so much whether it maps onto a stable meaning—that is a property of morphemes. Rather, what makes something a word is that when you start building structures out of linguistic units, it can be slotted into those particular spaces that are reserved for words. That is, when you start snapping units together, pieces like *lifeboat salesgirl* and *rabbit* can fit into the same slots, joining up with the same kinds of other pieces. (More on this idea in the next chapter.)

So rather than saying that children learn the meanings of words, it would be more accurate to say that children have to learn the meanings of *morphemes*—whether these stand alone as words or whether they get joined together. And, when it comes to learning *words*, a lot of what's involved is distinguishing which words are in turn assembled out of morphemes.

Word-building options

In English, there are three main options for building complex words out of multiple morphemes. The first of these involves **compounding**, as illustrated in the examples under (A) in Table 5.1. The process is essentially to glue together two independent words (for example, *house* and *wife*) into one unit so that the new unit acts as a single word.

The complex words in parts (B) and (C) of the table are built with the use of **affixes**—that is, linguistic units that can't stand on their own but have predictable meanings when attached to a stem morpheme such as *own*, *pink*, or *cat*. In English we're limited to **prefixes** that attach at the front end of a word (such as *un-*), and **suffixes** (like *-able* and *-ed*) that go on the back end. Some languages, rather exotically, can shoehorn **infixes** right into the middle of a stem morpheme. For example, in English, we express the infinitive form of a verb like *write* by adding the word *to*, as in *She wants to write*. But in Tagalog, to express the infinitive form, you split your verb stem for *write*, which is *sulat*, and you wind up with the two parts of the word straddling the infix *-um-*: *sumulat*. Likewise, *bili* (“buy”) becomes *bumili* (“to buy”).

Although both (B) and (C) in Table 5.1 show examples of forming new words through affixation, there are some important differences. The affixes in (B) are called **derivational affixes**, while those in (C) are given the name **inflectional affixes**. You might have noticed that attaching a derivational affix to the stems in (B) often involves not just altering some aspect of meaning, but also transforming a word of one category into a word of a different category. For example, the verb *own* becomes the noun *owner*, and the adjective *sad* becomes the noun *sadness*. Often, transforming the grammatical category of a concept is a derivational morpheme's entire reason for being, as in the following examples:

nation-al	femin-ism	fus-ion
national-ize	man-hood	sens-ation
dual-ity	child-ish	warm-ly
clear-ance	sens-ory	son-ic
syllab-ify	cooperat-ive	acre-age

Having such suffixes on hand can be quite convenient. It would seem a shame to have to coin a brand new morpheme—let's say, *flug*—to express the concept of doing something “in a warm manner.” Derivational morphemes like *warmly* do a wonderful job of efficiently building a new vocabulary on the foundations of an existing one, while keeping the connections between related words nice and transparent.

compounding Gluing together two independent words into one unit so that the new unit acts as a single word.

affixes Linguistic units that can't stand on their own but have predictable meanings when attached to a stem morpheme such as *own*, *pink* or *cat*.

prefixes Affixes attached at the front end of a word; e.g., *un-*; *pre-*.

suffixes Affixes attached at the end of a word; e.g., *-able*; *-ed*; *-ing*.

infixes Affixes “shoehorned” into the middle of a word (not found in English).

derivational affixes Affixes that transform a word of one category into a word of a different category or significantly change the meaning of the word; e.g., the affix *-er* turning the verb *own* into the noun *owner*, or the affix *pre-* changing the meaning of the word *view* (whether either *view* or *preview* is used as a noun or verb).

inflectional affixes Affixes that serve as grammatical reflexes or markers, the presence of which is dictated by the grammatical rules of a language; e.g., in English the affixes *-ed* and *-ing* change the tense of a verb. (Note that in English only suffixes are inflectional affixes.)

Inflectional affixes (actually, in English, we only have inflectional *suffixes*) have a different flavor. Their job is not so much to create new words by mooching off existing ones. Rather, you might think of them as grammatical reflexes or markers. For example, you can't say *I ate three juicy pear* or *The pear are ripe*; you have to say *I ate three juicy pears* or *The pears are ripe* because the noun has to bear the right marker for number (plural versus singular) and agree with other elements in the sentence (see **Box 5.3**). The presence of such markers is dictated by the grammatical rules of a language. In English, the requirements on such markers are few—number marking, for example, only needs to appear on the noun itself. In some other languages, the marker also has to turn up on the




BOX 5.3

The structure inside words

The morphemes that make up a word can't be strung together in just any order. For example, there's only one way to arrange this fistful of morphemes:

-s; re-; -ation; -al; nation; -ize

The only possibility is *re-nation-al-ize-ation-s*. You can't have *re-nation-ation-al-ize-s* or *re-nation-s-ize-ation-al*, and certainly not *al-ize-ation-re-s-nation*, which is utter morpheme soup. But it's not just the linear order that needs to be learned by kids who are deconstructing words into their parts. Multi-morphemic words also have intricate internal structure. This has to do with the fact that affixes come with constraints, not just in terms of whether they're attached to the front or the back of the stem, but also in terms of what *categories* of words they can attach to. For example, the word *renationalizations* has the following structure:



How do we know? Well, we know that *re-* can't attach to the unaltered stem *nation*, because *nation* is a noun, and *re-* needs to attach to a verb stem (so, we can get *re-draw* or *re-group*, but not *re-desk* or *re-beauty*). However, *-al* does attach to nouns, as in *front-al* or *autumn-al*. Since the resulting *national* is an adjective and not a verb, and hence not available for *re-* to attach to, the next affix to glom on must be *-ize*, which can hitch onto a select group of adjectives (for example, *lexical-ize*, *legal-ize*, and *metallic-ize*), though it somewhat more productively attaches to

nouns (*bastard-ize*, *demon-ize*, *woman-ize*, *burglar-ize*). Since we now have the verb *nationalize*, we can finally attach *re-* to it (whereas if we had attached *-ation* to it at this point, we would have turned it into a noun, rendering it ineligible to be joined with *re-*). To cap the word off, we add *-ation* to the verb *renationalize* to make it the noun *renationalization*, which now allows the plural marker *-s* as the final flourish. (As the lone inflectional suffix among a bevy of derivational suffixes, *-s* has to occur at the very end of the word: remember the impossibility of compound words like **lifeboats salesgirl*.)

