Submitted

Burnett, R. & Medin, D.L. (2008). Reasoning across cultures. In L. Rips & J. Adler (Eds.), *Reasoning: Studies of human inference and its foundations* (pp. 934-955). Cambridge, MA: Cambridge University Press.

# Reasoning Across Cultures

Russell C. Burnett and Douglas L. Medin

7 June 2006

Please send correspondence to Russ Burnett at rburnett@mscc.huji.ac.il.

# Reasoning Across Cultures

In 1931 A. R. Luria traveled to rural Uzbekistan with a question: Is thinking influenced by its social and cultural environment? The fieldwork he did there (Luria, 1976) was likely the first attempt to answer this question using methods of experimental psychology. Rural folk in Uzbekistan were at that time in the midst of a socioeconomic transition that involved, among other things, the collectivization of agriculture and the growth of schools and literacy. Luria saw in this a natural experiment. He found groups of participants with different levels of involvement in this transition—including different levels of schooling and literacy—and gave them a variety of cognitive tasks. Some of these were designed to elicit deductive inferences from premises like "precious metals don't rust" and "gold is a precious metal." Luria found that subjects with no formal schooling often balked at such problems, rejecting the premises and saying, for example, that "one can speak only of what one has seen." In contrast, subjects who had been to school were more likely to use the hypothetical premises to draw conclusions (e.g., "gold doesn't rust") with no obvious basis in personal experience. Luria also exercised some experimental control, manipulating whether the content of an argument was familiar or unfamiliar to participants. Participants who had not been to school treated the two kinds of content differently, drawing conclusions from premises more often when the content was familiar than when it was unfamiliar. Schooled participants tended to reason from the premises regardless of content.

We revisit Luria's work because it illustrates two challenges in comparative cultural research. First, there is the natural confounding of numerous factors that might constitute "culture," and various other factors besides. Luria favored the theory that literacy enables "verbal-logical" reasoning, but his groups differed in countless other ways, including general schooling, "practical activities," "modes of communication," "cultural outlooks," "access to a technological culture," "social relations," and "life principles" (Luria, 1976, p. 15). Which of these, or what other factors, caused the groups to respond differently? Second, even if the practical problem of confounding could be solved (and nature may sometimes disentangle such factors for us), there is a theoretical problem: Which factors should be understood as constituting culture? Or should culture be thought of as an irreducible construct, something that would remain even after one has controlled for "practical activities," "social relations," and so on? Any empirical demonstration of cultural differences presupposes some definition, however vague, of culture, but without a suitably specific theory of culture comparative research has no basis for causal analysis and is capable of little more than cataloging phenomena or disproving their universality.

In evidence of these difficulties, there has been a good deal of disagreement over which social or cultural factors were responsible for the group difference Luria (1976) observed. Cole and Scribner (1974) suggested that this difference was due not to literacy but to schooling or involvement in "complex acts of social planning." This interpretation found some support in work by Scribner and Cole (1981), who capitalized on a partial disentangling of literacy and schooling among the Vai, a people of Liberia who employed a system of writing apart from formal education. Yet Scribner and Cole's own findings have been interpreted as due to literacy, where literacy is defined in functional context (e.g., literacy for formal education; Greenfield, 1983). Distinguishing between these theories would require finding a case of formal schooling without literacy, which seems unlikely.

The cognitive causes of the group difference are also unclear. Cole and Scribner (1974) argued against a qualitative difference in methods of reasoning: "There is no evidence for different *kinds* of reasoning processes such as the old classic theories alleged—we have no evidence for a ‘primitive’ logic” (p. 170). Scribner (1977) proposed that schooling promotes a shift from an “empirical bias,” or a bias to draw on personal experience, to a more “theoretical” approach to reasoning that allows greater use of hypothetical premises. More recently, Dias, Roazzi, and Harris (2005) found that whereas both schooled and unschooled participants could be prompted to reason from unfamiliar premises (by a suggestion that the premises described a distant planet), there was a persistent gap between schooled and unschooled participants, as in Luria’s studies. This suggests that reasoning from novel or hypothetical information does not require schooling, but also that schooling (or one of its correlates) does promote something like a stable orientation or stance that facilitates such reasoning (see also Harris, 2000).

In this chapter we will review recent research on reasoning and culture that differs from Luria’s (1976) work in an important respect. Whereas Luria’s probes were designed to elicit abstract reasoning, guided by content-free principles like class inclusion and entailment, the work we will describe focuses on reasoning that makes greater use of a knowledge base. We will suggest that culture-related variations in knowledge are an important source of differences in reasoning. Of course, differences attributable to the mere presence or absence of relevant knowledge might be fairly uninteresting from a theoretical point of view. Our story will be more interesting. We will suggest that knowledge is often organized according to culture-related framework theories or expectations and that graded differences in the organization or accessibility of knowledge are reflected in reasoning.

That culture influences reasoning by way of such things as framework theories and expectations is consistent with the epidemiological view of culture, or the idea that culture can be understood as socially distributed mental representations, as well as expressions of these representations and behaviors associated with them in given ecological contexts (Atran, Medin

& Ross, 2005; Sperber, 1996). The epidemiological approach leads naturally to treating within-culture variation not as noise, but rather as important information that may be used to identify paths of cultural transmission and relevant correlates of within- and between-culture differences (Medin, Ross, & Cox, in press). To be sure, we are far from a full understanding of the sources of the kinds of representations we will describe in this chapter and the modes by which they are transmitted—we will return to this topic later. Nonetheless, the work reviewed in this chapter benefits from a theory of culture that is specific enough to move us beyond a simple catalog of group differences in reasoning, to hypotheses about proximal causes of these differences in individual minds.

We begin with a distinction that is slippery but still useful. Reasoning can be a tool for understanding cultural differences, or cultural comparisons can be a tool for understanding reasoning. As an example of the former, Choi, Nisbett, and Smith (1997) investigated the suggestion that Westerners (undergraduates at a university in the United States) are more likely to encode examples into categories than Easterners (undergraduates at a Korean university) by giving participants inductive reasoning tests where they might spontaneously generate categories, according to the similarity-coverage model of Osherson, Smith, Wilkie, López, and Shafir (1990). We will describe this model in a moment, but for now the point is that the reasoning task was used as a tool to make observations about a cultural difference in propensity for categorizing. As an example of the latter, López, Atran, Coley, Medin, and Smith (1997) gave Itza' Maya agro-foresters and University of Michigan undergraduates the same sorts of reasoning probes as a test of the generality of the similarity-coverage model. They failed to find evidence for one of the reasoning phenomena predicted by the model, and in this way the cross-cultural comparison revealed something about the limits of the model.

Of course, no strict principle distinguishes these two types of studies. Choi et al. (1997) were also, implicitly or explicitly, testing the generality of the similarity-coverage model, and a different pattern of results could have led to changes in the model. Furthermore, as we have said, an observed cultural difference tells us little about reasoning if we do not understand the source of the difference. Nonetheless, the distinction is important for present purposes, because our focus will be on what comparative research tells us about theories of reasoning and not vice versa.

We will focus on two kinds of reasoning. The first is inductive reasoning about categories and their properties (what is often called category-based induction), especially in the biological domain. Cultural research has shown the importance of framework theories and the organization of knowledge to this kind of reasoning. The second is causal reasoning, where interesting cross-cultural research is being done and where, at the same time, a promising new body of theory has been adopted by cognitive scientists. In this case, the cross-cultural findings and their implications are less clear, but we can begin to see how they might inform the new theory. In trading breadth for depth, we will not discuss some other kinds of reasoning that have been topics of recent cross-cultural research. For recent reviews of other research on culture and cognition, see Cohen (2001); Medin and Atran (2004); Medin, Unsworth, and Hirschfeld (in press); Nisbett and Norenzayan (2002); and Norenzayan and Heine (2005).

## Inductive Generalization of Properties over Categories

One form of reasoning studied extensively in cultural research involves the inductive generalization of properties over objects or, more typically, over categories of objects. A fisherman who learns that brown trout are affected by a certain disease might infer, with some degree of confidence, that rainbow trout, or all trout, or catfish, or fish in rivers, or even all fish are affected by this disease. Each of these inferences can be thought of as an argument in which the premise doesn't guarantee the truth of the conclusion but provides some support for it. From this perspective, understanding knowledge generalization involves understanding how reasoners judge the support that a premise like "brown trout have a certain disease" lends to a conclusion like "rainbow trout have this disease."

*1. Reasoning from Similarity and Taxonomic Relationships*

Rips (1975) gave undergraduates at Stanford University premises like "all of the robins on an island have a certain disease" and asked them to judge what proportion of, say, the geese on the island are affected by the disease. It was found that judgments were well explained as a function of two constructs. The first was the similarity between the premise category and the conclusion category (the similarity between robins and geese); all else equal, diseases associated with robins generalized to sparrows more strongly than to geese. The second construct was the typicality of the premise category with respect to a salient category that included both the premise and the conclusion (the typicality of robins with respect to the bird category); for generalizing to other birds, robins were a better premise category than were geese, all else equal. As Rips noted, participants reasoned as if they expected the novel property (the disease) to be distributed over categories in a way that mirrored the distributions of known properties. That robins and sparrows are similar is a consequence (or restatement) of the fact that they share many known properties, and because they share many known properties they are likely to share a novel one, too.

These ideas were elaborated by Osherson et al. (1990) in their *similarity-coverage model*. Consider the following inferences, each of which involves projecting a novel property from a specific category to a more general, inclusive category.

Sparrows have a certain enzyme; therefore, all birds have it. (1)

Penguins have a certain enzyme; therefore, all birds have it. (2)

Osherson et al. and many others have found a *typicality effect* in inferences like these. Sparrows are deemed more typical of birds than are penguins, and (1) is deemed a stronger inference than (2). The similarity-coverage model explains this effect and others as due to (a) similarities among categories and (b) the degree to which small categories "cover" inclusive categories. (1) is better than (2) because sparrows provide better coverage of the bird category, in the sense that they have higher average similarity to other birds. The model also predicts the *similarity effect* found by Rips (1975), in which similar kinds are better bases for inferences about each other than are dissimilar kinds, and a *diversity effect*, in which multiple kinds jointly form a better basis for inferences about an inclusive category if they are more dissimilar and therefore more diverse (less redundant) in their coverage. For example:

Robins and sparrows have a certain enzyme; therefore, all birds have it. (3)

Robins and hawks have a certain enzyme; therefore, all birds have it. (4)

According to the similarity-coverage model, (4) is better than (3) because robins and hawks are more diverse in their coverage of the bird category than are robins and sparrows.

