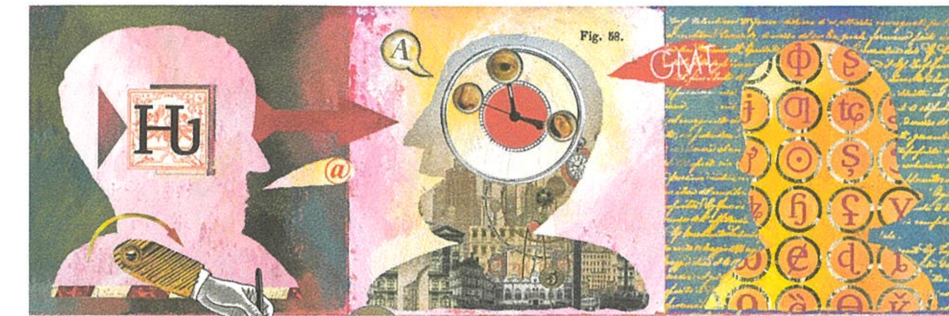

Language in Mind

An Introduction to Psycholinguistics

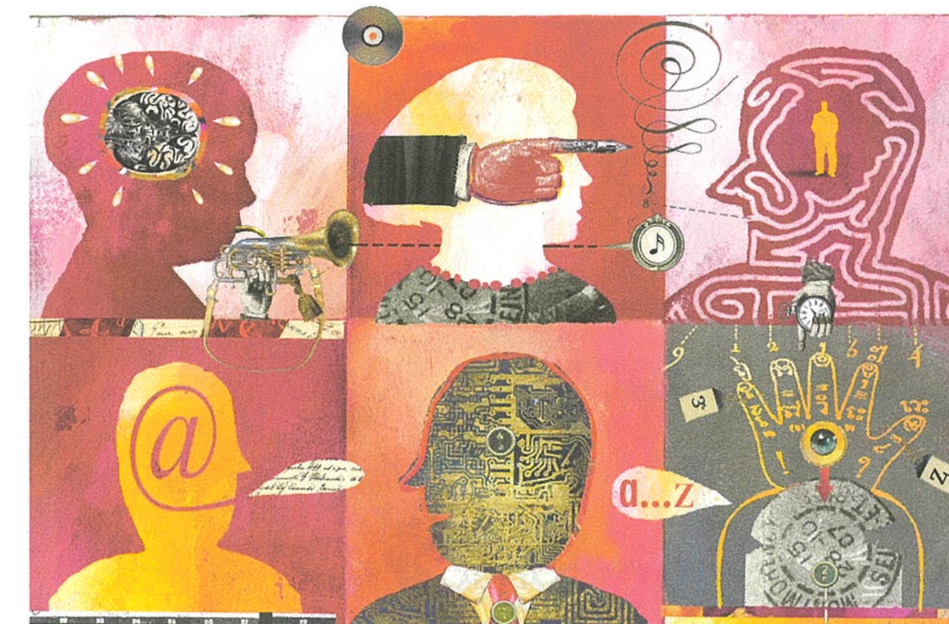

JULIE SEDIVY *University of Calgary*

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

5 Learning Words

At roughly 12 months of age, children begin to utter their first words. After that much-celebrated milestone event, they manage to add a smattering of new words over the next several months until they have a stash of about 50 or so. And then their word learning starts to accelerate until they reach, by conservative estimates, about 60,000 words in their total vocabulary by the time they graduate from high school. This translates into about 10 new words per day.

To put this achievement into perspective, think back to the days when you were learning "math facts"—memorizing the relationships between numbers and arithmetic operations, such as 2 + 5, or 4 × 7. For many kids, it can take weeks of sweating over practice sheets, drills, and flash cards to solidly commit to memory the multiplication tables—a modest set of several dozen facts. And yet, these same kids seem to be able to pluck words out of the air without any effort, sometimes after hearing a word only once. (If you're in doubt about one-shot word learning, see what happens when you inadvertently utter a swear word within earshot of a toddler.) How do children manage to do this? Presumably, the process of learning words is quite unlike learning math facts.

When you learn words as an adult in a foreign-language classroom, of course, what you do is memorize vocabulary lists that are either translated into words of your first language, or linked to pictures and videos describing scenes and events. You might think that parents helpfully simulate a version of this teaching environment by pointing out objects and events and describing them for their children. But when they talk to their kids, parents typically act more

duality of patterning The concept that language works at two general levels, with units of sound combining into meaningful units (usually words) and these meaningful units combine into a larger pattern of meaningful syntactic units.

like conversational partners than like language instructors. You rarely see parents repetitively putting objects into their children's hands and barking out single-word statements like: "Cup." "Peas." "Spoon." Instead, they're more likely to offer a running commentary such as, "Honey, don't throw your cup on the floor" or, "Boy, do I have some delicious peas for you!" This would be a bit like having your foreign language teacher talk at you in sentences, leaving you to sort out how many words she's uttered and what they all mean in the context. Add to that the fact that parents often talk about things that aren't in the here and now: "Did you have a nice nap?" or, "Grandma and Grandpa are coming tonight." Or, they might comment on a child's state of mind in a way that is only very loosely related to reality: "I know you want these peas!"

Even if parents were to describe only objects and events in the immediate environment, there would still be a lot of explaining to do about exactly how kids make the connections between the stream of sounds they hear and their interpretation of the scenes that these sounds accompany. Researchers of language development like to describe the child's task of mapping utterances onto meanings by invoking a famous example from the philosopher Willard Quine. In his example, an anthropologist is trying to characterize the language of a small indigenous tribe previously unknown to Western scientists by noting the relationship between what speakers say and the contexts in which they say it (which also seems like a reasonable strategy for word learning by children). When a rabbit scurries by, a tribesman points to it and exclaims, "Gavagai!" The anthropologist might then assume that *gavagai* corresponds to the English word *rabbit*. But how can he be sure? It might also mean any one of the following: "There goes a rabbit," "I see a rabbit," "That thing's a pesky varmint," "white," "furry," "hopping," "lunch," "rabbit parts," "a mammal," "a living thing," "a thing that's good to eat," or "a thing that's either good to eat or useful for its skin." Or, for that matter, the tribesman may not be commenting on anything to do with the rabbit; he might be pointing out the lovely color of the newly sprung leaves on the trees, or the fact that the sun is about to set. The logical possibilities for choosing a meaning to go with the utterance are fairly vast.

Of course, we have the sense that not all of these logical possibilities are equally likely—for example, if the tribesman is pointing to the rabbit, it would be rather strange for him to be describing the approaching sunset, and it's a good bet that he'd be more likely to be referring to the rabbit itself than its various parts, or some very general category of living things. But do *babies* know these things? We can't assume that just because we as adults can rely on certain inferences or biases to figure out new meanings that these would also be available to your average toddler. Ultimately, the child needs to have some way of constraining the enormous space of possibilities when it comes to linking sounds with meanings. In this chapter, we'll explore what some of these constraints might look like, and how the child might come to have them in place.

But learning words goes beyond just figuring out which meanings to attach to bundles of sounds. In Chapter 2, I introduced the idea that human natural languages show **duality of patterning** (see Box 2.1). That is, language operates at two very general levels of combination. At one level, meaningless units of sound combine to make meaningful units (for example, words), and at the second level, meaningful units (words) combine with each other to make larger meaningful syntactic units. Words, therefore, are the pivot points between the system of sound and the system of syntactic structure. So, in addition to learning the meanings of words, children also have to learn how words interface with the sound system on the one hand and syntactic structure on the other. This chapter will also deal with how individual words become connected with these two distinct levels. As was the case in Chapter 4, what might intuitively *feel* like a straightforward learning task is anything but.