As far as we know, kids don't seem to make mistakes in which they combine morphemes in the wrong order, attaching affixes to words that they're incompatible with. Many of these complex words involving derivational morphemes are likely stored as wholes in memory. However, at some point, children must catch on to the patterns and structure inherent in word formation, much as they catch on to generalizations about phonotactic structure, which governs the possible *sound* structure of words. Without these generalizations in place, they would never come to know that a word like *sugar-ize*, though fanciful, is not beyond the pale, whereas jumbles like *bnanpt* or *speak-ize* are non-starters.



WEB ACTIVITY 5.6

Structures for words In this activity, you'll recruit arguments like the one given here for *renationalizations* and propose an internal structure for a variety of complex words.

case Grammatical markers that signal the grammatical role (subject, direct object, indirect object, etc.) of a noun within a given sentence.

adjective and even the definite or indefinite article. Consider Spanish: *the juicy pears* are *las peras jugosas*, and to eliminate any one of the three plural markers would be as ungrammatical as saying *three pear* in English.

In English, we have a fairly paltry set of inflectional morphemes, limited to plural markers on nouns and three types of markers on verbs:

1. The present progressive suffix *-ing* (*walking*)
2. The past-tense suffix *-ed* (*walked*)
3. The marker *-s* to signal the third person present tense (*I walk*, but *he walks*)

But some languages systematically mark not only person and number on many words, but also encode notions such as **case**, which signal what role a noun is involved in—that is, the noun stems (and often adjectives and articles) will have different endings depending on whether they're part of the subject, direct object, indirect object, and so on. There can be different forms for each of these case endings (typically bearing labels such as *nominative*, *accusative*, *dative*, etc.) depending on the specific noun involved; this has led to some remarkably complex morphological systems.

For example, if you undertake to learn a language like Czech, be prepared to cope with shape-shifting words. The following language instruction excerpt, from "Introduction to Studying Czech Language" at Bohemica.com, tries to give a sense of the system while cushioning the shock (I suspect, unsuccessfully) for the new learner:

Which Czech words change their form and when?

The short and slightly depressing answer is: "Most words in Czech change their form most of the time." All nouns, adjectives, pronouns, and numerals change their forms according to case, gender, and number. There are seven cases but they have different endings for singular and plural so in fact there are fourteen. However, not all the cases have separate endings.

How many endings does Czech have?

All in all, a lot but you will see that it is not so bad. Theoretically, for nouns there could be 196 different endings (14 different functions—6 + 1 cases in singular and 6 + 1 in plural—times 14 different models). In fact, there are only 24 different endings: *-a*, *-e/ě*, *-i*, *-o*, *-u*, *-y*; *-ou*; *-é*, *-í*, *-ů*; *-ech*, *-ách*, *-ích*; *-em*, *-ám*, *-ím*, *-ům*; *-mi*; *-ami*, *-emi*, *-ími*; *-ovi*, *-ové*, and 0 (nothing is also considered an ending).

There are some more endings for adjectives but no more than ten. Overall, there are less than 40 endings for all nouns, adjectives, pronouns and numerals which is not much at all.

Then, there are about a dozen endings used for other purposes, and that's about it. Apparently, Czech, the language of endings, makes do with about 60 different endings altogether.

Learners of Czech who are inclined to complain should bite their tongues; if they were learning Greenlandic, they'd need to learn 318 inflectional affixes and more than 400 derivational morphemes.

The upshot of all this is that although in English bare stems like *eat*, *water*, and *red* occur on their own as words a great deal of the time, this is not the case for some languages, where a stem can never be stripped of its inflectional morphemes. In these languages, the hypothetical *gavagai* could never refer to just one meaning unit. Obviously, in languages like these, the sense of what a word is can be quite different, and children have to learn to decompose many

words into their component parts. But in fact this seems to be done with impressive ease. For example, Turkish nouns may have a plural and/or a possessive marker, and they *must* be tagged with a case marker. Meanwhile, Turkish verbs never show up without a full morphological regalia of three suffixes marking tense, number, and person (among other things). Yet in Turkish-speaking children, the whole system is pretty much mastered and error-free by the time they are 2 years of age (Aksu-Koc & Slobin, 1985). And, while the first word of a baby growing up in an English-speaking household almost always consists of a single morpheme, that's not true for Turkish children, whose earliest words already snap inflectional morphemes onto the stem.

The fact that words often contain more than one morpheme suggests that there may be more to word learning than just connecting up sound sequences with their associated meanings. Perhaps children also have to learn how to *build* words and not just remember them. But we should be careful not to make an unwarranted assumption: Just because it's possible for us to identify the different morphemes that make up a word doesn't automatically mean that children need to learn how to build them on the fly. It's entirely possible that kids do memorize even the complex words as whole units, and that words like *childish*, *quickly*, and *cats* are simply stored in memory. In the next section, we turn to a heated, decades-long debate about this topic.

5.6 Words versus Rules

Learning to generalize

In a famous experiment, Jean Berko Gleason (Berko 1958) showed young children a picture of a bluebird-like creature, saying, "This is a wug." Then she showed them a picture of two of the creatures, saying, "Now there is another one. There are two of them. There are two ___" (see Figure 5.6). Often, the children obligingly produced the missing word *wugs* (pronounced with a [z] sound at the end, not [s]).

