Available online at www.sciencedirect.com
SCIENCE DIRECT
COGNITION
Cognition 91 (2004) 113–153
www.elsevier.com/locate/COGNIT

Category coherence and category-based property induction

Bob Rehder[a,*], Reid Hastie[b]

[a]Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA
[b]Center for Decision Research, Graduate School of Business, University of Chicago, Chicago, IL 60637, USA

Received 17 April 2002; revised 6 June 2003; accepted 23 July 2003

Abstract

One important property of human object categories is that they define the sets of exemplars to which newly observed properties are generalized. We manipulated the causal knowledge associated with novel categories and assessed the resulting strength of property inductions. We found that the theoretical coherence afforded to a category by inter-feature causal relationships strengthened inductive projections. However, this effect depended on the degree to which the exemplar with the to-be-projected predicate manifested or satisfied its category's causal laws. That is, the coherence that supports inductive generalizations is a property of individual *category members* rather than *categories*. Moreover, we found that an exemplar's coherence was mediated by its degree of category membership. These results were obtained across a variety of causal network topologies and kinds of categories, including biological kinds, non-living natural kinds, and artifacts.

Keywords: Categories; Classification; Coherence; Category-based property induction

1. Introduction

Why do humans categorize the world the way they do and not otherwise? Intuitions inform us that there is something natural and sensible about categories like *birds*, *apples*, and *diamonds*, but something unnatural and strange about categories like *weighs less than a ton*, *striped with more than one leg*, and *small blue triangles*. Indeed, the latter examples do not seem to be categories at all. These intuitions are not based on mere experience with the categories, as *unicorns* and *death stars* seem like good categories even though they

* Corresponding author.
E-mail address: bob.rehder@nyu.edu (B. Rehder).

doi:10.1016/S0010-0277(03)00167-7

don't exist, and we experience objects that weigh less than a ton everyday. Nor are they based on the presence of a single linguistic term denoting them: *New World Squirrel Monkey* is a separate species and *Apple Titanium Laptop Computer* is a common everyday object.

Psychological research has enumerated some of the ways in which good or natural categories differ from arbitrary collections of objects. When asked to identify a new object, people will typically respond with a category label rather than an arbitrary class descriptor (Brown, 1958). For example, when a questioner points at a koala and asks "What's that?", people are likely to respond with "koala" but never with "weighs less than a ton". Perhaps most important, good categories also support confident inductive projections. Upon learning that a category member possesses a novel property, a category supports an inductive projection to the extent that the observer concludes that many or all of the members of the category also possess the new property (Goodman, 1983; Kornblith, 1993; Shipley, 1993). For example, learning that one koala has a three-chambered heart would provide support for the conclusion that all koalas have three-chambered hearts, but no support at all for the conclusion that all things less than a ton have three-chambered hearts. This is the case even though the koala is a thing that weighs less than a ton, as much as it is a koala.

Properties of categories such as preferential classification and property induction are two important ways in which humans use conceptual systems in order to "go beyond the information given", to use their past experience to respond effectively to new situations. Classification allows individuals to reason from the general to the specific, by allowing unobserved features to be inferred on the basis of category membership. Property induction supports reasoning from the specific to the general by allowing observed features to be projected to the category as a whole and then to other individual category members. Establishing the origins of the human categories that support preferential classification and strong inductive projections has been and remains an important practical and theoretical problem (Wisniewski, 2002).

Perhaps the simplest answer to the question of the origins of human categories, and the one that has dominated categorization research, is that people learn categories directly from what they observe. On this account, categories correspond to the information-rich clusters of features that occur naturally in the environment. This *family resemblance view* has been described as one where boundaries between categories are those that "cut nature at its joints" (Rosch & Mervis, 1975; see also Malt, 1995). The family resemblance view explains not only why koalas are classified as koalas and not "mammals" or "brown koalas" (Corter & Gluck, 1992; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976), but also why they are not classified as "weighs less than a ton". Objects that weigh less than a ton do not exhibit a family resemblance system of information-rich correlated features that define a natural category.

More recently the family resemblance approach has been elaborated by a *theory-based view* of categorical structure. Instead of observed clusters of features, the theory-based view emphasizes the importance of *category coherence*, the extent to which the features of a category "make sense" or "go together" in light of a person's prior knowledge (Murphy & Medin, 1985). The importance of category coherence has been established in a variety of category construction and learning tasks. For example, research participants asked to sort multi-dimensional stimuli exhibiting a family resemblance structure group those items into the family resemblance categories only when the categories' features can be related to a common theme on the basis of participants' prior knowledge. In the absence of such knowledge, objects are sorted on the basis of a single dimension (Ahn & Medin, 1992; Kaplan & Murphy, 1999; Medin, Wattenmaker, & Hampson, 1987; Regehr & Brooks, 1995; Smith, 1981; Spalding & Murphy, 1996). In addition, the rapid acquisition of new concepts that occurs throughout our lives does not seem to be reflected in the slow, labor-intensive process that characterizes learning of experimental categories employing meaningless stimuli (Murphy, 2002). Instead, drastically accelerated category learning occurs only when the to-be-learned category is coherent in light of prior knowledge (Heit & Bott, 2000; Kaplan & Murphy, 2000; Murphy & Allopenna, 1994; Rehder & Ross, 2001; Wattenmaker, Dewey, Murphy, & Medin, 1986). Considered more generally, prior knowledge has also been found to have effects on classification behavior itself, both when prior knowledge in long-term memory is prompted (Heit & Bott, 2000; Rehder & Ross, 2001; Wisniewski, 1995) and when it is taught during an experimental session (Ahn, 1998; Ahn, Kim, Lassaline, & Dennis, 2000; Rehder, 2003, in press; Rehder & Hastie, 2001; Sloman, Love, & Ahn, 1998; Waldmann, Holyoak, & Fratianne, 1995).

In contrast to research on category construction and learning however, less is known about the effects of category coherence on the inductive projection of novel properties (for reviews see Heit, 2000; Rips, 2001). However, in one suggestive study, Gelman (1988) (see also Gelman & O'Reilly, 1988) found that second graders, but not preschoolers, were more likely to generalize new properties of natural kinds than of artifacts. She attributed this result to folk theories about the coherence of natural kinds acquired by the older children. On this account, the coherence provided by emerging theories of biological kinds led the older children to expect such kinds to be more structured and constrained, and hence more homogenous. Novel properties of natural kinds are generalized more strongly because this expectation of homogeneity extends to new properties in addition to existing ones (see also Rips, 2001; Shipley, 1993).

Although suggestive, these studies are correlational in nature, and thus leave unanswered whether the changing pattern of property inductions in children is due to the influence of developing theories, or to factors that affect both theories and induction (e.g. the development of general cognitive mechanisms; Heit, 2000). In contrast, in the current study we adopt a strictly experimental approach: we manipulate category coherence by manipulating the knowledge associated with novel, experimental categories and ascertain the effects of those manipulations on the resulting pattern of property inductions.

We consider three specific hypotheses regarding the nature of the influence of category coherence on property induction. Each hypothesis makes different predictions regarding how the structure, or topology, of the knowledge that links category features lends the category coherence and promotes the projection of novel properties. According to the first hypothesis, people assume that many categories (especially biological kinds) possess an underlying *essence* that makes an object the kind of thing it is (Rips, 2001 provides a perceptive analysis of the notion of necessity in categories). The essence is presumed to occur in all members of a category and to generate many of the properties of natural kinds that people observe (Atran, 1990; Coley, Medin, & Atran, 1997; Gelman, Coley,

Table 1
Features, common-cause causal relationships, and blank properties for Lake Victoria Shrimp

	Features
F_1	High amounts of ACh neurotransmitter.
F_2	Long-lasting flight response.
F_3	Accelerated sleep cycle.
F_4	High body weight.
	Causal relationships
$F_1 \rightarrow F_2$	A high quantity of ACh neurotransmitter causes a long-lasting flight response. The duration of the electrical signal to the muscles is longer because of the excess amount of neurotransmitter.
$F_1 \rightarrow F_3$	A high quantity of ACh neurotransmitter causes an accelerated sleep cycle. The neurotransmitter speeds up all neural activity, including the internal "clock" which puts the shrimp to sleep on a regular cycle.
$F_1 \rightarrow F_4$	A high quantity of ACh neurotransmitter causes a high body weight. The neurotransmitter stimulates greater feeding behavior, which results in more food ingestion and more body weight.
	Blank properties
	has mucus that is slightly acidic
	engages in cannibalistic eating behavior
	is infected with a fungal retrovirus that attacks its reproductive system

& Gottfried, 1994; Gelman & Kalish, 1993; Hirschfeld & Gelman, 1994; Mayr, 1988; Medin & Ortony, 1989; Rehder & Hastie, 2001). The existence of such an essence explains not only many observed correlations between category features, it also leads to the expectation that there are more yet-to-be-discovered features in common (Kornblith, 1993; Shipley, 1993). That is, the presence of a causally-potent essence may promote stronger inductive generalizations because the essence is expected to generate many category features (Gelman & Coley, 1991; Gelman et al., 1994).

Table 1 presents an example of the features and inter-feature causal relationships for one of the six novel categories used in the present research. Lake Victoria Shrimp were described to participants as possessing four distinctive binary features and three inter-feature causal relationships. In this example, the causal links formed a *common-cause causal schema* in which one feature (F_1) was the common cause of the three remaining features (F_2, F_3, and F_4). The knowledge associated with categories such as Lake Victoria Shrimp was intended to be a simplified analog of the essentialist knowledge associated with real-world natural kinds, for example the belief that shrimp DNA (for many people, the essence of animal kinds) causes shrimp features like having fins, locomoting by swimming, living in the sea, having a carapace, and so on.

To determine whether categories with a common-cause schema support stronger inductive inferences, participants in the present experiments performed a series of property induction trials. They were told about one category member that possessed a novel property, and were asked to estimate the proportion of other category members that possessed the novel property. Because such inductive inferences are also sensitive to the type of property, and the interaction between the type of property and category

(Heit & Rubinstein, 1994; Lassaline, 1996; Lopez, Atran, Coley, Medin, & Smith, 1997; Nisbett, Krantz, Jepson, & Kunda, 1983; Proffitt, Coley, & Medin, 2000; Smith, Shafir, & Osherson, 1993; Thagard & Nisbett, 1982), we use *blank properties* that are unfamiliar to the participants (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990). Table 1 also includes the blank properties for Lake Victoria Shrimp. For example, participants are posed the hypothetical, "Suppose one Lake Victoria Shrimp has been found that has mucus that is slightly acidic", and then asked, "What proportion of all Lake Victoria Shrimp have mucus that is slightly acidic?" The predicate "mucus is slightly acidic" was deliberately chosen to be unfamiliar or "blank" to participants. According to our first hypothesis, categories with a common-cause schema may promote inductions because it prompts categorizers to assume that new features are also likely to be produced by the common cause. For example, when instructed on the Lake Victoria Shrimp category of Table 1 and a single category member that has acidic mucus, participants might reason that, "High amounts of the ACh neurotransmitter cause many properties of Lake Victoria Shrimp, it might cause acidic mucus as well." To test this prediction, performance on the property induction task by participants who were provided with a common-cause causal schema is compared to a control group that was taught the novel category but without the associated knowledge.

Our first hypothesis ties the strength of property induction to the specific pattern of causal links exhibited by a common-cause schema. A finding that a common-cause schema promotes inductive generalization will be important, because such a knowledge structure is presumed to be a common one for many natural categories (Gelman et al., 1994; Gelman & Wellman, 1991; Keil, 1989; Rips, 1989). However, the common-cause schema may promote inductive generalizations for reasons other than the presence of its single generative cause. For example, according to our second hypothesis, categories support stronger inductive generalization to the extent they are organized around any central feature or theme, regardless of whether that feature is a cause or an effect. Indeed, research reviewed earlier suggests that categories organized around a central theme have special properties regarding both category construction (e.g. Medin et al., 1987) and learning (e.g. Heit & Bott, 2000; Murphy & Allopenna, 1994). In addition, the importance of a central feature or theme has also been suggested by prior work that has manipulated experimentally-provided causal knowledge. For example, Rehder and Hastie (2001) found that features occupying a central position in a causal network dominated classification decisions. And, Ahn (1999) found that common-cause and common-effect schemas elicit family-resemblance sorting, but not chain schemas in which features were arranged in a causal chain. To test the possibility that any category with a central feature or theme will promote inductions, we also tested categories exhibiting a *common-effect schema* in which one feature (F_4) is caused by the other three features (F_1, F_2, and F_3) (see Fig. 1). Whereas only the common-cause schema should strengthen inductive generalizations if such inductions are promoted by a single underlying generative cause, both the common-cause and common-effect schemas will do so if they are promoted by a central theme, regardless of the direction of causality.