WEB ACTIVITY 5.1

Inferring language meaning In this exercise, you'll get a direct feel for the task of inferring language meaning without the benefit of any explicit teaching. You'll see some short videos, accompanied by unintelligible speech. Your job will be to generate plausible (and implausible) guesses at the meanings conveyed by the speech.

5.1 Words and Their Interface to Sound

Which sounds to attach to meanings?

When it comes to learning the meanings of words, a baby's first job is to figure out *which* blobs of speech, extracted from an ongoing stream of sounds, are linked to stable meanings. For example, it would be helpful to know if *gavagai* really consists of one word or three—and if it's more than one, where the word breaks are. No problem, you might think, having read Chapter 4. In that chapter, we saw evidence that infants can segment units out of running speech well before their first birthday, relying on a variety of cues, including statistical regularities. So, by the time babies learn the meanings of their first words, they've already been busy prepackaging sounds into many individual bundles. The idea would be that these bundles of sound are sitting in a young child's mental store, just waiting for meanings to be attached to them.

For example, let's suppose that in our hypothetical language, *gavagai* breaks down into two words: *gav agai*. Assuming they've heard each of these words in speech often enough, babies who've been exposed to this language from birth should be able to recognize *gav* and *agai* as cohesive linguistic packages, but they'd treat the sound sequences *avag* or *vagai* as just a collection of random sounds. As seen from the experiments with artificial languages in which "words" have been completely disembodied from meanings, babies should be able to segment these units based entirely on their knowledge of sound patterns and their regularities in running speech. In theory, then, when trying to figure out the *meaning* of *gavagai*, they should know that they're looking to attach meanings to the sound bundles *gav* and *agai*. If so, this would make some of the challenges of learning the meanings of words a bit more manageable.

But, once again, never take the workings of the infant mind for granted. We've been referring to the segmented units of speech as *word* units. This is convenient because as adults, we recognize that these sequences of sounds correspond to meaningful words. But babies don't necessarily know that. Just because they can slice these units out of a continuous speech stream doesn't mean that they understand that these particular groupings of sounds are used as units for *meaning*. They might simply be treating them as recurring patterns of sound that clump together. This might seem odd to you, since you're so used to thinking of speech sounds as carrying meaning. But consider this: studies of statistical learning have shown that both babies and adults can learn recurring groupings of musical tones in just the same ways that they learn word units in an artificial language. It might be easier for you to imagine how someone might recognize a group of tones as a recurring and cohesive bundle of sounds, without having the expectation that there's any meaning attached to them. It's entirely possible that for babies, this is exactly how they treat the "word" units they pull out from speech. They've figured out that certain sounds systematically go together, but they haven't yet figured out what these bundles of sounds are *for*.

A number of researchers have suggested that something very much like this is going on in the minds of babies at the earliest stages of word learning. The idea is that they've gotten as far as memorizing a bunch of meaningless sound bundles, but the next phase of linking sounds with meaning is an entirely new step. When they first start linking up sounds to meanings, they don't simply dip into their bag of stored sound sequences; rather, they start from scratch in building up brand new sound-based representations. The evidence for this argument goes as follows:

Remember that when babies segment units from speech, these units contain an exquisite amount of phonetic detail. Babies are able to attend to even very subtle sound variations, such as aspiration, and can use them as segmentation cues. On the other hand, a number of researchers have claimed that the units of sound that babies attach meaning to are much coarser in their phonetic detail than the strings that they segment out of speech. They've suggested that the mental representations that result from slicing up the speech stream aren't in fact the *same* representations that underlie a child's early meaningful words. In other words, let's suppose that a baby has managed to successfully segment the sound sequence for the word *dog* based on sound regularities alone—that is, she has the sense that this is a non-random, recurring bundle of sounds that clump together. Her representation of this bundle of sounds would look quite different from her representation of the sounds for the *meaningful* symbol *dog* when she starts to figure out that *dog* refers to the furry family pet. It's even been proposed (see, for example, Hallé and de Boysson-Bardies, 1996) that the representations that serve as containers for meaning aren't made of strings of individual sounds at all, unlike the representations for segmented speech. Rather, they're general and quite fuzzy holistic *impressions* of sounds.

This might seem like an odd proposal to make. After all, why would babies go to the trouble of segmenting speech into units if they weren't going to use these units as the basis for meanings? But there's some familiarity to the idea that the mental representations for words might be to some extent separate from the detailed representations of the sounds that make up those same words. If you remember, back in Chapter 3, I summarized some arguments for the separation of linguistic knowledge into dorsal and ventral streams in the brain. These arguments included evidence that some patients do well on word-recognition tasks even though they struggle with basic sound-discrimination tasks; and conversely, some patients are able to make fine sound discriminations but have a hard time recognizing familiar words. This suggests there's some potential slippage between the two systems; perhaps it takes babies some time to fully establish the connections between them.

However plausible the notion might be, though, in order to understand the sound representations that babies first map meanings onto, we need to take a close look at the mental life of these early word-learners.

Studying how children map sounds to meaning: The switch task

switch task A simple word-mapping test in which infants are exposed to a visual representation of an object paired with an auditory stimulus during a habituation phase. During the subsequent test phase, the infants hear either the same object–word pairing, or they hear a new word paired with the familiar object. A difference in looking times between the novel and familiar pairings is taken as evidence that the child had mapped the original auditory stimulus to the familiar object.

habituation Decreased response to a stimulus after repeated presentations of that stimulus.

Our best bet for studying how babies map sounds onto meaning is through carefully designed lab studies that look at the precise conditions under which these mappings are made. One way to do this is through an association test known as the **switch task**. Stager and Werker (1997) used this technique to test whether children pay attention to fine details of sound in learning new object-word mappings (see **Figure 5.1**).

In this task, there's a *habituation phase* in which two objects are each paired up with a novel word that is not an actual English word but obeys its phonotactic rules (for example, *lif* and *neem*; see Figure 5.1A). To learn these associations, babies watch pictures of objects one at a time, accompanied by their associated labels spoken through a loudspeaker over a number of trials during this phase. The trials are repeated until the babies show signs of **habituation**—that is, of

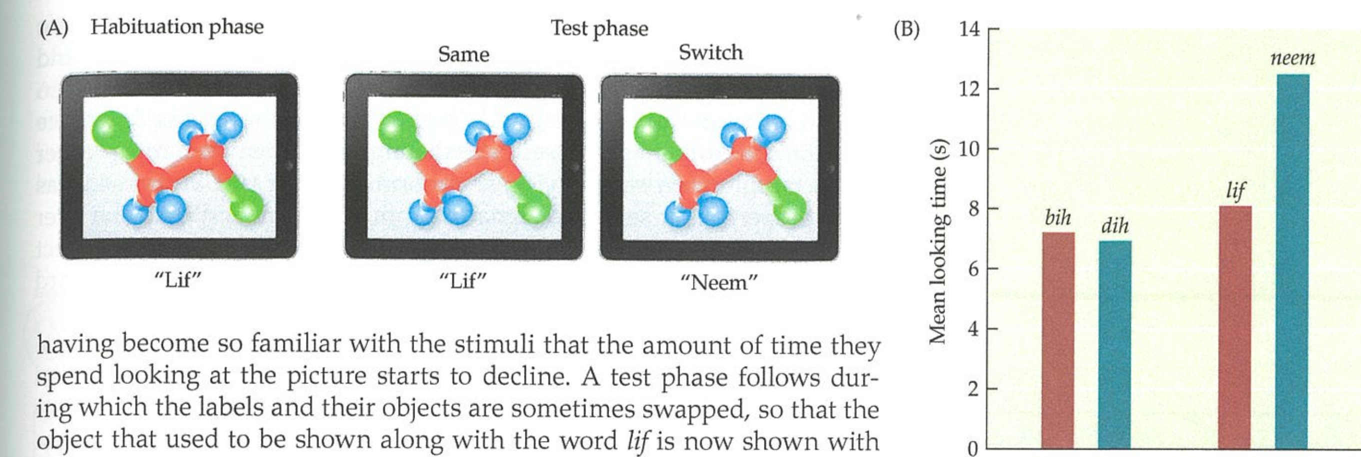

Figure 5.1 (A) A switch task using the highly distinct stimuli *lif* and *neem*. During the habituation phase, children heard *lif* paired with the visual image. In the test phase, children either heard *lif* again or a new word (*neem*) paired with the original visual image. In a different version of the same experiment, the highly similar sound stimuli *bih* and *dih* were used. (B) Mean looking times for 14-month-old participants. Babies readily distinguished between *lif* and *neem*, but not between *bih* and *dih*. (Adapted from Stager & Werker, 1997.)