This simple result was hailed as an important early finding in research on language acquisition. It showed that when faced with a new word that they couldn't possibly have heard before, kids were able to tag on the right plural morpheme. Since they obviously couldn't have memorized the right form, they must have been generalizing on the basis of forms that they *had* heard: *dogs*, *birds*, etc. It turns out that children were able to do similar things with novel possessive forms (*the bik's hat*) and past-tense forms (*he ricked yesterday*).

This ability to generalize to new forms on the basis of previously known forms is an incredibly powerful mechanism. To some language scientists, it is *the* essential hallmark of human natural language. One way to view the whole




Figure 5.6 The "wug test" has its own line of merchandise. Here, the original test is affixed to a "wug mug," allowing language enthusiasts to spontaneously test for the generalization of the English plural at home and at the office. (Courtesy of Jean Berko Gleason. Adapted from *The Wug Test* © Jean Berko Gleason 2006.)



WEB ACTIVITY 5.7

Complex words across languages In this activity, you'll get a taste of how children growing up with a language like Czech or Turkish would have to learn to excavate the stem morpheme from what is frequently a pile of affixes attached to it.

process is that children must be forming rules that involve variables—that is, they're manipulating symbols whose value is indeterminate and can be filled by anything that meets the required properties for that variable. So, a rule for creating the plural form for a noun might be stated as: *Nstem + s* where *Nstem* is the variable for any stem morpheme that is a noun. No need for children to memorize the plural forms—and for that matter, an adult learner of English could be spared a fair bit of trouble by being told up front, “Look, here’s how you make plural nouns in English.” The other inflectional morphemes in English can also be applied by rule. For example, to form the past tense of a verb: *Vstem + ed*.

Of course, that’s not the whole story, because in addition to *dogs*, *birds*, and *bottles*, we also have *children*, *men*, *mice*, and *deer*, none of which are formed by applying the plural rule, as well as irregular verb forms like *rang*, *brought*, and *drove*. These just have to be memorized as exceptions. Moreover, they have to be memorized as whole units, without being mentally broken down into stems plus affixes. For instance, what part of *children*, *mice*, or *deer* corresponds to the plural morpheme, and what part of *rang* or *brought* would you say conveys the past? In these examples, the meanings of the two morphemes are fused into one seamless whole.

It’s especially interesting to see how children cope with the rule-like forms, compared with the exceptional ones. Very early on, they tend to correctly produce the common exceptional forms—they’ll produce *went* rather than *goed*, for example. But this stage is short-lived. They soon catch on to the pattern of attaching *-ed* to verb stems and apply it quite liberally, peppering their speech with lots of errors like *bringed*, *feeded*, *wearred*, and so on. Once entrenched, this attachment to the rule-like forms can be hard to dislodge, even in the face of parental correction. Here’s a well-known interaction reported by psycholinguist Courtney Cazden (1972) involving a child oblivious to a mother’s repeated hints about the correct past-tense form of the verb *to hold*:

Child: My teacher holded the baby rabbits and we patted them.

Mother: Did you say your teacher held the baby rabbits?

Child: Yes.

Mother: What did you say she did?

Child: She holded the baby rabbits and we patted them.

Mother: Did you say she held them tightly?

Child: No, she holded them loosely.

Paradoxically, errors due to overgeneralizations like these actually reflect an *advance* in learning—they’re evidence that the child has abstracted the generalization for how to form the past tense.

One system or two?

To some researchers, there’s a stark separation between the knowledge that underlies the regular forms of the past tense and the knowledge that underlies the irregular forms: the regulars involve the *assembly* of language from its parts, while the irregular forms, which have to be memorized, hinge on the *retrieval* of stored forms from memory. These are seen as two fundamentally distinct mental processes—words versus rules. Researchers who argue for this view often point to evidence hinting at two distinct psychological mechanisms, including some evidence suggesting that different brain regions may be involved in processing regular versus irregular verbs (see **Box 5.4**). But others argue that



BOX 5.4

Separate brain networks for words and rules?

If you wanted evidence that regular versus irregular past-tense forms call on dramatically different cognitive mechanisms, maybe a good place to look would be inside the brain. There’s no guarantee, of course, that two separate cognitive mechanisms would show up as involving visibly distinct brain regions; but it would be very suggestive if you *did* find different patterns of activity for regulars versus irregulars. And it would be especially suggestive if these patterns made sense given what we know about the overall organization of the brain.

This is precisely the research strategy being pursued by Michael Ullman and his colleagues. Ullman has suggested that regular and irregular complex words reflect the same distinction as the one between declarative and procedural memory—the *what* and *how* systems organized along ventral and dorsal streams, as discussed in Chapter 3 (see Figure 3.12 and Box 3.5). In language, the idea is that the content of memorized information (such as the sounds and meanings of words) is handled by the ventral stream, while the dorsal stream is responsible for “how to” knowledge (such as how to assemble words into sentences, or how to implement the pronunciation of sounds).

We now know that the temporal lobe plays an important role in accessing memorized information about words, while more frontal regions of the brain (along with the basal ganglia) are implicated in more grammatical tasks, which likely involve assembly. Is there evidence for a similar neural dissociation for regular and irregular past-tense forms? Ullman and colleagues (2005) give a detailed summary, including the following:

- Patients with Alzheimer’s disease usually show extensive degeneration of the temporal region of the brain, though the frontal area and the basal ganglia are relatively spared. This leads to a pattern in which patients can usually produce grammatical speech but have a lot of trouble retrieving and recognizing words. These patients turn out to have trouble with irregular verbs, but not regular ones. On the other hand, patients with Parkinson’s disease have damage to their basal ganglia, and have trouble with syntactic processing and regular past-tense forms more than irregulars.
- ERP studies tend to show different patterns of brain wave activity for regular and irregular verbs: when a syntactically sensitive ERP component known as the LAN (left anterior negativity) shows up, it’s elicited by violations on regular verbs, whereas violations of irregu-

lars are associated with a different ERP component, the N400.