The similarity-coverage model, like Rips's (1975) account, assumes two kinds of knowledge: (a) knowledge of similarity relationships among categories and (b) knowledge of class-inclusion relationships—for example, knowledge that bird is the relevant category that includes robins and geese, so that robin's coverage of the bird category (or the average similarity of robins to other birds, or the typicality of robins with respect to the bird category) can be assessed. Another well-known model, Sloman's (1993) *feature-based model*, does away with the latter component and works just on the similarities between premises and conclusions. This buys some flexibility, as the model applies even in domains that lack clear class-inclusion relationships, but the driving principle is the same: A novel property is likely to cluster with known properties, so categories are likely to share a novel property to the extent that they are similar in their known properties. This model makes many of the same predictions as the similarity-coverage model, including typicality, similarity, and diversity effects.

Insofar as these models work on the general assumption that novel properties mirror known properties, they are *similarity-based*.[1] Studies of cognitive psychology's standard participants (undergraduates at research universities) have found that similarity-based reasoning—including the similarity, typicality, and diversity effects—is quite robust (López et al., 1997; Osherson et al., 1990; Rips, 1975; Sloman, 1993; and many others).

Studies of culture and expertise have revealed a very different picture. First, reasoners with even moderate knowledge of the domain often employ other strategies; only research subjects with impoverished knowledge consistently reason in line with the similarity-based models. Second, even among reasoners with equal knowledge of the domain, cultural differences in the organization of this knowledge may influence the strategies that are used. Third, even when reasoning is similarity based, culture-related beliefs or theories may influence how broadly properties are generalized from different premise categories. In the sections that follow we first

---

[1] Heit (1998, 2000) has proposed a Bayesian model of inductive reasoning based on this same idea. The Bayesian model, however, is somewhat more flexible in its application of the idea that novel properties mirror known properties—as we will discuss later.

describe these findings in some detail and then explore their implications for models of reasoning.

*2. A Focus on Biology*

Before proceeding, we should explain that most studies of this sort of reasoning have focused on biology. Participants are usually asked to generalize properties like a disease, an enzyme, or “sesamoid bones” over kinds of animals. Biology is an especially informative domain for several reasons. First, it is a domain in which all cultures have some knowledge, and so it constitutes a common ground for cultural comparisons. Second, as a domain, biology has an abstract structure that all cultures seem to recognize (Atran, 1998; Berlin, 1992; Medin & Atran, 2004). This abstract structure involves clusters of properties that travel together (for example, things that have wings tend to fly, to build nests, and so on) and are distributed systematically through downward-branching taxonomies in which nonoverlapping categories are nested under higher-level categories. Importantly, and not accidentally, this structure is just what is required by the similarity-coverage model. Clusters of co-occurring properties tend to provide clear similarity relationships among categories, and taxonomies yield clear class-inclusion relationships. Note also that the mere fact of property clustering fits the very principle behind similarity-based reasoning: All else equal, a novel property is likely to be distributed along with known properties. Since what the feature-based model requires is just a subset of what is required by the similarity-coverage model (namely, similarity relationships among categories), the abstract structure of the biological domain suits it, as well. In short, if similarity-based models apply anywhere, they should apply in biology. Moreover, since the understood structure of biology is largely invariant across cultures, these models tend to predict invariance in reasoning.

Third, however, biological entities are often the objects of people's goals, theories, beliefs, and practices. If reasoning is sensitive to such factors, then we might expect differences in these factors to be reflected in reasoning. Fourth, biology is an information-rich domain where a lot is hidden even from experts, and so there is room for biases or framework theories to guide the interpretation of experience (Keil, 1995; Keil, Levin, Richman, & Gutheil, 1999), highlight certain types of information over others, and promote certain inferences over others. In short, people's relationships with plants and animals often involve many of the factors that one might take to constitute or be related to culture. Focusing on biology allows us to test hypotheses about whether and how such factors influence reasoning.

To begin to see how the complexity of biology might play into reasoning, consider that not all of the properties that one might want to reason about participate in the abstract similarity-based structure of the domain, and that similarity-based reasoning might therefore work better for some properties than for others (cf. Heit & Rubinstein, 1994). If one interprets the abstract domain structure as a consequence of progressive speciation via natural selection, then similarity-based reasoning should tend to work for properties that "load" heavily on genes. Other properties are likely to be distributed in ways that are less closely related, or perhaps completely unrelated, to the similarity-and-taxonomy structure of the domain. Take, for example, diseases, which have been used as novel properties in many studies. Some diseases may have little basis in genes or innate potential. Whether an organism or a population is affected by such a disease may have mostly to do with whether it is exposed to a pathogen, and intrinsic susceptibility to the pathogen may sometimes be relatively independent of genes, at least in the population of organisms under consideration. (Of course, many other diseases are species specific.) Other properties might have some basis in innate potential, but it might also be clear that these properties don't travel with other large bundles of features. For example, flying-related properties in bats do not generalize outward through taxonomic space or its associated similarity space very well. Simply put, there is reason to expect that similarity-based models work only for certain properties. Indeed, Osherson et al. (1990) emphasized that the similarity-coverage model is meant to apply to "blank" properties, or properties for which the reasoner has little prior knowledge about their distribution. In practice, experimenters have tended to use properties, like "sesamoid bones," that have a "biological flavor" (Rips, 2001) and more than a hint of genes or innate potential.

*3. Reasoning from Causal-Ecological Knowledge*

The first finding from studies of culture and expertise is that even in biology, with its structure involving similarity and taxonomic relationships, reasoners often prefer inductive strategies that have little to do with this structure—that is, strategies that are not similarity based at all. Knowledgeable reasoners often prefer to project properties like diseases and enzymes on the basis of specific causal mechanisms by which these properties might have arisen in both the premise and the conclusion, and these causal mechanisms often involve ecological interactions between members of the different categories. For example, in judging which of two kinds of trees was more likely to share a disease with all other trees, three different groups of Chicago-area tree experts preferred the tree with wider geographic distribution or greater intrinsic susceptibility to disease (Proffitt, Coley, & Medin, 2000). In the former case, the idea is that trees with wider geographic distribution have greater potential to pass the disease to other species. In the latter case, the rationale is that the disease will spread more easily among trees of the susceptible species, which renders the disease widely distributed and more likely to spread to other species. These experts tended not to invoke similarity or taxonomic strategies.

Similarly, in reasoning about diseases and enzymes (or "little things inside") in birds, both North American birdwatchers and Itza' Maya farmers in Guatemala tended to base their inferences on causal-ecological interactions, often focusing on geographic distribution (Bailenson, Shum, Atran, Medin, & Coley, 2002). They tended to prefer, as premises, birds that were rich in known ecological associations with other birds. Similarly, in a study that involved reasoning about diseases among mammals, Itza' Maya tended to focus on geographic range and ecological diversity (López et al., 1997).

In a recent study, conducted in collaboration with Norbert Ross, we asked fishermen of two cultural groups in northern Wisconsin to reason about diseases and enzymes in fish. Most of the probes lent themselves to reasoning by typicality, similarity, or diversity. For example, in one item river shiners were said to have one enzyme (or disease), and sunfish another. Participants were asked which enzyme was more likely to be found in (or which disease was more likely to affect) smallmouth bass. Sunfish are more similar to smallmouth bass than are river shiners, and in taxonomies reported by members of the same populations (including many of the same participants) smallmouth bass tended to be closer to sunfish than to river shiners. Similarity-based reasoning thus predicts that sunfish are the better premise. Nonetheless, a great majority of participants chose river shiners and explained this choice by saying either that smallmouth bass eat river shiners or that smallmouth bass and river shiners are found in the same waters. More generally, over 20 items, participants tended to focus on ecological interactions or associations through which the novel property might be transmitted among fish. In many cases participants reasoned upward through the foodchain from premise to conclusion; that is, they chose the premise fish that was more likely to transmit the property to the conclusion fish by being eaten (for a similar result, see Shafto & Coley, 2003). But foodchain knowledge was used in other ways, too. Participants sometimes reasoned downward through the foodchain, choosing the premise that was more likely to have “caught” the property by eating the conclusion fish. In still other cases, participants chose the premise that shared with the conclusion a common food source, the idea being that both fish might have gotten the property from this source. Overall, only 9% of inferences were based on similarity and/or taxonomic relationships. Fully 90% were based on transmission of the property through ecological interactions (Burnett, Medin, & Ross, 2004).

These various groups of participants differ in many ways, of course, and we will turn to cultural differences in a moment. For now, a reasonable generalization is that participants with normal levels of experience with plants and animals often invoke causal knowledge and reason about the mechanisms by which properties might come to be distributed in different ways across categories. For properties like diseases and enzymes (even enzymes), these causal mechanisms often involve ecological interactions that are more or less unrelated to the similarity-and-taxonomy structure of the domain.

*4. Flexibility in Reasoning*

The above findings on causal-ecological reasoning notwithstanding, knowledgeable reasoners almost surely prefer similarity-based strategies for properties that (they believe) participate in the similarity-based structure of the domain. What knowledge provides is flexibility, in the form of a variety of strategies that allow the reasoner to project different properties in different ways. In evidence of this flexibility, Shafto and Coley (2003) found that whereas fishermen projected diseases according to food chain relations among marine animals, they projected more abstract or ambiguous properties like “a property called sarca” among the same animals according to similarity or taxonomic relatedness, as did domain novices.

Even novices show some flexibility when stimuli tap their knowledge. Using ecological contrasts that even undergraduates often know (e.g., jungle creatures versus desert creatures), Shafto, Coley, and Baldwin (2005) found that, against a background preference for reasoning according to taxonomic relatedness, undergraduates showed a tendency to distinguish between properties, such that participants with greater knowledge of the ecological groups were more likely to project diseases and toxins, but not abstract properties like "a property called sarca," among ecologically related animals. (For other examples of relative novices basing inferences on causal considerations instead of, or in addition to, similarity, see Gelman & Markman, 1986; Hadjichristidis, Sloman, Stevenson, & Over, 2004; Heit & Rubinstein, 1994; and Ross & Murphy, 1999. For reviews see Heit, 2000, and Rips, 2001.)