having become so familiar with the stimuli that the amount of time they spend looking at the picture starts to decline. A test phase follows during which the labels and their objects are sometimes swapped, so that the object that used to be shown along with the word *lif* is now shown with the word *neem*, and vice versa. These "switched" trials are compared with "same" trials, in which the words accompany their original pictures.

If babies have linked the correct label to the picture of each object, they should register surprise when the wrong label is paired with the familiar picture; that is, they should look longer at the "switch" trials than at the "same" trials. This result is exactly what we see for babies who are 14 months of age or older—shortly after the average age at which babies gurgle their first recognizable words (Stager & Werker, 1997). Or rather, the effect shows up for two words that are very different from each other, such as *lif* and *neem*. But babies at this age *don't* notice the switch if the words are very similar to each other—for example, *bih* and *dih*. Instead, they act as if these were just variants of the same word, linked to the same meaning (see Figure 5.1B). This is interesting, because infants can clearly hear the difference between /b/ and /d/ in categorical perception tasks that *don't* require them to link meanings with sounds, but they seem to ignore this difference when it comes to linking up sounds with simple pictures.

These results raise the possibility that babies' representations of meaningful words don't actually contain all the phonetic detail of the strings of sounds that they segment from speech. But why would this be? One idea is that there's a difference between information that babies can pay attention to while processing language in the here and now, and the information they commit to long-term memory in the form of a stable **lexical representation** by which sound and meaning properties are recorded. The thinking might go like this: babies (and maybe adults too) are somewhat miserly about their long-term memory space, and only keep as much detail about sound in the lexical representation of a word as they need for the purpose of distinguishing that word from other words in their long-term memory. But adults and older children have large vocabularies of thousands of words, so they need a lot of sound detail to keep their lexical representations separate from one another.

For example, older individuals really need to encode the difference between /b/ and /d/ so they can distinguish between word pairs like *big/dig, bean/dean, bad/Dad, bought/dot*, and so on. But put yourself in the shoes of an infant who only knows a smattering of meaningful words. Take the following meager collection:

baby	mommy	daddy
hat	shoe	diaper
yummy	milk	juice
spoon	bottle	blanket

lexical representation Information that is committed to long-term memory about the sound and meaning properties of words, and certain constraints on their syntactic combination.

None of these words are very similar to each other, so there's no need to clutter up lexical representations with unnecessary details. However, with time and a burgeoning vocabulary, the infant might notice not only that she needs to distinguish between words like *big* and *dig*, but also that the voicing difference between /b/ and /d/ is mighty useful for distinguishing between many other words. She would then start paying attention to voicing (see Section 4.3) as a general property that serves to signal meaning differences. When she later comes across a new word such as *pill*, her lexical representation would reflect the fact that its first sound is voiceless, even though she may never have heard the contrasting word *bill*.

This is a plausible account. But it's not clear that babies' failure to notice the difference between *bih* and *dih* in the switch task really reflects less detailed lexical representations. Another possible explanation is that babies are just very inefficient and prone to error when it comes to *retrieving* words from memory. So, they might confuse words that are similar to each other during the retrieval process, even though sound distinctions between them are actually there in their lexical representations. For example, it's possible that they do in fact have separate lexical representations for *bih* and *dih*, but they might easily confuse them because of their similarity—just as you might mistakenly pull a bottle of dried basil out your cupboard while searching for oregano. You know they're different herbs, but their similar appearance has momentarily confused you.

This alternative retrieval account gains support from evidence that even adults are less sensitive to detailed sound differences when they're trying to match newly learned words with their meanings. Katherine White and her colleagues (2013) showed that when adult subjects were taught a new "language" in which made-up words were paired with abstract geometric objects, they showed a tendency to confuse novel words like *blook* with similar-sounding words like *klook* if they'd heard the word only once before. But this confusion lessened if they'd heard the novel word multiple times. These findings suggest that, like those of adults, children's lexical representations might well be linked to rich and detailed phonetic representations, but be vulnerable to sound-based confusion when words are unfamiliar. And in fact, 14-month-olds seem perfectly capable of distinguishing familiar words like *dog* from very similar ones like *bog* (Fennell & Werker, 2003). Together, these two experiments argue against the notion that the sound representations that children map onto meanings are different in nature from the sound representations they segment out of running speech.

We can find even more direct evidence that the units babies pull out of the speech stream are put to good use in the process of attaching meanings to clumps of sounds. In a study led by Katharine Graf Estes (2007), 17-month-old babies first heard a 2.5-minute stream of an artificial language. Immediately after that, the babies were exposed to a novel word-learning task using the switch paradigm. The added twist was that half of the new words in the switch task corresponded to word units in the artificial language the babies had just heard. The other half were sequences of sounds that had occurred just as often in the artificial language but that straddled word boundaries ("part-words"), as seen in the experiments in Chapter 4. Babies were able to learn the associations between pictures and words, but only for the sound units that represented word units in the artificial language, showing that they were applying the results of their segmentation strategies to the problem of mapping sounds to meaning.

Results like these suggest that, indeed, very small children are able to draw on their stores of segmented units to help them in the difficult task of matching words and meanings. In other words, if you're a baby and the transitional probabilities among sounds lead you to treat *agai* as a coherent and recurring clump of sounds, you might have a leg up in figuring out what is meant by "Gavagai!" All of this suggests that word learning relies heavily on the statistical experience that babies have with language. As it turns out, the *amount* of exposure to language that children get in their daily lives can vary quite dramatically, leading to some striking consequences for their early word learning (see **Box 5.1**).

BOX 5.1
The 30-million-word gap

Small children can learn new words after just one exposure, which might lead you to think that they can acquire a decent vocabulary even if they hear just a modest amount of language. But there's good reason to believe that the greater the quantity of input, the more solid their learning of language will be.

In Chapter 4, we saw that children can track the statistical patterns in their speech. While even 2 minutes of a very simple sample was enough for them to be able to parse the speech stream, in real life, the language sample that children hear is much messier and more complex. Remember that in the speech segmentation experiments, for example, babies only heard samples of *four* frequently repeated "words," hardly a realistic reflection of real-world speech. In order to be able to notice the statistical patterns in everyday speech, children no doubt need a much heftier amount of input.

We've also seen that children gloss over many details of the sounds of unfamiliar words. This too suggests that they might benefit from many repetitions of the same word, in order to solidify their representations of words and their ability to pull them efficiently from memory.

Are there repercussions for children who hear less language? In a well-known study published in 1995, researchers Todd Risley and Betty Hart found that there were massive differences in the amount of speech that children heard in their homes; the average number of parental words directed at the babies ranged from 200 to 3000 words per hour. The amount of exposure to language was strongly correlated with the size of the children's vocabulary (Hart & Risley, 1995).

