

Interpreting Asymmetries of Projection in Children's Inductive Reasoning

Douglas L. Medin and Sandra Waxman
Northwestern University

I. Introduction

Like adults, children use categories as a basis for inductive inference. Having learned that some property is true of some individual (e.g., "My dog, Magic, likes marshmallows"), a child might assume both that other members of the same category (dogs) share this property and that members of other, similar categories (e.g. cats) might also share this property. In short, inductive inference may be guided by and reflect categorical relationships, hence the term, category –based induction or CBI. Cognitive and developmental researchers have used this paradigm to not only to study the use of categories in reasoning but also to draw inferences from patterns of reasoning about the nature of conceptual structures themselves (Carey, 1985, Keil, Inagaki & Hatano, 2002; Hatano & Inagaki, 1994; Johnson & Carey, 1998; Gutheil, et al, 2004; Atran et al, 2001, Ross et al 2003.; Gelman & Markman, 1986; Waxman et al., 1997).

One of the most influential examples of the use of pattern of inductive project to draw inferences about conceptual organization comes from research in the domain of folkbiology involving categories of living things, including humans, non-human animals, and plants. Developmental evidence has revealed, in particular, certain systematic *asymmetries* in inductive strength among these categories. For example, researchers have found that children are more willing to project properties from humans to dogs than from dogs to humans. As we will see, one interpretation of this result is that children's biological knowledge is organized around humans as the prototype and that it is more natural to generalize or project from the prototype to its variants than from variants to the prototype (more about this later). There is even some evidence that young, urban children may violate the principle of similarity by generalizing more from humans to bugs than from bees to bugs (Carey, 1985).

In this chapter, our focus will be on asymmetries in children's inductive projections. We begin by describing Carey's studies and outlining the framework within which she interpreted her findings. Subsequent work has questioned the generality of Carey's findings across populations and procedures but relatively little attention has been directed as the theoretical basis for the asymmetries themselves. Although the interpretation that Carey offered is consistent with the observed asymmetries, there are several other alternative interpretations that are also consistent with asymmetries. We describe a number of these alternatives and their associated claims about conceptual organization and processing principles. The upshot is that human-animal asymmetries underdetermine theoretical accounts, and that further empirical constraints are needed.

We next continue to examine asymmetries but broaden their range to include asymmetries involving humans, nonhuman mammals, plants and insects. We also expand the empirical base by drawing category-based

induction in children from a range of cultures. The data show certain asymmetries that appear to hold across populations. These additional sources of evidence offer considerable leverage for evaluating alternative interpretations of asymmetries. To foreshadow, from this broader view, we find that the asymmetries in inductive inference are not well characterized as knowledge or typicality effects favoring humans, as Carey has suggested. Instead, these asymmetries seem to reflect the distinctive categories and features that become activated by the comparison processes associated with inductive inferences. On this account, a major source of human animal asymmetries is the ambiguous status of humans as members of the category animal *and* as members of a category human that is contrastive with animal. The distinctive features/categories account leads to further predictions that are supported by children's open-ended justifications of their responses. We close with a discussion of implications of this newer interpretation of asymmetries for the understanding of category structure and conceptual organization.

II. Asymmetries in Category-Based Induction:

An early source of evidence for asymmetries in young children's inductive reasoning comes from Carey's (1985) CBI task. She used human, dog and bee for a base (the kind to which some property is attributed) and multiple living and nonliving kinds as targets (candidate kinds to which the learned property might be projected). The now-classic finding was that children from 4 to 6 years of age willingly project novel properties from humans to non-human animals, but are reluctant to make the converse generalization from non-human animals to humans. More specifically, young children who are told that people have a little green thing inside them, called an omentum, are quite willing to assert that dogs also have an omentum. In contrast, when children are told that dogs have an omentum, they are reluctant to infer that people also have an omentum. In general, the youngest children were very reluctant to generalize from any base other than humans to other kinds. One striking consequence of this tendency is that the youngest children were more likely to generalize from humans to bugs than from bees to bugs. Carey took this asymmetry as a reflection of the underlying conceptual structure of the young child's mind. Let's take a closer look.

Carey interpreted her results within a framework that argues that children's development can be understood in terms of domain-specific competences. At a minimum, one can distinguish between a naïve physics - understandings of how things in the world work (e.g. Spelke, 1990; Baillargeon, 2000) and theory of mind or naïve psychology - understandings of beliefs, desires and intentions (e.g. Wellman & Gelman, 1992; Leslie, 1984). A third candidate domain is naïve biology - understandings of plants and animals. One issue is whether these domain competencies are present at birth and become elaborated with development, or whether they are wholly acquired. Carey and her colleagues argued that the pattern of findings on biological induction supports the view that young children do not have a distinct naïve biology, but rather their reasoning reflects a naïve psychology where humans are the prototypical psychological entity. Only later, when children are roughly 7 to 10 years of age, do they develop a distinct biology in which humans are seen as one animal

among many. Older children do generalize from bases like dog and bee and, at this developmental point, the asymmetries no longer are there. Notice that on this view, there is radical conceptual change, as children move from an intuitive theory of psychology to a naïve biology.

There are, however, lingering questions concerning Carey's interpretation of the asymmetries in inductive reasoning. One set of questions concerns the conditions under which these asymmetries arise. We touch on this only briefly in the current chapter (see Ross et al, 2003, Gutheil, et al, 2000, and Atran et al, 2001 for a broader analysis of the generality of human animal asymmetries). More central to this chapter are the questions concerning the underlying basis of the asymmetries, whenever they do arise. In the next section, we outline a broad range of possibilities, focusing initially on the observed human - non-human mammal asymmetries. As will be seen, it is straightforward for most of these alternative interpretations to account for the human - non-human mammal asymmetries. Ultimately, then, it will be necessary to broaden our focus to include asymmetries among other categories of living things, including mammals, insects and plants, and to use these additional sets of asymmetries to select among the alternatives.

