Chapter 6

Concepts and Categories

Imagine a language with only proper names. A new word that names a dog must refer to that particular dog and nothing else. Learning this language would require the ability to track individuals over time, but it wouldn't require any ability to generalize, to recognize how collies are different from terriers or how dogs are different from tables.

Such languages don't exist, of course. Children who hear a word that refers to a dog have to cope with the possibility that it can also refer to other individuals that belong to the same category or kind. It could be a common noun, such as *dog.* To learn such words, they need some grasp of the conditions underlying category membership, some understanding of what is and is not a dog. Within psychology, this understanding is usually described as a concept (the concept of dog), and the concept that is associated with a word is usually described as the word's meaning (the meaning of *dog*).

This chapter discusses the concepts expressed by common nouns. It begins by asking why such concepts exist at all and goes on to defend an essentialist theory of their nature, contrasting this with the view that children's words are generalized solely on the basis of perceptual properties. It concludes with a discussion of the relationship between essentialized concepts and naive theories.

What Are Concepts For?

Why do we have concepts at all? What is the value of treating dogs as members of a category, instead of seeing them just as distinct individuals? John Locke (1690/1964, bk. 2) has an insightful discussion of this issue, in the course of asking why *general terms* (what we would call *common nouns*) exist in natural language.

He has three proposals. The first is often found in the contemporary psychological literature. This is the idea that thinking about the world only in terms of individuals is just too hard for us (p. 14): "it is beyond the power of human capacity to frame and retain distinct ideas of all the particular things we meet with: every bird and beast men saw; every tree and plant that affected the senses, could not find a place in the most capacious understanding.''

But this is a weak argument for several reasons. First, even if Locke is right about our capacity limitations (and he might not be; Pinker, 1997), it doesn't actually explain what concepts are good for. It is like answering the question ''Why is that person carrying a duck?'' by saying: ''Well, he isn't strong enough to carry a *car!*'' You cannot explain the existence of X (carrying a duck, having concepts) just by stating the impossibility of Y (carrying a car, storing all individual instances).

Second, in some cases we form concepts even though we also store the individual instances, and so categorization is not necessarily an alternative to individuation. For instance, I know all the members of my immediate family, but I also recognize that they fall into different categories (males, females, siblings, parents, and so on). For reasons that we get to below, even if an animal had a perfect memory and lived in a world with a small number of objects, it would still benefit from categorizing these objects.

Finally, categorization is not always useful. Consider the following categorization: all objects bigger than a bread box fall into one category; everything else falls into another. This would lead to great savings in memory—but presumably it isn't what Locke had in mind. Perhaps it isn't merely possessing concepts that is advantageous; it is possessing the *right* concepts.

Here is Locke's second proposal (p. 15): ''If it were possible [to name all individual instances], it would be useless; because it would not serve to the chief end of language. . . . Men learn names, and use them in talk with others, only that they may be understood. . . . This cannot be done by names applied to particular things; whereof I alone having the ideas in my mind, the names of them could not be significant or intelligible to another, who was not acquainted with all those very particular things which had fallen under my notice.''

This is reasonable. In a language with only proper names, I couldn't tell you about anything unless you were also directly acquainted with it, and this would severely limit the language's efficacy as a communication system. But this explanation applies specifically to words; it does not extend to concepts more generally (nor was it intended to). Since at least some concepts exist independently of words (see chapter 10), the fact that general terms are necessary for communication about novel entities does not explain the function of concepts more generally.

I think Locke's third proposal, however, is exactly right, and applies to both words and concepts (p. 15): ''a distinct name for every particular thing would not be of any great use for the improvement of knowledge: which, though founded in particular things, enlarges itself by general views; to which things reduced into sorts, under general names, are properly subservient."

In other words, we need general terms—and the concepts that underlie them—for the purpose of generalization. Generalization is essential for inductive learning, and successful induction is the stuff of life.

You drink orange juice, and you like it. You drink oil, and you don't. A chair supports you when you sit on it; when you sit on a cardboard box, the box collapses, and you fall to the floor. If you pat the cat, it rubs against you and purrs; if you pat the dog, it snarls. These events provide valuable lessons—about the joys of drinking juice and the hazards of drinking oil, about what you can and cannot sit on, and what creatures it pays to interact with. But you can learn from these events only if you have some mental representation of the relevant kinds. To learn from the juice episode, it is not enough to know that *this* liquid at *this* time is tasty; you have to be able to generalize to other liquids. A creature without concepts would be unable to learn and would be at a severe disadvantage relative to creatures that did have these sorts of mental representations.

Not all potential concepts are of equal value. Objects in the world are not randomly distributed with regard to the properties they possess. Instead, there is what Tolman and Brunswik (1935) call "the causal texture of the environment." Objects fall into categories because they are the products of physical law, biological evolution, and intentional design. And concepts might be useful only insofar as they correspond to categories that have many relevant properties in common.

This is an old idea. For instance, John Stuart Mill (1843, p. 153) notes that "some classes have little or nothing to characterize them by, except precisely what is connoted by the name." As an example of this, he says that all that white things have in common is that they are white, nothing else, and hence *white things* would be a poor category for inductive purposes. (And this might be why no language has a noun that refers only to white things; Markman, 1989.) Similarly, Jorge Luis Borges imagines a dictionary that divides animals into different categories, one of which is "those that have just broken a flower vase," and a children's joke goes as follows:

Question: What do a donkey and a refrigerator have in common?
Answer: Neither one of them is an elephant.

It is not a thigh-slapper, but it counts as a joke because of the manifest unfairness of the answer. The category "not an elephant" is profoundly unnatural from a psychological standpoint because there is nothing that nonelephants have in common that distinguishes them from elephants—other than the fact that they are not elephants, of course. Real categories—such as dogs, chairs, and water—have members that tend to share certain properties; once you know that something belongs to such a category, you know further facts about it that are not true of things that don't belong to that category.

Considerations of inductive utility help explain the privileged status of certain concepts. Roger Brown (1958a) noted that some names are more frequently used than others when talking about objects. For instance, we usually call Fido "a dog," not "an animal" or "a terrier." Eleanor Rosch and her colleagues (e.g., Rosch & Mervis, 1975; Rosch et al., 1976) have confirmed the importance of such basic-level concepts. They found that the basic level is the most inclusive level in which objects are judged to have many features in common, that people tend to interact with categories at the basic level using the same bodily movements, and that members of a basic-level category tend to share a common shape. And names for such categories are among the first common nouns learned by children.