A more troubling result was that kids from the poorest families generally heard much less talk than children from more economically privileged homes—and were significantly behind in their vocabulary growth as well (see **Figure 5.2**). The authors of the study estimated that by the age of three, lower-income kids would have heard an average of 30 million fewer words than their wealthier counterparts. This raised a fair bit of concern, since the size of a child's vocabulary upon entering school is a good predictor of how easily he'll learn to read. It's well known that school-age children from lower-income homes lag behind higher-income children in reading skills.

The study by Hart and Risley focused attention on socioeconomic differences, but a broader lesson is the relationship between the quantity of language input and the child's word-learning trajectory. These days, educators and language researchers raise concerns about whether the constant presence of electronic devices and social media in many households might be cutting into a family's conversation time.

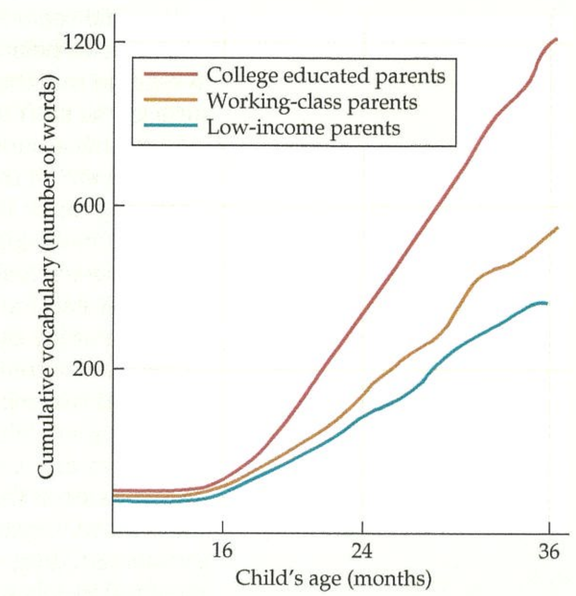

Figure 5.2 Hart and Risley's data show that disparities emerge very early in life among children living with college-educated professional parents, working-class parents, and low-income parents living on welfare. (Adapted from Hart & Risley, 1995.)

5.2 Reference and Concepts

Words and objects

A young child's sense that certain bundles of sound are more likely than others to be linked to meaning is a start. But in order to explain how sound–meaning pairings come about, we also need to look at the other side of the equation, and know a bit about which *meanings* a child considers to be good candidates for linguistic expression. It would be a startling child indeed whose first word turned out to be *stockbroker* or *mitochondria*—even if her parents held down jobs as genetics researcher and financial analyst and uttered such words on a daily basis at home. It's simply implausible that these *concepts* would be part of a toddler's mental repertoire. These are complex notions that are probably impossible to grasp without the benefit of some already fairly sophisticated language to explain them. Most children's first words instead refer to tangible objects or creatures that can be experienced more directly through the senses: *rabbit, bottle, milk, baby, shoe.*

But how is it that a child knows (or comes to learn) that *rabbit* is more likely to be referring to the whole squirming animal than to its fur or, even more implausibly, three of its four limbs? It *feels* obvious to us. Our intuitions tell us that some concepts—such as whole objects—are simply more psychologically privileged than others, almost crying out to be named. But, as always, we can't assume that what feels obvious to us would also feel obvious to a very young baby. Our own intuitions may be the result of a great deal of learning about the world and about how languages use words to describe it. (Think back, for example, to how phonemic categories which seem so deeply ingrained to us in fact had to be *learned* early in childhood.) Have intuitions about "obvious" candidates for word meanings been with us since babyhood, guiding our acquisition of words throughout our lives?

It would seem that they have. Many studies have shown, for example, that when babies hear a new word in the context of a salient object, they're likely to assume that the word refers to the whole thing, and not its parts, color, or surface, the stuff it's made of, or the action it's involved in. Researchers often refer to this assumption as the **whole-object bias** in word learning. This bias doesn't seem all that surprising when you consider the landscape of very early infant cognition. Even as young as 3 months of age, babies clearly organize the jumble of lines, colors, and textures that make up their visual world into a set of distinct *objects*, and they have robust expectations that objects in the world will behave in stable and predictable ways.

In one study by Philip Kellman and Elizabeth Spelke (1983), 3-month-old babies saw a screen with a stick visible at its top and another stick at the bottom. If the two sticks moved simultaneously, the babies assumed that they were joined to a single object, and they were surprised if the screen was removed to reveal two disconnected objects. ("Surprise" was measured by how long the babies stared at the scene once it was revealed.) On the other hand, if the two sticks moved separately, the babies were unfazed to find that two separate objects were hiding behind the screen. Aside from knowing that objects usually act as indivisible wholes, young babies also seem to know that objects can't disappear at one point and reappear at another, that they can't pass through other objects, and that inanimate objects can't move unless they come into contact with other objects.

Given that babies can clearly parse the world into whole objects long before they can parse speech into word-like units, it would seem natural that once children figure out that words are used to *refer*, they then take for granted that

whole-object bias The (theoretical) assumption by babies that a new word heard in the context of a salient object refers to the whole thing and not to its parts, color, surface, substance, or the action the object is involved in.

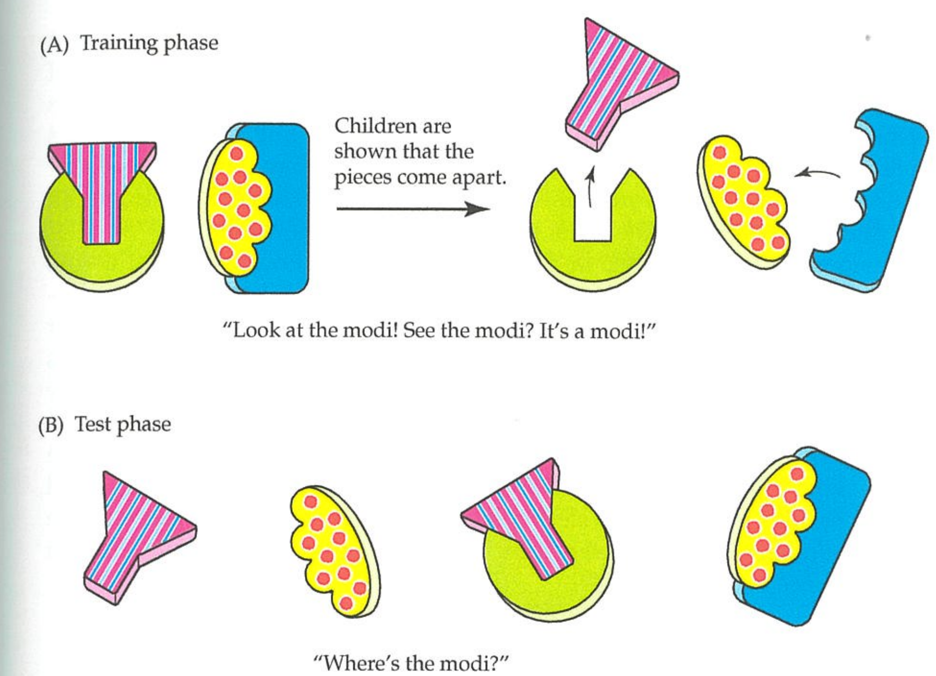

Figure 5.3 (A) Two of the novel objects used in experiments by Hollich et al. Experimenters demonstrated to young children (ages 12 and 19 months) how the objects could be separated into two parts. (B) During the test phase, children saw a display showing both the entire assembled object and the more colorful of its parts. (Adapted from Hollich et al., 2007.)

objects are great things to refer to. At one level, this might be simply because objects are perceptually important, so when babies hear a word, they easily slap that word onto whatever happens to be most prominent in their attention. But in fact the relationship appears to go deeper. It's not just that whole objects have a tendency to draw babies' attention; it appears that babies have a sense of what kinds of things *words* attach to.