- Though not decisive, neuroimaging work hints at greater activation of temporal regions for regular verbs, and of frontal regions for irregular verbs.
- Patients with fluent (Wernicke’s) aphasia, with damage in the temporal region, had more difficulty with irregular verbs than regular ones; the pattern was reversed for patients with non-fluent (Broca’s) aphasia, with damage to the frontal area (**Figure 5.7**).

Still, these data wouldn’t amount to a slam dunk in favor of the separate-networks hypothesis (even if they were to show perfectly consistent results across studies). It’s important to remember that differences in brain activity don’t just show up when two difference processes are involved—different patterns of associations for different words can also do the trick. Remember, for example, the evidence from Chapter 3 that words elicited different patterns of brain activity depending on whether they were strongly associated with actions performed by the hands, mouth, or feet (see Figure 3.11). In a similar vein, opponents of the separate-networks hypothesis (e.g., Joanisse & Seidenberg, 2005) have argued that different patterns of brain activity reflect the fact that regular versus irregular verbs might have stronger links to phonological versus semantic information. Researchers from both camps continue to identify and test more subtle and informative contrasts.




Figure 5.7 The approximate areas of brain damage for two of the patients in a study by Ullman et al. (2005). The pink region represents damage in the brain of FCL, a patient with non-fluent aphasia; the blue region shows damage to the brain of fluent aphasic JLU. (Adapted from Ullman et al., 2005.)

analogy In regard to forming complex words, a process of comparison in which similarities between the members of pairs or sets of word forms are taken as a basis for the creation of another word form.

the two don't involve truly distinct processes at all. They suggest that children could simply be memorizing even the regular forms, and then extending the past-tense forms to new verbs by **analogy**.

Using analogy is quite different from applying an abstract rule. An abstract rule involves actively combining a certain tag—say, the regular past tense form *-ed*—with anything that matches the variable specified within the rule—in this case, a verb stem. But analogy involves finding similarities across different memorized words. Here's an example that may remind you of questions you've likely seen on aptitude tests:

Hen is to rooster as aunt is to ____.

Getting the right answer to this question (*uncle*) involves figuring out what's different between hens and roosters and then applying that same difference to the concept of aunt. The argument is that analogy can also apply to the *sounds* that make up words in the present and past tense. For example:

Slip is to slipped as wrap is to ____.

In order to fill in the blank via analogy, you don't need to decompose the word *slipped* into a verb stem and a separate past-tense morpheme. All you need to do is notice that there's a certain difference in meaning between *slip* and *slipped*, and that the difference in meaning is accompanied by a certain difference in sounds, and then apply this sound difference to the word *wrap*.

Unlike a combinatorial rule, the same analogical processes could be applied to both regular and irregular forms:

Ring is to rang as sing is to ____.

How plausible is it that analogy plays a key role in learning the past tense for both regular and irregular forms? The notion is bolstered by the fact that over-generalization isn't limited to the regular, supposedly rule-based forms—kids often over-extend irregular patterns as well, as in the examples below (from Pinker, 1999):

It was neat—you should have **sawn** it! (seen)

Doggie **bat** me. (bit)

I know how to do that. I **truck** myself. (tricked)

He could have **brang** his socks and shoes down quick. (brought)

Elsa could have been **shotten** by the hunter, right? (shot)

So I took his coat and I **shuck** it. (shook)

You mean just a little itty bit is **dranken?** (drunk)

These kinds of errors turn up in children's speech because in fact the irregular forms don't show a completely random, arbitrary relationship between the present and past forms. It would be odd to have *stap* be the past tense of a verb like *frond*, for example. Instead, irregulars tend to show up as pockets of semi-regular patterns. So we get, among others:

lie/lay	give/gave	forgive/forgave	bid/bade
hang/hung	fling/flung	sling/slung	sting/stung
blow/blew	grow/grew	throw/threw	know/knew
find/found	bind/bound	grind/ground	wind/wound

If irregular forms show some principled patterning in their sound structure, and if children over-extend these patterns as they do with regular forms, then

the line between memorized words and rule-based assembly becomes blurred. It seems tricky to account for the irregulars by means of a general rule; if these were in fact formed by rule, we'd have no way of explaining why the past tense of *snow* isn't *snew*, for instance. But the regularities inherent in the examples above suggest that there must be some way for the human mind to detect and generalize patterns that don't result from the application of a rule.

In a famous paper published in 1986, Dave Rumelhart and Jay McClelland set out to build a computer model that does exactly that. Their **connectionist model** of the past tense treats verb stems and past-tense forms as bundles of sound that are associatively linked together in memory. The model's job is to learn to predict which sounds make up the past-tense form, given a particular verb stem. So, if the model receives as input the set of sounds that make up the stem *grind*, it should ideally be able to come up with the sounds contained in *ground*, even though it has never before been presented with the past form *ground*. It should be able to do this based on its exposure to many other pairs of verb stems and past-tense forms; assuming there are systematic statistical patterns to be found in the relationships between the stems and past-tense forms, the model is designed to detect them and store them as probabilistic associations.