An interesting question is how reasoning by causal mechanisms or ecological interactions relates to similarity-based reasoning. On one hand, there are reasons to think of similarity-based reasoning as simpler or more basic. Similarity and taxonomic relationships are often available even to extreme novices. They seem to be at the core of folk biological knowledge, robust even under devolution in knowledge (Medin & Atran, 2004). Furthermore, reasoners who employ ecological knowledge under normal conditions may abandon this knowledge and fall back on similarity under time pressure (Shafto et al., 2005; see also Coley, Shafto, Stepanova, & Baraff, 2005).

On the other hand, causal reasoning and similarity-based reasoning are not mutually exclusive, and reasoning based on similarity and taxonomic relatedness may itself involve or interact with causal considerations. Hadjichristidis et al. (2004) found that similarity had a greater influence on inference when the property in question was more causally central to the category in which it appeared (i.e., when more of the category's other properties depended on it).[2] This suggests that similarity-based reasoning is invoked to the degree that causal considerations support it—that is, to the degree that the property in question is involved in the causal mechanisms that give rise to the similarity-and-taxonomy structure of the domain in the first place.

Or consider a related task: inferring whether an individual category member has some property that is associated with the category. Rehder and Burnett (2005) found that inferences were stronger when the individual was known to have other category-associated features. Importantly, results suggested that this was not because the known features made the individual a better or more typical member of the category, as a similarity-based account might suggest, but because they indicated that the individual possessed some underlying cause or causes of the category's other features. Though known features did tend to boost the individual's typicality, their importance for reasoning lay not in this, but in their relevance to causal considerations.

Of course, this point is meaningful only if people do actually have beliefs about causal structure even in cases where they appear to invoke similarity-based strategies. Other work has suggested that they often do. Even young children and domain novices have intuitions about what properties, or what kinds of properties, of biological organisms are due to intrinsic causes or innate potential (Keil, Smith, Simons, & Levin, 1998; Medin & Atran, 2004). As for the finding that reasoners fall back on similarity under time pressure, this effect has been most clearly demonstrated in relative novices (Shafto et al., 2005), and it is possible that more knowledgeable reasoners are often as fluent with causal strategies as with similarity.

---

[2] Hadjichristidis et al. distinguish between causal structure and more generic "dependency structure" and draw their conclusions about the latter, but for simplicity we focus here on causal structure.

In general and in short, it is difficult to distinguish (a) true similarity-based reasoning from (b) causal reasoning based on causes that participate in the similarity-and-taxonomy structure of the domain. For example, consider two interpretations of Shafto and Coley's (2003) finding that fishermen reason (as if) from similarity or taxonomic relatedness for properties like "sarca" but from foodchain relationships for diseases. One possibility is that, given an ambiguous or abstract property, these experts fell back on similarity-based reasoning as a default strategy. The other is that these experts assumed that ambiguous properties were causally related to innate potential—for example, properties grounded in genes—and projected them accordingly. As Rips (2001) observed, studies showing that relative novices override similarity in cases where causal considerations provide better bases for reasoning (e.g., Gelman & Markman, 1986; Heit & Rubinstein, 1994; Ross & Murphy, 1999) suggest that, even in novices, inferences consistent with similarity may conceal an influence of causal considerations.

*5. Cultural Factors*

We have explained the first of three findings: Knowledgeable reasoners use similarity and taxonomic relationships only some of the time; even for properties like enzymes, which are plausibly related to the similarity-and-taxonomy structure of the domain, knowledgeable reasoners often invoke causal-ecological strategies instead. In itself this finding concerns the importance of the reasoner's knowledge base, but it also sets the stage for culture. We turn now to two ways in which cultural factors might influence reasoning.

*A. Differences in Organization or Accessibility of Knowledge*

Knowledgeable reasoners often have many bases for generalization available to them. Causal-ecological reasoning is not one strategy but rather a potentially large set of strategies that draw on various causal mechanisms and ecological interactions known by the reasoner. In generalizing diseases, reasoners with knowledge of trees sometimes prefer premise categories with great geographic range or ecological diversity, the idea being that widely or diversely distributed trees will have greater opportunity to transmit the disease to other trees. Sometimes they prefer premises with greater intrinsic susceptibility to disease; here the rationale is that the disease will spread more easily among trees of the susceptible kind, which renders the disease widely distributed and more likely to spread to other kinds. Sometimes they seem to go in the opposite direction, preferring a premise with greater intrinsic resistance to diseases, the idea being that the disease itself must be highly infectious if it managed to spread to an intrinsically resistant kind of tree. Of course, such reasoners also have similarity-based strategies available to them.

Cultural beliefs, values, goals, and attitudes might influence the contents of a person's knowledge base by constraining the practices that he or she undertakes with respect to plants and animals. If a group abhorred trout and avoided catching or touching them, then members of this group would likely know less about trout than if they regarded trout as ideal or desirable. Here we focus on a somewhat more interesting possibility: that even among reasoners with the same knowledge, cultural factors may influence how this knowledge is organized, or the relative accessibility of different pieces of the knowledge base (e.g., Hong, Morris, Chiu, & Martinez, 2000; Medin, Ross, Atran, Burnett, & Blok, 2002; Medin, Ross, Atran, Cox, Coley, Proffitt, & Blok, in press). Differences in organization or accessibility may, in turn, be reflected in how reasoning strategies are derived from the knowledge base. That is, a single knowledge base may be organized so as to make different reasoning strategies more or less fluent (see also Higgins, 1996). In what follows we will first describe recent studies of cultural differences in the organization or accessibility of folkbiological knowledge and then consider how these differences play into reasoning.

One way of assessing organization of knowledge is to ask participants to sort biological kinds into hierarchies that have the structure characteristic of folkbiological taxonomies. The participant is presented with a set of cards. On each card is printed the name of a kind (e.g., “rainbow trout”), and the participant is asked to sort these kinds into groups “that go together by nature.” This instruction is also accompanied with telling informants that they should create groups that make sense to them. After this initial sort, the participant is given opportunities to join these groups into progressively more inclusive categories, and then to split the groups into progressively smaller, more specific categories. The result is a downward-branching hierarchy of biological kinds.

This task has revealed differences among various groups of participants in various subdomains of biology (e.g., Bailenson et al., 2002; López et al., 1997; Medin, Lynch, Coley, & Atran, 1997; Proffitt et al., 2000; Shafto & Coley, 2003). One recent study is especially interesting in that it has shown cultural differences in the organization of knowledge that cannot be reduced to differences in raw experience or practices, but rather seem to arise from differences in what might be called framework theories (Medin, Ross, Atran, et al., in press). This case involves fishermen of two groups in northern Wisconsin: Native American Menominee Indians and a nearby majority-culture (European-American) community. These groups allow a close comparison because they are similar in many ways. They fish in similar waters and are familiar with the same fish. Though they do differ slightly in some of their practices—for example, Menominee put somewhat more emphasis on fishing for food (versus for sport) and are less likely to practice catch and release—these differences in themselves have small, if any, consequences for knowledge acquired through experience.

Whereas these groups showed substantial overall agreement in their sortings of local fish, they also showed some reliable differences. Multidimensional scaling revealed that Menominee sorts, but not majority-culture sorts, tended to express a dimension related to habitat or ecological niche. This difference was also evident in participants' explanations of their sorts. Majority-culture participants gave many taxonomic or morphological justifications like "bass" and few ecological justifications like "lake fish" or "fish you find in cool, fast-moving water." In contrast, Menominee participants gave fewer taxonomic or morphological justifications and more ecological ones. In short, Menominee participants seem to organize their knowledge somewhat more around ecological considerations, whereas majority-culture participants seem to focus more on taxonomic and morphological characteristics of fish. This was also reflected in a "species interactions" task, where participants were asked to say how various kinds of fish affect one another. In this task, Menominee participants reported more causal interactions among kinds of fish than did majority-culture participants (Medin, Ross, Atran, et al., in press).

Just as these differences in organization cannot be explained by group differences in practices, neither can they be explained by differences in mere possession of knowledge. When members of the same groups were asked to sort fish according to ecological relatedness, there were no significant group differences. Likewise, when the stimuli used in the "species interactions" task were pared down so that participants spent more time thinking about each response, group differences disappeared (Medin, Ross, Atran, et al., in press). In short, it is not that Menominee participants know more ecological relations; rather, ecological relations seem to play a greater role in organizing their knowledge of fish and are more accessible. The differential importance of ecological relations in organizing knowledge of biological kinds has also been seen in studies with other cultural groups (e.g., Atran et al., 1999, 2002; López et al., 1997).

A significant challenge is to understand how these differences arise. Ross, Medin, Coley, and Atran (2003) reported parallel differences between young rural Euro-American children and young Menominee children—Menominee children were more likely to give ecological justifications on a reasoning task—and so the difference in emphasis on ecological relationships seems to be present early. One possibility is that the mediating factor is cultural differences in skeletal principles or framework theories. Several Menominee participants commented that "every fish has a role to play," and in the "species interaction" task several Menominee participants made explicit mention of the idea that, in general, any two fish are likely to affect each other somehow. In interviews, majority-culture parents often said that they wanted their children to learn to take care of nature, whereas Menominee parents said they wanted their children to see themselves as part of nature (Bang, Medin, Unsworth, & Townsend, 2005). Such ideas might function like framework theories to guide the interpretation of experience, but just how this happens remains a challenging question.