A study by George Hollich and colleagues (2007) explored whether the act of *naming* by an experimenter affected how babies visually examined objects and their parts. Babies age 12 and 19 months were shown objects made of two parts—a single-colored "base" piece and a more exciting, colorfully patterned piece that inserted into the base (see **Figure 5.3**). During a training phase, the experimenter labeled the object with a nonsense word ("Look at the modi! See the modi? It's a modi"). The experimenter did nothing to hint at whether the word referred to the whole object or just one of its parts but did emphasize that the object came apart, by repeatedly pulling it apart and putting it back together again. Then, during the testing phase, the experimenter put the whole object and an exact copy of the colorful part on a board in front of the child and asked the child, "Where's the modi? Can you find the modi?" Since babies reliably look longer at an object that's been named than at other objects, it's possible to measure looking times to infer whether the child thinks that *modi* refers to the whole object or to the colorful part; the child's eye gaze should rest longer on whichever object he thinks the word refers to.

The results of the study showed that babies looked longer at the whole object than at the object part, suggesting that they had taken the word *modi* to refer to the whole object rather than to one of its parts. But before we draw this conclusion with any certainty, we want to make sure that kids weren't just looking at the whole object because it was more interesting or visually important, but that their eye gaze really reflected something about making a link between the object and the word *modi*. To check for this, during an earlier phase of the experiment, even before the experimenter labeled the object, babies' free-ranging eye movements to the whole object and the isolated part were also measured. Before hearing the novel name, babies spent roughly the same amount of time looking at the colorful part as the whole object. So it wasn't just that the babies found the whole object more visually compelling; it was specifically the act of naming that drew their attention to it. In other words, the whole-object bias seems to be about something more specific than just finding whole objects more interesting, and even at 12 months of age babies have some ideas about what kinds of meanings words are likely to convey. This kind of knowledge can really help narrow down the set of possible meanings for words.

Obviously, though, people *do* talk about parts of an object as well as the whole thing, and they can also describe an object's color, texture, temperature, function, taste, ingredients, location, origin, the material it's made of, or any actions or events it's involved in. A whole-object bias might help babies break into an early vocabulary, but sticking to it too rigidly would obviously be pretty limiting. For example, researcher John Macnamara (1972) describes a scenario where a child thought the word *hot* referred to the kitchen stove. (You can imagine a parent warning "Don't touch that! It's hot!") Luckily, as we'll see shortly, kids make use of a slew of helpful cues in figuring out that you can talk about other things besides objects, and that words can map onto different kinds of concepts.

Categories large and small

So far, I've talked about object names as if they refer to specific objects, actual things that the child can see or touch. But that's actually the wrong way to think about words like *rabbit* or *bottle* or *blanket*. These nouns don't just apply to any particular object—they apply to the entire *categories* of rabbits, bottles, and blankets. It turns out that only a small subset of nouns—the ones we call proper nouns, like *Dave, Betty, Marilyn Monroe,* or *Cleveland*—refer to particular entities or individuals. And, when you think about it, language would be only very narrowly useful if words didn't generalize beyond specific referents. We'd like to be able to talk to our kids not just about *this* family dog ("Honey, don't tease the dog"), but also about dogs more generally ("Don't get too close to a strange dog"; "Some dogs bite"; "Dogs need to go outside for walks"; and, eventually, "The dog is descended from the wolf"). Presumably, one of the huge benefits of language for humans is its ability to convey more general information along these lines.

Mapping words onto categories is even more complex than mapping words onto specific referents. In principle, any single object could fall into an infinite number of categories. Along with categories like "dogs," "furniture," or "reggae music," people could (and sometimes do) talk about categories such as "foods that give you heartburn"; "professions that pay poorly but give great satisfaction"; "things I like to do but my ex never wanted me to do"; or even "objects that broke at exactly 2 PM on March 22, 2006." But categories like these rarely get their own words—in the scheme of categoryhood, they just don't make the cut. How do children figure out which categories are most likely to have words bestowed upon them?

WEB ACTIVITY 5.2

Mapping word meanings In this activity, you'll explore whether, as an adult, you show evidence of any biases in the way that you map the meanings of novel words onto various possible categories.

Just as whole objects seem to be more natural candidates for reference than their properties or their parts, we also have intuitions about which categories seem the best candidates for reference. Even if we look at just those categories that *have* been awarded their own words, though, it becomes clear that people are inclined to talk about some kinds of categories more than others. For instance, unless you're a dog breeder, you probably use the word *dog* in your daily communication much more often than you do the

more specific word *Dalmatian*. What's more, you probably also use it more often than the more general terms *mammal* and *animal*, even though these broader categories obviously encompass a larger number of creatures. This is generally true of categories at a mid-level degree of specificity—for example, chairs get talked about more often than recliners or furniture, and the same holds for apples versus Cortlands or fruit. These privileged midlevel categories are called **basic-level categories**, in contrast with the more general **superordinate-level categories** and the more specific **subordinate-level categories**.

When you look closely at their informational value, the basic-level categories seem to be especially useful for capturing generalities. For instance, if you know that something is a dog, you know a lot of things that are likely to be true about it: you know that it probably barks, needs exercise, eats meat, has a pack psychology, marks its territory, has sharp teeth, and so on. On the other hand, if you only know that something is an animal, the number of things you know about it is much smaller and vaguer. So, being told "Fido is a dog" allows you to infer many more things about Fido than being told "Fido is an animal." At the same time, basic-level categories (say, dogs versus cats or birds) also tend to be fairly distinct from one another, unlike subordinate-level categories (for example, Dalmatians versus collies). Often it makes sense to talk about what the category members have in common that differentiate them from other categories. So, if you're told "Dogs need to be walked," this is a more useful piece of information than being told "Dalmatians need to be walked." It may be that basic-level categories are favored by language users exactly because they strike this balance between similarity among members and distinctiveness from other categories.

In learning their early words, then, kids have to figure out how to map the words onto just the right categories, including how to hit exactly the right level of specificity. They sometimes show evidence of **under-extension** of category names—for example, if a child first learned the word *flower* in the context of referring to a carnation, she might not realize that daisies can be referred to by the same word. But you wouldn't necessarily notice a child's vocabulary under-extensions without some explicit probing—she might just not call a daisy *anything*, rather than use an obviously wrong word for it. **Over-extension**, in which a child uses a word for an inappropriately general category—for instance, referring to all animals as *doggie*—is more easily observed. Errors of over-extension may well be more common in children's minds, although it's a bit hard to tell based only on how they label things in the world.

Despite these missteps, children's early speech (just like adult language) contains a disproportionate number of basic-level words—a fact that raises the question, is this simply because basic-level words are more common in their parents' speech, or is it because young children are especially attuned to words in this middle layer of categories? It might be helpful to young children if they showed up at the word-learning job with some preconceptions about not just what aspects of the physical environment are most likely to be talked about (whole objects more than parts), but also what categories make for good conversation. To answer this question, we need to look at experiments that specifically probe for the links young children make between new words and their meanings.