III. Alternative Interpretations of Asymmetries.

We will consider three broad classes of interpretations of the observed asymmetries in inductive inference. Within each class there are important variations. Our analysis begins with Carey's original interpretation and moves on to consider two others.

A. Typicality effects.

The explanations included in this section are consistent with the view espoused by Carey. In essence, the argument is that the asymmetries in inductive inference reflect an underlying conceptual structure in which humans are considered the prototypical entity. There are two variants of this general view, as described below.

1. Inductive confidence displays a typicality effect. Asymmetries in inductive confidence, which have been documented in adults, have often been attributed to differences in typicality (Osherson, et al, 1990). On this view, inferences from typical to atypical category examples are stronger than the reverse direction of comparison. This typicality view was endorsed by Carey (1985) to account for the human – non-human-animal asymmetries displayed by young children. Specifically, Carey suggests that young children treat humans as the prototypic animal, and that the asymmetries are just typicality effects.

To bolster this view, one would need independent, converging evidence that young children treat humans as the prototypical animal. It would also be beneficial to provide an account of how differences in typicality lead to asymmetries. With respect to the former, if we allow frequency of exposure to examples of the concept to bias typicality (see Barsalou, 1985, for evidence that frequency of instantiation as a member of the category is the more relevant variable), then humans SHOULD be the prototype, at least for urban children. With respect to the latter, let's look at how the Osherson et al, similarity-coverage model

(SCM) accounts for typicality effects.

True to its name, the model has two components. The first is similarity. The greater the similarity between the base and the target, the greater the confidence that some novel property true of the base will be true of the target. The SCM assumes that similarity is symmetrical (the similarity of a to b is the same as the similarity of b to a) and therefore, asymmetries do NOT arise from the similarity component. Instead, it is the coverage component that leads to typicality effects.

The coverage component works as follows: In addition to considering the similarity of the base and target categories, the model assumes that participants generate examples of the lowest level category that encompasses both the base and target, and then compute their similarity (of this parent category) to the base. Because participants tend to generate typical examples (even if sampling is random, because a body of other work suggests that natural object categories have a typicality structure—see Rosch and Mervis, 1975, Smith, Shoben and Rips (1974) for evidence and Smith and Medin, 1981 for a review), they will be more similar to a typical than to an atypical base. Because inductive confidence is assumed to increase with similarity, typical bases will support stronger inductive inferences since they have greater similarity than atypical bases to the examples generated by the coverage component process.

Although this coverage component accounts for asymmetries at a theoretical level, it is hard to see how it would account for the human asymmetries observed in young children. More specifically, it is difficult to see why humans would have better coverage than other mammals (e.g., dog, wolf), since humans are not especially similar to other mammals. For this reason it may make more sense to conceptualize humans as a prototype in the ideal, rather than similarity sense, and to use something other than the similarity-coverage model to account for the asymmetries.

2. Asymmetries may represent ideal-based prototypicality effects. An alternative view is that asymmetries derive from typicality effects but that these typicality effects are based on ideals rather than central tendency (See Medin and Atran, 2004 for a review). To get on with developing this position, two related assumptions are needed and both of them are drawn from Carey (1985). One is her assumption that young children do not have a distinct naïve biology and that biological induction draws on an undifferentiated module (that is, naïve psychology and naïve biology are not distinct) that is organized primarily by naïve psychology. The second and uncontroversial assumption is that humans are the prototypic psychological entity. If we adopt these assumptions, we can appeal to other research (e.g. Gleitman et al, 1996, Palmer, 1978) suggesting that it is more natural to compare the variant to the standard than vice versa. For example, when faced with a picture of a person, we notice that the picture has a likeness to the person rather than the person having a likeness to the picture. That is, what is relevant is that the picture has some aspects of, or similitude to, the person. Practically speaking, this means that if two entities share certain aspects, it is much more natural to compare the variant to the standard than vice versa. Although this sounds a bit like a “P implies Q” therefore “Q implies P” argument, the claim is that in making an inductive inference, we act as if a comparison in the natural direction is more likely to

involve a shared property than a comparison in the unnatural direction. (See Medin et al, 2003 for a relevance framework that seems generally compatible with this view)

Specifically, the idea is that comparisons go from the target to the base. Hence when a non-human animal is a target and a human is the base, the non-human gets compared to the (prototypic – ideal) human and it is natural to assume that the property is shared. In the reverse direction, when the standard is being compared to the variant, differences become more accessible. This position has some testable consequences. First of all, if humans aren't the prototype (as they may not be in some populations), then asymmetries should disappear. All we need is some independent index of prototypicality to tie things down. Once we have determined a typicality ordering, we should be able to predict a range of asymmetries. In summary, from this perspective, there is nothing special about humans with respect to generating asymmetries---the results follow from the (ideal-based) typicality of humans.

B. Knowledge and experience effects.

The explanations included in this section focus on participants' knowledge about the base and target, focusing mainly on the base. In essence, the argument is that willingness to generalize from a given base to a given target depends on how much one knows about that base. We describe three main variants of this view.

1. Knowledge increases projection. The first notion is that the more a participant knows about some base, the more likely they are to project from it to some target. This idea is endorsed by Kayoko Inagaki and Giyoo Hatano (2001) and fits their observation that children who have experience raising goldfish use goldfish as well as humans as a good base for inductive generalization. The paradigm used by Inagaki and Hatano is somewhat different from that employed by Carey, so one has to be careful about their inter-relationships. Inagaki and Hatano have shown that children who have knowledge of certain biological properties will extend them to near neighbors. In a sense, this amounts to affirming the existence of similarity effects in inductive reasoning. We do not know if we would find asymmetries if a novel property were attributed to goldfish versus, say a turtle. That is, we don't know if an inference from goldfish to turtle would be stronger than an inference from turtle to goldfish. If such an asymmetry were found, it would provide strong support for knowledge of a base driving induction.