Are differences in organization or accessibility of knowledge reflected in reasoning? One fairly direct way of addressing this question is to measure reasoning and organization of knowledge separately and compare the results. López et al. (1997) asked Itza' Maya farmers in Guatemala and undergraduates at the University of Michigan to both sort and reason about mammals. Not surprisingly, the groups differed in their sorts. Itza' tended to draw finer distinctions among mammals based on more detailed knowledge of morphology, behavior, and ecological associations. Some of the reasoning probes tapped specific differences in the two groups' sorts, and in these cases the groups diverged in reasoning in ways that mirrored their differences in sorting. For example, in their sorts the undergraduates tended to group foxes with dogs, whereas the Itza' grouped foxes with cats. Both groups were given a forced-choice reasoning item in which foxes were said to have some disease, and the task was to generalize this disease to either dogs or cats. Consistent with their respective sorts, undergraduates preferred dogs, and Itza' preferred cats. In short, each group presumably used its relevant knowledge, and the knowledge differences led to group differences in induction. Also, as we have seen, there is some evidence of a relationship between reasoning and organization of knowledge among Menominee and majority-culture participants in Wisconsin, though in this case the relevant pieces of evidence come from different age groups. In sorting, Menominee adults tend to emphasize ecological relationships more than do majority-culture adults. In reasoning, Menominee children invoke ecological mechanisms more than do majority-culture children.

The previous method is correlational, and it would be nice to show the influence of accessibility on reasoning by controlled experimentation. Accessibility itself is difficult to manipulate (it is unclear how to prime, say, ecological relations without introducing a task demand to use these relations in reasoning), but an indirect way to get at accessibility is to manipulate the amount of time the reasoner has to access knowledge. Shafto et al. (2005) observed that undergraduate participants knew certain ecological associations (e.g., jungle animals, desert animals) but, when forced to choose between ecological and taxonomic relations, tended to generalize novel properties according to taxonomic relatedness. (This was true even for diseases.) To test whether this was due to poor accessibility of ecological relations, they ran speeded and unspeeded versions of another task, in which undergraduates did show evidence of reasoning from the ecological relations. In this task a novel property was attributed to a premise category (e.g., tigers) and the participant rated the likelihood that the property was also present in a conclusion category (e.g., anacondas). In the unspeeded version of the task, ratings were highest for taxonomically related animals, intermediate for ecologically related animals, and lowest for animals that had no close taxonomic or ecological relation. Under time pressure, however, ratings for ecologically related animals dropped to roughly equal those given to unrelated animals. Ratings given to taxonomically related animals were unaffected. This is consistent with the idea that graded differences in accessibility are reflected in reasoning; present but relatively inaccessible knowledge is sometimes invoked in reasoning but in a way that is sensitive to processing costs (Shafto et al., 2005; see also Coley et al., 2005).

In this case the relevant distinction was between similarity-based reasoning and causal-ecological reasoning, but causal-ecological reasoning consists of a potentially large set of strategies that draw on various causal mechanisms and ecological interactions known by the reasoner. Do different kinds of causal-ecological knowledge vary in accessibility, and are these variations reflected in reasoning? Bailenson et al. (2002) found that in reasoning U.S. experts relied on geographic distribution, whereas Itza' used specific causal-ecological interactions as well as geographic distribution. It remains to be seen whether this difference reflects knowledge accessibility.

*B. Asymmetries and Other Differences in Breadth of Generalization*

Similarity-based reasoning itself is subject to variability in how broadly one generalizes from a given premise category. Similarity and taxonomic relationships specify ordinal relationships among categories—a property should generalize outward like a ripple through similarity space or through the taxonomy—but these ordinal relationships do not in themselves say how far the ripple should travel.

Breadth of generalization has been a focus of developmental research since Carey (1985) argued that young children's understanding of biology is organized around humans. She gave children a reasoning task in which a novel property was said to be true of one biological kind (e.g., "Humans have a little green thing inside them called an omentum") and the child was asked whether that property was true of other biological kinds (e.g., "Do you think that dogs also have an omentum?"). Her participants tended to treat humans as a privileged base for inferences. In general, inferences from humans to nonhuman biological kinds were stronger than (a) from these same nonhuman kinds to humans and (b) from nonhuman kinds to other nonhuman kinds. In some cases this meant that children violated similarity; inferences from humans to bugs were stronger than from bees to bugs. In short, reasoning was anthropocentric.

One interpretation is that Carey's (1985) participants—mostly urban children—knew more about humans than about other kinds, and that better known or more richly represented categories are better premises for generalization. In support of this interpretation, Inagaki (1990) compared two groups of urban children, one group who had raised goldfish at home and another group who had not. Children who had raised goldfish were more likely to draw inferences from both goldfish and humans than were children who had not raised goldfish; the latter group reasoned more like Carey's participants, treating humans as a uniquely privileged premise category. Atran et al. (2001) studied Yukatek Maya children and adults in southern Mexico and found no evidence of systematic anthropocentrism. Instead they found a pattern of age- and gender-related differences that were consistent with familiarity effects. For example, girls knew less about the peccary than did boys and also treated the peccary as a weaker premise category. In related work, Tarlowski (in press) has demonstrated that children's patterns of inductive generalization are influenced both by urban versus rural status and by parents' knowledge of the domain.

The finding that better-known kinds make for better premises is problematic but not fatal for the similarity-based models; they can accommodate the finding but have trouble explaining it. For example, the similarity-coverage model can handle the finding by (a) putting much greater weight on the coverage component (e.g., the degree to which 'human' covers 'living thing') than on the similarity component (e.g., the similarity of 'human' to 'bug') and (b) assuming that the representations of the various categories are such that the privileged premise category ('human') covers the relevant inclusive category ('living thing') better than do other, nonprivileged categories. Still, the model provides no reason why coverage should be weighted so much more heavily than similarity (Rips, 2001).

A bigger problem comes from recent cross-cultural work which suggests that richness of knowledge of the premise category is not all that matters. Ross et al. (2003; see also Medin & Atran, 2004) studied three groups of children: Menominee children in rural Wisconsin, majority-culture children in rural Wisconsin, and majority-culture children in (urban) Boston. Rural children of both cultural groups have similar levels of experience with animals and plants, and so if amount of knowledge were the only strong determinant of breadth of generalization, then we would expect these two groups to be similar. However, Ross et al. found different developmental trajectories in all three groups.

The rural majority-culture children treated humans as a privileged base at an early age; this anthropocentrism waned with age. When they declined to generalize from nonhuman animals to humans, children in this group (at all ages) often explained that "people aren't animals." In contrast, Menominee children showed no reliable anthropocentrism at any age.

Also, they showed less differentiation between "higher" animals and "lower" animals when generalizing from humans, as if humans are intimately related to all other animals.[3] In interpreting this finding, Medin and Atran (2004) note that "the Menominee origin myth has people coming from the bear, and even the youngest children are familiar with the animal-based clan system. In short, there is cultural support for a symmetrical relation between humans and other animals" (p. 967).

Interpreting asymmetries is difficult. Medin and Waxman (in press) propose that they often reflect not just richness of knowledge of the premise kind but also the distinctive features of both premise and conclusion kinds and also the higher-level categories that these kinds belong to. In many cases the properties and higher-level categories of the *conclusion* kind seem to matter more than properties and categories of the *premise* kind. Thus, when children fail to generalize a property of peccaries to a target like humans, this may be not because peccaries are unfamiliar or atypical but rather because humans have distinctive properties that limit generalization to them. In other cases, asymmetries may be due to ecological reasoning, where some kinds are seen as more active ecological agents than others. Until we have a better understanding of the cognitive mechanisms that underlie asymmetries, interpretations should be made with caution. For now, a reasonable conclusion is that at least some asymmetries are due to (culture-related) knowledge of the properties of the relevant kinds and also to (culture-related) tendencies to think of certain kinds as belonging to certain higher-level categories (e.g., a

---

[3] Another finding of Ross et al. (2003) and Atran et al. (2001) was that some groups of children—especially Yukatek Maya and Menominee children—showed causal-ecological reasoning even at early ages. Since our focus here is on breadth of generalization by similarity and/or taxonomic relationships, we have disregarded causal-ecological generalizations in our description of the data.

tendency to think of humans as a kind of animal or as a kind distinct from other animals; Anggoro, Waxman, & Medin, 2005). Such tendencies are problematic for the similarity-based models.

*6. Implications for Models of Reasoning*

To summarize, cross-cultural and other studies have revealed the following.

(1) Reasoners with domain knowledge flexibly invoke a variety of strategies—including, but not limited to, the similarity-based strategies described by the similarity-coverage model (Osherson et al., 1990) and the feature-based model (Sloman, 1993).

(2) Reasoning favors causal knowledge. Of course, the strategies we have called causal-ecological involve causal knowledge directly (e.g., reasoning about how a property might be transmitted through ecological interactions), and knowledgeable reasoners seem to prefer these to similarity-based strategies even for properties that might plausibly be generalized according to similarity (e.g., enzymes). But even similarity-based reasoning might sometimes involve causal considerations indirectly. Though it may sometimes serve as a mere fallback or default strategy, at other times similarity is invoked just because the property in question is understood to be involved in the causal mechanisms that underlie the similarity-and-taxonomy structure of the domain (Hadjichristidis et al., 2004; Rips, 2001).

(3) Taxonomic relations matter, and they may differ, at least in salience, across cultures. For example, some cultures do not have a superordinate term for animals, and cultures that do have such a term do not always include humans in it (see Anggoro et al., 2005, for evidence and implications). This may mean that, for example, generalization from a

given premise category will be broader when the superordinate category that supports broad generalization is more salient (say, because it is named rather than covert).

(4) When various similarity-based and causal-ecological strategies are available to the reasoner, which strategy is invoked depends in part upon the organization or relative accessibility of different pieces of the reasoner's knowledge base (e.g., Shafto et al., 2005). Or, to put it differently, a single knowledge base can be organized in different ways that render different reasoning strategies more or less fluent. The organization of the knowledge base, in turn, is sensitive not just to the practices or experiences through which people acquire domain knowledge but also to cultural milieu (perhaps in the form of framework theories) (e.g., Medin et al., 2002; Medin, Ross, Atran, et al., in press).

(5) Even similarity-based reasoning, which one might expect to be well constrained by the intrinsic structure of the domain (taxonomic relationships, clusters of related features, and so on), is sensitive to cultural milieu. Even at early ages, cultural beliefs or framework theories seem to influence how broadly reasoners generalize from different premise categories (Ross et al., 2003; see also Medin & Atran, 2004).