One way to do this is with a slightly modified version of Web Activity 5.2, in which you had to guess the possible referents for the made-up word *zav*. If children see a novel word like *zav* used to refer to a Dalmatian, and they assume that it applies to members of the same basic-level category, they'll also accept collies and terriers as possible referents of the word and will gladly apply the word *zav* to these kinds of dogs as well. On the other hand, if their guesses

basic-level categories The favored midlevel category of words that strike a balance between similarity among members of the category and distinctiveness from members of other categories; e.g., of the words *dog, Dalmatian,* and *animal, dog* would fall into the basic-level category.

superordinate-level categories The more general categories of words that encompass a wide range of referents; e.g., *animal.*

subordinate-level categories More specific categories comprising words that encompass a narrow range of referrents; e.g., *Dalmatian.*

under-extension Mapping new words into categories that are too specific; e.g, referring to a carnation, but not a daisy, as *flower.*

over-extension Mapping new words into categories that are too general—for instance, referring to all animals as *doggie.*

are more conservative, they might limit their choices to only Dalmatians and reject other kinds of dogs as falling under the meaning of *zav*. Or, they might guess more liberally, and extend the word to apply to all animals, including cows and cats.

Several studies (e.g., Callanan et al., 1994; Xu & Tenenbaum, 2007) have found that young children are clearly reluctant to extend a new word to a broad category such as animals, as are adults. But, whereas adults seem to have fairly strong assumptions that *zav* extends to a basic-level category, preschoolers seem a bit more conservative, sticking more often with the subordinate-level category as the appropriate meaning of the word. However, these early assumptions about categories seem to be fairly fluid. For example, the study by Fei Xu and Josh Tenenbaum (2007) found that preschoolers became willing to shift their interpretation of a new word to a basic-level category if they'd heard just

LANGUAGE AT LARGE 5.1

How different languages cut up the concept pie

Babies seem to approach the task of word learning with some reasonable starting assumptions. These line up fairly decently with the way language cuts up conceptual space into word-sized bites. Many of these assumptions mirror facts about the natural world, and capture the sorts of properties that define or distinguish categories of things in the world. For example, no language that I know of has a word that groups together roses and rabbits but excludes all other flowers and mammals. We have the feeling that such a language would be deeply alien. Concepts, then, seem to offer themselves up in natural classes—much as the sounds of speech tend to pattern together in natural classes, as we saw in Chapter 4.

The starting assumptions that children have about word meanings may well help them avoid unnatural hypotheses about word meanings. But they only take kids so far. After all, in addition to talking about the natural world, we also use language to talk about some very abstract notions and relationships. And once we move into the space of abstract concepts, it's less obvious which concepts cry out for their own words. As a result, languages can vary quite widely in how they match up words and meanings.

For instance, consider words for family members, such as *mother*, *brother*, *cousin*, *aunt*, *grandfather*, *sister-in-law*, *stepfather*, etc. Some of the dimensions that English enshrines in distinct words include information about gender, generation, and whether the relationship is by blood or marriage. This may seem very natural to you. But you don't have to go very far to find dramatically different ways of referring to relatives. For example, many languages (such as Bengali, Tamil, Korean, and Hungarian) also feel the need to encode the relative age of family members, so there would be a different word for an older brother than for a younger brother. In Hawaiian, you'd use the same word for your mother and your aunt, and there would be no way to distinguish siblings from cousins. On the other hand, Sudanese children have to learn different words to distinguish their father's brother's children from their father's sister's children, from their mother's brother's children, and their mother's sister's children; English speakers are satisfied with the single word *cousin*. Other languages can get even more complicated, having certain distinctions among relatives made only on the mother's side, but not the father's side.

No language has a word to cover every concept, and it's not always obvious why languages make the choices they do. Ever wondered, for example, why English has the word *hum*, but no word for what you do when you sing the melody of a song but substitute "la la la" for the song's actual words? (In the Pirahã language, the situation is reversed.) And how come we make no distinction between addressing one person or many with our pronoun *you* (though a previous version of English did, with the two words *thee* and *ye*)?

In some cases, lexical gaps can be filled when speakers of a language encounter a handy word in another language and simply appropriate it, rather than inventing their own. In this way, we've imported nifty words like *Schadenfreude* (from German), *déjà vu* (from French), *bozo* (from West African), *robot* (from Czech), *aperitif* (from French), *tsunami* (from Japanese), *pogrom* (from Russian via Yiddish), *bonanza* (from Spanish), and *schlep* (from Yiddish) among many others.

three examples of the word being applied to members of the same basic-level category. So, hearing that both a collie and a terrier could also be called a *zav* would lead children to assume that *zav* could also be used to name a poodle or a Labrador retriever, even though they'd never heard the word applied to those particular kinds of dogs. And, if they'd heard *zav* being used three times *only* in reference to a Dalmatian, they became even more convinced than they'd been at the outset that this word applied only to members within this subordinate-level category.

This study suggests that youngsters don't necessarily match up words with the right categories the first time they hear them, but that they're not set in their ways when it comes to their first hypotheses about word meanings. Instead, they gather evidence that either confirms or disconfirms their preliminary guesses and fine-tune these guesses accordingly. Of course, as outside observers we might never become aware of their internally shifting word meanings unless we specifically choose to test for them.

LANGUAGE AT LARGE 5.1 *(continued)*

Still, many potentially useful words remain unimported from other languages, so numerous wordless gaps remain even in a language as porous as English. Author Howard Rheingold (2000) has collected hundreds of examples in his book *They Have a Word for It*. In it, he offers numerous gems, including the following:

> *rasa* (Sanskrit): The mood or sentiment that is evoked by a work of art
>
> *razbliuto* (Russian): The feeling a person has for someone he or she once loved but now does not
>
> *fucha* (Polish): Using company time and resources for your own ends
>
> *mokita* (Kiriwina, New Guinea): Truth everybody knows but nobody speaks
>
> *koro* (Chinese): The hysterical belief that one's penis is shrinking

The above examples are easily grasped, and one wonders why English speakers haven't been clever enough to think of them. But some untranslatable words have astonishingly nebulous meanings. This is what the Czech novelist Milan Kundera (1980) says about the Czech word *lítost*:

> *Lítost* is a Czech word with no exact translation into any other language. It designates a feeling as infinite as an open accordion, a feeling that is the synthesis of many others: grief, sympathy, remorse, and an indefinable longing. The first syllable, which is long and stressed, sounds like the wail of an abandoned dog.
>
> Under certain circumstances, however, it can have a very narrow meaning, a meaning as definite, precise, and sharp as a well-honed cutting edge. I have never found an equivalent in other languages for this meaning either, though I do not see how anyone can understand the human soul without it.

There. Clear on the concept of *lítost*? Likely not. Which brings us back to the problem of word learning. How is anyone, child or adult, supposed to arrive at an accurate understanding of a word whose meaning can best be captured as "a feeling as infinite as an open accordion"? How do Czech speakers come to reliably learn its meaning?

In all likelihood, learning such subtle words requires being exposed to many situations in which it is appropriate, and inferring its nuanced meaning from the contexts of its use. In fact, Kundera uses exactly this approach to teach his non-Czech readers the meaning of the word, offering vignettes such as the following:

> Let me give an example. One day a student went swimming with his girlfriend. She was a top-notch athlete; he could barely keep afloat. He had trouble holding his breath underwater, and was forced to thrash his way forward, jerking his head back and forth above the surface. The girl was crazy about him and tactfully kept to his speed. But as their swim was coming to an end, she felt the need to give her sporting instincts free rein, and sprinted to the other shore. The student tried to pick up his tempo too, but swallowed many mouthfuls of water. He felt humiliated, exposed for the weakling he was; he felt the resentment, the special sorrow which can only be called lítost. He recalled his sickly childhood—no physical exercise, no friends, nothing but Mama's ever-watchful eye—and sank into utter, all-encompassing despair.