There is also some more direct evidence consistent with the idea that knowledge or familiarity with a base increases inductive projection (Atran et al., 2001) and in the limiting case of complete ignorance about the kind in question, it seems plausible. Nonetheless, an element is missing from the argument linking unfamiliarity and uncertainty to a tendency *not* to generalize. It's not obvious that this should be true, especially given classic evidence that experience in the form of discrimination training tends to sharpen generalization gradients.

2. Inductive confidence may be driven by a history of successes and failures of induction (Goodman, 1983, Shipley, 1993). The category of "blue things" doesn't permit much by way of inferences

beyond the fact that they are things that are blue, but the category of “bluejays” is rich in the sense that the members share a large number of properties besides being blue. In this view, both categories and properties (or predicates) become entrenched as result of their history of use, and therefore based on this history, some categories will come to support more inferences than others and some properties (or predicates) will become more widely projectable than others. In the variant that we consider here the focus is once again on the categories and properties of the bases in induction tasks.

Projections are like predictions, and to the extent that they are supported by data, they may be reinforced. Here the idea is that learners have made, seen, or heard previous successful inductive projections from humans to animals but are much less likely to have done so from other animals to animals. Success breeds success so children should be more willing to project a novel property from humans than from some other animal. One straightforward way to obtain evidence bearing on this hypothesis would be to examine the input (this could be written text or speech corpora) to ascertain whether humans are preferred over non-human animals as a base in the comparisons that are available to children. To take this idea one step further, distinguishing between comparisons involving physical (biological) versus psychological properties should permit us to make more detailed predictions. In any event, this view falls into the projectability category given that the history of inductive inference may give rise to greater or lesser willingness to generalize from a particular base.

3. Differences in prior odds or surprisingness of predicates. In virtually all studies of induction, researchers use “blank” properties; that is, properties for which participants have no strong a priori beliefs. This design feature is important because participants’ projection or generalization of the blank property cannot be guided by a simple application of their prior knowledge. Instead, the assumption is that their projection of a blank property must be guided by a genuine inductive inference. There is, however, a wrinkle in this assumption. Although it is usually assumed that blank properties have uniform prior odds, this may not be the case. Specifically, assuming that people are experts on people, they may have lower odds for novel (blank) properties attributed to people than for the same properties attributed to some other base. That is, participants should be more surprised to hear that humans have an omentum inside (where omentum is a novel property) than to hear that a raccoon does.

Two factors (at least) appear to support this intuition. The first involves a “lack of knowledge inference” (Collins and Michalski, 1986). This would run something like the following: “If humans had omenta inside then I would have heard about it, so it’s unlikely that they do.” The second and related idea is that there may be a bias for assuming that novel properties are associated with unfamiliar categories (Testa, 1975). If so, the asymmetries will follow directly for two reasons. First, the priors are lower for humans versus other mammals (by hypothesis) and even if the premise has no effect whatsoever, an inference from humans to (other) mammals should be higher than and inference from mammals to humans. Second, it’s generally assumed that the more surprising a premise is, the greater its effect. We suspect that it is more surprising to hear that humans have an omentum inside than to hear that some other

mammal does.

C. Distinctive categories and features of the target.

The previous two accounts have focused primarily on the ways in which properties of the bases lead to asymmetric patterns of induction. Our third class of explanations, which brings into sharper focus the contribution of the targets, represents a strong departure from prior work. Specifically, it proposes that when bases and targets are compared, inferences are limited by distinctive properties of targets, relative to the base.

1. Different patterns of category label activation. This explanation draws upon the observation that for most individuals living in the US, humans have a dual status: they are considered members of the category animal (as distinct from plants) and are also members of the category human (as distinct from non-human animals)¹. In contrast, non-human animals do not have this dual status. Fleshing out this idea (no pun intended), it is important to point out that the inductive strength of a given target is diminished by its distinctive categories. These observations are relevant to inductive inference in the following way: When a human is the base and a non-human animal is a target, the more inclusive category “animal” gets activated and this tends to prime the broad sense of animal that includes humans. In contrast, when a nonhuman animal is the base and a human is the target, the distinctive category of “human” gets activated in the target and, as we noted above, this diminishes inductive confidence to non-human animals.

What might count as evidence for this proposal? First, if this proposal is correct, it should be evident in participants’ justifications for their inductions. For example, participants should mention the distinctive feature of humans (e.g., “people aren’t animals”) as a justification much more frequently when a human is a target than when it is a base for induction. A second piece of evidence involves cross-cultural comparisons. We know that there are cultures and languages in which the contrast between human and non-human animals is more clear-cut than in urban and suburban US English-speaking communities. For example, in Indonesia, these categories are kept quite distinct and the Indonesian names are mutually exclusive (Anggoro, Waxman & Medin, 2005). That is, the Indonesian term *animal* cannot be applied to humans. In such cultural contexts, we would expect to find less inductive generalization between humans and non-human animals and, as a consequence, the asymmetries should be diminished. See Anggoro et al. (2005) for evidence that this is the case.