Finally, our third hypothesis is that categories support stronger inductive generalization as a result of the global coherence afforded by inter-feature relations, regardless of the specific arrangement of the causal links. To test this possibility, the final schema shown in Fig. 1, the *chain schema* in which feature F_1 causes F_2 which in turn causes F_3 which in

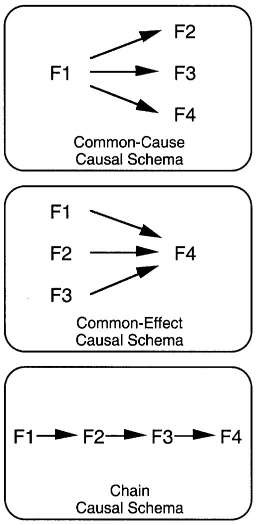

Fig. 1. Causal schemas tested in Experiments 1–3.

turn causes F_4, was also tested. According to this third hypothesis, all three of the schemas shown in Fig. 1 will promote inductions if they are a function of a category's global coherence, rather than a pattern of reasoning associated with one specific topology of knowledge relations among features.

Although the main focus of the current research is on the effect of a category's causal structure on property induction, some investigators have proposed that different cognitive mechanisms will be invoked depending on the kind of the category. For example, researchers have suggested that reasoning with biological kinds will be influenced by a domain-specific *folk biology* with which humans are innately endowed or which emerges early in development (Gelman & Kremer, 1991; Gopnik & Wellman, 1994). Keil (1995) has suggested that artifacts are also often understood in terms of the purpose they serve, a mode of understanding referred to as teleological-adaptive explanation (Bloom, 1996, 1998; Gould & Lewontin, 1978; Mayr, 1982; Rips, 1989). Indeed, Gelman's (1988) finding that elementary school children more often generalize new properties to natural kinds than to artifacts may have arisen from differences in the domain-specific mechanisms with which those kinds are processed, rather than differences in their background knowledge. To explore these questions, we manipulated the causal schema (common-cause, common-effect, or chain) and crossed that factor orthogonally with the kind of category. Specifically, we tested two biological kinds (Lake Victoria Shrimp, Kehoe Ants), two non-living natural kinds (Myastars and Meteoric Sodium Carbonate) and two artifacts (Romanian Rogos and Neptune Personal Computers), each with all four of the causal schema structures (common-cause, common-effect, chain, and no-cause control).

Finally, the current study also investigated whether the notion of "coherence" that promotes inductive generalizations is a characteristic of a category per se, or a property of the category exemplar that displays the new to-be-generalized property. That is, any given category exemplar may manifest, or instantiate, its category's theoretical laws to a greater or lesser extent, and it may be only the former that exhibit the coherence required to promote inductive generalizations. Indeed, prior research has found that the goodness of category membership of an exemplar is enhanced for those exemplars that instantiate a category's causal relationships (Rehder, 1999, 2003, in press; Rehder & Hastie, 2001). And, a well-established finding is that property inductions are strengthened when the new property is displayed by a typical vs. atypical category member (Rips, 1975; see also Osherson et al., 1990). To assess the possibility that it is coherent *category members* that promote inductive generalizations rather than coherent *categories*, the present experiments manipulated the extent to which the exemplars that displayed new to-be-generalized blank properties also manifested their category's causal laws, and collected judgments of the goodness of category membership of those exemplars.

2. Experiment 1

2.1. Method

2.1.1. Materials

The materials were based on those originally employed in Rehder and Hastie (2001). Six categories were constructed: two biological kinds (Kehoe Ants, Lake Victoria Shrimp), two non-living natural kinds (Myastars, Meteoric Sodium Carbonate), and two artifacts (Romanian Rogos, Neptune Personal Computers). The features and blank predicates for each category are presented in Appendix A. Each of the four binary features was described as contrasting with respect to a superordinate category. For example, "Some Lake Victoria Shrimp have high body weight whereas others have normal body weight." Hereafter the presence of a feature ("high body weight") is also denoted by "1" and its absence ("normal body weight") by "0".

Appendix A also presents for each category the six causal relationships needed to construct the common-cause, common-effect, and chain schemas of Fig. 1: $F_1 \rightarrow F_2$, $F_1 \rightarrow F_3$, $F_1 \rightarrow F_4$, $F_2 \rightarrow F_3$, $F_2 \rightarrow F_4$, and $F_3 \rightarrow F_4$. The description of each causal relationship consisted of one sentence indicating the cause and effect features (e.g. "A high quantity of ACh neurotransmitter causes a long-lasting flight response."), and then one or two sentences briefly describing the mechanism responsible for the causal relationship (e.g. "The duration of the electrical signal to the muscles is longer because of the excess amount of neurotransmitter.").

Finally, Appendix A presents the three blank properties used for inductive generalizations for each category. Because prior research has shown that the strengths of inductive generalizations vary as a function of the type of category and the type of property (e.g. more induction for compositional properties for natural kinds, and for functional properties for artifacts, Gelman, 1988; for a review of property effects see Heit, 2000), each category has a composition, a function-behavior, and a disease-malfunction blank property. In addition, a fourth blank property described as "X" was also tested with each category.

2.1.2. Procedure

Experimental sessions were conducted by computer. Participants first studied several screens of information about their assigned category at their own pace. All participants were presented with the cover story and the category's features. Participants in a schema condition were additionally presented with a description of three causal relationships: $F_1 \rightarrow F_2$, $F_1 \rightarrow F_3$, and $F_1 \rightarrow F_4$ for the common-cause group, $F_1 \rightarrow F_4$, $F_2 \rightarrow F_4$, and $F_3 \rightarrow F_4$ for the common-effect group, and $F_1 \rightarrow F_2$, $F_2 \rightarrow F_3$, and $F_3 \rightarrow F_4$ for the chain group. These participants also observed a diagram like those in Fig. 1 representing the causal schema. When ready, participants took a multiple-choice test that tested them on the knowledge they had studied. Participants could request help in which case the computer re-presented the information about the category. Participants were required to retake the test until they committed zero errors and made zero requests for help. For control participants the test consisted of seven questions that tested their knowledge of the category's features. For schema participants the test consisted of an additional 14 questions that tested their knowledge of the causal links (e.g. the direction of causality, the details of the causal mechanism involved, etc.).

Participants then performed three tasks counterbalanced for order: a property induction task, a classification task, and a similarity rating task. The results from the similarity task are unrelated to the theoretical issues raised in this report and are omitted. During the induction task participants were given 32 induction problems in which they were presented with a hypothetical category exemplar that possessed a new blank property and asked what proportion of all category members possessed the property. These 32 trials were divided into four blocks of eight trials each, with the same blank property presented on each trial of a block. Within a block the category exemplar that the blank property appeared with varied: 0000, 0001, 1000, 0101, 1010, 0111, 1110, and 1111. The order of presentation of blocks (i.e. blank properties) was randomized for each participant, except that the "X" blank property was always presented last. As with the other (blank) properties, on this trial participants were asked what proportion of category members have the property X. Within a block, the order of presentation of the eight exemplars was also randomized.

Participants entered their rating with a response bar. The response bar consisted of a computer-displayed scale and a tick mark (a "bar") placed on top of that scale. The left and right arrow keys were used to move the bar along the scale to a position which reflected their confidence that the exemplar was a category member. The left end of the scale was labeled "None" and the right end was labeled "All", indicating the proportion of category members that had the new property. The bar could be set to 21 distinct positions along the scale.

During the classification task, participants rated the category membership of 48 exemplars, consisting of all possible 16 objects that can be formed from four binary feature and the eight single-feature exemplars, each presented twice. For example, those participants assigned to learn the Lake Victoria Shrimp category were asked to classify a shrimp that possessed "High amounts of the ACh neurotransmitter", "A normal flight response", "Accelerated sleep cycle", and "Normal body weight". The features of each to-be-rated exemplar were listed in order (F_1 through F_4) on the computer screen. The list of attribute values for single-feature exemplars contained "???" for the three unknown attributes. The order of the 48 exemplars was randomized for each participant. Participants entered their rating with a 21-position response bar in which the left end of the scale was labeled "Definitely not an X" and the right end was labeled "Definitely an X", where X was the name of the category.

Responses from both the property induction and classification tasks were scaled into a number in the range 0–100. Experimental sessions lasted approximately 45 minutes.

2.1.3. Participants

One hundred and forty-four University of Illinois undergraduates received course credit or $5 for participating in this experiment. They were assigned in equal numbers to the common-cause, common-effect, chain, and control conditions, and to one of the six experimental categories.

2.1.4. Analyses

In all of the experiments reported in this article, initial ANOVAs were performed on the property induction ratings in which schema (common-cause, common-effect, chain, or control), category (six levels), predicate type (four levels), and exemplar displaying the new variable were entered as variables. Induction ratings did not generally vary as a function of whether the category was a biological kind, a non-living natural kind, or an artifact, or whether the predicate type was composition, function-behavior, disease-malfunction blank, or "X". Although there were occasional interactions between predicate type and category (indicating that certain predicates were generalized more strongly with some categories than others), those effects were not consistent across experiments. As a consequence, the results are collapsed over the category and predicate type variables in all experiments.

Because of this article's focus on the effects of causal theories on property induction, category membership ratings are reported only for those exemplars that were also involved in property induction trials: 0000, 0001, 1000, 0101, 1010, 0111, 1110, and 1111. See Rehder (2003, in press) for a complete analysis of the categorization results from Experiments 1 and 3.

2.2. Results

Property induction ratings averaged over categories, blank properties and participants are presented in Table 2 as a function of causal schema condition and the exemplar displaying the new property. A central question asked by the current experiment concerned the efficacy of causal schemas to promote the generalization of new properties to a category. In fact, average induction ratings did not vary as a function of causal schema: ratings were 46.8, 51.5, 46.4, and 51.0 in the common-cause, common-effect, chain, and control conditions, respectively. In a 4 (schema condition) $\times$ 8 (exemplar) mixed ANOVA on induction ratings, the effect of schema was non-significant ($F(3, 120) = 1.17$, $\mathrm{MSE} = 1791$, $P > 0.20$).

Although there was no overall effect of schema, the design of the current experiment also examines the possible role of the exemplar displaying the blank property in mediating the effect of causal schema on property induction. In fact, the interaction between

Table 2
Induction ratings from Experiment 1 (standard errors are shown in parentheses)

Exemplar	Common cause	Common effect	Chain	Control
0000	45.6	57.4	52.2	51.4
	(3.9)	(4.4)	(4.1)	(3.3)
0001	42.3	38.1	40.7	48.2
	(3.1)	(3.7)	(2.9)	(2.8)
1000	37.7	46.7	39.7	49.5
	(3.9)	(3.4)	(3.6)	(2.8)
0101	38.9	48.6	38.1	51.2
	(3.3)	(3.0)	(3.2)	(2.6)
1010	46.4	46.0	41.3	49.1
	(3.4)	(3.5)	(3.5)	(2.6)
0111	41.1	58.8	42.6	51.1
	(4.0)	(3.2)	(3.7)	(3.4)
1110	55.0	46.5	51.0	52.1
	(3.8)	(4.2)	(3.8)	(3.0)
1111	67.2	69.9	65.4	55.0
	(3.5)	(3.7)	(3.8)	(3.3)

exemplar and causal schema ($F(21, 840) = 4.17$, $\text{MSE} = 236.9$) was highly significant ($P < 0.0001$) indicating that the pattern of induction ratings across exemplars depended on the category's causal schema. The pattern of induction ratings in all three schema conditions differed from those in the control condition (i.e. the exemplar variable interacted with the common-cause/control contrast, the common-effect/control contrast, and the chain/control contrast, all Ps < 0.0001).

Earlier in this article we suggested that the presence of a causal schema might make certain category exemplars support stronger inductive generalizations because they manifest the category's causal laws. Fig. 2 presents the induction ratings as a function of causal schemas for the category exemplar 1111. Exemplar 1111 manifests the causal laws in each of the three causal schemas because in all three conditions cause and effect features are all present. As predicted, Fig. 2 indicates that the induction ratings associated with exemplar 1111 are higher in the causal schema conditions (67.2, 70.0, and 65.4 in the common-cause, common-effect, and chain conditions, respectively) as compared to the control condition (55.0); in each causal schema condition these ratings were significantly greater than in the control condition (Ps < 0.05). There were no differences in induction ratings associated with exemplar 1111 across the three causal schemas, indicating that the strength of inductive generalization was not stronger with the common-cause causal schema than with the common-effect or chain schemas.