Naturally, in order to map meanings onto categories, children need to first carve the world up into the right sorts of categories, which in itself is no small feat. In fact, having immature and still-unstable categories may well explain why kids are less eager than adults to map new words onto basic-level category members. Adults may simply be more certain than young children that collies and terriers *do* fall into the same category as Dalmatians, but that cats and pigs don't.

Cues for forming categories

Clearly, the process of mapping meanings onto words has to go hand-in-hand with the process of forming categories. So how do children learn that bananas belong in a category that excludes oranges? Or that dining room chairs, stools, and recliners all are the same *kind* of thing? Or that their own scribbled attempts at art have more in common with paintings hung in a museum than with splatters on the ground from accidentally spilled paint?

There's no single cue across all these categories that obviously groups category members together. For bananas versus oranges, it might be handy to pay attention to shape. But for different examples of chairs, what really seems to matter is that people sit on them. As for pictures versus accidental paint splatters, what's important is not so much how they look or even how people use them (a picture is still a picture even if someone can wash it off its surface), but whether someone *intended* to create them.

What are the essential properties that bind together the members of a category? The photographs in **Figure 5.4** give you an idea of some of the challenges children face in identifying both the features that category members have in common with each other and the features that distinguish them from members of other categories. While shape seems to be an especially useful cue in the early days of word learning, children as young as age two are able to tune in to all of these cues—shape, function, and the creator's intent—as a way of grouping objects into categories, paving the way for correctly naming a great variety of different categories.

It's easy to see that naming things accurately depends on having the right kinds of categories in place. What may be less obvious is that words themselves can serve as cues to forming categories. That is, just hearing someone name things can help children group these things into categories. An interesting line of research by Sandy Waxman and her colleagues (e.g., Waxman & Markow, 1995) reveals a very early connection in the minds of children between words and categories. The researchers used the following experimental setup: One-year-old babies were shown several objects of the same superordinate-level category (various animals, for instance). In the key experimental condition, all of the objects were named with the same made-up word during a familiarization phase—for example, "Look at the toma! Do you see the toma?" Afterward, the babies were shown either another animal or an object from a completely different superordinate category, such as a truck. The babies tended to be less interested in the new animal than in the truck—as if thinking, "I've already seen one of those—show me something different!" But what was interesting is that the novelty effect of the truck only showed up if all of the animals had been named by the same name. If each animal had been named with a different word, the babies paid as much attention to the new animal as the truck. But using the same word for all of the animals seems to have focused their attention on what the animals all had in common, so when they later saw another one, they were more indifferent to it.

In this version of the study, babies tuned in to the fact that when objects share a common label, this means they can be grouped into a common cate-

Figure 5.4 Do all of these images represent category members for the words on the left? Do you have any intuitions about the kinds of features that might be easiest for very young children to attend to?

gory. But the more general expectation that *language* communicates something about categories emerges at the tender age of 3 or 4 months. Alissa Ferry and colleagues (2010) found that after the familiarization phase, these tiny infants paid attention differently to objects of the same category than they did to objects of a completely new category, as the 1-year-olds did in the earlier study. But they did this only when the objects were accompanied by a phrase, such as, "Look at the toma! Do you see the toma?" When the same objects were accompanied by a musical tone instead, the babies didn't seem inclined to group the objects into categories.

Remember: at this age, babies haven't managed to segment words out of the speech stream yet. So it seems that well before they produce their first word, or possibly even pull out their first word-like unit from a torrent of speech, babies have the preconception that words are useful for talking about object categories. Happily, then, infants aren't left to flounder around among all meanings that words might theoretically convey. Once they begin learning the meanings of words, they can rely on certain built-in expectations about which meanings are the best candidates for words.

There's still some debate among researchers about the exact nature of some of these expectations. Are they innate? Are they specifically part of the language faculty? That is, expectations that act as constraints on word learning could amount to default assumptions that specifically link words to certain

BOX 5.2
Word learning in dogs

While Baby is busy learning the word *doggie*, the family dog may well be involved in its own word-learning project. Many a dog owner has boastfully claimed that his canine best friend can recognize many words of English, beyond the usual commands *sit*, *heel*, and *stay*. Research suggests that at least some of these dog owners may be telling the truth.

A number of studies over the past decade have put canine "wordiness" to the scientific test. Some dogs do indeed show evidence of being able to remember hundreds of words for different objects, demonstrating that they're able to retrieve a target object from among ten or more other objects, solely on the basis of the object's name. Dogs have also shown the capacity to learn new words after a single exposure, much as human infants can (Kaminski et al., 2004).

Naturally, this is fodder for those researchers who argue that language—or at least the capacity for fast word learning—is not a uniquely human trait, but one that we share to different degrees with other species. Nevertheless, showing that dogs can learn words doesn't necessarily mean that they get there by means of the same cognitive machinery.

For example, do dogs show the same kinds of word-learning biases that human babies do? One recent study investigated whether dogs, like children, accord special importance to shape as a diagnostic trait of objects that share the same name. This shape bias was first noted in children by Barbara Landau and her colleagues (1988), who demonstrated it as follows: Suppose you show a small child a novel object, saying, "This is a dax." Then you present her with a variety of objects that share features of color, size, texture, or shape with the original object and ask her, "Show me another dax." Does the child generalize to objects that are the same color as the first? The same size? There are a number of possible features to attend to as possible ways of defining the object category. It won't surprise you to learn that, for infants, object shape wins out over the other features. It seems clear that objects that have the same shape are more likely to be the same *kind* of object than those that merely happen to be of the same color or size.

But dogs may not show the same bias, according to a study by Emile van der Zee and colleagues (2012). These researchers studied a border collie named Gable who had solidly learned more than 40 words of English. They presented Gable with a word-learning conundrum similar to the one Landau's team used for children. They found that Gable had a tendency to extend the name *dax* to objects of the same size, ignoring their shape, color, or the other features. But even this was not a stable bias. Gable was allowed to take the "dax" home and play with it for several weeks while being repeatedly exposed to its name. When he was tested again after this period, he preferred to extend the word *dax* to objects that had a similar texture. This study suggests that dogs might have to explore a broader hypothesis space when it comes to generating guesses about the meanings of new words. Unlike children, they might not come into the world with pre-existing notions about how words and meanings are likely to line up.

meanings. Alternatively, they could be side effects of other kinds of cognitive and perceptual mechanisms that tend to organize the world in certain ways. For example, novice word learners often extend names on the basis of the shapes of objects. If they hear someone use the word *zav* to refer to a red ball and then they are asked whether *zav* also refers to an orange or to a red hat, they'll choose the object of the same shape (the orange) over the other object that has the same color as the original *zav* (see, e.g., Landau et al., 1988). Is this because there is a specific word-learning bias that directs them to equate shape with object labels? Or is it simply that they've figured out that shape is a more reliable cue than color for deciding which category the object falls into? (After all, grapes, apples, and shirts all come in different colors, so it's best to pay attention to their shapes in figuring out what kind of objects they are.)

We don't always know why young children make the assumptions they do about the alignment of words and meanings, but it's become clear that at various points in their early word learning, children have a whole repertoire of ways in which to constrain their possible hypotheses about word meanings. Given that the ability to learn words is not limited to humans, it would be interesting to know whether similar constraints apply to word learning in other species (see Box 5.2).