2. Different patterns of feature or property activation. On this view, the asymmetry of induction is no more and no less than the asymmetry of similarity (comparisons). The argument here is parallel to the previous one but the focus is on constituent properties, rather than categories. As we have suggested, in urban and suburban US communities, humans certainly can be considered ANIMALS, but they are considered to be a special case, as animals with more distinctive features than other animals (e.g., humans are special because they talk, have beliefs, construct buildings, cook their food, etc). The core idea, then, is that inductive inferences from non-human animals to humans should be diminished whenever

the distinctive features of humans are activated, and such features are activated more when humans are the target than when they are the base. By logic, the same would hold for any animal, human or non-human, whenever participants attribute more distinctive features to the target than to the base. In other words, there is nothing special about humans (save their distinctive features). If this is the case, then for any distinctive category (e.g., goldfish for Inagaki and Hatano's goldfish caretakers), generalizing a novel property (e.g., from goldfish to frogs) should be stronger than generalizations in the opposite direction (e.g., from frogs to goldfish). This would follow if the child goldfish caretakers knew more distinctive features about goldfish than about frogs.²

3. Relevance and distinctive features. This position can be seen as a modest variation on either or both of the first two explanations in this class. The idea is that asymmetries of induction are not based on overall similarity but rather on matches and mismatches of a subset of features that appear to be most relevant in the experimental context. The motivating principle is that similarity comparisons are mainly focused on how things are similar rather than how similar they are (Medin, Goldstone & Gentner, 1993).

This position is compatible with the two others in this section, as long as one comes down on the side of distinctive features of the target being more important than those of the base. In general, one would expect that if a salient property of some target is also a feature of the base, then inductive confidence should increase (and if it isn't, confidence should decrease). For example, an argument from skunks to zebras should be stronger than from zebras to skunks because the most distinctive property of zebras (being striped) is also shared by skunks (note that the reverse does not hold). This seems like a pretty strong (and perhaps incorrect) prediction in that a hypothesis-based model of induction (McDonald, Samuels, & Rispoli, 1986) would clearly make the opposite prediction (skunk should give rise to two hypotheses (odor, stripes) and zebras only one (stripes)). To evaluate this position it would be important to obtain an independent measure of distinctive features.

Of course, relevant features of both the base and the target might well matter. There is good evidence that there are at least two different strategies for comprehending conceptual combinations (see Wisniewski, 1997, for a review) and that they depend on the similarity (read, alignability) of the base and target. In the case of induction the idea would be that for a base and target within the same lifeform, the base and target are aligned and induction is therefore limited by distinctive features of the base. For bases and targets that are from different lifeforms (and for humans versus other animals), attention would shift to salient properties and categories of the target. This idea is certainly testable.

Summary. We will stop with these three main classes of explanation for asymmetries in induction, though there certainly are others.³ It should be clear that human – non-human animal asymmetries reported by Carey do not unambiguously implicate any one theoretical account over others.

III. Broadening the Analysis of Asymmetries

We suggest that the most straightforward means of obtaining some leverage in choosing among the alternative accounts is to look for and evaluate additional asymmetries in inductive reasoning. For example, if asymmetries are based on relative amounts of knowledge about bases and targets, then they should appear wherever there are demonstrable differences in knowledge, independent of whether humans are involved in the comparison. It seems uncontroversial that children being raised in urban and suburban communities in the US know more about mammals than they do about plants, so one should expect mammal-plant asymmetries in which mammal to plant inductions are stronger than plant to mammal inductions. In this section we will look at data from an expanded analysis of asymmetries. First, however, there are a few methodological and conceptual issues that we need to review.

A. Methodological cautions. Although a number of studies have examined the generality of Carey's results across populations and procedures (e.g. Atran, et al., 2001, Ross, et al., 2003), these have not followed Carey's procedure in every detail. First, although in Carey's design, the same kinds appear as both base and target⁴, in most recent instantiations of the category-based induction task, this is not always the case. This difference in design may have consequences on patterns of performance: a human – non-human animal asymmetry may be a consequence of differences in the similarity of the base vs. target animals to humans (e.g., if the animal used as a target bears a greater similarity to humans than does the animal used as a base). Ideally we would have studies where we equate for base and target similarity and have the same items as both base and target, at least across participants but ideally within participants. In our ongoing work, we have designed experiments in this way, but we do not yet have sufficient data to report on. Still, because we find that most of the asymmetries of interest consistently hold across studies, across populations, and across a variety of stimulus materials, the patterns of asymmetries reported below are sufficiently strong to support our main empirical claims.

There is a second difference between Carey's design and the more recent versions. In Carey's original version of the task, 4- and 6-year olds were trained and tested on separate days, though this apparently was not done for older children and adults. On the first day, children were taught about "spleen" or "omentum" within the context of a review/probe of other, presumably more entrenched properties like "has lungs" or "has a heart". Moreover, the teaching was fairly elaborate. Children saw a diagram that indicated the location of the omentum, but there was no explicit appeal to other kinds that might or might not have that property. They were then tested on the novel property (e.g., "Does a dog have an omentum?") within the context of the more familiar properties (e.g., "Does a dog have a heart?"). We suspect that as a result, children in Carey's study may have treated the probes regarding the novel property more as a test of knowledge than as an invitation to make an inductive inference. This raises the interesting possibility that there is a fundamental difference between spontaneous inference versus prompted inference. We note that Inagaki and Hatano's work may be conceptually related to Carey's. In their work, they provided a group of young children with extensive experience with goldfish, and later examined children's spontaneous inferences rather than explicit inferences (Inagaki and Hatano, 2002). A third difference between Carey's

design and the more recent versions pertains to the inclusion criteria. Carey restricted her analyses to those children who correctly attributed the novel property (e.g., omentum or spleen) to the base on which it was introduced.