What are the implications of these findings for models of reasoning? That reasoners often abandon similarity in favor of causal-ecological strategies (findings 1, 2, and 4) represents a serious limitation of the similarity-coverage and feature-based models, because they seem to predict that similarity-based reasoning will be universal. Of course, one might argue that these models are only meant to apply to truly blank properties—that is, properties for which the reasoner has absolutely no prior belief about their distribution—and that people with some domain knowledge interpret almost any property in such a way that, functionally, it is not blank. There are two problems with this counterargument. One is that it restricts the applicability of such models to the point of irrelevance. The second is that it is not the case that knowledgeable reasoners never show similarity-based reasoning; rather, similarity-based reasoning is one strategy among many.

Indeed, the similarity-based models seem to emerge from this analysis as good accounts of one strategy that is invoked when knowledge is sparse or when the reasoner believes that the property in question is related to other properties that determine the similarity-and-taxonomy structure of the domain. Still, how can models of reasoning handle findings 1, 2, and 4? We turn now to three alternative theories, each of which addresses some part of these findings.

McDonald, Samuels, and Rispoli (1996) presented a *hypothesis-assessment model* of inductive reasoning. On this model, generalization of properties over categories can be viewed as a form of hypothesis assessment in which the conclusion is the relevant hypothesis and the premise is some evidence for this hypothesis. A general prediction of this model is that inferences will be sensitive to the same sorts of factors that influence hypothesis assessment in other contexts—especially the number of competing hypotheses. McDonald et al. provide support for this prediction by asking people to generate hypotheses or explanations and showing that confidence in inferences decreases when there are competing hypotheses (i.e., other candidates for a conclusion category). The hypothesis-assessment model places no constraints on the hypotheses or explanations that a reasoner might consider, and so in principle it is consistent with a variety of strategies (finding 1). However, it has trouble with finding 2, in that it gives no special status to causal factors. As for knowledge accessibility (finding 4), the model as initially described has little to say, but McDonald et al. discussed this. In their data, judgments of argument strength seemed to be sensitive to how accessible the corresponding conclusion categories were (as measured by the number of participants who, given the premises, spontaneously generated the conclusion category). McDonald et al. noted that adding accessibility to the model as a predictor would be reasonable, given that the influence of accessibility on hypothesis strength has been shown in other contexts.

Heit (1998, 2000) proposed a *Bayesian model* that works more or less as follows. If the task is to generalize a novel property from cows to sheep, then various known properties of cows and sheep are called to mind. Of these, some are likely shared by cows and sheep, and others distinctive to cows. From the numbers of properties that are shared and distinctive, the reasoner computes the likelihood that the novel property belongs to one group or the other. (If known properties tend to be shared, then it's likely that the novel property is shared. If known properties tend to be distinctive, then it's likely that the novel property is distinctive.) To the extent that the novel property is likely to belong to the shared group, the inference from cows to sheep is strong. In one important respect this model is like the similarity-based models (and especially the feature-based model): It works on the assumption that the novel property is associated with known properties, and larger clusters of known properties carry more weight in reasoning. Yet the Bayesian model allows flexibility in just which of the known properties of sheep and cows are considered. In this way the model explains, for example, Heit and Rubinstein's (1994) finding that anatomical properties were generalized according to animals' anatomical similarity, whereas behavioral properties were generalized according to behavioral similarity. On the Bayesian model, this is because a novel anatomical property calls to mind known anatomical properties, and these dominate in the inference process. In contrast, a novel behavioral property calls to mind known behavioral properties.

Because the Bayesian model is similarity based (i.e., works on the assumption that the novel property is associated with known properties), it has trouble with findings 1 and (especially) 2. As for finding 4, the Bayesian model allows that different bits of knowledge are called to mind in different contexts, and this can be seen as reflecting context-specific differences in knowledge accessibility. Still, the model's ability to account for finding 4 in this way is limited. First, it relies on an independent theory to explain which bits of knowledge—that is, which known properties of the premise and conclusion categories—are called to mind. Second, and more importantly, the relevant bits of knowledge are always known properties of the premise and conclusion categories; they are not causal mechanisms by which properties are acquired or transmitted. This is another way of saying what has already been said, namely, that the model is similarity based and does not predict causal-ecological reasoning.[4]

A theory better suited to finding 4 is the *relevance framework* outlined by Medin, Coley, Storms, and Hayes (2003). One way to motivate this framework is to consider some responses that a tree expert gave to Proffitt et al. (2000). The expert was given probes such as the following: "Suppose we know that river birch get disease X and that white oaks get disease Y. Which disease do you think is more likely to affect all trees?" In this case, the expert said disease X, noting that river birches are very susceptible to disease (so that "if one gets it they all get it"). The very next probe involved the gingko tree, and the expert chose the disease associated with it as more likely to affect all trees on the grounds that "gingkos are so resistant to disease that if

[4] One might suggest that the Bayesian model handles causal-ecological strategies by computing probabilities over just those properties that are associated with a certain ecological interaction, but here the activation of just those properties is doing most of the explanatory work, and this requires an independent theory. Furthermore, this method of mimicking causal-ecological reasoning is limited to cases where the premise and conclusion categories share some property (e.g., habitat) that can serve as a proxy for the relevant causal mechanism—for example, an inference from one river fish to another might be predicted if 'lives in rivers' is called to mind as a shared property—but it is difficult to imagine what shared property might mimic an inference that cows get a property from grass by eating it.

they get it, it must be a very powerful disease." He then said that he felt as if he had just contradicted himself, but that nonetheless these seemed like the right answers.

Normatively, this expert's answers do not represent a contradiction. Instead, he appeared to be using the information that was most salient and accessible to guide his reasoning (birches are notoriously susceptible to, and gingkos notoriously resistant to, diseases). Simply put, the expert was using the knowledge that he considered most relevant. Medin et al. (2003) suggested that Sperber and Wilson's (1986) relevance theory provides a good framework for understanding these and related patterns of responding. One motivation for this view is the fact that experiments take place in a social context and participants reasonably infer that the experimenter is being relevant and informative with respect to the inductive argument forms (cf. Grice, 1975). Furthermore, this view leads to a number of novel predictions that contrast with those of other models.

In relevance theory, relevance is seen as a property of inputs to cognitive processes:

> An input is relevant to an individual as a certain time if processing this input yield cognitive effects. Examples of cognitive effects are the revision of previous beliefs, or the derivation of contextual conclusions, that is, conclusions that follow from the input taken together with previously available information. Such revisions or conclusions are particularly relevant when they answer questions that the individual had in mind (or in an experimental situation, was presented with)….Everything else being equal, the greater the cognitive effects achieved by processing an input, the greater its relevance. On the other hand, the greater the effort involved in processing an input, the lower the relevance.…One implication of the definition of relevance in terms of effect and effort is that salient information, everything else being equal, has greater relevance, given that accessing it requires less effort." (Van der Henst, Sperber, & Politzer, 2002, p. 4)

In support of this approach, Medin et al. (2003) experimentally manipulated effort and effect to determine whether they have the sorts of consequences predicted by relevance theory.

Undergraduates contrast with experts in having little background knowledge to bring to bear on these sorts of reasoning tasks, and consequently it is not surprising that they rely heavily on more abstract reasoning strategies. In their studies with undergraduates, Medin et al. were able to identify accessible background knowledge to bring out the effect side of relevance and manipulated the premise and conclusion categories to show consequences on the effort side. As an example of the former, they found that the argument “bananas have enzyme X, therefore monkeys have Enzyme X” was rated stronger than the argument “mice have enzyme X, therefore monkeys have Enzyme X.” In this case relevant background knowledge that monkeys like bananas leads to a violation of similarity.

As an example of varying effort, Medin et al. (2003) showed that undergraduates rate the inductive strength of the argument “grass has enzyme Y, therefore humans have enzyme Y” to be less strong than the argument “grass has enzyme Y, therefore cows and humans have enzyme Y.” (The arguments were not juxtaposed but rather were used in a between-subjects design.) In this case, the data yield a “conclusion conjunction fallacy” since, normatively, the former argument’s conclusion cannot be less likely than the conclusion of the latter argument. From a relevance perspective, the addition of cows to the conclusion made it easier for the participants to access a sensible causal pathway between grass and humans.

In other conditions, Medin et al. demonstrated premise nonmonotonicity, that is, a drop in argument strength with the addition of premises. For example, “white oaks get disease X, therefore sugar maples do” was rated stronger than “white oaks, red oaks, and burr oaks get disease X, therefore sugar maples do.” In this case the idea is that multiple premises involving oaks make “oaks” salient and relevant and reinforce the idea that disease X is specific to oaks.

Of the other models we have described, the relevance framework is most closely related to the hypothesis-assessment model, in that one could see the relevance framework as a basis for predicting which hypotheses people will tend to generate. Importantly, due to its emphasis on effort, the relevance theory explains differences due to knowledge accessibility (finding 4) in a natural way: Greater accessibility means less effort, and more accessible pieces of knowledge are therefore favored in reasoning. We readily concede that relevance theory seems vague, especially in relation to computational models like the similarity-coverage model. But a theory must also be judged on its ability to generate novel predictions, and relevance theory fares well by this standard.

Although there are numerous other studies looking at inductive reasoning in a cultural context, their primary focus is on induction informing culture rather than vice versa. We now turn to a more speculative consideration of how culture may inform theories of causal reasoning.

## Culture and Causal Reasoning

Psychologists have recently begun to explore a theory of learning and reasoning about systems of causal relationships. At the same time, cross-cultural work has revealed ways in which reasoning is guided by abstract expectations or understandings of causal structure. Both lines of work are young (consequently, this section will be shorter and more speculative than the previous one), yet both are far enough along that we can begin to see tensions between them and ways in which the cultural research might inform the new theory.

### *1. Causal Bayes net theory*

The use of causal knowledge to predict and control events is a form of reasoning that is ubiquitous both in everyday life and in more formal contexts like science and medicine. Causal knowledge enables us to discern with some precision whether, or to what degree, variables are relevant to predictions about one another, or to interventions to control one another. When we know causal relationships, we tend to base our predictions on factors that are causally relevant, and to focus our interventions on factors that will in fact transmit influence to the things we wish to control. For example, if we knew that two variables *X* and *Y* were associated just because they had a common cause *C*, and if we wanted to make a prediction about *Y*, we would rather base this prediction on *C* than on *X*. If we wished to exercise control over *Y* by manipulating one of the other variables, we would manipulate *C*, not *X*.