What makes them so special? Brown speculated that we describe things at the basic level "so as to categorize them in a maximally useful way" (p. 20). Subsequent work has led to more explicit formulations of this insight. Murphy and Lassaline (1997) propose that the basic level is an optimal compromise between informativeness and distinctiveness: you can infer many unobserved properties once you know which basic-level category something belongs to (informativeness), and it is also relatively easy to make this categorization (distinctiveness). For instance, if you know something is a dog, you can infer a lot about it (it barks, eats meat, is a pet), much more than if you just know that it belongs to a superordinate category such as animal. And it is fairly easy to distinguish dogs from members of other basic-level kinds, such as horses, much more so than distinguishing members of contrasting subordinate categories such as collies versus terriers. For these reasons, we are most prone to think about and name Fido as a dog than as an animal or terrier.

This analysis gets us only so far. To say that good concepts divide the world into categories that are informative and distinctive is a bit like saying that good investment strategies make a lot of money: it might be true, but it is a poor guide to practical action. Humans have no direct access to the causal texture of the environment. So although we have an answer to the question of *why* we have concepts—having the right concepts underlies adaptive generalization—we are left with another, harder question: *How* is it that we come to possess these useful concepts?

Perception, Properties, and Essences

Regardless of how smart babies are, all they have access to when categorizing novel entities is the perceptual properties of these entities—what they look like, what sounds they make, and so on. Under one view, humans possess an "animal similarity space" that guides our initial learning about the world (Quine, 1960). We have evolved to see one dog as looking more like another dog than like a tree, and this perceptual similarity serves as the basis for early concept formation (see Keil, 1989, for discussion).

But a single invariant similarity space runs afoul of the fact that different properties are relevant for the adaptive categorization of different entities. For rigid objects, for instance, shape is highly relevant; this is how we typically distinguish tables and chairs. But for substances, color and texture are what matter: a circle of white paste is likely to have the same unobservable properties (such as taste) as a square of white paste, not a circle of red foam. Some animals undergo radical shape transformations, such as snakes; others don't, such as starfish. And the same entity might be categorized in different ways depending on the sort of induction one needs to make. For instance, different properties are relevant for determining whether something is poisonous versus whether it floats.

Young children are appropriately flexible in their categorization; the properties they attend to when categorizing a novel entity depend on whether it is an object versus a nonsolid substance (Soja, Carey & Spelke, 1991), a plant versus a rock (Keil, 1994), a real monkey versus a toy monkey (Carey, 1985), or an animal versus a tool (Becker & Ward, 1991). A shift in the properties that underlie categorization can be caused by quite subtle cues, as when eyes are added to simple geometrical shapes, giving them the appearance of being snakelike animals (Jones, Smith & Landau, 1991), and can also be motivated simply by changes in how an entity is described (e.g., Keil, 1994). Even babies show different patterns of generalization depending on the nature of what they are observing: the movements of objects are categorized on the basis of their trajectories, while the movements of animate beings are categorized on the basis of inferred goals (Spelke, Phillips & Woodward, 1995; Woodward, 1998).

Perhaps multiple innate similarity spaces are triggered by different stimuli, such as rigid objects, substances, and animate beings. In addition to this, similarity is a flexible notion (Jones & Smith, 1993). People can learn to categorize objects as falling into distinct categories on the basis of arbitrary contrasts along one or more dimensions, and such learning affects their subsequent intuitions of object similarity

(Goldstone, 1994). If you view objects as points in a multidimensional similarity space, and categories as clusters of objects, this modification can be viewed as the shrinking or stretching of different dimensions, bringing objects closer together or pulling them apart (Nosofsky, 1988).

If we were to stop here, we would have a minimalist proposal about the nature of human concepts. People are born with one or more perceptual similarity spaces that can be modified through experience with specific domains. Objects are seen as belonging to the same category to the extent that they cluster together in that space, and a concept is a representation of that clustering, as captured by exemplar (Nosofsky, 1988) or prototype (Hampton, 1995) models.

There are two reasons, however, to believe that this proposal is too minimal, even for young children. The first has to do with novel properties; the second concerns intuitions about essences.

Properties

Consider more abstract categories. We categorize people as, among other things, *stockbrokers* and *atheists,* objects as *dollars* and *menus,* actions as *swearing* and *weaning.* Under the theory that entities are categorized on the basis of the properties they possess (encoded either as features in a prototype representation or dimensions in a similarity space), the questions arise: What are the relevant properties in these cases, and where does our knowledge of them come from?

It is not a serious proposal that they emerge directly from our perceptual systems; we do not categorize people as stockbrokers solely on the basis of what they look like. It is also unlikely that these properties are abstract innate primitives. The existence of stockbrokers is unlikely to have been anticipated by biological evolution. And if stockbroker is an innate concept (or equivalently if stockbrokerhood is an innate property), then some way is needed for it to make contact with the external world, to connect (however indirectly) to our perceptual and motoric systems. This problem of "embodiment" arises for all putatively innate categories, and for many of them—such as *noun, cause, object,* and *person*—there exist plausible theories as to how the problem can be solved (e.g., Carey & Spelke, 1994; Pinker, 1984). There are no such proposals for stockbroker.

The traditional solution to the problem of abstract properties is empiricist. Humans start off with a perceptually based similarity space, but somehow—through perceptual and linguistic experience—this space comes to respond to increasingly more abstract notions. It starts off being able to categorize on the basis of features such as color and shape, and, through a sensitivity to the statistical properties of the environment, it becomes sensitive to the presence of stockbrokers and menus. Abstract notions are built up from perceptual ones. Explaining precisely how this occurs has long been a central project in empiricist philosophy, but the extent of progress has not been impressive. As Fodor (1981) points out, despite thousands of years of trying, *no* cases can be found in which any natural concept has been shown to emerge from a perceptually based similarity space.

There might be a principled reason for this failure. Perhaps concepts are not statistical abstractions from perceptual experience. Instead, they might be constituted, at least in part, in terms of their role in naive theories of the world (e.g., Carey, 1985, 1988; Gopnik & Meltzoff, 1997; Keil, 1989; Murphy & Medin, 1985). Our concept of stockbroker isn't a vector in some multidimensional perceptual state-space, then; it is instead rooted in our implicit understanding of society, money, jobs, and so on.

This proposal—sometimes called the *theory theory*—comes in different strengths. In its mildest form, the role of theories guides the choice of which of a set of available features will be relevant in any given situation—color for food, shape for artifacts, and so on. Under a stronger version, theories actually lead to the creation of new features (e.g., Wisniewski & Medin, 1994). The most radical version rejects the very idea that category membership is computed on the basis of similarity of features. Murphy and Medin (1985) give the example of seeing someone at a party jump fully dressed into a swimming pool and categorizing him as being drunk. This categorization isn't made because the person is perceptually similar to other drunks; it is made because categorizing him this way provides the most plausible explanation for his behavior. The proposal that causal and explanatory considerations underlie categorization is something that I return to at the end of this chapter.

Essences

A related issue concerns the relationship between the superficial properties of an object and how it is categorized. This relationship might be direct. That is, if objects resemble one another to a sufficient extent, they fall into the same category.