Finally, and perhaps most interestingly, there are differences in the types of participants included in the original Carey task as compared to more recent work. Several recent studies have moved beyond the urban and suburban populations to examine how children raised in rural environments, who enjoy a more intimate contact with nature, reason about the biological world. We have found that children raised in rural environments, sometimes appeal to ‘ecological reasoning’, focusing on potential mechanisms of transmission of a property from one biological kind to another. For example, when they told that “bees have sacra inside them”, these children often infer that bears must also have sacra, and they justify their responses by mentioning that sacra could be transmitted when a bee stings a bear or when a bear eats honey (Ross et al, 2003).⁵ The problem, however, is that when ecological reasoning is combined with procedures where base and target items are not counter-balanced, the problems loom larger. For example, “bee” may be a base that supports ecological reasoning, but “fly” might not. The bottom line is that skeptics might want to defer judgments about the claims we make in the next section. Nonetheless, because we find that most of the asymmetries consistently hold across studies, across populations, and across stimulus materials, we are confident that the patterns of asymmetries reported below are strong enough to build upon.

B. Empirical asymmetries in induction.

In this section we focus on the patterns of asymmetries produced by children and adults from a range of communities, in urban and rural USA settings and rural Mexico. The primary data of interest, taken from Ross et al (2003) and Atran et al, (2001) data⁶, appear in Table 1. We refer the reader to the original articles for methodological details, but include the tables here to permit readers to check our assertions about patterns of asymmetries. A number of the individual comparisons no doubt fall short of statistical reliability, but the trends are generally quite robust across populations.

Finding 1. Human to non-human mammal asymmetries. As we noted before, studies on this topic have tended to use different mammals as base and target; humans are essentially the only kind that have appeared as both base and target. However, despite this nontrivial concern, it nonetheless appears as if the human to non-human mammal asymmetry is fairly robust. Although the evidence for this asymmetry is stronger in Ross et al, (2003), than in Atran et al, (2001)⁷, the key finding is that this asymmetry never reverses. In short, although there may be cultural factors moderating the relationship, we suggest that the human –non-human mammal asymmetry is strong.

Finding 2. Human to mammal vs mammal to mammal asymmetries. In Carey’s view, young children lack a theory of folk-biology and rely instead on a theory of folk-psychology in which humans serve as the prototype. If this were the case, then inferences from humans to other non-human mammals should be stronger than from one non-human mammal (in our case: wolf, dog, peccary) to another. To the best of our

knowledge, Carey is the only researcher who has found this pattern of results. The results of Table 1 reveal the opposite pattern. Although the absolute level of human to mammal inferences ranges across groups, the (other) mammal to mammal inferences are always higher. This pattern of reverse typicality, apparent at each age and in all of our populations, strongly undermines the argument that humans serve as the prototypic animal for either young children or adults. On the contrary, the reverse typicality pattern evident in our data is consistent with the notion that humans are not the prototypical animal. Instead, the data suggest that throughout development, humans retain a somewhat atypical status within the animal category, and that it is something of a conceptual achievement for young children to realize that humans are, in fact, also animals.

Finding 3. Human to plant asymmetries. In contrast to the previous result, here the asymmetries appear to vary across populations, and as a result, no clear overall pattern emerges. Urban children show a strong asymmetry favoring humans over plants as a base, rural children show little generalization and no asymmetry, and Menominee children show a reverse asymmetry. One possibility is that this pattern is a consequence of a shift to increasing use of ecological/causal reasoning across groups, but additional research is required to tie this down.

Finding 4. Mammal to plant asymmetries. A review of Table 1 reveals that the rate of generalization between mammals and plants is limited. However, we do see a consistent trend in favor of reverse asymmetries, with more inductive inferences from plants to mammals than from mammals to plants. This reverse asymmetry pattern holds for all but one (7-9 year old urban children show no difference) of the age and population samples.

Finding 5. Human to insect asymmetries. This set of comparisons is currently limited because the data are quite variable and in most cases, include only one target insect (fly). In the Ross et al data, asymmetries seem to be present for the urban children, but they are weak and in a few cases reversed for the rural and Menominee children. In the Atran et al data, the asymmetries diminish with age and are sharply reversed in adults. This is consistent with the possibility that ecological reasoning increases with presumed intimacy of contact with nature. For more recent data with native and exotic bases intermixed, where we appear to have eliminated ecological reasoning, there is a pretty consistent asymmetry. This asymmetry appears to be larger for the exotic set of bases than the native set (bee versus tarantula as the base) for all but the Yukatek children.

Finding 6. Mammal to insect asymmetries. There's no reliable trend across populations and there are no really notable differences. Comparing human-insect with mammal-insect the main difference is that insect to mammal inferences are higher. Note also that humans to insects is consistently lower than mammals to insects.

Finding 7. Insect to plant asymmetries. In the Ross et al. (2003) data, where the trends are a bit variable, there doesn't seem to be any clear overall asymmetry. For the newer data with mixed native and exotic bases there is a consistent pattern of reverse asymmetries, though in several cases the effects are

small because there is little overall generalization between insects and plants (especially compared to the generalization seen in Ross et al., 2003).

Other asymmetries. A careful examination of other asymmetries in Table 1 reveals no consistent pattern across populations or age groups. As we noted earlier, we are suspicious of comparisons involving insects because base (bee) and target (fly) were not counter-balanced and may differ substantially in the availability of children's ecological/relational knowledge. In some follow-up studies we have unwittingly more or less eliminated ecological reasoning and these induction data provide still stronger asymmetry constraints.

IV. What is the underlying basis for the asymmetries in induction?

Let's turn now to the implications of these results for choosing among the alternative theoretical accounts of asymmetries in induction. With respect to potential for asymmetries, we have an overall ordering of human, plant > insect, mammal (and insect > mammal when ecological reasoning is eliminated). That is, asymmetries are observed favoring the base when the premise category is higher on this ordering than the conclusion or target category. In addition, we consistently find that mammal to mammal inferences are stronger than human to mammal inferences. As we'll see, these orderings place strong constraints on theories of induction.