A detailed account of how causal knowledge can be used to make predictions and exercise control has been developed in philosophy, statistics, and computer science (Pearl, 2000; Spirtes, Glymour, & Scheines, 2000) and is known as causal Bayes net theory, or causal graphical model theory. There has recently been a good deal of interest in using this theory as a basis of psychological accounts (e.g., Glymour, 2001; Gopnik et al., 2004; Lagnado & Sloman, 2004; Rehder & Hastie, 2001; Sloman & Lagnado, 2005; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann & Martignon, 1998). Briefly, causal Bayes net theory specifies mappings between causation and correlation (where we use *correlation* broadly, to refer to statistical relationship of any form). Causal knowledge is represented as graphs in which variables (events, states, and so on) appear as nodes, and causal relationships as directed links between nodes. The heavy explanatory work is done by a principle called the *causal Markov assumption*: Any variable is uncorrelated with, or statistically independent of, all variables that are not its descendants in causal structure—that is, not its direct or indirect effects—conditional on its immediate cause(s).

Given a causal graph, the Markov assumption says just which variables are relevant to any instance of prediction or control. For example, suppose that a certain virus, $V$, causes two symptoms, $A$ and $B$, by different mechanisms. In graphical form: $A \leftarrow V \rightarrow B$. According to the Markov assumption, $A$ and $B$ should be correlated in general but uncorrelated conditional on $V$, that is, when $V$ is controlled for. (Neither is a descendant of the other, and each has $V$ as its sole immediate cause.) A reasoner who respects the Markov assumption should treat $A$ as predictive of $B$ except when there is information about $V$, in which case a prediction about $B$ should be based on $V$ alone. If it is known that a patient has (or does not have) the virus, then the appearance of symptom $A$ should not be predictive of the appearance of $B$. Interventions to control variables are represented as surgeries on graphs (Pearl, 2000), in which the links leading into an intervened-upon variable are broken, and the intervention itself becomes the sole cause of the variable. The Markov assumption can then be applied to determine how relevance flows through the new graph. For example, intervening to control symptom $B$ (by means of a perfectly effective drug, say) breaks the influence of the virus $V$ on this symptom; consequently, $B$ is no longer predictive of $V$ or $A$.

The causation-correlation mappings specified by the Markov assumption can also be applied in the other direction, to learn causal models from correlational evidence. If $X$ and $Y$ are correlated in general but uncorrelated when $Z$ is controlled for, then, given a few additional assumptions, there are three ways in which these variables might be related: $X \rightarrow Z \rightarrow Y$, $X \leftarrow Z \leftarrow Y$, and $X \leftarrow Z \rightarrow Y$. Other models, like $Z \rightarrow X \rightarrow Y$, can be ruled out. Though the details are beyond the scope of this chapter, causal Bayes net theory also provides leverage for understanding some seemingly more complicated tasks, including (a) learning causal models by intervening to manipulate variables and (b) inferring the presence of hidden causes.[5]

Several features of causal Bayes net theory are pertinent:

(1) The theory requires that causal knowledge be *complete*, in two respects. First, there can be no unknown cause of any combination of known variables; this requirement has been called *causal sufficiency* (Spirtes et al., 2000). Second, there can be no unknown paths of influence between variables. If causal knowledge is not complete in these ways, then the Markov assumption does not apply. To see why, suppose there were an unknown common cause of symptoms *A* and *B* that did not act via the virus *V*, or an unknown path of influence between *A* and *B* (perhaps *A* promotes *B* independently of their joint dependence on *V*). In these cases *A* and *B* would be relevant to predictions about one another (and perhaps to interventions to control one another) even conditional on *V*.

(2) Causal knowledge is *concrete*, in two senses. First, causal graphs are made of relationships among specific variables, by which we mean variables encoded with enough specificity to map onto the events and states that we make predictions about, base predictions on, manipulate, and so on. There is (as yet) no place in the formalism for more abstract notions about what *kinds* of variables might be related or relevant to one another. Second, the mechanisms underlying causal relationships must often be understood in some detail if causal knowledge is to be properly complete (Hausman & Woodward, 1999). For example, in cases of a single cause

[5] For applications of these ideas to psychology, see... On reasoning about interventions: Sloman and Lagnado (2005), Waldmann and Hagmayer (2005). On learning causal models by intervention: Gopnik et al. (2004), Lagnado and Sloman (2004); Steyvers et al. (2003). On inferring hidden variables: Gopnik et al. (2004); Kushnir, Gopnik, Schulz, & Danks (2003).

with multiple effects, it is often difficult to tell whether the effects come about by truly distinct mechanisms or by way of a single unseen mechanism. If virus *V* gave rise to symptoms *A* and *B* by way of a single hidden variable (say, lack of a certain enzyme), then *A* and *B* would have an unknown common cause, completeness would be violated, and the Markov assumption would not apply. Such arrangements have been called *interactive forks* (Salmon, 1984; on the failure of the Markov assumption to apply in these cases, see Sober, 1988).

(3) There are two general approaches to learning causal models from correlational evidence: (a) a bottom-up or "constraint-based" approach, in which learning is primarily a matter of using the causation-correlation mappings specified by the Markov assumption to work backwards, from patterns of correlation to causal models consistent with these patterns (e.g., Spirtes et al., 2000); and (b) a top-down Bayesian approach, in which learning begins with a set of candidate causal models and evaluates these models for their likelihood of having generated the correlational evidence (e.g., Heckerman, Meek, & Cooper, 1999). The bottom-up view ties causal structure closely to the correlational input. This view suggests that, in cases where causal knowledge is acquired from correlational evidence, there should be little cultural variation in the understood causal structure of the world (unless, of course, there is corresponding variation in the correlational evidence available to learners in different cultures). If there *were* cultural variation in the interpretation of correlational evidence, this would favor the top-down approach, where culture-related expectations of causal structure might guide the construction of the initial set of candidate structures.

*2. Culture-related understandings of causation*

Recent cross-cultural research suggests that causal reasoning is often guided by abstract expectations or understandings of the causal structure of the world. Consider the theory that

Easterners and Westerners have holistic and analytic theories, respectively, about causation (Nisbett, 2003; Nisbett, Peng, Choi, & Norenzayan, 2001). The idea is that Westerners (roughly, Europeans and European Americans) think of the world as partitionable into causally unrelated objects. When reasoning about the behavior of an object, they tend to look for causes in the intrinsic disposition or attributes of the object itself (e.g., an object falls to earth because of its weight). In contrast, Easterners (roughly, east and central Asians) are thought to have a more holistic view of the causal structure of the world, which can be characterized (or perhaps caricaturized) as “everything affects everything else.” They tend to consider multiple, perhaps interactive causes and, when reasoning about the behavior of an object, to look outside of the object to situational or environmental influences (e.g., an object falls because of an external force) (Choi, Nisbett, & Norenzayan, 1999; Morris & Peng, 1994; Nisbett, 2003; Nisbett et al., 2001; Peng & Knowles, 2003).

The evidence for this view is mostly in the social domain (for a review see Choi et al., 1999). Miller (1984) found that, in explaining a person’s behavior, participants in India invoked more contextual or situational factors than did Americans, who spoke to a greater extent about the person’s disposition. Similarly, Morris and Peng (1994) asked graduate students to explain the behavior of a murderer and found that Americans favored dispositional factors and Chinese situational factors. These findings have been echoed in studies of attributions made in print by American and Chinese journalists (Lee, Hallahan, & Herzog, 1996; Morris & Peng, 1994). When asked to explain the behavior of a cartoon fish which moved in various ways relative to a group of other fish, Americans showed a greater tendency than did Chinese to rate internal causes as more important than the other fish (Morris & Peng, 1994).

Predictions in other domains have found less support. Morris and Peng (1994) asked participants about the causes of movements of physical objects (in animated depictions). Participants were asked to rate the relative importance of internal and external causes. This revealed only one difference between American and Chinese participants, and it was in the direction opposite what one might have predicted. Given an animated depiction of an "entraining" event in which one object collides with another and then continues on its path, pushing the second object before it, Americans seemed to favor external over internal causes for the second object's motion, whereas Chinese showed no preference.

Ji, Peng, and Nisbett (2000) used a contingency learning task in which the goal was to judge the strength of a relationship between which of two abstract shapes appeared on the left side of a screen and which of two other shapes appeared on the right. Whereas there were differences in the judgments made by American and Chinese participants, these differences can be explained largely as reflecting a group difference in response bias.[6]

Peng and Knowles (2003) found some suggestive differences in how American and Chinese participants explained abstract, seemingly physical events. When participants' explanations were classified as either dispositional (e.g., referring to a target object's weight or composition) or contextual (e.g., referring to another object or a surrounding medium), American college students showed a bias, relative to Chinese college students, toward dispositional explanations on a few items. Still, this trend was not statistically reliable over the whole set of

---

[6] Ji et al. also reported that Americans but not Chinese showed a primacy effect, in which the judged strength of the relationship was more sensitive to cases presented early in the learning phase than to those presented later. This might have been due to a group difference in level of engagement in the task.

items, and large differences between items suggest that a considerable part of the story is yet unexplained.

Our belief is that there is something right about the idea that cultures differ in understandings of causal structure, but that these understandings have sometimes been theorized too abstractly.[7] They are probably bound up with beliefs, habits, goals, and ideals in particular domains. On this view, the robust findings in the social domain are the result not of domain-general stances like holism but of more specific theories of personal behavior and social interaction. Such theories were measured directly by Norenzayan, Choi, and Nisbett (2002), who asked Korean and American participants to rate their agreement with each of three general explanations of human behavior: a dispositional theory, a situational theory, and an interactionist theory (according to which dispositional and situational factors interact to yield behavior). Compared to Americans, Koreans reported reliably greater agreement with situationism in two studies and with interactionism in one. We suspect that theories at this level of abstraction are likely guides of causal reasoning.[8]

Research in folkbiology has revealed other cases in which cultural groups seem to differ in (domain-specific) understandings of causal structure. This work involves the "species interaction" task described earlier. Atran et al. (1999, 2002) asked participants in three cultural

---

[7] As a consequence, the mapping between an abstract principle and a particular task typically involves a series of assumptions, often implicit, that themselves may not be straightforward.