Many creatures might categorize objects this way, but humans do not. We see commonalities that transcend appearance. We categorize a caterpillar and the butterfly it later becomes, and a squalling infant and a grown man, as the same individuals, despite the radical changes in appearance. A hummingbird, ostrich, and falcon belong to the same category (birds) and are distinct from other perceptually similar animals such as bats. Whales are not fish, fool's gold is not gold, marsupial mice are closely related to kangaroos, and glass is actually a liquid. At least for educated adults, the category that a thing belongs to is not merely a matter of what it resembles.

The term for what is maintained through the transformation from caterpillar to chrysalis to butterfly, and for what hummingbirds, ostriches, and falcons share, is *essence.* Locke (1690/1964, p. 26) defines this as follows: "Essences may be taken for the very being of anything, whereby it is what it is. And thus the real internal, but generally . . . unknown constitution of things, whereon their discoverable qualities depend, may be called their essence."

People's naive intuition that certain categories have essences is called *naive essentialism* or *psychological essentialism* (Medin & Ortony, 1989). This proposal does not entail that people actually know what the essences are. For instance, to be an essentialist about water does not require that you know the internal properties that make something water (presumably being H_2O), just that you believe some such properties exist. Hence an essentialist should be able to entertain the possibility that something might resemble water but not actually be water (because it lacks the essence) or not resemble water but be water nonetheless (because it has the essence). It is possible that people were essentialists about water before the development of modern science; in fact, the belief that certain entities have essences might be what motivates scientific inquiry in the first place.

There is a difference, of course, between the claim that people believe that categories have essences and the claim that such essences actually exist. In an important review, Susan Gelman and Larry Hirschfeld (1999, p. 405) suggest that although naive essentialism might be innate and universal, it nonetheless "may yield little insight about the nature of the world." In particular, as Gelman and Hirschfeld point out, contemporary biologists are adamant that species actually have no essences. This is nicely summed up in the title of a classic article by David Hull (1965): "The Effect of Essentialism on Taxonomy: Two Thousand Years of Stasis."

But I don't think the human propensity toward essentialism is actually a mistake (see also Pinker, 1997). In modern biology, species are viewed as populations that evolve, with no sharp boundaries in space or time, and so biologists reject the Aristotelian notion that species are unchangeable ideal types with no intermediate forms; something either is a bird or isn't; there's nothing in between (e.g., Mayr, 1982). It is *this* sort of "essentialism" that is mistaken. But this is much stronger than Lockean essentialism, under which the superficial features of entities are the result of deeper causal properties. Essentialism in this more general form is simply a belief that reasons exist as to why things fall into certain categories: birds are not merely objects that resemble each other but instead have deeper properties in common. This sort of essentialism is rampant in current biological thought, both in the sciences and in the folk theories of different cultures (Atran, 1998).

Essentialism is an adaptive way of looking at the world; it is adaptive because it is true. In biology, animals are hierarchically grouped into species, families, classes, and so on because of their evolutionary history. One can usually classify an animal on the basis of its appearance, but when there is doubt, the most reliable indicators of an animal's category are properties such as embryonic features and genetic structure. (This is how we know, for instance, that chimpanzees are more related to humans than they are to any other primate). The scientific notion that a hummingbird, ostrich, and falcon all belong to a different category than a bat or that a peacock and a peahen belong to the same species are biological insights that emerge from an essentialist world view (see Pinker, 1997, for discussion).[1]

The idea that essentialism is an adaptive stance toward the material and biological world raises the possibility that it is unlearned. This is a controversial proposal. Many scholars would argue instead that essentialism is a cultural construct. Jerry Fodor (1998, p. 155), for instance, suggests that it emerged only with the modern rise of science, and so "*of course* Homer had no notion that water has a hidden essence, or a characteristic microstructure (or that anything else does); a fortiori, he had no notion that the hidden essence of water is causally responsible for its phenomenal properties."

Fodor notes that we have to be cautious with regard to what counts as evidence that children, or adults from other cultures, are naive essentialists. It is not enough to show that they have categories that are not purely perceptual (surely Homer knew about jokes), or that they distinguish between appearance and reality, or even that they believe that things have hidden properties. To be an essentialist in the Lockean sense, you must believe that these hidden properties are causally responsible for the superficial properties of an entity and determine the category that it belongs to.

Fodor is right that there is no proof that essentialism is either innate or universal. But quite a bit of evidence is highly suggestive. Some of this comes from cross-cultural research (e.g., Atran, 1998), but I want to focus here on studies of young children.

Preschoolers use category membership to infer deeper properties that animals have. For instance, in one set of experiments (Gelman & Markman, 1986, 1987), children were told that a brontosaurus has one property (cold blood) and rhinoceros has another (warm blood) and were then asked which property a triceratops had. A triceratops looks more like a rhinoceros than a brontosaurus, and so if children's inductions are solely based on perception, they should guess that the triceratops has warm blood like the rhinoceros. But when both the triceratops and the brontosaurus are described as ''dinosaurs,'' children infer that the triceratops has cold blood. Such a finding can be explained by children's use of the sameness of label as a cue that the triceratops and the brontosaurus fall into the same category and their belief that members of the same category share the same deep properties, even if they don't look alike (see also chapter 10).

Young children also know that membership in a category is not solely determined by appearance. A porcupine that has been transformed so that it looks like a cactus is still a porcupine; a tiger that is put into a lion suit is still a tiger (Keil, 1989). Four-year-olds know that if you remove the insides of a dog (its blood and bones), it is no longer a dog and cannot do typical dog things such as bark and eat dog food, but if you remove the outside of a dog (its fur), it remains a dog, retaining these dog properties (Gelman & Wellman, 1991). They believe that the skin color of a human child is determined by the biological parents, not the people who raised the child (Hirschfeld, 1996).

Finally, four-year-olds are more likely to accept a common label for animals described as sharing internal properties (''the same sort of stuff inside, the same kind of bones, blood, muscles, and brain'') than superficial properties (''lives in the same kind of zoo and the same kind of cage''). This holds for both middle-class children raised in Western societies (Diesendruck, Gelman & Lebowitz, 1998) and Brazilian children from shanty-towns, with little formal education, and with limited access to books and television programs (Diesendruck, under review).

Early emergence of essentialist ideas does not entail their innateness, of course. It is conceivable that children somehow pick up this perspective from the adults around them, but there is little support for this view. Even highly educated parents in university towns rarely talk to their children about insides and essences (Gelman et al., in press a), and working-class parents are considerably less likely to do so (Heath, 1986).