- A. Typicality effects. Neither variant of typicality theory can explain these patterns of asymmetries. If humans are the prototype, then inferences from human to (nonhuman) mammal should be stronger than those from mammal to mammal, and this was not the case. Instead, the ordering with respect to asymmetries appears to conform to a reverse typicality gradient, suggesting that mammals are the most prototypical biological entity and humans the least. If typicality alone cannot account for the underlying basis for asymmetries, perhaps an account in which typicality is considered in conjunction with similarity would fare better. The proposal would be that humans are the ideal animal but, at the same time, they are not especially similar to other mammals. But this account seems to add a new parameter for every data point it attempts to explain. Moreover, it provides no account of the plant to mammal asymmetry (and the simultaneous absence of a human to plant asymmetry). In short, the evidence from the broadened range of observations undermines the view that typicality effects can account for the human-animal asymmetries.
- B. Knowledge/familiarity effects. The idea that inductive strength increases with knowledge about the base is also challenged by our observations. It is plausible to suggest that children know more about humans than other biological kinds, but it is extremely unlikely that they know more about plants than insects and more about insects than mammals. Therefore, none of the knowledge-based accounts would predict the observed ordering.

C. Distinctive features and categories. The final class of explanations provides a nice account of the Ross et al data. One can explain most of the results by claiming that humans are the most distinctive (either in terms of their features or in terms of the category itself), followed closely by plants, followed by insects and mammals. Specifically, according to this account, when a base premise is given and a target suggested, participants compare the target to the base. Distinctive features or category memberships of the target reduce inductive confidence more than distinctive features of the base. So if the ordering on number of distinctive features is: humans equal to or slightly higher than plants and both much higher than insects or mammals, the following implications are straightforward:

1. Human to mammal should be stronger than mammal to human
2. Human to mammal should be weaker than other mammal to mammal
3. Plant to mammal should be stronger than mammal to plant
4. Plant to insect should be stronger than insect to plant
5. Human to insect should be stronger than insect to human
6. No clear human/plant asymmetry (trends should favor human to plant) or insect/mammal asymmetries (trends should favor insect to mammal).

This pattern conforms closely with the data. So far this looks very good for the distinctive features/categories position. Still, because it would be good to get some converging evidence to bolster this position, in the next section we present justification data.

D. Justification data. Children's justifications for their responses in the inductive inference task provide converging evidence for the distinctive features/categories view. If the distinctive features/categories position is correct, then failures to generalize from a base to a target should be accompanied by justifications focusing on the distinctive features and categories of targets rather than bases. For example, a child should justify her failure to generalize from humans to raccoons by saying that raccoons have a tail or raccoons are animals. In contrast, a failure to generalize from squirrels to humans should bring with it comments that humans are not animals or that humans have two legs (and squirrels four).

At this point in our investigation, we only have extensive justifications from rural majority culture children in Ross et al, 2003 (with a few additional justifications from Menominee children). However, these justifications provide strong support the distinctive features/categories account. The pattern about to be described holds for human-mammal and mammal-insect. Both groups of rural children show the same qualitative trends (though Menominee children don't spontaneously offer justifications with sufficient frequency to justify reporting them). Furthermore, the younger and older rural majority culture children show the same pattern, though the younger children focus relatively less on distinctive categories and relatively more on

distinctive features.

First consider the extent to which justifications focus on common versus distinctive features and categories. For human to mammal inferences, the ratio of common to distinctive features is 0.64 for young children and 0.68 for older children. The corresponding ratio for mammal to human is 0.25 and 0.19 for the younger and older ages, respectively. The ratio of common to distinctive categories for human to mammal is 0.23 and 0.77 for the younger and older ages, respectively. When the comparison shifts to mammal to human, these proportions drop to 0.12 and 0.43, respectively. For mammal to insect inferences, exactly the same qualitative trends are observed in both age groups. For example, the young children show a common to distinctive features ratio of 0.22 in the insect to mammal case and a ratio of 0.00 in the reverse direction. The corresponding proportions for common to distinctive categories are 0.38 and 0.10, respectively. This pattern is noticeably absent for the human-insect justification data where the young children only give distinctive features and distinctive categories regardless of the direction of comparison. Older children show a very slight tendency to have more common features and categories in the human to insect direction.

We can also examine justifications to see if they focus on distinctive features/categories of the base versus those of the target. The answer is quite clear: the justifications overwhelmingly (95% or better) focus on distinctive features and categories of the target object.⁸

E. Summary. Overall, the induction data coupled with the justifications provide detailed support for the distinctive features/categories account. Again, the claim is that humans have the most distinctive features/categories, followed by plants, then *insects and* finally mammals. Note that other explanations for the pattern of asymmetries can, in principle, describe some of these trends, but they do only by imposing some implausible assumptions. For example, consider the alternative view that differences in knowledge underlie the asymmetries. To capture the full range of data, this view would have to make the dubious claim that children have more knowledge about plants than mammals. The same difficulty holds for accounts that appeal to ideals, histories of induction, or prior odds. Only the distinctive features/categories account can account for the ordering of asymmetries that emerge in the data.

V. Implications.

What are the important implications of these findings? First, we suggest that the human-animal asymmetries that have been observed in category-based induction tasks do not bear on the status of bases as being either more familiar or better examples than targets. More generally, we suggest that accounts which focus primarily on the base are theoretically and empirically insufficient, for they cannot account for the range of asymmetries that emerge in children's

inductive inferences about the biological world.