[8] To illustrate how theories and goals might vary by domain, a Korean colleague has suggested that Korean explanations of mental illness or deviant behavior might actually be less holistic than American explanations, precisely because social interconnectedness is valued in Korean culture. A holistic explanation (e.g., society drove him mad) would implicate society, and to avoid this Koreans might prefer analytic explanations (e.g., bad traits or genes).

groups in Guatemala to say whether and how various kinds of plants and animals affect each other. It was found that the groups differed in the numbers of interactions they reported and, in particular, in the numbers of helpful effects of animals on plants. In short, the groups seemed to have different understandings of the kinds of causal relationships in the ecosystem. As we noted earlier, Medin, Ross, Atran, et al. (in press) found that Menominee and majority-culture fishermen in Wisconsin also have different views of causal structure, such that reciprocal relationships among fish are more salient to Menominee participants. Informally, Menominee participants have sometimes articulated framework theories of causal interaction like “every fish has a role to play” and “all in all, living things affect each other.”

In sum, cross-cultural research suggests that causal reasoning is often guided by understandings of causal structure that are more abstract than causal Bayes nets (though probably not as abstract as domain-general stances like holism). It is too early to say just what these understandings are—perhaps they are skeletal framework theories (Keil, 2003a, 2003b) or “causal grammars” (Tenenbaum, Griffiths, & Niyogi, in press)—but in at least some cases they seem to vary along a dimension that runs from *autonomy* (things in a domain tend not to affect each other) to *influence* (things tend to affect each other).

*3. Bringing the cultural research to bear on causal Bayes net theory*

These two lines of work are young, and the tasks they involve are quite different—for example, learning whether a certain object causes a machine to light up (Gopnik et al., 2004) versus explaining a murder (Morris & Peng, 1994). Yet we can begin to see tensions between them that are similar to the tensions between similarity-based models of inductive reasoning and cross-cultural research on reasoning about plants and animals.

Recall that the principal way in which causal Bayes net theory explains prediction and control is to specify the relevance of variables to one another, based on their relative positions in causal structure. It does this by applying a general principle, the causal Markov assumption, to causal models that are complete and concrete. There are no representational tools for accommodating abstract beliefs about relevance in a domain, and the Markov assumption fails if causal knowledge is incomplete in certain ways. These limitations should be taken seriously in building psychological accounts of reasoning on causal Bayes net theory, because natural causal knowledge, knowledge that supports reasoning, is often abstract and incomplete (Keil, 2003b; Rozenblit & Keil, 2002). Indeed, it is often abstract and incomplete in just the ways that invalidate the causal Markov assumption (interactive forks, etc.; Hausman & Woodward, 1999). In these cases (and they may be the norm), the lack of complete, concrete knowledge makes room for abstract understandings of relevance and causal relatedness to guide learning and reasoning. Cultural research shows that learning and reasoning are indeed guided by abstract expectations—and, more specifically, expectations that may vary along a continuum from autonomy to influence.

Expectations of autonomy are roughly consistent with the causal Markov assumption, which has the form “variables are irrelevant to one another unless proven otherwise.” Proving otherwise requires knowing just how variables are related in a complete, concrete causal graph, with the result that relevance is assigned to variables conservatively. In contrast, expectations of influence seem to run against the Markov assumption, as they assign relevance more liberally. For a reasoner who knows the model $A \rightarrow B \rightarrow C$, the causal Markov assumption dictates that a prediction about $C$, given information about $A$ and $B$, should be based on $B$ alone. But if this reasoner has an expectation of influence, then their prediction might be based also on $A$. This can be seen as allowing for the possibility that *A* and *C* are related in some unknown way. We know of no cross-cultural research on this question, but within-culture work has shown that both undergraduates and domain experts do tend to deviate from the Markov assumption in this way (Burnett, 2005; Rehder & Burnett, 2005). There is also some tentative evidence that scientists in different fields override the causal Markov assumption to different degrees, consistent with abstract expectations in their respective fields (Burnett, 2004). The clear prediction for cultural comparisons is that reasoners with a greater expectation of influence in a given domain will go farther in assigning relevance or predictive value to variables which, according to the causal Markov assumption, should be irrelevant or nonpredictive. We might make similar predictions about active intervention. Consider again a reasoner who knows the model $A \rightarrow B \rightarrow C$. If *B* is manipulated as a way of controlling *C*, the Markov assumption says there is no additional benefit to manipulating *A*. Reasoners with an expectation of influence might see additional benefit and, when given the option, might manipulate *A* as well as *B* in order to control *C*.

Cultural research also has implications for the learning of causal systems from correlational evidence. In the bottom-up or constraint-based approach, where learning is driven mainly by the Markov assumption, there is no room for abstract expectations of influence or autonomy. Furthermore, because an infinity of causal structures involving latent variables are (according to the Markov assumption) consistent with any pattern of correlational evidence, the bottom-up approach requires additional assumptions that favor causal structures that are parsimonious in some respect(s) (e.g., having the fewest latent variables). Cultural research suggests that learning is guided by abstract expectations about causal structures in different domains, and this argues in favor of the top-down approach, where abstract knowledge may influence the causal models that a learner considers (e.g., Tenenbaum & Griffiths, 2003;

Tenenbaum et al., in press). Furthermore, when expectations favor influence instead of autonomy, learning may deviate from the parsimony assumptions used in the bottom-up approach. Learners with expectations of influence may infer causal structures that are more than minimally elaborate, perhaps with more latent variables than are necessary to account for observed correlations.

## Conclusion

In both of the cases we have discussed, models have been proposed which ground reasoning in concrete knowledge. In the case of inductive reasoning about biological properties, the similarity-based models explain reasoning in terms of that aspect of the domain that is most readily detected: clusters of correlated features and the taxonomies that form around these clusters. In the case of causal reasoning, causal Bayes net theory explains reasoning (i.e., prediction and control) in terms of graphical models of causal systems that are complete and concrete. In both cases, one can at least imagine how the relevant knowledge might be acquired by an individual learner exploring the world independently of any cultural influence.

In both cases, cultural research suggests that reasoning is often guided by more abstract knowledge that is less constrained by the observable structure of the world and more culturally variable. Reasoning about biological properties draws preferentially on an aspect of the domain that is less easily detected than feature clusters and taxonomies, namely, causal mechanisms by which properties arise and by which properties are transmitted among categories. Furthermore, the knowledge base that supports reasoning—which may include feature clusters, taxonomic relationships, causal-ecological relationships, and more—may be organized according to different framework theories, so as to render different kinds of knowledge more or less accessible. These differences are reflected in reasoning. As for causal reasoning, in at least some domains it draws not only on concrete knowledge but also on abstract theories or expectations of causal structure.

This is not to say that framework theories and expectations of the kinds described in this chapter are the only paths through which culture influences reasoning. It is possible, for example, that cultural factors sometimes promote certain modes of reasoning over others. Futhermore, some cultural differences related to reasoning—for example, differences in typicality among members of biological categories—seem better explained by ideals and goals than by theories and expectations (Burnett, Medin, Ross, & Blok, 2005; Lynch, Coley, & Medin, 2000).

Given that we take the proximal mechanisms of culture to be cognitive constructs like theories and expectations, one might wonder what is to be gained by studying culture itself. There are several things to be gained. First, in many cases it would be difficult or impossible to identify the relevant cognitive constructs without knowing their sources in culture. Second, the cognitive constructs that depend on culture may be stable or entrenched in ways that cannot be mimicked or modified in laboratory experiments. Finally, the cultural sources of cognitive constructs are interesting in their own right, especially if one wants to predict reasoning in real-world settings.

We have said little about what exactly the operative framework theories and expectations are, beyond suggesting that they are somewhat domain specific—like the theories of human behavior described by Norenzayan et al. (2002) and the idea that "every fish has a role to play," mentioned by some Menominee fishermen (Medin et al., 2002). We have said even less about how they are acquired and transmitted. Possibilities include imitation, inference from other people's behavior, explicit communication of abstract principles (e.g., one Menominee participant explained his bias to see interactions among species by saying that his grandmother had taught him that all living things affect each other), explicit communication of more concrete facts from which principles can be abstracted, communication of goals and ideals that lead indirectly to differences in knowledge, and so on. As an example of the work that can be done in this area, Atran et al. (2002) employed social network analysis to trace the transmission of mental models of the ecosystem among members of different cultural groups in Guatemala. As this work illustrates, collaboration between psychology and anthropology has the potential to reveal determinants of reasoning that are neither easily detected in the structure of the world nor easily identified without considering the cultural milieus in which individual minds develop and function.

References

Anggoro, F., Waxman, S., & Medin, D. (2005). The effect of naming practices on children's understanding of living things. *Proceedings of the 27th Annual Conference of the Cognitive Science Society.*

Atran, S. (1998). Folk biology and the anthropology of science: Cognitive universals and cultural particulars. *Behavioral and Brain Sciences, 21,* 547–609.

Atran, S., Medin, D., Lynch, E., Vapnarsky, V., Ucan Ek', E., & Sousa, P. (2001). Folkbiology doesn’t come from folk psychology: Evidence from Yukatek Maya in cross-cultural perspective. *Journal of Cognition and Culture, 1,* 3–42.

Atran, S., Medin, D. L., & Ross, N. O. (2005). The cultural mind: Environmental decision making and cultural modeling within and across populations. *Psychological Review, 112,* 744–776.

Atran, S., Medin, D., Ross, N., Lynch, E., Coley, J., Ucan Ek', E., & Vapnarsky, V. (1999). Folk ecology and commons management in the Maya Lowlands. *Proceedings of the National Academy of Sciences of the United States of America, 96,* 7598–7603.

Atran, S., Medin, D., Ross, N., Lynch, E., Vapnarsky, V., Ek', E. U., Coley, J., Timura, C., & Baran, M. (2002). Folk ecology, cultural epidemiology, and the spirit of the commons: A garden experiment in the Maya lowlands, 1991–2001. *Current Anthropology, 43,* 421–441.

Bailenson, J. N., Shum, M. S., Atran, S., Medin, D. L., & Coley, J. D. (2002). A bird’s eye view: Biological categorization and reasoning within and across cultures. *Cognition, 84,* 1–53.

Bang, M., Medin, D., Unsworth, S., & Townsend, J. (2005, April). *Cultural models of nature and their relevance to science education.* Paper presented at the annual meeting of the American Educational Research Association, Montreal.