The Importance of Shape

Even if children do have an understanding of essences, this might play little role in their learning and use of words (e.g., Malt, 1991, 1994). It is sometimes suggested that children's naming is done entirely on the basis of superficial properties that entities possess. The strongest and most interesting version of this claim has been defended by Linda Smith, Susan Jones, and Barbara Landau (1992, pp. 145–146):

> In learning language, children repeatedly experience specific linguistic contexts (e.g., "This is a ____" or "This is some ____") with attention to specific object properties and clusters of properties (e.g., shape or color plus texture). Thus, by this view, these linguistic contexts come to serve as cues that automatically command attention. The evidence on children's novel word interpretations also suggests that the momentary salience of object properties interacts with contextual cues and influences children's generalization of a novel word (e.g., Smith et al., 1992). All in all, the data from artificial word-learning experiments are consistent with the idea that dumb forces on selective attentions—that is, associative connections and direct stimulus pulls—underlie the seeming smartness of children's novel word interpretations.

As they discuss, evidence suggests that children rely on perceptual properties—and specifically shape—when generalizing words. In particular, when given a new count noun that refers to a rigid object, children will typically extend that noun to other rigid objects of the same shape, not to those of the same size, color, or texture (e.g., Baldwin, 1989; Jones, Smith & Landau, 1991; Landau, Smith & Jones, 1988, 1998; Smith, Jones & Landau, 1992, 1996). Following Landau, Smith, and Jones (1988), we can call this the *shape bias.*

There are two theories about the nature of this bias. One is outlined in the quote above: shape is important because children observe that words used in the context "This is a ____" that refer to rigid objects are generalized on the basis of shape. Because of this, when children first hear "This is a dog" used to refer to a dog, they know that it should be extended to other similarly shaped objects. We can call this the *brute-shape theory.*

An alternative, which is more consistent with an essentialist perspective on concepts, is that shape is important because it is seen as a cue to category membership. Children know that count nouns can refer to kinds of objects, and they believe that an object's shape is highly related to the kind it belongs to (e.g., Bloom, 1996a; Gelman & Diesendruck, in press; Keil, 1994). We can call this the *shape-as-cue theory.*

Where would this belief about the importance of shape come from? There are several possibilities. It could be that the importance of shape for object categories is unlearned (e.g., Biederman, 1987; Landau & Jackendoff, 1993). It could be that shape is important because it is seen as highly nonrandom (see Leyton, 1992), and therefore if two objects are of the same shape, children infer that they most likely share deeper properties. Or it could be that children learn through experience that entities of the same kind tend to share the same shape. Under this view, the shape bias is learned, just as in the brute-shape view.

These two theories differ in a couple of substantial ways. First, brute shape predicts that the shape bias should be limited to words, while shape-as-cue predicts that it should apply to categorization more generally (see Ward, Becker, Hass & Vela, 1991, for discussion). Second, brute shape implies that shape *determines* naming, while shape-as-cue predicts that shape is a *cue* to naming. This raises a clear prediction: If brute shape is right, then children should sometimes generalize object names on the basis of shape, using them to refer to entities that belong to different kinds, while if shape-as-cue is right, children should sometimes generalize object names on the basis of kind, using them to refer to entities that have different shapes.

Familiar Categories

One way to address this issue is by studying how children learn new words that refer to familiar categories. Some researchers have used a procedure in which they show children a target object, give it a novel name (such as "This is a dax"), and test whether children extend the word to either a same-shaped item that belongs to a different-kind, or to a different-shaped item that belongs to the same kind (e.g., Baldwin, 1992; Golinkoff et al., 1995; Imai, Gentner & Uchida, 1994). An example of this, from Baldwin's study, is shown in figure 6.1.

Such studies typically find that children are biased to extend the word to the same-shape item, despite the difference in kind. For instance, preschool children who are told that the egg is "a dax" think that the football, and not the loaf of bread, is also a dax. This is said to support the view that young children assume words refer to categories that share a common shape, not a common kind, in support of brute shape.

But there is something worrying about the design of these experiments. As mentioned above, if two objects are the same shape, they are likely to belong to the same basic-level kind (Rosch et al., 1976),

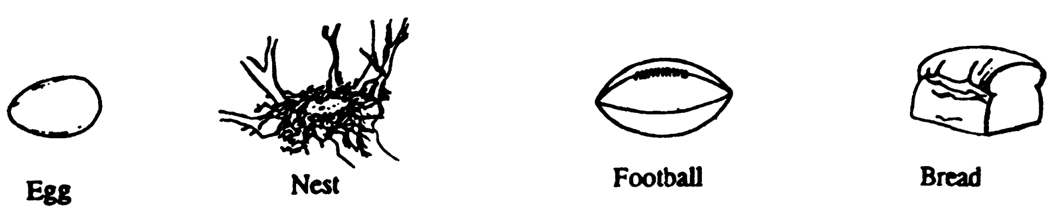

Figure 6.1
Stimuli from generalization study (from Baldwin, 1992)

and so, to contrast shape and kind, the stimuli are designed so that the sameness of kind is at the *superordinate* level. For instance, in the example above, the egg and the loaf of bread are both foods.

This makes generalization on the basis of kind particularly difficult for children. For one thing, if an object is simply labeled "This is an X," *X* is typically a basic-level term. It would be odd to point to an egg and say "That is food" (e.g., Brown, 1957; Horton & Markman, 1980; Mervis & Crisafi, 1982; see chapter 2). In addition, categorizing on the basis of a superordinate kind is harder than categorizing on the basis of a basic-level kind (e.g., Rosch & Mervis, 1975; Markman, 1989). For these reasons, children might fail to consider the superordinate kind when it comes to extending the name and instead might fall back on choosing the more perceptually similar item (Gelman & Diesendruck, in press).

If this is so, then one would expect quite another result if the different-shaped item belonged to the same *basic-level* kind. Golinkoff et al. (1995) contrasted a target item (such as a banana) with another object of the same shape but a different kind (the moon) as well as an object of the same basic-level kind (another banana). This shift to a basic-level kind made a dramatic difference: when children were given a name for the target object, they tended to generalize this word to the object of the same kind, not the one of the same shape.

This result could be taken as supporting shape-as-cue. But unfortunately the Golinkoff et al. study has another problem, one that bedevils all experiments that use familiar categories. Children may interpret the words they hear as synonyms for existing category labels. There are no pragmatic cues to the contrary and no other plausible meanings the words could have. In the example above, then, they might have construed the word as *banana,* and then, when asked to extend the word, they simply looked for the other banana.

Even if this is the case, this study still tells us something interesting: it shows that children know some object names that they do not extend solely on the basis of shape. That is, for three- and four-year-olds, *banana* does *not* merely mean "an object that has such-and-so shape." What does this understanding tell us about the nature of the shape bias?