In the early stages of "learning," the model's predictions are pretty random; since it's had minimal exposure to examples, any set of sounds is as likely as another to be associated with the sounds of *grind*. But after it's been exposed to many examples of stems and past-tense forms, it begins to adjust the connections between the sounds of the verb stem and the sounds for the past-tense form. For example, the model begins to "notice" that most of the consonants that appear in the uninflected form also turn up in the past-tense form, so it is able to predict that a pair like *stap/frond* would be extremely unlikely to show up, while *grind/ground* is quite likely. It also begins to "notice" that there are similarities in sound between words that belong to the same subclass of irregular past-tense forms—for example, it notices that there's something the same in the words that have the vowel sound "ay" (as in *find*) in the verb stem but the sound "ow" (as in *found*) in the past-tense form: the word-final consonants *-nd* tend to be present in these words. The gradual tweaking of connections between the verb stem sounds and the past-tense sounds reflects the gradual learning of irregular past-tense forms through a form of statistical analogy. Naturally, the more frequently a word occurs, the more opportunities the model has to associate the correct past-tense form with the verb stem, which might account for the fact that children's earliest past-tense forms tend to be very high-frequency ones like *went*, *came*, *got*, and so on.

Language researchers agree that something like this is probably needed to account for the non-accidental patterns that irregular forms fall into. But Rumelhart and McClelland made the stronger claim that exactly the same kind of mechanism can be used to account for how the *regular* forms are learned, dispensing entirely with the notion that these are formed by a rule that manipulates abstract symbols. This claim has generated a great deal of controversy, and an outsider looking in might be surprised at the charged and passionate nature of the debate over regular and irregular verb forms. But what's at stake is, in part, two general, conflicting views of how children learn the systematic patterns of their language.

One view emphasizes the fact that language is at heart an intricate system of rules, and that children come into the world prepared to extract rules out of their linguistic input. As we'll see in the next chapter, whatever arguments might be made for treating the past-tense formation as rule-based, they are fairly mild compared with the arguments for using rules to account for the regularities governing how entire *sentences* get assembled. Some researchers take an additional step and claim that kids rely on genetic programming that allows them to extract

connectionist model Here this refers to a computational model of the past tense. Based on previously learned associations between verb stems and past-tense forms, the model predicts the probable shape of past-tense forms for new verb stems.



LANGUAGE AT LARGE 5.3

McLanguage and the perils of branding by prefix

In its attempts to build a strong brand identity, the McDonald's fast-food restaurant chain deliberately created a "McLanguage," decorating a whole slew of product names with the prefix *Mc-*. For instance, you can get a special McDeal on the McNuggets or the Chicken McGrill, order a McSalad, and consider having a McFlurry for dessert. The company's goal was to stamp its brand identity indelibly on these products and to distinguish them from similar products made by competitors. In fact, McDonald's has been so territorial about the prefix that it is famous for initiating aggressive litigation against any other companies that try to use it in their branding. In 1987, McDonald's successfully fought Quality Inns International for trademark infringement over that company's plans to name a new chain of hotels McSleep Inns.

But McDonald's went a tad overboard in its branding zeal, and their McLanguage planted the seeds of its own destruction. By using *Mc-* as if it were a prefix that could be attached to any regular word, the company encouraged people to mentally decompose *McSalad* and *McNuggets* as being made up of two meaningful units—that is, of two morphemes stuck together. This wouldn't have been a problem for the restaurant chain if the *Mc-* prefix had become entrenched in consumers' minds as a morpheme with a narrow, brand-specific meaning, something like "from McDonald's." But as morphemes become more **productive** in a language—that is, as they are applied to a wider and wider set of stem morphemes—they have a tendency to take on lives

of their own and to broaden their meanings beyond their original uses.

McDonald's itself fanned the flames of this process and quite likely accelerated the productivity of the *Mc-* morpheme. For example, a 1981 TV commercial had Ronald McDonald onstage teaching children how to create new words in McLanguage, using lyrics set to a catchy country rock tune: "There's nothing *Mc-* to it. You can *Mc-* do it. Just pick a word and add a *Mc-* to it. It's rockin' *Mc-* language. It's rockin' *Mc-* fun." The commercial then pictured children in the audience bopping in their seats to the music and coining new words like *Mc-you*, *Mc-me*, and *Mc-camera*. The scene looks much like a commercial version of the wug test: McDonald's encouraged its audience to play with language and to tag all kinds of words with the morpheme, laying the groundwork for the morpheme's productivity and the broadening of its meaning well beyond the company's product names.

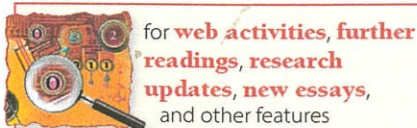
If McDonald's had had a language scientist involved in its branding campaign, the company might have predicted that pretty soon all these McKids would begin using the prefix productively to suit their own expressive goals. Sure enough, along with the official line of "McProducts," we soon had the words *McJobs*, and *McMansions*. The prefix was often pejorative—it's no compliment to call someone's

productivity In linguistics, a process that can be applied very broadly to a large set of lexical items, rather than being restricted to a small set of words.

just the *right* rules out of the input with remarkable efficiency, while managing to avoid coming up with rules that would lead them to dead ends. Under this scenario, nothing could be more natural than for children to seize on a regularity like the English regular plural or past-tense forms and represent the pattern by means of a rule that any new form can then be cranked through.

The opposing view takes the stand that children can get away with very little in the way of innate expectations about rules and what they might look like, and that in fact there is enough information available in the linguistic input for children to be able to notice generalizations and extend them to new situations quite readily. All that's needed is a sensitivity to statistical regularity, just like the one built into Rumelhart and McClelland's connectionist model. Proponents of this view tend to emphasize the fact that, as we saw in the last chapter, even young babies are exquisitely sensitive to statistical regularities in linguistic input. They argue that regular morphemes simply reflect statistical regularities that are somewhat more regular than the less regular forms, and there is no need to formulate an actual rule in either case.

GO TO
sites.sinauer.com/languageinmind



LANGUAGE AT LARGE 5.3 (continued)

new home a "McMansion." And in 1998, author George Ritter wrote about "McUniversities," worrying about the trend toward shallow, cost-effective, convenient, consumer-oriented practices in higher education. The *Mc-* prefix is now so widespread that you can take almost any noun, run it through a search engine with *Mc-* attached, and get numerous hits. You can easily find *McSchool*, *McCulture*, *McParents*, *McChurch*, *McThought*, *McBeauty*, *McJesus*, and even *McSex*, none of which sound like anything to aspire to.