We have argued that the most promising explanation for the observed asymmetries in inductive inference is one that focuses on the target, the base, and the relation between them. More specifically, we suggest that such an account must incorporate not only the categories and features of the target, but also the distinctive categories and features that emerge when the target is compared with the base. This account correctly describes the full pattern of asymmetries, the observation that nonhuman mammal to nonhuman mammal inferences are stronger than human to nonhuman mammal inferences, and the patterns of justification provided by children in category-based induction tasks. When it comes to categories, it appears that mammals are animals and little else, other than a specific kind of animal. Specific plants may be both plants (as a contrast with animals) and living things. Humans both contrast with and constitute animals. These categorical distinctions are paralleled by corresponding differences in distinctive features as a function of comparison direction.

A review of the literature reveals that, as is often the case, the seeds of this account are evident in previous work. For example, the idea that distinctive features of the target may be weighted differently from distinctive features of the base is not new (e.g. Tversky, 1977), though in the domain of similarity judgments the claim has been that distinctive features of the base get more weight than distinctive features of the target. As another example, we point out that at least one earlier model of induction, the similarity-coverage model (Osherson et al., 1990), also drew upon categories and relations among categories to account for patterns of inductive inference. However, the account that we have offered here, the distinctive categories/features account, goes beyond its progenitors in (at least) three ways. First, it broadens the range of categories that are relevant for induction. Second, it considers the relation between the coordinate categories in the target and base, rather than relying on inclusion relations between bases and their superordinates. Third, our account highlights a role for language. We suggest that category names may play a more instrumental role in biological induction than previous accounts would suggest (see Waxman, 2005 for an amplification of this point and Angorro, Waxman and Medin, 2005 for an example involving Indonesian). Waxman (2005) has described the various inclusion and contrast classes associated with different biological kinds and outlined some of the ways in which these categories may influence children's learning about biological properties. Her results and ours suggest that the reasoning task provides a window on how different categories are conceptualized.⁹

In summary, we think that asymmetries in induction are theoretically significant, despite the fact that there are many contending, plausible accounts of human-animal asymmetries. As is often the case, broadening the set of empirical constraints allows us to gain leverage on these alternative theoretical explanations of asymmetries. The main theoretical implication of our work

is that researchers should shift attention away from the inductive base alone in favor of examining the relationship between target and base. This relationship gives rise to distinctive features and categories and suggests new avenues for research on the interplay between categories and inductive reasoning.

- Anggoro, F, Waxman, S., & Medin, D. (2005). The Effect of Naming Practices on Children's Understanding of Living Things. Paper presented at Cognitive Science 2005, Stresa, Italy.
- Atran, S., Medin, D., Lynch, E., Vapnarsky, V., Ucan Ek', E. & Sousa, P. (2001). Folkbiology doesn't come from folkpsychology: Evidence from Yukatek Maya in cross-cultural perspective. Journal of Cognition and Culture 1:3-42.
- Baillargeon, R. (1994). How do infants learn about the physical world? Current Directions in Psychological Science, 3(5), 133-140.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. Journal of Experimental Psychology: Learning, Memory, and Cognition, 11, 629-654.
- Carey, S. (1985). Conceptual change in childhood. Cambridge, MA: MIT Press.
- Collins, A., & Michalski, R. (1989). The logic of plausible reasoning: A core theory. *Cognitive Science*, 13, 1--49
- Gelman, S.A., and Markman, E.M. (1986). Categories and induction in young children. *Cognition*, 23: 183–209.
- Gleitman, L.R., Gleitman, H., Miller, C., * & Ostrin, R. (1996). Similar, and similar concepts. *Cognition*, 58 (3), 321-376.
- Goodman, N. (1983). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Gutheil, G., Bloom, P., Valderrama, N., & Freedman, R. (2004). The role of historical intuitions in children's and adult's naming of artifacts. *Cognition*, 91, 23–42.
- Hatano, G & Inagaki, K. (1994). Young children's naive theory of biology. Cognition, 50, 171-188.
- Hatano, G. & Inagaki, K. (2003) The formation of culture in mind : A sociocultural approach to cognitive development. In J. Mehler, S. Carey & L.L. Bonatti (Eds.) Cognitive development and conceptual change. Cambridge, MA :MIT Press.
- Inagaki, K. & Hatano, G. (2001). Children's understanding of mind-body relationships. In M. Siegler and C. Peterson (Eds.), *Children's Understanding of Biology and Health*. Cambridge, UK: Cambridge University Press.
- Inagaki, K., & Hatano, G. (2002). *Young children's naive thinking about the biological world*. New York: Psychology Press.

- Johnson, S. C., Slaughter, V., & Carey, S. (1998). Whose gaze would infants follow: The elicitation of gaze following in 12-month-olds. *Developmental Science*, *1*, 233–238.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237-251.
- Leslie, A. M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception*, *13*(3), 287-305.
- Markman, A. B. and Gentner, D. (1997) The Effects of Alignability on Memory. *Psychological Science* *8*(4):pp. 363-367.
- McDonald, J., Samuels, M., & Rispoli, J. (1996). A hypothesis-assessment model of categorical argument strength. *Cognition*, *59*, 199-217.
- Medin, D., & Atran, S. (2004). The Native Mind: Biological categorization and reasoning in development and across cultures. *Psychological Review*, *111*(4), 960-983.
- 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., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, *100*, 254-278.
- Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category Based Induction. *Psychological Review*, *97*, 185-200.
- Palmer, S. E. (1978) Fundamental aspects of cognitive representation. In E. Rosch & B. L. Lloyd (Eds.), *Cognition and categorization* (pp. 259-302). Hillsdale, N.J.: Erlbaum.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573-605.
- Ross, N., Medin, D., Coley, J.D., & Atran, S. (2003). Cultural and experiential differences in the development of folkbiological induction.
- Shipley, E. F. (1993). Categories, hierarchies, and induction. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 30, pp. 265–301). New York: Academic.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, E. E., Shoben, E. J., & Rips, L. J. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, *81*(3), 214-241.
- Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*(1), 29-56.