Familiar Words

Consider a child who uses the word *clock* to refer to both a grandfather clock and a digital clock. A proponent of shape-as-cue might say that this reflects the child's understanding that these different-shape objects belong to the same kind. But there is a reasonable brute-shape reply: since the grandfather clock and the digital clock really are both clocks, the child might have heard other people name each of them this way and is simply parroting back this usage. A similar response could be made about the non-shape-based responses in the Golinkoff et al. study.

A more revealing source of information may be children's errors. Suppose, for instance, that the child calls a sandal "a boot." Since adults do not use the word in this manner, this sort of error should give us some insight into how children make sense of the meanings of words.[2]

Many of children's errors are uninformative, given that perceptual similarity and sameness of kind are highly correlated in real-world object categories. If a child calls a cat "a dog," this could be because she thinks that the cat belongs to the same kind as other dogs she has seen (in accord with shape-as-cue), or it could be because it is similar in shape to other dogs she has seen (in accord with brute shape).

There are potentially more telling errors, but these are also difficult to interpret. For instance, as mentioned in chapter 2, at 20 months, Max put a piece of yellow pepper on his head at dinnertime and said "hat." It is tempting to view this as evidence against brute shape, since the pepper was not shaped like a typical hat. On the other hand, at 21 months, he called a penis-shape ice cream cone "pee-pee"—which seems like a serious problem for shape-as-cue. But it is hard to know what to make of either case; perhaps Max might have been trying to express his opinions that these items *were similar to* a hat and a penis, not that they actually were members of these categories. (Or he could have been goofing around.)

Experimental studies have problems as well (Gelman et al., in press b). Suppose, for instance, that two-year-olds are shown a picture of a collie and a picture of a horse and are asked to "Point to the dog." They are likely to choose the collie. But this doesn't show that they have an adult understanding of *dog*, since the children might actually believe that both of the pictures denote dogs. It is just that the collie is a more typical dog, and so it is a better choice. On the other hand, if you were to show the children a display without a dog, such as a picture of a book and a picture of a horse, and ask the same question, they might choose the horse. But this would not necessarily show that they believe that a horse is a dog, it might be because the horse most resembles a dog and the children felt pressured to choose *something*. Preferential-looking tasks (e.g., Naigles & Gelman, 1995) suffer from the related problem that direction of gaze does not necessarily reflect intuitions about word meanings. If I am shown a picture of a horse and a picture of a book and hear the word *dog*, I might again look at the horse, not because I believe it is a dog but because it looks most like a dog.

Dromi (1987) points out that a robust test of children's understanding is a procedure that allows children to refuse to choose a picture if no correct referent is available. This is precisely what Gelman et al. (in press b) devised, based on a methodology originally used by Hutchinson and Herman (1991). They presented children with two pictures. One was visible, and the other was hidden behind a cardboard screen. The children were asked questions such as "Where's the dog?"

In one condition, the visible picture shows a dog, and children who know the word should point to this picture; they should not choose the hidden one. In another condition, the visible picture shows a cow. If children know that a cow is not a dog, then they should choose the hidden picture when asked for "the dog." If they do pick the cow (and if they show other indications of being able to understand the task), this suggests that they really do think that the word *dog* extends to cows.

Gelman et al. tested children on two nouns—*apple* and *dog*. Even two-year-olds were usually correct (they did not tend to confuse apples and oranges, for instance). On the occasions that overgeneralizations did occur, the children were just as likely to generalize to other objects of the same (superordinate) kind but of a different shape (picking a banana as "an apple") as they were to choose objects of a different kind but the same shape (picking a baseball as "an apple"). Gelman et al. concluded from this that shape has no privileged status in children's errors.

Gelman et al. went on to discuss why other investigators so often find shape-based errors in production and comprehension. Such errors tend to be most frequent with items that are unfamiliar for children, which follows from shape-as-cue. If you know little or nothing about an object, it is a reasonable assumption that shape is relevant to kind membership—particularly in an experiment where all that is available is a line drawing. Also, some overgeneralizations on the basis of shape might occur because the same-shape object is seen as a *representation* of the item being named. For instance, Gelman et al. found that even four-year-olds often selected a pink spherical candle as "an apple," perhaps because the children, not knowing that the item was a candle, thought it was a representation of an apple, such as a toy apple. The naming of representations is an issue we return to in the next chapter.

New Words for New Categories

When learning a new word for a novel category, how do children generalize this word? It depends. One consideration is the word's syntactic category. For instance, nouns are treated differently from adjectives. If children hear an unfamiliar object named as "a zav," they will tend to generalize the word on the basis of shape, but if they hear it described as "a zavish one," they are more prone to generalize on the basis of color or texture (e.g., Hall, Waxman & Hurwitz, 1993; Smith, Jones & Landau, 1992; Taylor & Gelman, 1988; see chapter 8).

It also depends on the kind of entity. While names for rigid objects are typically generalized on the basis of shape, even two-year-olds generalize names for nonsolid substances, such as jelly, paste, or instant coffee, on the basis of texture and color (Macario, 1991; Soja, Carey & Spelke, 1991). If a named object looks like an animal that can undergo postural change, the shape bias again goes away (Becker & Ward, 1991; Jones, Smith & Landau, 1991).

Finally, a word's interpretation can depend on how the entity is described. Keil (1994) finds that three- and four-year-olds tend to generalize names for objects described as animals on the basis of shape ("This is my hydrax. It's a kind of animal"), but if the same objects are described as nonliving natural kinds ("This is my malachite. It's a kind of rock"), they ignore shape and focus on color. There is the same effect of being told that something is a kind of animal versus a kind of food (Becker & Ward, 1991).

These studies tell us that children are not limited to generalizing words—even count nouns that refer to rigid objects—on the basis of shape. They are appropriately sensitive to other considerations, such as color and texture, which accords with the results of Gelman et al.'s study of children's overgeneralizations.

This is all consistent with shape-as-cue, but it is hardly decisive evidence for it, let alone for naive essentialism more generally. The results reviewed in this section are easily accounted for by the theory that perceptual properties directly underlie categorization, so long as one makes the reasonable assumption that the relevant features (or weights of the relevant dimensions) can change as a function of experience (e.g., Jones & Smith, 1993). Under such a view, for instance, it is *not* that the presence of eyes on an object tells children that the object is a snakelike animal (one that can bend and twist), and this is why they extend a name for this object to other objects of different shapes. Instead, children have simply observed that words that refer to objects with eyes are generalized on the basis of texture and color, not shape. None of these studies, then, resolves the question of whether children's naming of objects is based on an essentialist understanding of categories.

Artifacts

One domain of considerable interest is the learning of names for artifact categories—for human-made entities such as chairs and clocks.