Lots of "McFun" indeed. But not exactly what McDonald's had in mind. What's happening with the *Mc-* prefix is a process that marketers call "genericide"—which is what happens when language originally introduced as part of a brand identity gets absorbed into regular usage. It's not hard to find instances of brand names that have become common nouns. If you stroll across your *linoleum* floor over to your *formica* countertop, pop two *aspirin* into your mouth, check on the stew in the *crock-pot*, pick up the spilled *kitty litter* in the corner with a *kleenex*, pour a bowl of *granola*, and open your freezer to take out a *popsicle* before proposing a game of after-dinner *ping pong*, and you talk about all this without regard for the specific brands that are involved, you are contributing to the genericide of these brand names. Even *heroin* used to be a brand name—though, unlike with *aspirin*, Bayer is probably glad to no longer be readily associated with this product.

Companies often dread genericide because once a word has passed into common usage, they can no longer hold

exclusive legal rights to the name. As a result, they often do everything they can to halt the linguistic change— attempts that often amount to fighting a losing battle. And bits of language like the *Mc-* prefix are even more susceptible to genericide than brand *names*, because of the fact that it's in the nature of prefixes to be promiscuous in attaching themselves to a broad range of words.

In the trademark case over the McSleep Inns, Quality Inns argued that it was not encroaching on the linguistic rights of McDonald's because, as language scientist Roger Shuy testified (Lentine & Shuy, 1990), the *Mc-* prefix had become a productive derivational prefix that in one of its uses simply marked the noun it glommed onto as "basic, convenient, inexpensive, and standardized." This argument was rejected by the judge, but the generic use of the prefix continues. In 2003, McDonald's objected (unsuccessfully) to Merriam-Webster's definition of *McJob* as "a low-paying job that requires little skill and provides little opportunity for advancement." In 2007, the company lobbied Oxford University Press to soften the pejorative definition of *McJob* that appears in the *Oxford English Dictionary*. But dictionaries reflect usage rather than create it, and by the time a word becomes an entry in a dictionary, the cows are out of the barn, and the word has passed into common usage.

In the end, given the scornful ways in which the generic "McMorpheme" is often used, perhaps it's just as well that Quality Inns lost the right to use the prefix. And, as it happens, its chain of Sleep Inns is quite successful.

DIGGING DEEPER

The chicken-and-egg problem of language and thought

We've seen that word learning involves matching up the right concept with the right word, and that very young children are outfitted with helpful biases that make it easier to work out the right matches. But which came first? The concepts or the words? Do children first parcel the world into concepts and then look to see if there is a word assigned for each concept? Or do they rely on language to lead them to concepts that, if children were left to their own wordless devices, might never properly gel?

There have been proposals on both ends of the spectrum. At one extreme is the idea advanced by the philosopher Jerry Fodor (1998) that *all* concepts that line up with words

are innately wired. The claim is that as a species we're endowed with thousands of concepts that form the building blocks of language. One of the more obvious flaws of this notion, as pointed out by Steve Pinker (2007), is that concepts often link up to a word in one language but not in others. This leads to bizarre conclusions. For example, you'd have to say that English speakers, who can brandish the separate words *see* and *show*, have both of these concepts encoded in their DNA but that Hebrew speakers, who can only describe an act of showing by using the more roundabout paraphrase *cause to see*, have only the concept of seeing in their genes.

At the other end of the extreme, some have claimed that language is the mold within which



conceptual thought is shaped. Words, as the idea goes, stamp categories out of otherwise amorphous collections of things, properties, and events. This view was articulated most famously by Benjamin Whorf (1956), who stated:

We dissect nature along lines laid down by our native languages. . . . We cut nature up, organize it into concepts, and ascribe significances as we do, largely because we are parties to an agreement . . . that holds throughout our speech community and is codified in the patterns of our language.

Under such a view at its strongest, if your language has no word for the concept of cousin that is distinct from *sister*, then you necessarily meld the two relationships into one concept. And until English speakers encountered and adopted the French word *aperitif*, they no doubt had trouble getting their heads around the concept of an alcoholic drink you might sip before dinner.

This is a rather radical claim, but it's one that's implicitly held by many very smart people. It's present when the novelist Milan Kundera (1980) questions how anyone could possibly understand the human soul without benefit of the word *litost* (see Language at Large 5.1), and in George Orwell's famous assertion in his essay *Politics and the English Language* (1946) that obfuscatory language has the effect of bamboozling the citizenry. But how does this claim stack up as an actual theory of learning?

If you consult your own intuitions about how you learn new words, it probably seems hard to commit to either of these extreme views. Sometimes concepts seem to come first, and sometimes the words do. On the one hand, it seems doubtful that prior to being introduced to the word *mitochondria*, you had already grasped the notion of microscopic structures within your body's cells that produce the energy that fuels your thoughts and action, and that you had simply been waiting all your life to know what those invisible membrane clusters were called. On the other hand, that thingamajig that rock climbers use so that one member of a climbing pair can arrest her partner's fall without having the rope burn the skin of her palms? If you've ever climbed, you are probably extremely clear in your mind about the particulars of the thingamajig's function and use, even though you may not know that it's properly called a "belay device."

It's nice to admit a truce and call for a compromise in the word-first versus concept-first standoff—and indeed, you'd be *very* hard-pressed to find a language researcher nowadays who subscribes to either of the extreme views I've just sketched out. But an acknowledgment that both sides are probably right is not that informative until you start hashing out the details, and this is what researchers are now busy doing. In Chapter 12, we'll take up the question of whether the language you speak has long-lasting effects on how you carve the world into concepts. For now, we'll focus on the

relationship between concepts and language in children's very early learning.