In addition, Fig. 2 indicates that in each schema condition certain exemplars were *less* likely to support the inductive generalization of new properties compared to the control condition: ratings associated with exemplars 1000, 0001, and 0101 were 37.7, 38.1, and 38.1 in the common-cause, common-effect, and chain conditions, respectively, as compared to 49.5, 48.2, and 51.2 in the control condition (all three Ps < 0.05). The likely explanation for these lower ratings is that exemplars 1000, 0001, and 0101 are poor category members because they each possess violations of the three causal laws in their

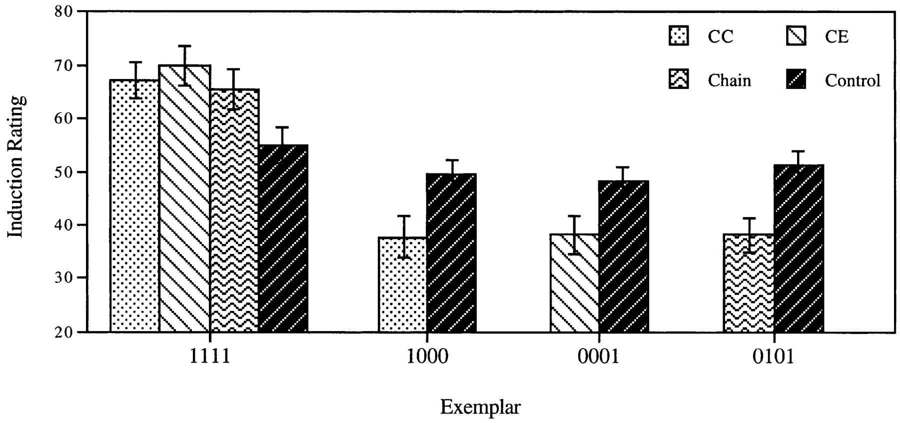

Fig. 2. Induction ratings from Experiment 1 for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

respective conditions, where a violation occurs when a cause feature is present but the effect feature is absent, or when a cause feature is absent but the effect feature is present. For example, in exemplar 1000 in the common-cause condition the common cause (F_1) is present even though all three of its effects (F_2, F_3, and F_4) are absent. In exemplar 0001 in the common-effect condition the common effect (F_4) is present even though all three of its causes (F_1, F_2, and F_3) are absent. Finally, in exemplar 0101 in the chain condition F_2 and F_4 are present even though F_1 and F_3 are absent, respectively, and F_3 is absent even though F_2 is present.

The classification ratings gathered from each participant enable a direct test of our conjecture that the effect of schemas on exemplars' tendency to support inductive generalizations is mediated by their goodness of category membership. Average category membership ratings for the same exemplars shown in Fig. 2 are presented in Fig. 3. As for

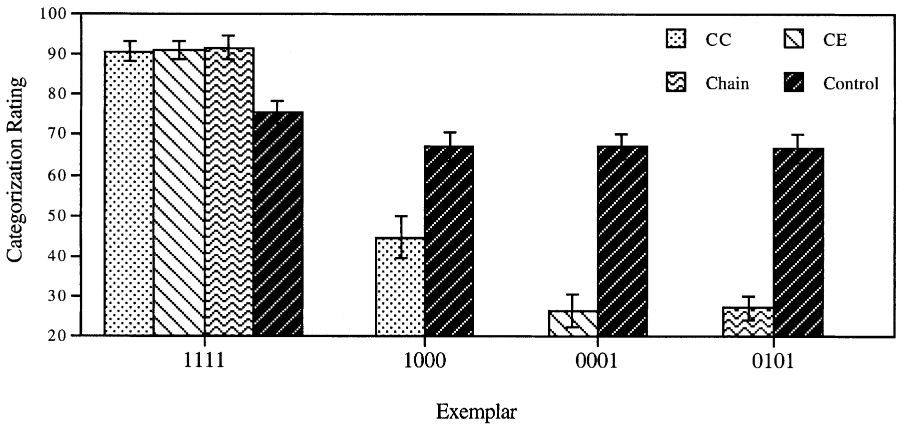

Fig. 3. Categorization ratings from Experiment 1 for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

induction ratings, the pattern of exemplars' category membership ratings was influenced by causal schema (schema by exemplar interaction, $F(21, 840) = 13.71$, MSE $= 377.7$, $P < 0.0001$). Moreover, Fig. 3 shows that the changes brought about by the causal schemas to the induction ratings of those exemplars shown in Fig. 2 are mirrored by the changes to their degree of category membership. That is, in each schema condition exemplar 1111 was rated a better category member than in the control condition, and those exemplars that violated causal laws were rated as worse category members.

To examine this relationship between category membership and property induction ratings more closely, Fig. 4 plots the induction ratings of all eight exemplars as a function of their categorization rating in each causal schema condition. Regression lines indicate that exemplars' categorization ratings are an excellent predictor of their associated induction ratings. Categorization ratings explain more than 90% of the variance in induction ratings in all three causal schema conditions.

Fig. 4 also presents exemplars' categorization and induction ratings in the control condition. As expected, these ratings did not vary across exemplars because participants were not given any information that would lead them to treat one exemplar as a better category member than any other. However, Fig. 4 invites us to consider whether the causal schemas supported stronger inductive generalization controlling for exemplar typicality. That is, it is conceivable that causal schemas lead to stronger inductive generalizations all else being equal, but that in the causal schema conditions a greater proportion of the tested exemplars were viewed as poor category members (e.g. because they violated causal laws), and the result was that these two effects canceled one another out. Indeed, the regression lines of Fig. 4 suggest that induction ratings in the common-cause, common-effect, and chain conditions were somewhat higher than in the control condition at the same level of categorization rating. However, this hypothesis was not supported by a statistical analysis. Each participant's induction ratings were regressed onto his or her categorization ratings that were mean deviated around the control condition's average categorization rating of 70.4 (i.e. 70.4 was subtracted from each exemplar's categorization rating). The intercept from each regression is the participant's predicted induction rating for an exemplar that receives a category membership rating of 70.4. These intercepts were subjected to a one-way ANOVA with schema condition (four levels) as the single variable. The effect of schema condition was not significant ($F < 1$). That is, induction ratings were not influenced by the presence of a causal schema above and beyond its effect on the degree of category membership of the exemplar that displayed the new property.

2.3. Discussion

The purpose of Experiment 1 was to determine if the presence of inter-feature causal relations in novel experimental categories would promote the inductive projection of new properties from one member to the entire category. The results indicated that those category members that exemplified their category's causal laws supported stronger inductive generalizations of the new properties they displayed.

Exemplars were rated as worse category members when their combinations of features violated the causal laws that applied to the category (causes present but effects absent or vice versa), compared to members that confirmed the laws or members of categories that

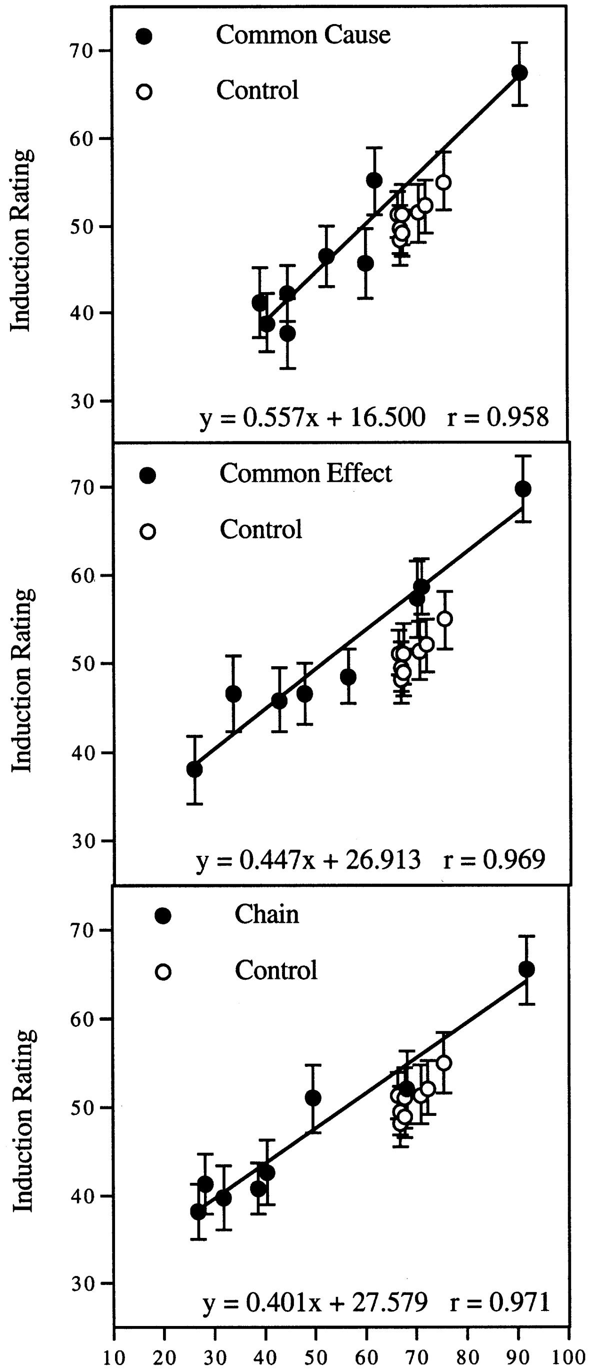

Fig. 4. Induction ratings as a function of categorization ratings in Experiment 1.

were associated with no causal laws. These findings replicate Rehder and Hastie's (2001) results regarding the effects of causal knowledge on category membership. Regression analyses performed in that study indicated that ratings of category membership increased with the number of category relationships confirmed (cause and effect features both present or both absent), and decreased with the number of category links violated (one present and the other absent). Moreover, features that were central to a causal network (i.e. the common cause in the common-cause schema, the common effect in the common-effect schema) had greater impact on category membership ratings than other features.

The use of several causal schemas in Experiment 1 enabled a test of whether the effect of coherent category members on property induction depended on the specific structure of the causal schema that participants were told the category possessed. For example, according to our first hypothesis, categorizers are likely to assume that a central cause of many known category features is responsible for newly-observed features as well, and hence conclude that many category members also possess the new feature. This principle predicts stronger inductive generalizations in the common-cause condition than in the common-effect and chain conditions. Alternatively, inductive generalization may be strengthened by the presence of any feature that functions as a theme around which other features can be organized, regardless of whether it is a cause or effect, thus predicting stronger inductions for the common-cause and common-effect schemas as compared to the chain schema. In fact, we found that the strength of inductive generalizations did not vary as a function of causal schema. Inductive strength depends on the coherence afforded a category by inter-feature relations, regardless of the specific arrangement of those relations.

Finally, the current experiment found no effect of the ontological kind of the category (biological kind, non-living natural kind, or artifact) on inductive generalizations, above and beyond the effect of causal schemas. That is, when the theoretical knowledge associated with natural kinds and artifacts is controlled, there are no systematic differences between natural kinds and artifacts regarding inductive inferences. This results suggests that inductive generalizations are influenced more by the knowledge that people possess than by broad category kind distinctions such as natural kinds vs. artifacts.

3. Experiment 2

There are two reasons why the causal schemas studied in Experiment 1 may not hold in many real-world categories. First, previous research has established that natural categories may be described as clusters of features probabilistically associated with a category (Hampton, 1979; Rosch, 1973; Rosch & Mervis, 1975; Smith & Medin, 1981). In contrast, although the categories in Experiment 1 possessed what we have described as "features", in fact participants were never informed about how often category members possessed those features. For example, participants were told that Lake Victoria Shrimp sometimes had a high amount of ACh neurotransmitter and sometimes a normal amount, but were never told whether Lake Victoria Shrimp typically have high or normal amounts. Given this absence of feature base rate information, it is not surprising that, in the control condition of Experiment 1, exemplars with all four of the "features" (1111) were assigned a category membership rating that was not substantially higher than exemplars missing all four features (0000, 76.9 vs. 71.5). That is, the experimental categories did not exhibit a normal family resemblance structure. Under these conditions, participants may have relied on the causal information only because they had nothing else on which to base their judgments.

Second, although the common-cause causal schema was intended as an analog of the knowledge of essential features associated with many natural kinds (e.g. bird DNA that causes having wings, singing, and looking for worms, etc.), we did not describe the common cause feature as present in all category members. Although such a probabilistic causal structure may correspond to some categories (e.g. *intelligence* is a causally-potent variable that varies over the category *human*; see Waldmann et al., 1995) for natural kinds the underlying essence of a category is assumed to be *fixed*, and to be present in all category members. The absence of any distinctive effects of the common-cause schema in Experiment 1 may have arisen from the use of a variable vs. a fixed common-cause feature value.

In Experiment 2, the categories used in Experiment 1 were modified to address both of these issues. First, to instantiate a more realistic essentialist common-cause schema, the common cause feature (F_1) was described as fixed and occurring in all category members. That is, the common-cause was *necessary* to category membership. The common-effect and chain causal schemas were likewise modified to have one necessary feature (F_4 and F_1, respectively). Second, to create a more realistic family resemblance structure, the remaining features were made *characteristic* of the category by describing them as occurring in 75% of category members.

Because the three causal schema conditions varied with respect to which feature was described as necessary, two control conditions were employed. The F_1-control condition served as a control for the common-cause and chain conditions by describing the feature base rates as 100%, 75%, 75% and 75% for features F_1, F_2, F_3, and F_4, respectively, but providing no information about inter-feature causal relations. Likewise, the F_4-control condition served as a control for the common-effect condition by describing the feature base rates as 75%, 75%, 75%, and 100% and providing no information about inter-feature causal relations.

The introduction of necessary and characteristic features in Experiment 2 enables a stronger test of our claim regarding the influence of theoretically coherent category exemplars on property induction, because participants were also provided with feature base rate information to make their decisions. For example, a crucial test in Experiment 2 will be whether exemplar 1111, which we refer to as the category prototype because it possesses all the features characteristic of its category, supports stronger inductive generalization when it is also coherent, that is, when it satisfies all the theoretical knowledge associated with a category.