Berlin, B. (1992). *Ethnobiological classification: Principles of categorization of plants and animals in traditional societies.* Princeton, NJ: Princeton University Press.

Burnett, R. C. (2004). *Inference from complex causal models.* Unpublished doctoral dissertation, Northwestern University.

Burnett, R. C. (2005). Close does count: Evidence of a proximity effect in inference from causal knowledge. *Proceedings of the 27th Annual Conference of the Cognitive Science Society*.

Burnett, R. C., Medin, D. L., & Ross, N. O. (2004). *Experts don't reason about fish out of water: Fishermen's inductive reasoning is based on causal-ecological relations* [Contribution to festschrift for Edward E. Smith]. Available in PsycEXTRA.

Burnett, R. C., Medin, D. L., Ross, N. O., & Blok, S. V. (2005). Ideal is typical. *Canadian Journal of Experimental Psychology, 59,* 5–10.

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: MIT Press.

Choi, I., Nisbett, R. E., & Norenzayan, A. (1999). Causal attribution across cultures: Variation and universality. *Psychological Bulletin, 125,* 47–63.

Choi, I., Nisbett, R. E., & Smith, E. E. (1997). Culture, category salience, and inductive reasoning. *Cognition, 65,* 15–32.

Cohen, D. (2001). Cultural variation: Considerations and implications. *Psychological Bulletin, 127,* 451–471.

Cole, M., & Scribner, S. (1974). *Culture and thought: A psychological introduction.* New York: John Wiley & Sons.

Coley, J. D., Shafto, P., Stepanova, O., & Baraff, E. (2005). Knowledge and category-based induction. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin.* Washington, DC: American Psychological Association.

Dias, M., Roazzi, A., & Harris, P. L. (2005). Reasoning from unfamiliar premises: A study with unschooled adults. *Psychological Science, 16,* 550–554.

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition, 23,* 183–209.

Glymour, C. (2001). *The mind's arrows: Bayes nets and graphical causal models in psychology.* Cambridge, MA: MIT Press.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111,* 3–32.

Greenfield, P. M. (1983). Review of the book *The psychology of literacy. Harvard Educational Review, 53,* 216–220.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41–58). New York: Academic Press.

Hadjichristidis, C., Sloman, S., Stevenson, R., & Over, D. (2004). Feature centrality and property induction. *Cognitive Science, 28,* 45–74.

Harris, P. L. (2000). *The work of the imagination.* Oxford, UK: Blackwell.

Hausman, D. M., & Woodward, J. (1999). Independence, invariance and the Causal Markov Condition. *British Journal for the Philosophy of Science, 50,* 521–583.

Heckerman, D., Meek, C., & Cooper, G. (1999). A Bayesian approach to causal discovery. In C. Glymour & G. F. Cooper (Eds)., *Computation, causation, and discovery* (pp. 141–165). Menlo Park, CA: AAAI Press.

Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition* (pp. 248–274). Oxford, UK: Oxford University Press.

Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review, 7,* 569–592.

Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 411–422.

Higgins, E. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. Higgins & A. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford.

Hong, Y., Morris, M.W., Chiu, C., & Martinez, V.B. (2000). Multicultural minds. *American Psychologist, 55,* 709–720.

Inagaki, K. (1990). The effects of raising animals on children's biological knowledge. *British Journal of Developmental Psychology, 8,* 119–129.

Ji, L.-J., Peng, K., & Nisbett, R. E. (2000). Culture, control, and perception of relationships in the environment. *Journal of Personality and Social Psychology, 78,* 943–955.

Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 234–262). New York: Oxford University Press.

Keil, F. C. (2003a). Categorisation, causation, and the limits of understanding. *Language and Cognitive Processes, 18,* 663–692.

Keil, F. C. (2003b). Folkscience: Coarse interpretations of a complex reality. *Trends in Cognitive Sciences, 7,* 368–373.

Keil, F. C., Levin, D. T., Richman, B. A., & Gutheil, G. (1999). Mechanism and explanation in the development of biological thought: The case of disease. In D. L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 285–319). Cambridge, MA: MIT Press.

Keil, F. C., Smith, W. C., Simons, D. J., & Levin, D. T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition, 65,* 103–135.

Kushnir, T., Gopnik, A., Schulz, L., & Danks, D. (2003). Inferring hidden causes. *Proceedings of the 25th annual conference of the Cognitive Science Society.*

Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 856–876.

Lee, F., Hallahan, M., & Herzog, T. (1996). Explaining real-life events: How culture and domain shape attributions. *Personality and Social Psychology Bulletin, 22,* 732–741.

López, A., Atran, S., Coley, J. D., Medin, D. L., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology, 32,* 251–295.

Luria, A. R. (1976). *Cognitive development: Its cultural and social foundations* (M. Lopez-Morillas & L. Solotaroff, trans.; M. Cole, ed.). Cambridge, MA: Harvard University Press.

Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory & Cognition, 28,* 41–50.

McDonald, J., Samuels, M., & Rispoli, J. (1996). A hypothesis-assessment model of categorical argument strength. *Cognition, 59,* 199–217.

Medin, D. L., & Atran, S. (2004). The native mind: Biological categorization and reasoning in development and across cultures. *Psychological Review, 111,* 960–983.

Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. K. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review, 10,* 517–532.

Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology, 32,* 49–96.

Medin, D. L., Ross, N., Atran, S., Cox, D., Coley, J., Proffitt, J. B., & Blok, S. (in press). Folkbiology of freshwater fish. *Cognition.*

Medin, D. L., Ross, N., Atran, S., Burnett, R. C., & Blok, S. V. (2002). Categorization and reasoning in relation to culture and expertise. In B. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 41, pp. 1–41). New York: Academic Press.

Medin, D. L., Ross, N. O. & Cox, D. (in press). *Culture and resource conflict: Why meanings matter.* New York: Russell Sage Foundation.

Medin, D., Unsworth, S., & Hirschfeld, L. (in press). Categorization and classification. In S. Kitayama & D. Cohen (Eds.), *Handbook of cultural psychology*. New York: Guilford.

Medin, D. L., & Waxman, S. R. (in press). Interpreting asymmetries of projection in children's inductive reasoning.

Miller, J. G. (1984). Culture and the development of everyday social explanation. *Journal of Personality and Social Psychology, 46,* 961–978.

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology, 67,* 949–971.

Nisbett, R. E. (2003). *The geography of thought: How Asians and Westerners think differently…and why.* New York: Free Press.

Nisbett, R. E., & Norenzayan, A. (2002). Culture and cognition. In D. Medin (Ed.), *Stevens' handbook of experimental psychology: Vol. 2. Memory and cognitive processes* (3rd ed.; H. Pashler, editor-in-chief) (pp. 561–597). New York: John Wiley & Sons.

Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review, 108,* 291–310.

Norenzayan, A., Choi, I., & Nisbett, R. E. (2002). Cultural similarities and differences in social inference: Evidence from behavioral predictions and lay theories of behavior. *Personality and Social Psychology Bulletin, 28,* 109–120.

Norenzayan, A., & Heine, S. J. (2005). Psychological universals: What are they and how can we know? *Psychological Bulletin, 131,* 763–784.

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97,* 185–200.

Pearl, J. (2000). *Causality: Models, reasoning, and inference.* Cambridge, UK: Cambridge University Press.

Peng, K., & Knowles, E. D. (2003). Culture, education, and the attribution of physical causality. *Personality and Social Psychology Bulletin, 29,* 1272–1284.

Proffitt, J. B., Coley, J. D., & Medin, D.L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory and Cognition, 26,* 811–828.

Rehder, B., & Burnett, R. C. (2005). Feature inference and the causal structure of categories. *Cognitive Psychology, 50,* 264–314.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General, 130,* 323–360.

Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior, 14,* 665–681.

Rips, L. J. (2001). Necessity and natural categories. *Psychological Bulletin, 127,* 827–852.

Ross, B. H., & Murphy, G. L. (1999). Food for thought. *Cognitive Psychology, 38,* 495–553.

Ross, N., Medin, D., Coley, J. D., & Atran, S. (2003). Cultural and experiential differences in the development of folkbiological induction. *Cognitive Development, 18,* 25–47.

Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26,* 521–562.

Salmon, W. (1984). *Scientific explanation and the causal structure of the world.* Princeton, NJ: Princeton University Press.

Scribner, S. (1977). Modes of thinking and ways of speaking: Culture and logic reconsidered. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking: Readings in cognitive science* (pp. 483–500). New York: Cambridge University Press.

Scribner, S., & Cole, M. (1981). *The psychology of literacy.* Cambridge, MA: Harvard University Press.

Shafto, P., & Coley, J. D. (2003). Development of categorization and reasoning in the natural world: Novices to experts, naive similarity to ecological knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 641–649.

Shafto, P., Coley, J. D., & Baldwin, D. (2005). *Knowledge-based induction: A matter of availability.* Manuscript submitted for publication.

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology, 25,* 231–280.

Sloman, S. A., & Lagnado, D. A. (2005). Do we "do"? *Cognitive Science, 29,* 5–39.

Sober, E. (1988). The principle of the common cause. In J. H. Fetzer (Ed.), *Probability and causality: Essays in honor of Wesley C. Salmon* (pp. 211–228). Dordrecht, Holland: Reidel.

Sperber, D. (1996). *Explaining culture: A naturalistic approach.* Oxford, UK: Blackwell.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition.* Oxford, UK: Blackwell.

Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science, 27,* 453–489.

Tarlowski, A. (in press). If it's an animal it has axons: Experience and culture in preschool children's reasoning about animates. *Cognitive Development.*

Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in Neural Information Processing Systems 15.* Cambridge, MIT Press.

Tenenbaum, J. B., Griffiths, T. L., & Niyogi, S. (in press). Intuitive theories as grammars for causal inference. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation.* Oxford University Press.

Van der Henst, J.-B., Sperber, D., & Politzer, G. (2002). When is a conclusion worth deriving? A relevance-based analysis of indeterminate relational problems. *Thinking & Reasoning, 8,* 1–20.

Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 216–227.

Waldmann, M. R., & Martignon, L. (1998). A Bayesian network model of causal learning. *Proceedings of the 20th Annual Conference of the Cognitive Science Society.*