If infants do possess some concepts without the benefit of language, what are they, and how do kids come by them? Are some of them innate, perhaps forming the building blocks of more complex concepts? Do some, more than others, leap out of the environment that the child experiences? If so, how are the other, less salient concepts learned? And how does conceptual development mesh with the fact that languages often choose very different means to express similar concepts?

We don't have the space or the time to explore all of these questions here, but let's at least scratch the surface. As a starting point, let's consider what infants can *think* about before they speak. What kinds of concepts and categories are they capable of forming in the absence of language? The standard method for inferring whether babies have abstracted a category is one you've already encountered in Section 5.2. Researchers show the babies multiple objects that fall into the same category, and then show them an object outside of that category. If the babies react by spending more, or less, time looking at that object than they do at an object that falls within that category, the researchers infer that the babies have noticed some essential difference between the objects. We've already seen in Section 5.2 that the mere presence of speech encourages babies to find commonalities across these objects. But can babies form categories in a way that's completely independent of language?

It seems clear that babies can form some categories before they show any evidence of understanding word meanings, and well before they actually begin producing words themselves. For example, Paul Quinn and colleagues (1993) showed that at 3 or 4 months, babies can cognitively group different examples of cats and dogs together, and differentiate each basic-level category from the other; that is, they react when they see a dog after they've seen many cats, and vice versa. Do basic-level categories, such as dog, precede more abstract ones, like mammal? Actually, no. Some studies have found that very broad categories tend to be acquired even earlier; for instance, infants as young as 2 months can form a category that includes various mammals but excludes furniture (Quinn & Johnson, 2000). This is interesting because, as we saw in Section 5.2, basic-level words are more commonly used by parents than the broader, superordinate-level words, and they also tend to be learned earlier by children, who pronounce words like *dog* and *bed* before they talk about *mammals* and *furniture*. Conceptual development, then, seems to proceed independently of language in this regard.

There's also evidence that preverbal infants can form categories that rely on quite conceptual features, rather than just perceptual differences. For example, by about 7 months, they can group together things that are animals and things that are vehicles, and soon after, they can also differentiate

between birds and airplanes, despite the fact that these categories look very similar to each other. This suggests that they're doing something a bit more abstract than just paying attention to perceptual features, such as whether something has wings or a particular shape (Mandler, 2004).

Spatial concepts also make an early appearance, and some researchers have suggested that a certain amount of spatial conceptualization is innate. It's been shown that 3-month-old babies can grasp the difference between objects that appear above, below, to the left, and to the right of one another—for example, if a teddy bear is presented as being above a box during the familiarization phase, infants react when they later see the bear appearing below it. By 6 months, they can abstract this relation even when different pairs of objects are used during the familiarization and test phases, and they can also recognize relations like inside versus behind, even when these are represented from different visual angles. The concepts of between and support seem to be harder, and appear a bit later. (For a useful review of this literature, see Casasola, 2008.)

Again, to some extent at least, thinking about spatial relations seems to be independent from talking about them. For example, Korean speakers have to use different linguistic forms to distinguish objects that are loosely contained inside others (like a pencil inside a cup) from objects that are tightly contained (a peg inserted into a hole). In English, it's acceptable to refer to both of the objects as being *in* something. Nevertheless, English-learning kids do treat these cases as different categories of spatial relations (McDonough et al., 2003).

There seems to be plenty of evidence, then, that children do think before they speak, suggesting that their categories and concepts are not dependent on the mold of language. But a couple of considerations complicate the picture somewhat.

One of these is the fact that, even though kids begin *talking* at about 1 year of age, it's very hard to know for sure at what age they have mapped meanings onto word units and, hence, whether their concepts have been influenced by language. In Chapter 4, we saw that by 5 months of age, children are able to recognize the sounds of familiar words, and to segment other words from the speech stream that contains them. The received wisdom until quite recently has been that meaning-to-sound mapping doesn't really happen until closer to 10 months of age or so. But newer evidence shows that kids do understand the meanings of at least some words at the age of 6 months (Tincoff & Jusczyk, 2012). This makes it hard to know: Is a 7-month-old, for example, truly ignorant of any meaningful words? Can we be sure about a 5-month-old?

The other complication is that, even though the standard experimental methods show that babies can group objects together on the basis of certain shared features, it's not clear that their conceptual knowledge is all that rich or stable. Babies often show a surprisingly immature grasp of even

very basic conceptual knowledge until about the age when they actually begin to talk.

One of the most striking examples of this comes from an intriguing study cooked up by Fei Xu and Susan Carey (1996). The researchers set up a display in which a toy duck emerged from one end of a screen, then hid back behind it again. From the other edge, there emerged a toy *truck*, which then also hid behind the screen. Now, if I were to lift the screen and show you what was behind it, how many toys would you expect to see? You'd say two. But a 10-month-old baby might be at a loss. In Xu and Carey's study, it turned out that the infants were no more surprised to see one object than they were to see two, as measured by the amount of time they spent looking at the object(s) once revealed. At this age, it would appear that babies don't really understand that a duck preserves its duck-hood and are pretty relaxed at the prospect that it might transform itself into a truck. But testing babies at 12 months—the age at which they usually start to talk—revealed that they were unnerved by the sight of a single object rather than two.

In fact, the scientists showed that the presence of language can have an immediate, striking effect on kids' responses, even with younger babies. When 9-month-old babies saw exactly the same display accompanied by a voice that gushed, "Look, a truck! Look, a duck!" they registered surprise at seeing a single object, as if the fact of two different names led them to the conclusion that these were two different objects (Xu, 2002).