What is the nature of adults' understanding of such categories? Chairs and clocks come in a range of shapes and sizes: there are beanbag chairs, basket chairs, deck chairs, chairs for dolls, chairs shaped like hands, and chairs suspended from ceilings on chains; there are grandfather clocks, digital clocks, clocks shaped like coke bottles, and clocks for the blind that tell the time at the press of a button. In the course of your life, you will be exposed to an extraordinary array of chairs and clocks, some emerging through technological advance, others that extend the boundaries of fashion or aesthetics, still others that exist in fiction, either historical, fantasy, or futuristic. What makes these things chairs and clocks plainly does not reduce to facts about their appearance.

What about function? Perhaps chairs are things we sit on; clocks are things that tell time. But this is also a nonstarter. One can sit on the floor, but this doesn't make it a chair. And a fragile chair that would break if you tried to sit on it is nonetheless still a chair. I can tell the time by looking at the shadow of a tree, but a shadow is not a clock, and if Big Ben stopped working—if it could no longer fulfill the function of telling time—it would not cease being a clock.

What about *intended* function? This is more promising. Perhaps a chair is something that was built with the intention that people sit on it; a clock is something that was built with the intention that it tell time. This is a better cue to artifact-kind membership than current function; in studies in which intended function and current function are pitted against each other, intended function wins out (Hall, 1995; Keil, 1989; Rips, 1989). A theory based on intended function also accounts for our intuitions that the floor isn't a chair and a shadow isn't a clock, but a broken chair and broken clock remain a chair and a clock.

But such a theory doesn't fully capture our intuitions about artifact category membership. Barbara Malt and Eric Johnson (1992) found that adults will often agree that something that looks very much like a boat *is* a boat, even if they are explicitly told that it wasn't created with the intent to serve the sorts of functions that boats typically perform. Similarly, it is true that chairs are usually designed for people to sit on—but benches, stools, and sofas are also designed for this purpose, and no unique function seems to distinguish chairs from these other categories. Finally, there is nothing incoherent about someone creating a chair without any desire that people sit on it. For instance, someone might build a chair as a prop in a play; it gets shattered with a sledgehammer in the first act. Still, it is a chair and everyone would call it one.

An alternative theory is that the categorization of artifacts is rooted in our intuitions about creator's intent and how it relates to the design of an object (Bloom, 1996a, 1998; Dennett, 1990; Keil, 1989; see also Malt & Johnson, 1998, for a critical discussion). We categorize something as being a member of an artifact kind if its current appearance and use are best explained in terms of the intent to create something that falls into that kind. Under this view, when we judge that something is a chair, we are not judging that it looks and functions like other chairs we have encountered; we are instead judging that it looks and functions as if it was created with the same intent as other chairs we have encountered.

Much of the time the relationship between appearance and intent is transparent. If something resembles a typical chair, for instance, then it is highly likely that it was created with the intent to be a chair. As Daniel Dennett (1990) notes, "There can be little doubt what an axe is, or what a telephone is for; we hardly need to consult Alexander Graham Bell's biography for clues about what he had in mind."

Categorization becomes more complicated, however, when dealing with atypical exemplars, and here the merits of an intention-based theory emerge. Malt and Johnson (1992) found that if an object was described as looking very different from typical members of a category, it was usually not judged to be a member of the category, regardless of its function. For instance, a rubber sphere hitched to a team of dolphins is not called *a boat,* even if it is described as being created for the function of carrying people across water. But there were a few exceptions to this finding. For instance, a small cube with an extendible rod and a digital display was categorized as a *ruler,* and a large disk suspended from the ceiling by cables, with a fold-down seat attached, was categorized as a *desk.*

What distinguishes the atypical entities that were judged as being category members from those that weren't? Malt and Johnson (1992) point out that for the atypical objects, "the unusual physical features . . . may be construed as more advanced or effective than the current features, and the descriptions may be interpreted as plausible futuristic versions of the articles" (p. 209). In other words, we can envision someone creating the cube and the disk with the intent to make a (more advanced, more effective) ruler and desk, while it is implausible that someone who intended to build an object belonging to the class of boats would make the sphere.

This is an essentialist account of artifacts. Applying essentialism to this domain is controversial. Philosophers such as Locke and Mill have proposed essentialism as a doctrine exclusively about categories in the natural world—about natural kinds—and many psychologists would argue that this holds as well for psychological essentialism: humans are naive essentialists about entities such as birds and water, but not about chairs and clocks. As Steven Schwartz (1978, p. 572) puts it, "The big difference between artifact kinds and [natural kinds] is that we do not propose that there is any underlying nature that makes something the kind of artifact that it is."

It is true that we don't think of artifacts as having internal essences in the sense of natural kinds. But this doesn't refute the view that artifacts are seen as having essences in the sense of having deeper causal properties that explain their superficial features and determine the categories they belong to. Instead of being biological and chemical, the essences of artifacts are social and psychological (see Keil, 1989; Putnam, 1975). In fact, the same essentialist intuitions we find for natural kinds apply to artifacts as well. Just as the superficial parts and properties of animals can be explained in part by their internal essence, the superficial parts and properties of artifacts can be explained in part by the intentions underlying their creation. Most of the time, we can categorize animals by their superficial properties, but in hard cases, experts might look at their internal structure. Most of the time we can categorize artifacts by their superficial properties, but in hard cases, experts—such as archaeologists, anthropologists, and historians—will attempt to figure out what they were intended to be. (Consider the debates over whether certain carved stones left by Neanderthals are tools, religious artifacts, or artwork.) This sort of *artifact hermeneutics,* to use Dennett's (1990) term, is rooted in essentialist assumptions about the nature of artifacts.

Naming by Children

Do children have this same essentialist understanding of artifacts? Few people have explored this issue. Most studies that have explored this domain have asked instead whether children extend these names on the basis of shape or on the basis of function. Implicit in this question is the notion that function determines how adults categorize artifact kinds (e.g., Miller & Johnson-Laird, 1976). Hence if children attend to function, they are being adultlike; if they attend to shape, they are more perception-bound.

Some evidence suggests an early understanding of function by children. Using a dishabituation paradigm, Kolstad and Baillargeon (1990) found that 10-month-olds are sensitive to a functional property—the ability to serve as a container—when categorizing objects and that this can override perceptual similarity. Similarly, Brown (1990) found that two-year-olds, having learned to retrieve an object with a rake, would attempt the same action with other objects of dissimilar shape but not with same-shape objects that lacked the length and rigidity to do this action—a finding that Hauser (1997) replicated with cotton-top tamarins. Other studies have found that two-year-olds are sensitive to correlations between the form of an animal and the "functions" that it can fulfill (McCarrell & Callanan, 1995).