3.1. Method

3.1.1. Materials

The same materials from Experiment 1 were used, with the exception that one feature was described as occurring in 100% of category members and the remaining features were described as occurring in 75% of category members. The necessary feature was F_1 in the common-cause, chain, and F_1-control conditions, and F_4 in the common-effect and F_4-control conditions. For example, for the category of Table 1 participants in the common-cause, chain, and F_1-control conditions were told that "100% of Lake Victoria Shrimp have high amounts of ACh neurotransmitter", and that F_2, F_3, and F_4 occurred 75% of the time (e.g. "75% of Lake Victoria Shrimp have a long-lasting flight response whereas 25% have a normal flight response." for F_2). Participants in the common-effect and F_4-control conditions were told that "100% of Lake Victoria Shrimp have a high body weight." and that F_1, F_2, and F_3 occurred 75% of the time.

3.1.2. Procedure

The procedure was identical to that employed in Experiment 1, with the exception that the eight exemplars that displayed blank properties during the property induction task always possessed a category's necessary feature. That is, in the common-cause, chain, and F_1-control conditions the eight induction exemplars each possessed feature F_1 (1000, 1001, 1010, 1100, 1011, 1101, 1110, and 1111), whereas in the common-effect and F_4-control conditions the eight induction exemplars each possessed feature F_4 (0001, 0011, 0101, 1001, 0111, 1011, 1101, and 1111).

3.1.3. Participants

One hundred and eighty University of Colorado undergraduates received course credit for participating in this experiment. They were assigned in equal numbers to the common-cause, common-effect, chain, F_1-control and F_4-control conditions, and to one of the six experimental categories.

3.2. Results

Property induction ratings averaged over categories, blank properties and participants are presented in Table 3 as a function of causal schema and exemplar. As in Experiment 1, the presence of a causal schema had no overall effect on average induction ratings. Mean induction ratings were 55.3, 54.9, 53.8, 54.4, and 51.2 in the common-cause, chain, common-effect, F_1-control, and F_4-control conditions, respectively. Because the common-cause, chain, and F_1-control conditions tested different exemplars than the common-effect and F_4-control conditions, these two sets of conditions were analyzed separately. In both a 3 (schema) × 8 (exemplar) ANOVA of the common-cause, chain, and F_1-control schema conditions, and in a 2 (schema) × 8 (exemplar) ANOVA of the common-effect and F_4-control schema conditions, the effect of schema on induction ratings was non-significant (both Fs < 1).

However, in both ANOVAs there were interactions between the schema and exemplar variables ($F(14,735) = 1.83$, $MSE = 131.2$ and $F(7,490) = 2.92$, $MSE = 178.0$, respectively, both Ps < 0.05). This result indicates that, as in Experiment 1, the pattern of induction ratings across exemplars depended on the category's causal schema. These effects are illustrated in Fig. 5 for high- and low-coherence exemplars. Fig. 5 shows that for the three causal schema conditions the induction ratings associated with the category prototype (1111) are higher than in the corresponding control condition (all three

Table 3
Induction ratings from Experiment 2 (standard errors are shown in parentheses)

Exemplar	Common cause	Chain	F_1 control	Exemplar	Common effect	F_4 control
1000	36.3	40.9	42.1	0001	33.0	39.8
	(2.7)	(4.0)	(3.4)		(3.6)	(3.7)
1001	49.0	47.0	49.9	1001	48.1	46.5
	(1.9)	(3.1)	(3.0)		(3.4)	(3.3)
1010	50.2	45.6	49.2	0101	47.2	46.3
	(2.5)	(3.2)	(2.8)		(3.6)	(3.0)
1100	48.1	47.8	47.6	0011	46.9	49.2
	(1.9)	(2.7)	(3.0)		(3.4)	(3.4)
1011	60.7	59.3	59.2	1101	59.1	53.6
	(2.5)	(2.9)	(2.8)		(3.2)	(3.4)
1101	59.0	59.4	59.2	1011	57.6	55.0
	(2.5)	(3.0)	(2.8)		(2.7)	(3.5)
1110	60.6	60.3	58.8	0111	58.1	53.4
	(2.8)	(2.7)	(2.4)		(3.0)	(3.3)
1111	78.2	79.1	69.4	1111	77.9	66.8
	(2.7)	(2.3)	(3.1)		(3.0)	(4.3)

*P*s > 0.05). In addition, Fig. 5 indicates that those exemplars that are likely to be viewed as poor category members because they violate all three causal relationships (1000, 0001, and 1010 in the common-cause, common-effect, and chain conditions, respectively) tend to result in lower property induction ratings than the ratings in the corresponding control conditions.

As in Experiment 1, our claim that exemplars' tendency to support inductive generalizations is a function of their degree of category membership is supported by the categorization ratings (Fig. 6). The pattern of exemplars' category membership ratings was influenced by causal schema (as indicated by a schema by exemplar interaction,

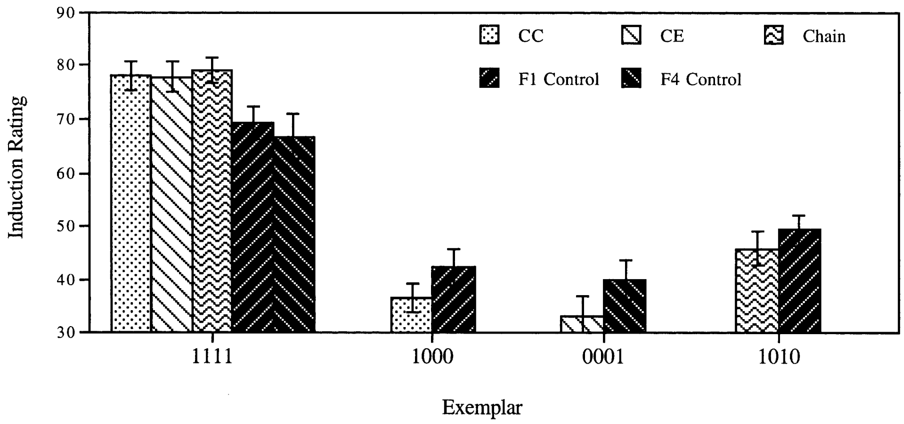

Fig. 5. Induction ratings from Experiment 2 for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

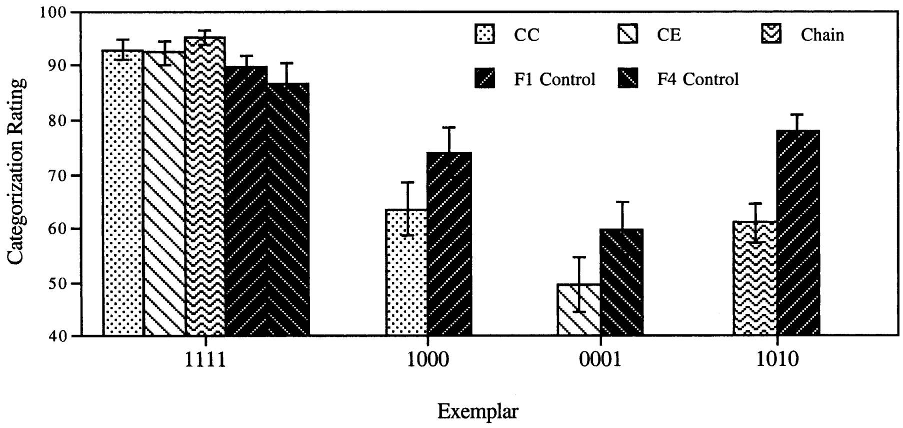

Fig. 6. Categorization ratings from Experiment 2 for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

$F(60, 26.25) = 43.86$, $\text{MSE} = 342.4$, $P < 0.0001$). Moreover, Fig. 6 shows that the changes brought about by the causal schemas to exemplars' induction ratings as illustrated in Fig. 5 are paralleled by changes to their degree of category membership. That is, the prototype 1111 supports stronger inductive generalizations apparently because it is viewed as a better category member in light of causal knowledge as compared to the control conditions. And, exemplars 1000, 0001, and 1010 support weaker inductive generalizations in the common-cause, common-effect, and chain conditions, respectively, apparently because they are viewed as worse category members in light of that knowledge.

Following Experiment 1, Fig. 7 plots each exemplar's induction rating as a function of its categorization ratings in each causal schema condition. The corresponding control condition is also included in each panel. Regression lines indicate that exemplars' categorization ratings are once again an excellent predictor of their associated induction ratings: categorization ratings explain more than 90% of the variance in induction ratings in all five conditions. In contrast to Experiment 1, Fig. 7 indicates that exemplars' categorization and induction ratings varied in the control conditions, a result that was expected on the basis of the 75% feature base rate information that was provided in Experiment 2 but not the first experiment.

It is notable that the regression line in each causal schema condition is higher than in the corresponding control condition, suggesting that causal schemas might support stronger inductive generalizations at the same level of exemplar typicality. To test this possibility, each participant's induction ratings were regressed onto his or her categorization ratings that were mean deviated around the control conditions' average categorization rating of 70.7. The intercept from each regression is the participant's predicted induction rating for an exemplar that receives a category membership rating of 70.7. These intercepts were subjected to an ANOVA with schema condition (five levels) as the single variable. The effect of schema condition did not approach significance ($F(4, 175) = 1.07$, $\text{MSE} = 0.584$, $P > 0.20$). That is, induction ratings were not strengthened by the presence

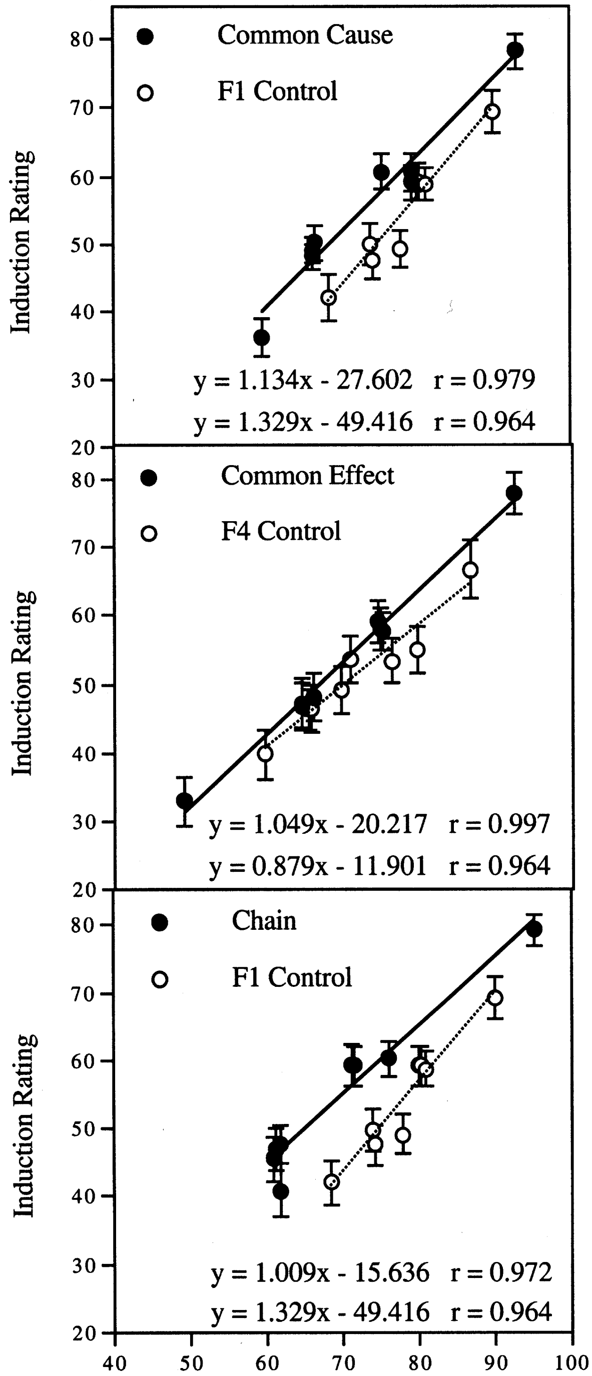

Fig. 7. Induction ratings as a function of categorization ratings in Experiment 2.

of causal schemas above and beyond their effect on the degree of category membership of the exemplar that displayed the new property.

3.3. Discussion

The results of Experiment 2 replicate those of Experiment 1 using more realistic category structures. First, category members supported stronger inductive generalizations of new properties when they were causally coherent category members, that is, when they instantiated inter-feature causal laws. This result was obtained even though the categories were described as possessing characteristic features along with one necessary feature. That is, even though participants in Experiment 2 had an opportunity to base their categorization and induction judgments solely on feature base rate information, exemplars that possessed all the characteristic features of a category received higher categorization and induction ratings when they also confirmed their category's causal laws as compared to control categories without those laws. This finding indicates that the extent to which an exemplar satisfies a category's causal laws influences categorization and induction above and beyond the extent to which it satisfies the category's empirical family resemblance information.