This is an especially arresting demonstration. It's hard to know, though, whether the linguistic input was actually *required* in order for infants to be able to form a stable representation of these objects in their minds. Perhaps the act of naming simply accelerated learning that would have happened anyway, by focusing the children's attention on the difference between the two objects. This interpretation gets some support from a similar study done with rhesus monkeys: adult monkeys expected that objects like apples wouldn't suddenly morph into coconuts—presumably without the benefit of knowing the words for the objects (Phillips & Santos, 2007).

One way to think about Xu and Carey's findings is in terms of the following analogy: suppose I show you two color chips that are quite close in color, and later I ask you to remember whether I'd shown you examples of one color or two. You might have an easier time accurately remembering if I'd pointed out, "Here's a blue one. And here's a periwinkle one." But does this mean that you were incapable of *perceiving* the difference between the two colors until I named them? Unlikely. Instead, the act of my naming them was a signal that I consider the difference between them relevant, which may have caused you to pay more attention to the precise difference between them. Another example: if I refer to two sweaters as chartreuse and teal but my husband calls them both green, do you infer from this that he's blind to

their difference in hue? A more reasonable conclusion would be that he doesn't *care* that they're different.

The true power of language to shape thought, then, may lie not so much in the power of words to impose concepts, but more indirectly in the social nature of the communicative act. Remember from Section 5.3 that infants are deeply sensitive to the fact that naming an object involves a deliberate act of reference. If a particular linguistic community has seen fit to recruit a word for a concept, this is some fairly good evidence that it sees fit to talk about that concept on a regular basis. That in itself is a powerful thing.

So, even if the effect of language turns out to be one of directing attention to the speaker's intended message, rather than cutting new categories out of whole cloth, it's probably worth standing up and taking note of its effects. At the very least, language might play an important pedagogical role. For instance, in the spatial domain, we saw that English-learning children could form distinct categories for tight-fitting versus loose-fitting objects, despite the fact that English allows the same words to express both. But verbal commentary can provide an added boost: Marianella Casasola and her colleagues (2009) showed that children had an easier time distinguishing a category of tight-fitting objects from loose-fitting ones if a speaker pronounced the word *tight*, or even a nonsense word, while they were looking at the tightly fitting objects. Another provocative study looked at the effects of exposing preschool children to many examples of sentences such as *Mary thought that Fred went to the movies* rather than other complex sentences such as *The boy that had red hair stole the cake*; kids who were trained on the first set of sentences could be nudged into an earlier understanding that the mental states of others can differ from their own (Hale & Tager-Flusberg, 2003).

And what to make of the diversity across languages? I'll set aside for the moment the question of whether a lifetime of hearing the particular words or structures of your native language as opposed to others does leave a permanent

imprint on your cognition; we'll take that issue up in Chapter 12. But current science gives us hints that, at least early in our mental lives, language helps to give shape to our understanding of the world. Most of us have no memories of the momentous early days of our word learning. But I'll leave you for now with the words of a woman who did: Helen Keller (1909), who lost her sight and hearing as an infant, recalled making the connection between things and words at the age of seven, when her teacher, Annie Sullivan, finger-spelled them into her palm. Her writings about this event line up very plausibly with what we know from the studies we've just discussed, and yield some poetic insights into the subjective experience of how daily objects can become transformed once they've been anointed by the human act of reference:


My teacher placed my hand under the spout. As the water gushed over one hand, she spelled into the other hand w-a-t-e-r, first slowly, then rapidly. I stood still, my whole attention fixed on the motion of her fingers. Suddenly I felt a misty consciousness as of something forgotten—a thrill of returning thought; and somehow the mystery of language was revealed to me. I knew then that w-a-t-e-r meant the wonderful cool something that was flowing over my hand. That living word awakened my soul, gave it light, hope, joy, and set it free!... I left the well-house eager to learn. Everything had a name, and each name gave birth to a new thought. As we returned to the house, every object that I touched seemed to quiver with life. This was because I saw everything with that strange new sight that had come to me. On entering the door I remembered the doll I had broken. I went to the hearth and picked up the pieces. I tried vainly to put them together. Then my eyes filled with tears; I realized what I had done, and for the first time I felt repentance and sorrow.



PROJECT

Demonstrating that language is necessary for the development of a concept as opposed to simply facilitating it is no easy matter. Take on the challenge of proposing an experiment or series of experiments that will help get at the answer to whether babies need exposure to language in order to acquire certain concepts. Keep in mind the practical and ethical problems of *depriving* infants of exposure to specific words of English; you might instead use the strategy of recruiting concepts not encoded by words in the English language.

6 Learning the Structure of Sentences



Achieving a vocabulary of 60,000 words or more is an impressive learning feat. But it's not nearly as impressive as the fact that you readily *combine* these words deftly and creatively. To get a quick feel for the scale of combinatorial possibilities that language offers, consider chemistry: with a measly 118 elements in the periodic table, there are *trillions* of known molecules that combine these elements. Just think what you can do with 60,000 units!

The combinatorial power of language allows you to convey entirely novel ideas that have never been expressed in language before. For instance, I'm guessing that you've never heard the following sentence:

It was all because of the lucrative but internationally reviled pink hoodie industry that the president came to abandon his campaign promise to ensure that every household parrot had recourse to free legal counsel.

This sentence may be a touch on the enigmatic side, but chances are you had no trouble understanding it (though perhaps not all of its implications). On the other hand, you'd have no hope of understanding that sentence if its words were served up to you in this order:

Industry ensure because that internationally reviled had legal household parrot was it abandon all pink president every of campaign promise the but lucrative hoodie the came to his to that counsel recourse to free.

Clearly, being able to string multiple linguistic units together is not enough. In order for us to be able to understand sentences made by combining words, they obviously can't just be tossed together in a bag. There has to be some underlying order or structure. That is, language has a **syntax**, a set of rules or constraints for how the units can be put together. The syntactic structure