Testa T.J. (1975). Effects of similarity of location and temporal intensity pattern of conditioned and unconditioned stimuli on the acquisition of conditioned suppression in rats. *Journal of Experimental Psychology: Animal Behavioral Processes*, 1: 114–21.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.

Waxman, S. R., Senghas, A., & Benveniste, S. (1997). A cross-linguistic examination of the noun-category bias: Its existence and specificity in French- and Spanish-speaking preschool-aged children. *Cognitive Psychology*, 43, 183-218.

Waxman, S. R. (2005). Why is the concept "Living Thing" so elusive? Concepts, languages, and the development of folkbiology. 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.

Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, 62, 1070-1090.

Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, 4(2), 167-183.

Table 1. Ross et al, 2003 Asymmetries

Ross et al. - Asymmetries									
Population	City			Rural			Menominee		
Age	7-9	9-11	5-7	7-9	9-11	5-7	7-9	9-11	
Human-Mammal	0.70	0.73	0.52	0.58	0.80	0.58	0.70	0.72	
Mammal-Human	0.33	0.35	0.24	0.16	0.47	0.42	0.47	0.52	
Mammal-Mammal	0.96	0.90	0.76	0.78	0.92	0.75	0.82	0.88	
Human-Insect	0.47	0.35	0.07	0.34	0.63	0.46	0.55	0.48	
Insect-Human	0.31	0.08	0.28	0.20	0.53	0.63	0.35	0.38	
Human-Plant	0.50	0.27	0.16	0.17	0.30	0.30	0.22	0.34	
Plant-Human	0.13	0.12	0.10	0.12	0.30	0.42	0.41	0.38	
Mammal-Insect	0.50	0.38	0.38	0.28	0.47	0.57	0.56	0.52	
Insect-Mammal	0.50	0.56	0.38	0.25	0.56	0.56	0.56	0.56	
Mammal-Plant	0.32	0.12	0.14	0.14	0.22	0.18	0.33	0.24	
Plant-Mammal	0.32	0.38	0.24	0.25	0.42	0.50	0.42	0.46	
Insect-Plant	0.40	0.44	0.26	0.30	0.58	0.46	0.40	0.42	
Plant-Insect	0.31	0.62	0.24	0.34	0.33	0.29	0.41	0.43	

¹ A further complication is that animal sometimes is treated as equivalent to “beast” or mammal or even quadruped.

² Note that, as in the assumption for category activation above, we are attributing to this view the claim that the distinctive features of the target affect induction more than do distinctive features of the base. This is the opposite of Tversky’s (1977) claim that it is the distinctive features of the base that receive more weight in asymmetries (in similarity judgment). This point is worth noting, but it is also worth noting that the induction paradigm is considerably different from a similarity judgment paradigm.

³ Here are two additional possibilities.

1. Causal associations produce the asymmetry. Many causal relations have directional implications for induction. If Grass has Enzyme x, we are more sure that Cows have it than we are sure that Grass has Enzyme Y given that Cows have it (Medin et al, 2003). For this argument to work, the causal associations should be more likely in the human to (other) animal case than in the (other) animal to human case. One major problem with this explanation is that among populations of children where we see this sort of reasoning (see Ross, et al, 2003 for details) the associations we see involving humans tend to have humans as an effect rather than a cause (e.g. bees stinging humans). If this position has value, it may be in explaining why human-animal asymmetries are *not* found in some cases.

2. Asymmetries are mediated by size differences between base and target. There are two motivations for this idea. One is that both Itza’ Maya and tree experts sometimes explicitly refer to relative size on induction tasks and argue that going from large to small (tree species) is more likely than going from small to large (Medin et al, 1997). The other impetus for worrying about size is that children may worry about it. Hatano and Inagaki (2003) observed that children sometimes deny that ants have hearts on grounds that they are too small to contain a heart. Note, however, that this example cuts the other way in that a property true of ants will fit into human dimensions but the converse is not true. Still, it may be that, in induction, there is a big-to-small bias

⁴ Humans are the exception to this generalization. Although humans have consistently been presented as both bases and targets, the specific other mammal, insect or plant that has been a base is rarely also a target.

⁵ We should also note that in the causal reasoning literature, there are some well-documented asymmetries in which it is easier to reason from cause to effect than vice versa (see Kahneman and Tversky, 1974 for asymmetries in causal judgment). In principle, at least some of the explanations that we have been considering might apply here as well.)

⁶ We also have considerable data from a study which follows up and expands on these studies with additional populations including Yukatek Maya children, majority culture and Menominee children from rural Wisconsin, and Chicago area urban and suburban children. In most cases we also have adult data but not always in sufficient quantity to be reported (yet). Although we do not present these data, they consistently support and expand the claims we make here from published data.

⁷ The asymmetry was clear for peccary as a base but essentially absent with dogs as a base. This provides some support for knowledge or familiarity affecting induction. Nonetheless, familiarity will not account for the other asymmetries even within the Atran et al, data.

⁸ Importantly, this is true both for what one might call “alignable features” (Markman and

Gentner, 1997) that reflect different values on a common dimension (e.g. number of legs) and for nonalignable features (e.g. presence versus absence of a tail). This not only points to the robustness of the focus on the target object but also undermines the idea that children might focus on the base for comparisons of alignable targets and the target for nonalignable targets. A further problem for this two-process account is that this same focus on the target category also holds for the mammal-mammal probes.

⁹ It may be possible to translate distinctive features/categories back into typicality language - in this view mammals are the typical biological organism, followed by *insects and then* plants and humans. Next you have to claim that comparing the variant to the standard yields a stronger induction than vice versa and that it's better (in terms of induction likelihood) to have distinctive properties in the base than in the target. Unless there's some obvious novel prediction that we're missing, this seems like a notational variation on the distinctive categories/features idea.