Second, as in Experiment 1, category membership ratings not only increased whenever an exemplar's combination of features confirmed a schema's causal laws, it also decreased (relative to the appropriate control condition) whenever those features violated those laws (causes present but effects absent or vice versa). Moreover, whenever categorization ratings decreased, so too did induction ratings. This result can be interpreted as supporting the previous finding that exemplar typicality is an important variable moderating the strength of inductive inferences (Osherson et al., 1990; Rips, 1975), with the qualification that "typicality" is partly determined by causal knowledge, not just by the number of features often observed in other category members.

Finally, as in Experiment 1 the common-cause schema did not support distinctively strong inductive generalizations as compared to the other two causal schemas. This result was obtained despite the fact that we used a more realistic fixed, "essential", common-cause, compared to Experiment 1's varying common cause. Once again we found that strong inductive inferences are based on those coherent category exemplars that instantiate a category's causal laws, regardless of the topology of the relevant causal schema.

4. Experiments 3A and 3B

Experiment 2 differed from Experiment 1 by describing the empirical base rate of the common cause feature as 100%, a change we introduced to make our common-cause experimental categories more like essentialist categories. However, the common-cause categories used in Experiment 2 still depart from "essentialized" real-world categories in at least two regards. First, a category essence is viewed as not only being possessed by all members of a category, but also by members of no other category. That is, essences are not just necessary to category membership, they are also sufficient. The materials used in Experiments 1 and 2 were modified to make F_1 the essential, or defining, property of the category in the common-cause and chain conditions, and to make F_4 the essential property in the common-effect condition, by describing those features as occurring in 100% of all category members and occurring in members of no other categories. In addition, the names of the categories were changed to emphasize the defining, essential nature of these features. For example, in the common-cause and chain conditions Lake Victoria Shrimp were renamed Acetylcholine Shrimp and participants were told that all Acetylcholine Shrimp have high amounts of acetylcholine and that no other kind of shrimp does. In the common-effect condition Lake Victoria Shrimp were renamed Heavy Shrimp and participants were told that all Heavy Shrimp have an usually high body weight that is not seen in any other kind of shrimp.

Second, when generating induction and category membership ratings in the first two experiments, participants were always given complete information regarding the presence or absence of each of the four binary features for the presented category exemplar. In contrast, in the real world we usually cannot directly observe underlying, essential features such as Bird DNA. To simulate this situation, the test exemplars that participants were presented with always listed the essential feature as unknown. That is, in the common-cause and chain conditions participants were presented with the eight three-feature exemplars with F_1 missing: x000, x001, x010, x100, x011, x101, x110, and x111 (where x denotes an unknown value). Likewise, in the common-effect condition participants were presented with the eight three-feature exemplars with F_4 missing: 000x, 001x, 010x, 100x, 011x, 101x, 110x, and 111x.

The presence of missing features in Experiment 3 provides a more stringent test of our proposal regarding the influence of coherent category exemplars on property induction. In the first two experiments, both categorization and induction ratings were strongly influenced by whether the exemplar confirmed or violated causal laws. Confirmations occurred when cause and effect were both present or both absent, and violations occurred when one was present and the other absent. In Experiment 3, participants never directly observe either the common-cause or the common-effect, and so the confirmation (and violation) of causal relationships is not explicit in the category exemplars. The third experiment tests whether exemplars are viewed as coherent (incoherent) in the absence of explicit confirmation (violation).

The common-cause and common-effect conditions were tested in Experiment 3A, and the chain condition was tested in Experiment 3B. Experiment 3A included two control conditions, one which served as control for the common-cause condition in which F_1 was described as the essential feature but no causal relationships were presented, and one which served as control for the common-effect condition in which F_4 was described as the essential feature but no causal relationships were presented. Experiment 3B included one control condition in which F_1 was described as the essential feature but no causal relationships were presented.

4.1. Method

4.1.1. Materials

The categories in Appendix A were renamed as follows in the common-cause and chain conditions and their associated F_1-control condition: Iron Sulfate Ants, Acetylcholine

Shrimp, Ionized Helium Stars, Radioactive Sodium Carbonate, Butanos, and Magnetic Computers. In those conditions each category feature F_1 was described as occurring 100% of the time in category members, and 0% of the time in members of other categories. For example, for the category of Table 1 participants were told that "100% of Acetylcholine Shrimp have high amounts of ACh neurotransmitter, and no other kind of shrimp does." In the common-effect condition and its F_4-control condition the categories were named: Fast Ants, Heavy Shrimp, Planet Stars, Reactive Sodium Carbonate, Carbonos, and Bright Computers, and F_4 was described as occurring 100% of the time in category members, and 0% of the time in members of other categories. For example, for the category of Table 1 participants were told that "100% of Heavy Shrimp have a high body weight, and no other kind of shrimp does." As in Experiment 1, no empirical base rate information was provided for non-essential features.

4.1.2. Procedure

The training procedure employed in Experiment 3 was identical to that in Experiment 1 except for the new information about F_1 (or F_4). After category learning, participants were presented with test exemplars where the presence or absence of the defining feature was always unknown. That is, in the common-cause, chain, and F_1-control conditions, participants produced category membership and induction ratings for the eight three-feature exemplars x000, x001, x010, x100, x011, x101, x110, and x111. In the common-effect and F_4-control conditions participants produced ratings for the eight three-feature exemplars 000x, 001x, 010x, 100x, 011x, 101x, 110x, and 111x.

4.1.3. Participants

For Experiment 3A, 108 New York University undergraduates received course credit or pay for their participation, with 36 participants each assigned to the common-cause and common-effect conditions, and 18 each to the F_1-control and F_4-control conditions. Experiment 3B consisted of 54 University of Illinois undergraduates with 36 and 18 participants randomly assigned to the chain and F_1-control conditions, respectively.

4.2. Results

Property induction ratings averaged over categories, blank properties and participants are presented in Table 4 as a function of causal schema and exemplar for both Experiments 3A and 3B. Three separate 2 (schema vs. control) × 8 (exemplar) ANOVAs were performed on the common-cause and F_1-control conditions, and on the common-effect and F_4-control conditions of Experiment 3A, and on the chain and F_1-control conditions of Experiment 3B. In all three analyses, there was no main effect of schema, but a large interaction between schema and exemplar ($F(7, 364) = 5.49$, $\mathrm{MSE} = 101.4$; $F(7, 364) = 7.93$, $\mathrm{MSE} = 142.7$; $F(7, 364) = 5.74$, $\mathrm{MSE} = 221.6$, in the common-cause, common-effect, and chain condition ANOVAs, respectively, all $P\text{s} < 0.0001$). That is, as in Experiments 1 and 2, the pattern of induction ratings across exemplars depended on the category's causal schema. These effects are illustrated in Fig. 8 for high- and low-coherence exemplars. The figure indicates that in the common-cause and chain conditions exemplar x111 received larger induction ratings than in the corresponding

Table 4
Induction ratings from Experiment 3 (standard errors are shown in parentheses)

	Experiment 3A		Experiment 3B			Experiment 3A	
Exemplar	Common cause	F_1 control	Chain	F_1 control	Exemplar	Common effect	F_4 control
x000	37.5	49.1	45.7	57.4	000x	30.8	46.2
	(3.2)	(3.1)	(4.7)	(5.4)		(3.8)	(3.1)
x001	48.8	52.6	40.4	53.8	100x	48.4	47.2
	(2.4)	(3.1)	(3.3)	(4.5)		(3.4)	(3.2)
x010	49.2	53.8	40.2	51.3	010x	46.7	47.6
	(2.3)	(2.8)	(3.3)	(4.5)		(3.3)	(3.1)
x100	48.2	50.6	43.0	54.3	001x	47.4	49.6
	(2.2)	(2.7)	(2.8)	(4.2)		(3.4)	(3.0)
x011	59.4	54.9	49.2	54.0	110x	59.3	50.3
	(1.9)	(3.0)	(3.4)	(4.4)		(3.2)	(3.5)
x101	58.2	53.1	50.0	55.8	101x	59.4	52.6
	(1.8)	(2.7)	(3.5)	(4.2)		(3.1)	(3.4)
x110	58.3	53.5	53.9	56.4	011x	59.7	49.6
	(1.9)	(3.2)	(3.2)	(4.4)		(3.0)	(3.4)
x111	69.2	60.3	74.8	56.5	111x	70.9	54.7
	(3.1)	(4.6)	(3.4)	(4.7)		(3.6)	(3.7)

F_1-control conditions (69.2 and 74.8 vs. 60.3 and 56.5, Ps < 0.15 and 0.05, respectively). In addition, 111x received larger ratings in the common-effect condition than in the F_4-control condition (70.9 vs. 54.7, $P < 0.05$). This result replicates those from Experiments 1 and 2 where those exemplars that confirmed causal laws were given higher induction ratings.

Also as in the first two experiments, those exemplars that violated causal laws were given lower induction ratings. In the common-cause condition x000 received lower ratings than in Experiment 3A's F_1-control condition (37.5 vs. 49.1, $P < 0.05$), in the common effect condition 000x received lower ratings than in the F_4-control condition (30.8 vs. 46.2,

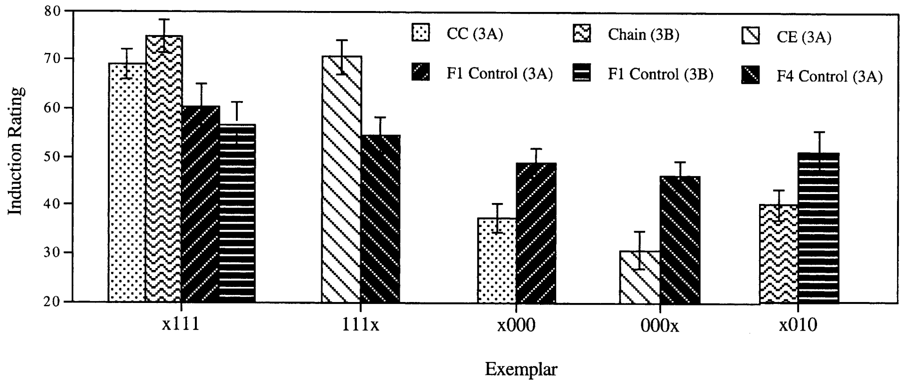

Fig. 8. Induction ratings from Experiments 3A and 3B for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

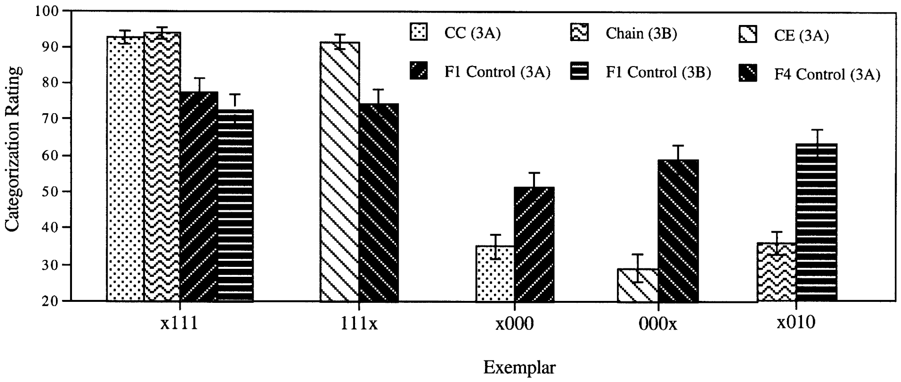

Fig. 9. Categorization ratings from Experiments 3A and 3B for high- and low-coherence exemplars as a function of causal schema. CC, common cause; CE, common effect.

$P < 0.05$), and in the chain condition x010 received lower ratings than in Experiment 3B's F_1-control conditions (40.2 vs. 51.3, $P < 0.10$).

The effect of causal schemas on induction ratings was reflected in categorization ratings (Fig. 9). In the three ANOVAs testing the effect of causal schemas against their corresponding control conditions there were large interactions between the schema and exemplar variables ($F(7, 364) = 10.86$, MSE = 115.4; $F(7, 364) = 13.87$, MSE = 173.3; $F(7, 364) = 13.37$, MSE = 220.9, for the common-cause, common-effect, and chain conditions, all Ps < 0.0001). A comparison of Figs. 8 and 9 indicates that the changes brought about by causal schemas to the pattern of induction ratings were closely mirrored by the corresponding changes in categorization ratings. This correspondence is shown explicitly in Fig. 10, which presents induction ratings as a function of categorization ratings for each of the eight exemplars in each condition. As Fig. 10 shows, the correlation between categorization ratings and induction ratings was >0.98 in all three causal schema conditions.

4.3. Discussion

The results of Experiment 3 replicated those of the previous two experiments. When causal laws were present, a category exemplar that confirmed (violated) those laws received higher (lower) induction and category membership ratings as compared to the same exemplar when the causal laws were absent. This result held despite the fact that the essential category feature could not be directly observed.

In this experiment, we "essentialized" the common-cause categories by giving them an unobservable necessary and sufficient property, but they still did not support stronger inductive generalizations as compared to essentialized common-effect and chain causal schemas. Once again, exemplars promote stronger inductive generalizations when they are coherent in light of causal laws *regardless of the topology of the relevant causal schema*.