But is function the basis for their generalization of novel artifact names? In a classic study, Dedre Gentner (1978) showed children and adults two novel objects. The "jiggy" was a square box with a cartoon face; operating a lever on the box made the face change expression. The "zimbo" was a modified gumball machine; operating a lever caused jellybeans to drop. Subjects were then shown a hybrid object that looked like a jiggy but dispensed jellybeans like a zimbo and were asked whether it was a jiggy or a zimbo. Two- to five-year-olds generalized the word on the basis of appearance (saying that it was a jiggy), while five- to 15-year-olds generalized on the basis of function (saying that it was a zimbo). Surprisingly, however, adults acted just like the younger children; they said it was a jiggy, generalizing on the basis of appearance.

Two recent studies that pit shape against function are those of Deborah Kemler-Nelson et al. (1995) and Landau, Smith, and Jones (1998).[3] Kemler-Nelson et al. showed three-, four-, and five-year-olds a novel artifact, such as a T-shaped metal object with attached brushes. They were told that the object was "a stennet." The experimenter modeled its function, such as dipping the stennet into paint and making a design, and the children were then encouraged to use the stennet to make their own designs. They were then shown other objects: some had the same overall shape, others had a different shape; some could do the same function, others could not. For each, they were asked, "Do you think this is a stennet or that this is not a stennet?" Children in all age groups tended to extend the word on the basis of function, not shape. In a recent set of studies with similar stimuli, Kemler-Nelson (in press) found that even two-year-olds generalized new artifact names on the basis of function.

Landau, Smith, and Jones (1998) used a similar design. Two-, three-, and five-year-olds and adults were shown a simple object and told, for instance, that it was "a rif." Then they were told "Rifs are made by a special company so they can do this," and a function, such as wiping up water, was demonstrated. Then the subjects were shown different objects and asked for each one: "Is this a rif?" Landau et al. obtained the opposite result from Kemler-Nelson. The word tended to be extended by children on the basis of shape, not function. Adults, in contrast, did extend the word on the basis of function.

How can one explain these different results? Both Kemler-Nelson et al. (1995) and Landau, Smith, and Jones (1998) have some suggestions, including the fact that the stimuli in the Kemler-Nelson et al. study were more complicated and realistic than those in the Landau et al. study and contained familiar object parts. Also, the subjects in the

Kemler-Nelson et al. study had more experience with the objects' functions, perhaps making function more salient.

As I noted earlier, one concern with these studies is that, at least for adults, artifacts aren't categorized exclusively on the basis of shape or function. It may be that these considerations are relevant only to the extent that they are cues to the intent underlying the creation of the object. From this perspective, we are in a position to speculate—in what is admittedly a post hoc fashion—as to why the two studies got such different results.

In the Kemler-Nelson et al. study, the functions were highly specific and reflected intentional design: being able to paint parallel lines isn't the sort of thing that an artifact can do by accident. This should motivate categorization on the basis of function; children (as well as adults) might reason that other objects would be able to do the same thing only if they were intended to fall into the same category.

In the Landau et al. study, the functions—such as wiping up water—were simple and dependent only on the substances that the artifacts were made of. There was no motivation to believe, then, that the objects were created with the express intent that they fulfill that function. This might motivate categorization on the basis of shape. After all, in the absence of any other reliable cue, shape is highly diagnostic as to creator's intent.

As I said, this is post hoc, but it does lead to a prediction. One can have a situation in which two objects have the same shape, there is some explanation for why they are that shape, but the explanation doesn't entail that they have been intended to belong to the same kind. If this hypothesis is right, the shape bias should go away.

To test this hypothesis, Lori Markson, Gil Diesendruck, and I exposed children to a situation like the one shown in figure 6.2, using stimuli based on those used by Landau et al. (Bloom, Markson & Diesendruck, 1998). We used a target object and two test objects—one of the same shape and a different material from the target, the other of a different shape and the same material.

We tested adults and four-year-olds. There were two conditions. In one, we simply put down the target object and named it, as in "This is a fendle." We then presented the other two objects and asked, "Which one of these is a fendle?" Not surprisingly, we got a shape bias, replicating Landau et al. (children: 83 percent, adults: 90 percent).

The second condition was identical, except that the same-shape object was used as a tight-fitting container for the target object—and, at the start of the study, the target object was removed from this container. We reasoned that this manipulation would indicate to our

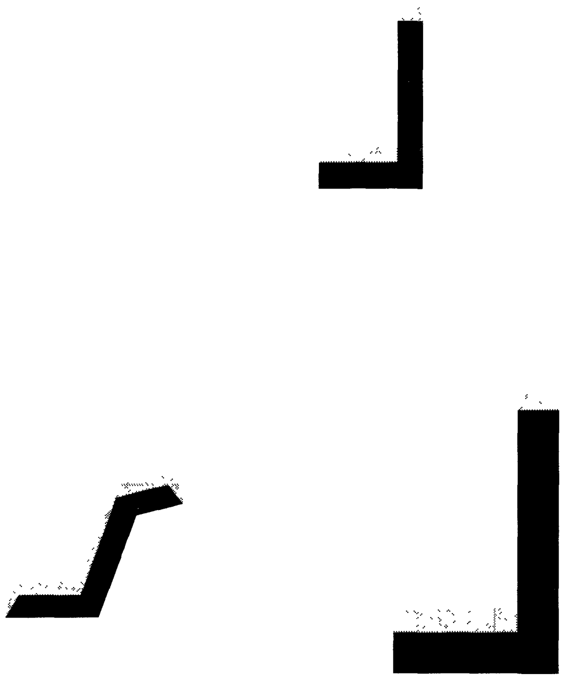

Figure 6.2
Stimuli from container study

subjects that the two same-shape objects were created for different purposes, while at the same time giving some intuitive rationale for why they were the same shape.[4] In this condition, the shape bias goes away (children: 55 percent, adults: 50 percent). This is exactly as predicted by shape-as-cue and, more generally, by an essentialist theory of artifact categories.

Several facts about children's generalizations still need to be explained: Why did the adults in the Gentner study behave differently from the older children and adolescents? Why did Landau et al. find a developmental shift while Kemler-Nelson did not? But I think that a change in perspective—moving away from pitting shape and function against each other and instead exploring the extent to which they serve as cues to underlying intent—might give us insight into some of these puzzles (see also Bloom, Markson & Diesendruck, 1998).

The debate over whether children are essentialists in their language and thought has focused almost exclusively on the domain of animals. Biological species have long been the paradigm case of essentialized categories within philosophy, and psychologists have followed suit. But by many accounts, young children's understanding of biology is quite limited (e.g., Carey, 1985; Keil, 1989), particularly in contrast with their rich understanding of, and interest in, goals and desires. If children really are essentialists, then, this might show up most clearly in a domain in which the essences have to do with these very mental states—the domain of artifacts.

The Structure of Concepts

If an essentialist perspective on concepts is correct, it would entail that many words correspond to concepts that do not exhaustively decompose into simpler notions. Although concepts might be associated with prototypes or sets of exemplars, they do not reduce to them.