5.3 Understanding Speakers' Intentions

Associations versus intentions

So far, we've looked at two aspects of word learning: the sound sequences that are the bearers of meaning, and the things in the world—whether objects or categories of objects—that are linked with these sound sequences. But there's an important dimension we've neglected so far. I've been talking as if words "have" meanings and as if objects or categories "have" names. But of course, it's not that objects *have* names in the same way that they *have* certain shapes, colors, or textures. It's that people *use* certain names to refer to these objects. As discussed in Chapter 2, word meanings are a fundamentally social phenomenon. We link words with objects only because we accept that as members of a particular linguistic community, we've entered into an implicit social agreement to use certain words to describe certain concepts. Linguistic meaning, then, is a social compact within a cultural community—much like our agreement to use paper money with certain markings, or numbers electronically associated with a specific bank account, as proxies for things of value that can be exchanged for other things like food or cars. Our ability to enter into these complex social agreements distinguishes us from other animals and is probably an important prerequisite for language.

As adults, we can recognize that linguistic labels aren't inherently associated with objects or categories, and that they're social constructs rather than properties of the natural world. When we interact with someone, we understand that their use of a word isn't automatically triggered by the presence of a certain object, but rather, that the other person is using the word in the service of some intentional message involving that object. But how prominently does this understanding figure in a child's early word learning? Do very young children have the same understanding of the social underpinnings of language, knowing that names are used by speakers with the intent to refer? And if they do, does this understanding help guide the process of learning what words mean? Or do they simply create associative links between words and objects without being preoccupied by what's in the mind of the speaker who used that word?

associative learning Learning process by which associations between two stimuli are made as ideas and experience reinforce one another.

It takes children quite a while to behave like adults in terms of their ability to fully appreciate the contents of another person's mind. For example, if you've ever played hide-and-seek with a toddler, you probably discovered that he can fall prey to the charming delusion that if he can't see you, *you* can't see *him*—as long as his face and eyes are covered, he may be unconcerned that the rest of his body is in plain view. And, until about the age of four, children have a hard time thoroughly understanding that something that they know could be unknown to another person. (In Chapter 11, we'll explore in more detail some of the ways in which children fail to take the perspective of others.)

Given that babies begin the process of word learning long before they have a completely formed sense of people's states of mind, it could well be the case that their early word learning is of a more simple associative variety, and that it's only over time that they begin to treat language as the social construct that it is, reasoning about speakers' motives and intentions. So, you might think of a child's early word learning as similar to the associations learned by Pavlov's dogs in the famous classical conditioning experiments in which a bell was rung every time dinner was served. The end result was that the dogs eventually began to salivate whenever a bell was rung, even if dinner wasn't forthcoming. It's rather doubtful that the dogs went through a process of reasoning that the bell was intentionally being rung by the human experimenter with the specific purpose of communicating that dinner was on its way. More likely, they just paired the two events—bell ringing and dinner—together in their minds. This kind of simple-minded **associative learning** is well within the learning capacities of animals with far less intelligence than dogs. So it stands to reason that as a learning mechanism, it would be readily available to infants in the early days of vocabulary mastery.

All of which makes it really interesting to consider the experimental results of Dare Baldwin and her colleagues (1996), in which young children *failed* to learn the meanings of words on the basis of paired association between an object and a particular word. In this study, 15- to 20-month-old toddlers were left in a room with a novel object. While they played with it, and were therefore obviously focusing their attention on the object, they heard a disembodied voice intone "Dawnoo! There's a dawnoo!" over a set of loudspeakers. But they didn't end up learning that *dawnoo* was the name for the object they were playing with. Since it's been well documented that kids can learn the meaning of a word after a single exposure, what gives? This seems to be the perfect scenario in which rapt attention to a single object is paired with hearing a single common noun—in other words, just the scenario in which an associative mechanism should produce results in word learning.

The trouble was, the researchers argued, that children *don't* learn words like Pavlov's dogs. Even when they're extremely young, they understand that language is rooted in highly social behavior. And there was nothing in the situation that clearly signaled to the babies that anyone was actually intending to use the word *dawnoo* to refer to that object. Without this evidence, the infants weren't willing to attach the word *dawnoo* to the object right in front of them merely because they heard the name simultaneously with focusing their attention on the object. Essentially, the kids were treating the simultaneous utterance of the word as a coincidence.

In other studies, Dare Baldwin (1993) demonstrated that when there was clear evidence of a speaker's intent to refer, young children were able to map the right word onto an object even if when they heard the word, they were paying attention to a *different* object than the one intended. For instance, in one study, 18-month-old infants played with toys on a tabletop, and while they were absorbed with a toy, the experimenter would utter, "It's a modi!" The infants would look up at the speaker, only to find that *her* attention was focused

on a different object, at which point they would direct their own attention to this new object. When later asked to "find the modi," they'd be more likely to choose the thing that the experimenter had been looking at than the toy they themselves had been playing with at the time they'd heard the word.

WEB ACTIVITY 5.3

Learning through joint attention In this activity, you'll watch video clips of how babies typically learn language "in the wild." Notice how both the child and adult are involved in establishing joint attention during the interaction.

By 2 years of age, toddlers can put together some pretty subtle clues about a speaker's referential intent in order to figure out what words are likely to mean. In one clever study by Michael Tomasello and Michelle Barton (1994), an experimenter suggested to the child, "Let's find the toma. Where's the toma?" The experimenter then proceeded to look inside five buckets in a row. For all but one of the buckets, he lifted an object out of the bucket, scowled, and put it back. One object elicited an excited "Ah!" before the experimenter put it back inside its bucket. For the most part, children guessed that *toma* referred to the object that seemed to satisfy the experimenter. In fact, when children's word learning in this condition was compared with a different condition in which the experimenter showed only one object to the child while uttering a satisfied "Ah," there was no difference in performance. The children had no trouble rejecting the other objects as possible referents, based solely on the reactions of the experimenter.

These studies show that very young children routinely fail to learn word/meaning pairs when they have reason to doubt that these pairings are intended by the speaker. Tots also look askance at word/meaning pairs that are produced by speakers who obviously intend them but whose labeling behavior has been demonstrably bizarre in the past. For instance, Melissa Koenig and Amanda Woodward (2010) had 2-year-olds interact with speakers who labeled various familiar objects. When a speaker named three familiar objects inaccurately (for example, calling a shoe a "duck"), the children were more reluctant to learn a new word (for example, *blicket*) that he applied to a novel object than when the speaker behaved normally, giving all the familiar objects their conventional names.

This body of works clearly establishes that even at a very young age, small kids don't just link words together with objects that happen to be in their attention when they hear the word. But associative learning mechanisms are extremely powerful, and it seems far-fetched that children would be totally unable or unwilling to use them in the process of learning the meanings of words. What's more plausible is that on their own, associations are taken to be weakish evidence of word meanings, as compared with witnessing direct evidence of referential intent by a speaker who's abiding by normal linguistic conventions (see **Method 5.1**). Barring clear signs of the speaker's intent, children may need to hear multiple pairings before settling on the meaning of a word, or they may have a much more fragile memory trace of the word. And, solid evidence of a competent speaker's intent may trump associative processes when the two are in conflict.

This story would suggest that words could still be learned to some extent if, for some reason, kids didn't have a deep appreciation of the fact that words are used by speakers to satisfy a communicative goal, or if they lacked the ability to make inferences about what those goals are. And this seems to be largely true for children with **autism spectrum disorder (ASD)**, a neurological condition that impairs the ability to coordinate attention with another person, or to make inferences about someone else's state of mind. Autism researcher Simon Baron-Cohen teamed up with Dare Baldwin and Mary Crowson (1997) to replicate Baldwin's experiments in which the speaker uttered a word while looking at a different object than the one the child was gazing at. The researchers found dramatically different results between typical kids and kids with autism. The children with autism did end up mapping novel words onto

autism spectrum disorder (ASD) A neurological condition that impairs the ability to coordinate attention with another person, or to make inferences about someone else's state of mind.