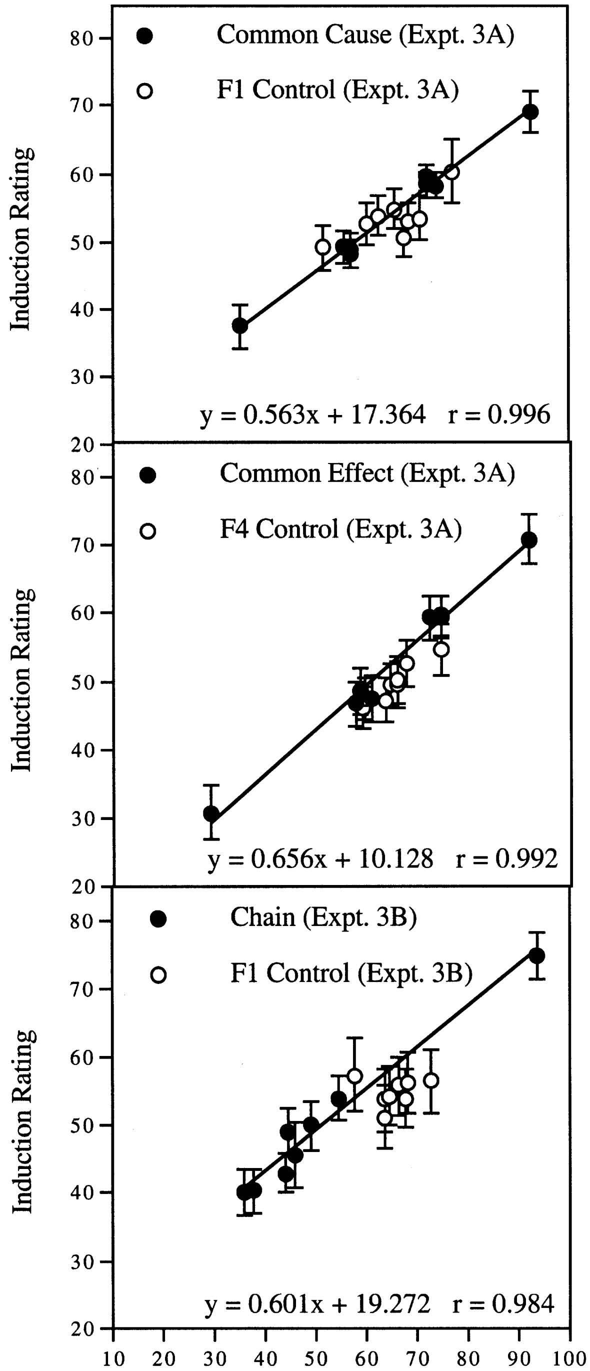

Fig. 10. Induction ratings as a function of categorization ratings in Experiments 3A and 3B.

5. General discussion

We began this research with a question: What distinguishes everyday human conceptual categories from arbitrary collections of objects? One answer to this question is that everyday categories exhibit theoretical coherence and are organized by causal schemas, that arbitrary collections of objects lack. The experiments reported in this paper show that the theoretical coherence afforded to "good exemplars" supports the projection of novel properties to other members of the same category. We found that categories organized around a causally-potent essence, a common-cause schema in which one category feature was the cause of all other distinctive features, were inductively strong. But, we also found that other causal schemas, including common-effect and chain structures, supported strong inductive inferences. Thus, a quick summary of our findings would be that structure mattered; the more causal structure, the more coherence, and the more inferential strength.

5.1. Essentialism, coherence, and category-based property induction

One goal of the present research was to examine inductive reasoning within a category whose known features were believed to be generated by a common cause. Many investigators have suggested that the special properties of real-world categories, especially biological kinds, stem from the belief that everyday categories possess causally-potent essences (Gelman et al., 1994; Gelman & Kalish, 1993; Medin & Ortony, 1989; Rehder & Hastie, 2001). Consistent with this proposal, the common-cause schema supported stronger property induction from coherent category members when no feature base rate information was provided (Experiment 1), when the common-cause was described as necessary to category membership (Experiment 2), and when it was described as defining (i.e. necessary and sufficient) of category membership (Experiment 3). The purpose of describing the common-cause as a defining feature in Experiment 3 was to mimic those real-world categories that people believe are organized around an essence that all members of the category possess and that is possessed by the members of no other category.

The present research demonstrates that *once a common-cause takes hold* in a person's mental representation of a category, stronger inductive generalizations are likely to follow. But, we have not addressed the prior question of why essentialist beliefs are so common, outside of our experiments, and especially in everyday beliefs about biological phenomena. According to some accounts, people may have an innate tendency to assume the existence of an essence for many categories, especially natural kinds (Hirschfeld, 1996; Medin & Ortony, 1989). Alternatively, essences may be the product of inductive reasoning processes that search for causal explanations for the properties and behaviors of kinds that one observes (Gelman et al., 1994). Or perhaps assumptions of essence may simply derive from the use of individuating category labels (Coley et al., 1997).

But, we also found that enhanced confidence in inductive inferences is not unique to the essentialist, common-cause schema, because any of the causal schemas we tested conferred more inductive strength than a-theoretical, zero-causal-link control concepts. That is, the stronger inductions do not appear to arise from a pattern of reasoning in which a causally-potent essence is assumed to cause a new feature because it already causes most other features. Rather, stronger inductions arise from the general coherence that theories provide by linking category features with meaningful relations.

Our use of several novel categories allowed us also to manipulate the ontological kind of the category, that is, whether it was a biological kind, a nonbiological natural kind, or an artifact. One possibility we considered was that different patterns of inductive reasoning may be evoked by natural kinds and artifacts. In fact, we found no systematic differences between kinds regarding inductive generalizations: both natural kinds and artifacts exhibited stronger inductive generalizations when they were accompanied by inter-feature theoretical knowledge as compared to when such knowledge was lacking. This result is consistent with Gelman's (1988) contention that elementary school children's stronger inductions for natural kinds are a result of the greater knowledge associated with those kinds, rather than something special about natural kinds per se. We also observed no affinities between kind and causal schema. For example, given their presumed affinity for an essence-based representation, we might have expected that the common-cause schema would be especially potent when making inductive judgments about biological kinds. No such kind-by-schema dependencies were observed however.

Another question we asked concerned the locus of the effect of theoretical coherence on property induction. We considered two possibilities. First, coherence might be a property of the category such that any category member that possesses a new property will have that property more strongly generalized. Instead, we found that inductions were strengthened only for those *category members* that were coherent, that is, those that instantiated or manifested the causal laws associated with their category. Incoherent category members that violated or disconfirmed their category's laws supported weaker generalizations.

To understand why the coherence of an individual category member affects its support for inductive inferences we considered the possibility that this effect was mediated by the exemplar's degree of category membership. Indeed, according to the well-known *Similarity-coverage Model*, a novel blank property is more likely to be generalized to a category when it appears with typical vs. atypical category members (Osherson et al., 1990; Rips, 1975). Consistent with this framework provided by the Similarity-coverage Model, we found that exemplars that were rated better category members also supported stronger inductive inferences, and those that were rated worse category members supported weaker inferences. Across nine causal schema conditions in four experiments, the average correlation between group-level induction and category membership ratings was 0.98.

On the one hand then, the relationship we found between induction and category membership may be viewed as an extension of previous findings. However, according to the Similarity-coverage Model, an effect of typicality arises because typical members are more similar to other members of a category than are atypical members, and stronger induction arises to the extent that the category's features are more "covered" by the premises of the inductive argument (see also Sloman, 1993). In the current experiments however, "typicality" was not simply determined by feature overlap, but rather by the extent to which exemplars satisfied the theoretical knowledge associated with a category. Thus, although our findings may be viewed as a replication of previous studies regarding the relationships between property induction and category membership, the basis for determining typicality is different than that usually assumed by those studies.

5.2. Coherence, ideals, and goodness of category membership

A critical question raised by the present findings therefore concerns why the same exemplar was viewed as a better category member when it manifested the causal laws associated with its category as compared to when no causal laws were provided. We suggest that one useful way to view this result is as a generalization of Barsalou's (1985) notion of "ideals" in determining category membership. According to Barsalou, exemplars become better category members not only when they possess features that appear frequently among category members (and rarely among non-members), but also when they fulfill their category's goals. For example, Barsalou found that despite the fact that people on diets often eat foods with many calories, "foods with zero calories" were rated as good members of the category "things to eat on a diet", a result he attributed to the fact that such foods satisfy the category's goals.

We propose that exemplars are not only ideal when they satisfy human goals, but also to the extent that they are *theoretically ideal*, that is, to the extent they satisfy the complex causal and explanatory knowledge that people have about a kind. Just as we might judge a person who is healthy and physically fit to be a good exemplar of the category human beings, we judge other exemplars whose causal systems are in an ideal state of effective operation to be typical or representative of their categories. Past research supports the importance of theoretical ideals in determining category membership. For example, Lynch, Coley, and Medin (2000) found that the "weediness" of a tree contributed (negatively) to how good a member of the tree category a given species of tree was rated to be, and did so more strongly than a measure of feature overlap or central tendency. According to Lynch et al., weediness was a negative ideal associated with trees because weedy trees fail to satisfy human goals (they have weak wood and a tendency to create messes and grow where they are not wanted). However, Lynch et al. (2000) also found that height was a strong predictor of category membership ratings above and beyond central tendency, but expressed uncertainty regarding what human goal tall trees fulfill.

Tallness is an ideal for trees because tallness is taken as diagnostic of a "well functioning" tree, that is, one whose underlying causal mechanisms are operating as they (ideally) should be. Stunted trees, in contrast, often suffer from disease or malnourishment. In support of this idea, height was a significant predictor of category membership only for people in the Lynch et al. study who would be likely to have theoretical knowledge of trees (taxonomists and landscapers), but not for respondents unlikely to possess such knowledge (undergraduates and maintenance workers).[1] In other words, we agree that "experts' categories may be organized around ideals more than novices'" (Lynch et al., 2000, p. 48) but suggest that ideals are determined not just by human goals but also by the greater theoretical knowledge that experts possess.

In the current research coherent or ideal category members were those that satisfied inter-feature causal relationships. However, many real-world categories not only exhibit

[1] Lynch et al. (2000) also consider the possibility that tallness influences typicality because it distinguishes trees from other organisms (e.g. shrubs). However, on this account the typicality ratings of undergraduates and maintenance workers should be influenced as much by height as those of taxonomists and landscapers, since all groups are equally aware of the distinctive height of trees.

internal causal link but also participate in causal relationships with other categories, and we suggest that ideal category members are ones that satisfy those external relationships as well. For example, people understand natural kinds to be involved in important relations with external entities (e.g. predator/prey relationships for animals, sources of nutrients for plants, habitat relationships for all living things), including complex environmental interactions that have taken place in the past (e.g. the evolution of particular adaptive mechanisms, the formation of non-biological kinds such as coal and diamonds). Rips (2001) has emphasized the importance of members of natural kinds satisfying causal relationships with other kinds, rather than just exhibiting certain critical features. According to his *interactional view* of such kinds, "an object's membership in a natural kind depends on whether the object instantiates the laws for that kind" (p. 846). In contrast, according to an *intrinsic view*, category membership depends on exhibiting essential (intrinsic) features. The current results that all three causal schemas tested in the current research produced better category members when they satisfied causal relationships are consistent with an interactional view of natural kinds, as long as the causal relationships that are required to be satisfied for kind membership are construed to include not only those between a kind and other kinds, but also those internal relationships between a kind's own features.

Satisfying relationships, and especially external relationships, will be important for members of artifact categories as well. Of course, artifacts enter into important external relationships insofar as they satisfy human goals. But artifacts participate in other external relationships as well. First, people understand that how an artifact functions is a result of interactions between it and other entities, and that whether an object functions as needed to warrant membership in an artifact category requires the success of those interactions. For example, Wisniewski (1995, Experiment 4) found that certain artifacts were better examples of the category "captures animals" when they possessed certain novel combinations of features (e.g. "contains peanuts" and "caught a squirrel") but not others ("contains acorns" and "caught an elephant"), a distinction presumably based on participants' understanding that an animal catching device will only function properly when it utilizes bait appropriate for the animal. Similarly, Rehder and Ross (2001) showed that artifacts were considered better examples of a category of pollution cleaning devices when they possessed a gathering instrument that was appropriate to the type of pollution being gathered (e.g. "has a metal pole with a sharpened end" and "works to gather discarded paper") but not otherwise ("has a magnet" and "removes mosquitoes"). Second, some researchers have argued that our conception of artifacts must also include a notion of the *original intentions* of the artifacts' designer and that the presence or absence of such intentions influences judgments of category membership (Bloom, 1998; Keil, 1995; Rips, 1989). Matan and Carey (2001) comment on the complexity of human artifact categories by observing that children must work out the causal relations between artifacts' functions, designers' intentions, and human goals to reach a mature understanding of artifacts.

Finally, the importance of satisfying relationships even extends to kinds usually considered to be nominal, that is, based on conventional definitions. For example, Lakoff (1987) suggested that concepts amount to *idealized cognitive models* and that category membership is affected by the extent to which an exemplar instantiates such a model. On this account, Catholic priests are poor examples of bachelors, not because they lack the features of the concept of bachelor (unmarried adult male), but because they fail to satisfy cultural background assumptions regarding marriageability on which the concept of bachelor depends.