I have focused here on common nouns, but the same considerations apply to the meanings of proper names. When we first meet someone and learn that he is named *Bob*, we form a representation of that person and associate it with that proper name. This representation might include Bob's appearance and the tone of his voice. But we nonetheless assume that the proper name corresponds to *that individual*, tracking him over space and time; we encode this other information as contingent facts about the individual. Because of this, it is possible for us to question whether Gödel was the man who did a certain mathematical proof (even if the only thing we know about Gödel was that he did this proof), whether Shakespeare wrote *Hamlet, Macbeth,* etc. (even if all we know about Shakespeare is that he wrote *Hamlet, Macbeth,* etc.), whether Moses really crossed the Red Sea, and so on (Kripke, 1980).

This is a plausible way for the human mind to have evolved: we would want the capacity to think about, recognize, and track people, not merely properties that are associated with people. Similarly, we would want the capacity to refer to members of categories that occur in the physical and social world (things like birds and chairs), to reason about them, to have desires that implicate them, and so on—but to also understand that things can be members of these kinds even if they appear different from typical members and that they might not be members of these kinds even if they appear similar to typical members.

Essentialism Lite

It is sometimes argued that humans have innate knowledge of space, time, causality, and number. A set of innate naive theories or stances includes intuitive physics, theory of mind, and naive biology. Long before they start to speak, babies think about people as intentional agents, about rabbits as biological entities, about chairs as intentionally created entities. The abstract understanding of the world that educated adult scientists possess is not a radical departure from this initial innate understanding; it is merely an extension of it.

This is one perspective on the structure of the mind (see Pinker, 1997), one I am sympathetic to. But suppose it is wrong, every word of it. Psychological essentialism could still be right. Imagine, following Quine, that humans start with nothing more than an animal similarity space: entities are thought about as similar and hence are thought more likely to fall into the same kind, depending on their relative locations in that space. To give rise to essentialism, all one needs to add is the notion that a distinction exists between, roughly, *looking* the same and *being* the same.

In particular, it is not the case that to have an essentialized concept you need to have a theory encompassing the domain that the category belongs to. And so the essentialism proposal does not entail the theory theory—the view that concepts are characterized by the role they play within naive theories. The theory theory works well for many concepts, particularly those that are part of bona fide scientific theories—concepts such as electricity and heat. Block (1986) notes that many such concepts are acquired not through ostensive naming but instead in the course of acquiring the relevant theories. And in some cases, a set of concepts might be learned all at once (force, mass, acceleration; gene, codon, chromosome), just by being connected to one another within a larger explanatory system.

But this approach does not naturally apply to all concepts. Laboratory studies have found that people can acquire concepts that correspond to arbitrary correlations of features (e.g., Nosofsky, 1988; Posner & Keele, 1968). While top-down expectations do affect the sorts of properties people attend to in these studies (Kelemen & Bloom, 1994), the *specific* concepts that they come to learn are defined entirely in terms of statistical information. And it is implausible that two-year-olds who know the meanings of the words *milk* and *juice* distinguish the two in terms of their explanatory roles within naive theories. More likely, they do so on the basis of color and taste.

Fodor (1998) makes a similar point using the example of water. He notes that we know the properties of water—that it is liquid, transparent, drinkable—and we know that these properties are correlated. But even most adults don't know *why* they are correlated. More generally, the notion that a concept is held together by an understanding of why it has the properties it does "set[s] the conditions for concept possession impossibly high" (p. 118).

The essentialist view doesn't have problems with such cases. It simply entails that children will have the implicit assumption that some deeper fact relates to the matter—that these properties of water are the result of some deeper essence. In such cases, children aren't like scientists who have theories; they are like scientists *before* they have theories, trying to make sense of some domain they know little about. More generally, before developing a theory as to why there is a correlation between certain features, you have to first notice that this correlation exists, and hence there is a symbiotic relationship between statistical learning and theory-based inference (Keil, 1994).

The studies reviewed in this chapter support the view that, early in development, children do possess intuitive theories of natural kinds and artifacts. A word that names a novel animal, such as an aardvark, will be extended in a different way than a word that names a novel artifact, such as a modem. But children also possess a more general intuition—one that can apply in the absence of intuitive theories. This is the intuition that concepts—and hence words—correspond to how things really are and not just how they appear to be.

Notes

1. This isn't to say that essentialism is always correct. Humans can *over*essentialize and infer causes when none exist. We are, at heart, conspiracy theorists. The most pernicious example of this is race. Human groups differ in properties such as skin color and facial morphology, and to a biologist these are literally skin deep—typically superficial adaptations to climate differences. But even young children tend to think of races as deep and immutable categories (Hirschfeld, 1996).
2. Underextension, in which a child has a narrower extension of the word than the adult does—refusing to call a pelican "a bird," for instance—could provide similar insights, but underextensions are harder to observe in spontaneous speech and harder to study experimentally.
3. Another study using a similar design was reported in Tomikawa and Dodd (1980), but those stimuli were described to the children as characters "who have various exciting adventures," and not as artifacts, and so the results from this study do not bear directly on the question of how children treat artifact names.
4. For some real-world examples of this, consider a glove (hand-shaped), a sheath (knife-shaped), and a violin case (violin-shaped). We do *not* call these "hands," "knives," or "violins," presumably because we know that there is another explanation for the sameness of shape, one that does not entail that they were intended to be—or represent—hands, knives, or violins.

BLANK PAGE

Chapter 7
Naming Representations

Modern humans live in a world of visual representations. William Ittelson (1996, p. 171) gives some examples of this:

> As I sit here at my breakfast table, my morning newspaper has printing on it; it has a graph telling me how the national budget will be spent, a map trying to tell me something about the weather; a table of baseball statistics, an engineering drawing with which I can build a garden chair, photographs of distant places and people, a caricature expressing what the editor thinks of a political figure, and an artist's rendition of what the city will look like 20 years from now. . . . On the wall in front of me hangs an abstract painting. Next to that, there is a calendar. Above the calendar is a clock. All this and more, and I haven't even turned on the TV or the computer.

Children live in the same world. While young ones do not read, they often look at picture books. They draw, paint, sculpt, and observe the artistic endeavors of others. In many homes, they watch cartoons, videos, documentaries, and other television programs. Much of what children—and adults—know about the world is the result of exposure to visual representations. In some instances, children's appreciation of pictures is parasitic on an understanding of the external world (as when a child looks at a photograph of his mother and knows who it represents), but often it is the other way around: the understanding of the representation comes first. Most children will see a picture of a gorilla before seeing an actual gorilla, and much of the mature understanding of everything from planets to popes comes not from experience with the actual entities but through experience with visual representations.

We talk about these representations using the same words we use for talking about real things. (The terminology I am using here is informal, since, of course, pictures are also real things.) We might point to a painting of a chair and say "Look at the chair" or to a drawing of a