In each of these domains, categorization may be viewed as a mental activity that involves judging whether an object's observable features are representative of the causal mechanisms that he or she believes give rise to a kind. In this light, object classification may be viewed as an instance of "reasoning to the best explanation" (Harman, 1986; see also Murphy & Medin, 1985; Rips, 1989). We suggest that people often strive to explain the objects they perceive, and that classification of an object into a category occurs when the category's causal mechanisms are implicated in the best explanation for that object. Theoretically ideal category exemplars are those that provide evidence for most or all of those mechanisms.

The emphasis on explanation and the satisfaction of causal relationships in the role of establishing category membership places the current categorization results beyond the reach of traditional models of classification that assume that categories merely reflect the statistical regularities of category features in the environment. For example, according to prototype models, the mental representation of a category consists of those features that appear in many category members and in few non-members (Hampton, 1979; Rosch & Mervis, 1975). And, according to exemplar models, it consists of memory traces of previously-observed category members (Medin & Schaffer, 1978; Nosofsky, 1986). On their own, such models are unable to account for the role of knowledge in producing category members that are not just prototypical, but are also coherent and ideal. However, more recent proposals have begun to address the effects of theoretical knowledge on classification. For example, one proposal that accounts for these findings is *causal-model theory* (Rehder, 1999, 2003, in press; see also Waldmann et al., 1995). Causal-model theory specifies both a representation of causal knowledge and a theory of categorization in light of that knowledge, and predicts that an exemplar's degree of category membership is a function of the probability that the exemplar was generated by a category's causal model. On this account, prototypical category members will be viewed as even better category members when a category also has theoretical knowledge (such as the common-cause, common-effect, or chain schemas used here), because that knowledge constrains the generation of category members such that the prototype is especially likely to be generated. Conversely, it predicts that exemplars that violate many causal laws will be less likely to be generated and hence will be granted a lower degree of category membership (Rehder, 2003, in press). The current results extend these findings by showing how an exemplar manifesting a category's causal laws influences not only its perceived degree of category membership, but also its propensity to support the generalization of new properties. With suitable modifications, causal-model theory could account for the importance of satisfying not only internal (i.e. feature-to-feature) relationships, but also external links with other concepts.

5.3. Domain-specific causal reasoning and category-based property induction

Finally, another notable characteristic of our experiments was their use of "blank predicates" about which participants have little prior knowledge. Blank predicates were used to determine the effect of theoretical coherence on inductive generalization in isolation of any effect of knowledge that directly links the category and the to-be-generalized property. However, in many real-world examples of category-based property induction, reasoners will possess knowledge linking the category and the property, and there is now considerable experimental evidence that people will also use that knowledge when reasoning about meaningful, non-blank properties. For example, Heit and Rubinstein (1994) found that the species-to-species inductive inference of a property was strengthened when the species shared its underlying cause. For example, when first told that bears had a new biological property P_1 (e.g. a liver with two chambers) participants were more likely to judge that whales also had P_1 as compared to when first told that tuna had P_1, presumably because whales are more likely to share biological mechanisms with other mammals (e.g. bears) than fish (e.g. tuna). In contrast, they were more likely to judge that whales had a behavioral property P_2 (e.g. travel in a zig-zag path) when first told that tuna rather than bears had P_2, because whales are more likely to share a survival behavior with other prey animals in the same ecology (e.g. tuna) than a non-prey animal in a different ecology (e.g. bears). Similarly, Lassaline (1996) found that undergraduates were more likely to project a new property when the target category included the property's purported cause (see also Wu & Gentner, 1998).

Ranging further afield, Lopez et al. (1997) studied category-based inductions among the Itzaj Maya, an indigenous population in central Guatemala. They found that the Itzaj frequently appealed to specific causal mechanisms when judging inductions (e.g. their estimates of the prevalence of a disease among members of a biological category depended on their beliefs about the mechanism whereby the disease might spread among category members). Lopez et al. argued that the greater reliance of the Itzaj on causal mechanism arguments, compared to American undergraduates, was due to the greater practical knowledge possessed by the Itzaj. Proffitt et al. (2000) found similar differences between groups that varied in their amount of domain knowledge (for additional examples of causal reasoning during induction, see Osherson, Smith, Myers, & Stob, 1994; Sloman, 1993; Smith et al., 1993).

These results implicate the importance of causal reasoning in category-based property induction, and point to the need for psychological models of induction that include the influence of reasoners' intuitive theories. Our own hypothesis is that inductive inferences will depend more directly on causal explanations when non-blank predicates are involved, in domains where the reasoner has some, perhaps superficial, inkling of the causal preconditions or "recipes" for the presence of the novel property. We believe that a useful theoretical approach is to describe these explanations for the novel property in terms of a Bayesian network mechanism (see Rehder, 2003, in press). The plausibility of that causal model and its explanation for the novel property would then be the major determinants of confidence in inductive projection inferences (cf. Sloman, 1993). If we think that a house down the street was burgled because it was surrounded by trees and ground cover, we project the predicate "burglar-izable" to other houses that are obscured by trees (Heit, 2000; Rips, 2001). As Rips (2001, p. 833) puts it, "We sometimes also need to understand why or how the premise is supposed to be true."

We close by noting that the importance of causal linkages in the induction of non-blank properties demonstrated by the studies just reviewed provides one clue for why the induction of *blank* properties is so closely tied to an exemplar's typicality. We speculate that, despite their "blankness", reasoners often assume that blank properties are causally linked to one or more of the exemplar's existing features. As a result, the prevalence of the new property will be a function of how often the exemplar's existing feature(s) appear in other category members. On this account, typical category members support stronger inferences because their features are more prevalent than atypical ones. Ideal category members support the strongest inductions of all, because their features are constrained to be generated by underlying theories, and hence are especially likely to be present in most category members.

5.4. Conclusion

The present research examined the relationship between category-based inductive generalization and theoretical coherence. Our primary finding is that providing causal knowledge that connected the known features of the category promoted the induction of new properties to the category as a whole. We found that induction was mediated by the degree of category membership of the exemplar possessing the new property such that category membership and inductive inferences were both stronger when the exemplar confirmed the causal laws, and both weaker when it violated them. We suggested that maximally coherent category members represent "theoretical ideals" which arise when an exemplar satisfies not only human goals, but also any of the complex causal and explanatory knowledge that people have about a kind.

Acknowledgements

We thank Evan Heit for his valuable suggestions and comments on this research. Support for this research was provided by funds from the National Science Foundation (Grant Numbers SBR 98-16458 and SBR 97-20304) and from the National Institute of Mental Health (Grant Number R01 MH58362).

Appendix A. Materials

Descriptions of the cover story, attributes, attribute values, causal relationships, and blank properties for each of the six categories are presented below.

A.1. Kehoe Ants

On the volcanic island of Kehoe, in the western Pacific Ocean near Guam, there is a species of ant called Kehoe Ants. For food, Kehoe Ants consume vegetation rich in iron and sulfur.

A.1.1. Features

(F_1) Some Kehoe Ants have blood that is very high in iron sulfate. Others have blood that has normal levels of iron sulfate.

(F_2) Some Kehoe Ants have an immune system that is hyperactive. Others have a normal immune system.

(F_3) Some Kehoe Ants have blood that is very thick. Others have blood of normal thickness.

(F_4) Kehoe Ants build their nests by secreting a sticky fluid that then hardens. Some Kehoe Ants are able to build their nests quickly. Others build their nests at a normal rate.

A.1.2. Causal relationships

($F_1 \rightarrow F_2$). Blood high in iron sulfate causes a hyperactive immune system. The iron sulfate molecules are detected as foreign by the immune system, and the immune system is highly active as a result.

($F_1 \rightarrow F_3$). Blood high in iron sulfate causes thick blood. Iron sulfate provides the extra iron that the ant uses to produce extra red blood cells. The extra red blood cells thicken the blood.

($F_1 \rightarrow F_4$). Blood high in iron sulfate causes faster nest building. The iron sulfate stimulates the enzymes responsible for manufacturing the nest-building secretions, and an ant can build its nest faster with more secretions.

($F_2 \rightarrow F_3$). A hyperactive immune system causes thick blood. A hyperactive immune system produces a large number of white blood cells, which results in the blood being thicker.

($F_2 \rightarrow F_4$). A hyperactive immune system causes faster nest building. The ants eliminate toxins through the secretion of the nest-building fluid. A hyperactive immune system accelerates the production of nest-building secretions in order to eliminate toxins.

($F_3 \rightarrow F_4$). Thick blood causes faster nest building. The secreted fluid is manufactured from the ant's blood, and thicker blood means thicker secretions. Thicker secretions mean that each new section of the nest can be built with fewer applications of the fluid, increasing the overall rate of nest building.

A.1.3. Blank predicates

Composition:	has an exoskeleton coated with ionized calcium.
Disease:	has choroidal parasites attached to its eyes.
Behavior:	attacks poisonous Juniper insects.

A.2. Lake Victoria Shrimp

Lake Victoria Shrimp are found in Lake Victoria, Africa. The concentration of algae that are rich in choline is unusually high in some parts of Lake Victoria.

A.2.1. Features

(F_1) Lake Victoria Shrimp use acetylcholine (ACh) as a brain neurotransmitter. Some Lake Victoria Shrimp have an unusually high amount of ACh. Others have a normal amount of ACh.

(F_2) Lake Victoria Shrimp have a "flight" response in which they flee from potential predators. The flight response consists of an electrical signal sent to the muscles which propel the shrimp away from a predator. Some Lake Victoria Shrimp have a flight response which is long-lasting. Others have a normal flight response.

(F_3) Some Lake Victoria Shrimp have an accelerated sleep cycle (4 hours sleep, 4 hours awake). Others have a normal sleep cycle (12 hours sleep, 12 hours awake).

(F_4) Some Lake Victoria Shrimp have high body weight. Others have a normal body weight.

A.2.2. Causal relationships

($F_1 \rightarrow F_2$). A high quantity of ACh neurotransmitter causes a long-lasting flight response. The duration of the electrical signal to the muscles is longer because of the excess amount of neurotransmitter.

($F_1 \rightarrow F_3$). A high quantity of ACh neurotransmitter causes an accelerated sleep cycle. The neurotransmitter speeds up all neural activity, including the internal "clock" which puts the shrimp to sleep on a regular cycle.

($F_1 \rightarrow F_4$). A high quantity of ACh neurotransmitter causes a high body weight. The neurotransmitter stimulates greater feeding behavior, which results in more food ingestion and more body weight.

($F_2 \rightarrow F_3$). A long-lasting flight response causes an accelerated sleep cycle. The long-lasting flight response causes the muscles to be fatigued, and this fatigue triggers the shrimp's sleep center.

($F_2 \rightarrow F_4$). A long-lasting flight response causes a high body weight. The shrimp are propelled over a greater area of the lake, and find more new food sources as a result.

($F_3 \rightarrow F_4$). An accelerated sleep cycle causes a high body weight. Shrimp habitually feed after waking, and shrimp on an accelerated sleep cycle wake three times a day instead of once.

A.2.3. Blank predicates

Composition:	has mucus that is slightly acidic.
Disease:	is infected with a fungal retrovirus that attacks its reproductive system.
Behavior:	engages in cannibalistic eating behavior.

A.3. Myastars

In certain parts of the known universe there exists a large number of stars called Myastars. Myastars are formed from clouds of helium.

A.3.1. Features

(F_1) Some Myastars are constructed from ionized helium. Others are constructed from normal helium.

(F_2) Some Myastars are very hot. Others have a normal temperature.

(F_3) Some Myastars are extremely dense. Others have a normal density.

(F_4) Some Myastars have a large number of planets. Others have a normal number of planets.

A.3.2. Causal relationships

($F_1 \rightarrow F_2$). Ionized helium causes the star to be very hot. Ionized helium participates in nuclear reactions that release more energy than the nuclear reactions of normal hydrogen-based stars, and the star is hotter as a result.

($F_1 \rightarrow F_3$). Ionized helium causes the star to have high density. Ionized helium is stripped of electrons, and helium nuclei without surrounding electrons can be packed together more tightly.

($F_1 \rightarrow F_4$). Ionized helium causes the star to have a large number of planets. Because helium is a heavier element than hydrogen, a star based on helium produces a greater quantity of the heavier elements necessary for planet formation (e.g. carbon, iron) than one based on hydrogen.

($F_2 \rightarrow F_3$). A hot temperature causes the star to have high density. At unusually high temperatures heavy elements (such as uranium and plutonium) become ionized (lose their electrons), and the resulting free electrons and nuclei can be packed together more tightly.

($F_2 \rightarrow F_4$). A hot temperature causes the star to have a large number of planets. The heat provides the extra energy required for planets to coalesce from the gas in orbit around the star.

($F_3 \rightarrow F_4$). High density causes the star to have a large number of planets. Helium, which cannot be compressed into a small area, is spun off the star, and serves as the raw material for many planets.

A.3.3. Blank predicates

Composition:	has adiabatic solar whirlpools on its surface.
Disease:	rotates around an asymmetrical gravitational axis that will lead to it breaking apart.
Behavior:	emits coherent neutrinos.

A.4. Meteoric Sodium Carbonate

A special form of sodium carbonate (Na_2CO_2) is found in meteors that land on earth. Molecules of "meteoric" sodium carbonate differ from molecules of normal sodium carbonate that are found on earth in that they have been exposed to intense X-rays in space.

A.4.1. Features

(F_1) Some meteoric sodium carbonate molecules are radioactive, i.e. theta particles get emitted from the nuclei of the sodium (Na) atoms. Other meteoric sodium carbonate molecules are non-radioactive.

(F_2) Some molecules of meteoric sodium carbonate have their five atoms arranged in an eight-bond pyramid (four atoms at the base of the pyramid, and one at the "peak"). Other molecules of meteoric sodium carbonate have their five atoms arranged in a normal five-bond ring, as in normal sodium carbonate found on earth.

(F_3) Some molecules of meteoric sodium carbonate are positively charged. Others have no charge.

(F_4) Some molecules of meteoric sodium carbonate are very reactive (tend to enter into chemical reactions). Others react at normal levels.

A.4.2. Causal relationships

($F_1 \rightarrow F_2$). Radioactivity causes the molecule to take on a pyramid structure. Theta particles provide the extra energy required to form the additional atom-to-atom bonds required for the pyramid.

($F_1 \rightarrow F_3$). Radioactivity causes the molecule to have a positive charge. Theta particles are negatively charged, and so leave the molecule with a positive charge after they are emitted.

($F_2 \rightarrow F_3$). The pyramid structure causes the molecule to have a positive charge. Atoms are packed close together in the pyramid structure, and so are able to share electrons. Because they can be shared, the molecule has fewer negatively-charged electrons, and hence the molecule has an overall positive charge.

($F_1 \rightarrow F_4$). Radioactivity causes the molecule to be reactive. Theta particles break up surrounding molecules and hence accelerate the natural rate of chemical reactions.

($F_2 \rightarrow F_4$). The pyramid structure causes the molecule to be reactive. Once one atom of the pyramid is involved in a chemical reaction, the remaining atoms break apart, providing the raw material for further reactions.

($F_3 \rightarrow F_4$). Having a positive charge causes the molecule to be reactive. The molecule attracts negatively-charged subparts of other molecules, which breaks up the other molecules, and causes chemical reactions.

A.4.3. Blank predicates

Composition:	has a high concentration of light tachyon particles.
Disease:	has an oscillating quantum field that will lead to its dissolution.
Behavior:	behaves as a catalyst for rare earth reactions.

A.5. Romanian Rogos

The Romanian Motor Company, located in Bucharest, Romania, manufactures an automobile called a Rogo which is designed to run on fuel refined locally in Romania. Depending on where it is refined, the fuel may or may not have butane (C_4H_{10}), a naturally-occurring hydrocarbon, blended in with the gasoline.

A.5.1. Features

(F_1) Some Rogos are filled with gasoline laden with butane. Other Rogos are filled with gasoline with no butane.

(F_2) The fuel filters of Rogos have gaskets. Some Rogos have fuel filter gaskets that are extra loose. Other have normal fuel filter gaskets.

(F_3) Some Rogos have a hot engine temperature. Others have a normal engine temperature.

(F_4) Some Rogos have a high amount of carbon monoxide in their exhaust. Others have a normal amount of carbon monoxide in their exhaust.

A.5.2. Causal relationships

($F_1 \rightarrow F_2$). Butane-laden fuel causes loose fuel filter gaskets. The butane tends to corrode the rubber out of which gaskets are made, and so the gaskets do not fit tightly.

($F_1 \rightarrow F_3$). Butane-laden fuel causes hot engine temperature. The butane in the fuel burns at a hotter temperature than normal gasoline.

($F_1 \rightarrow F_4$). Butane-laden fuel causes high amounts of carbon monoxide in the exhaust. Butane contains more carbon than normal gasoline, and so more carbon is available to bind with oxygen to form carbon monoxide.

($F_2 \rightarrow F_3$). A loose fuel filter gasket causes hot engine temperature. Loose gaskets allow more air to be mixed in with the fuel, meaning that the gas is more fully burned, resulting in the engine running hotter than normal.

($F_2 \rightarrow F_4$). A loose fuel filter gasket causes high amounts of carbon monoxide in the exhaust. Loose fuel filters allow more air into the gas-air mixture, providing the oxygen which binds with carbon to form carbon monoxide.

($F_3 \rightarrow F_4$). Hot engine temperature causes high amounts of carbon monoxide in the exhaust. The heat provides the energy required for the carbon to bind with the oxygen.

A.5.3. Blank predicates

Composition:	has asbestos in its muffler.
Disease:	loses its brakes at high temperatures because of the loss of brake fluid viscosity.
Behavior:	distributes weight transversely while taking fast turns.

A.6. Neptune Military Personal Computers

The power supplies for the Neptune Military Personal Computers are made from tungsten mined in southern Utah, some samples of which are magnetic.

A.6.1. Features

(F_1) Some Neptune Personal Computers have a power supply that is magnetic and extends a magnetic field. Others have a normal non-magnetic power supply that extends no magnetic field.

(F_2) Neptune Personal Computers have an internal clock based on a crystal oscillator that determines how fast the computer runs. Some Neptune Personal Computers have a clock speed that is too fast. Others have a normal clock speed.

(F_3) Some Neptune Personal Computers run at an unusually high temperature. Others run at a normal temperature.

(F_4) Some Neptune Personal Computers have a screen image that is unusually bright. Other have a screen image of normal brightness.

A.6.2. Causal relationships

($F_1 \rightarrow F_2$). Magnetic power supplies cause the computer to have a fast clock speed. The magnetic field interferes with the natural phase transitions of the crystal oscillator, the result being that the crystal oscillator emits square waves at a faster rate.

($F_1 \rightarrow F_3$). Magnetic power supplies cause the computer to run at a hot temperature. The magnetic field influences the copper atoms in electrical wire to orient themselves perpendicularly to the flow of electricity, increasing the resistance, and resulting in more heat being generated.

($F_1 \rightarrow F_4$). Magnetic power supplies cause the computer to display a bright image. The magnetic field concentrates the electron beam which strikes the phosphor on the computer screen, leading to an image that is slightly smaller but brighter.

($F_2 \rightarrow F_3$). A fast clock speed causes the computer to run at a hot temperature. With a faster clock speed the computer runs faster, performs more operations, and generates more heat as a result.

($F_2 \rightarrow F_4$). A fast clock speed causes the computer to display a bright image. The clock controls how fast the image is "repainted" on the screen. A faster clock means that the phosphors on the screen's surface are being irradiated with electrons more often, leading to a brighter image.

($F_3 \rightarrow F_4$). Hot temperature causes the computer to display a bright image. Heat increases the efficiency of the cathode ray tube, leading to a more energized electron beam and a brighter screen.

A.6.3. Blank predicates

Composition:	has disk drives that store a large electro-static charge.
Disease:	commits rounding errors during Fourier floating-point operations.
Behavior:	survives unexpected power surges during lightning storms.

References

Ahn, W. (1998). Why are different features central for natural kinds and artifacts? the role of causal status in determining feature centrality. *Cognition*, *69*, 135– 178.

Ahn, W. (1999). Effect of causal structure on category construction. *Memory and Cognition*, *27*, 1008–1023.

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, *41*, 361–416.

Ahn, W., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, *16*, 81–121.

Atran, S. (1990). *Cognitive foundations of natural history*. New York: Cambridge University Press.

Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experiment Psychology: Learning, Memory, and Cognition*, *11*, 629–654.

Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition*, *60*, 1–29.

Bloom, P. (1998). Theories of artifact categorization. *Cognition*, *66*, 87–93.

Brown, R. (1958). How shall a thing be called? *Psychological Review*, *65*, 14–21.

Coley, J. D., Medin, D. L., & Atran, S. (1997). Does rank have its privilege? Inductive inference within folkbiological taxonomies. *Cognition*, *64*, 73–112.

Corter, J. E., & Gluck, M. A. (1992). Explaining basic categories: feature predictability and information. *Psychological Bulletin*, *111*, 291–303.

Gelman, S. A. (1988). The development of induction within natural kinds and artifact categories. *Cognitive Psychology*, *20*, 65–95.

Gelman, S. A., & Coley, J. D. (1991). Language and categorization: the acquisition of natural kind terms. In S. A. Gelman, & J. P. Byrnes (Eds.), *Perspectives on language and cognition: interrelations in development* (pp. 146–196). Cambridge: Cambridge University Press.

Gelman, S. A., Coley, J. D., & Gottfried, G. M. (1994). Essentialist beliefs in children: the acquisition of concepts and theories. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the mind* (pp. 341–367). Cambridge: Cambridge University Press.

Gelman, S. A., & Kalish, C. W. (1993). Categories and causality. In R. Pasnak, & M. L. Howe (Eds.), *Emerging theories in cognitive development* (pp. 3–32). New York: Springer Verlag.

Gelman, S. A., & Kremer, K. E. (1991). Understanding natural causes: children's explanations of how objects and their properties originate. *Child Development*, *57*, 396–404.

Gelman, S. A., & O'Reilly, A. W. (1988). Children's inductive inferences with superordinate categories: the role of language and category structure. *Child Development*, *59*, 876–887.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: early understandings of the nonobvious. *Cognition*, *38*, 213–244.

Goodman, N. (1983). *Fact, fiction, and forecast* (4th ed.). Cambridge, MA: Harvard University Press.

Gopnik, A., & Wellman, H. M. (1994). The "theory theory". In L. Hirschfeld, & S. Gelman (Eds.), *Mapping the mind: domain specificity in culture and cognition* (pp. 257–293). New York: Cambridge University Press.

Gould, S. J., & Lewontin, R. C. (1978). The spandrels of the San Marco and the Panglossian paradigm. *Proceedings of the Royal Society of London*, *205*, 581–598.

Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *18*, 441–461.

Harman, G. (1986). *Change in view: principles of reasoning*. Cambridge, MA: MIT Press.

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

Heit, E., & Bott, L. (2000). Knowledge selection in category learning. In D. L. Medin (Ed.), *The psychology of learning and motivation* (pp. 163–199). San Diego, CA: Academic Press.

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

Hirschfeld, L. A. (1996). *Race in the making: cognition, culture, and the child's construction of human kinds*. London: MIT Press.

Hirschfeld, L. A., & Gelman, S. A. (1994). Toward a topology of mind: an introduction to domain specificity. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the mind* (pp. 3–36). Cambridge: Cambridge University Press.

Kaplan, A. S., & Murphy, G. L. (1999). The acquisition of category structure in unsupervised learning. *Memory and Cognition*, *27*, 699–712.

Kaplan, A. S., & Murphy, G. L. (2000). Category learning with minimal prior knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 829–846.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

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 approach* (pp. 234–262). Oxford: Clarendon Press.

Kornblith, H. (1993). *Inductive inference and its natural ground: an essay in naturalistic epistemology*. Cambridge, MA: MIT Press.

Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago, IL: University of Chicago Press.

Lassaline, M. E. (1996). Structural alignment in induction and similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 754–770.

Lopez, 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.

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 and Cognition*, *28*, 41–50.

Malt, B. C. (1995). Category coherence in cross-cultural perspective. *Cognitive Psychology*, *29*, 85–148.

Matan, A., & Carey, S. (2001). Developmental changes within the core of artifact concepts. *Cognition*, *78*, 1–26.

Mayr, E. (1982). *The growth of biological thought*. Cambridge, MA: Harvard University Press.

Mayr, E. (1988). *Toward a new philosophy of biology*. Cambridge, MA: Harvard University Press.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179–196). Cambridge, MA: Cambridge University Press.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207–238.

Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, *19*, 242–279.

Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Murphy, G. L., & Allopenna, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 904–919.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Nisbett, R. E., Krantz, D. H., Jepson, C. J., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*, 339–363.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology*, *115*, 39–57.

Osherson, D., Smith, E. E., Myers, T. S., & Stob, M. (1994). Extrapolating human probability judgment. *Theory and Decision*, *36*, 103–126.

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

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.

Regehr, G., & Brooks, L. R. (1995). Category organization in free classification: the organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 347–363.

Rehder, B. (1999). A causal model theory of categorization (pp. 595–600). *Proceedings of the 21st annual meeting of the Cognitive Science Society*, Mahwah, NJ: Lawrence Erlbaum Associates.

Rehder, B (2003). Categorization as causal reasoning. *Cognitive Science*, 27, 709–748.

Rehder, B (in press). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

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.

Rehder, B., & Ross, B. H. (2001). Abstract coherent categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1261–1275.

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

Rips, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21–59). New York: Cambridge University Press.

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

Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, *4*, 328–350.

Rosch, E. H., & Mervis, C. B. (1975). Family resemblance: studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Rosch, E. H., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.

Shipley, E. F. (1993). Categories, hierarchies, and induction. In D. Medin (Ed.), *The psychology of learning and motivation* (pp. 265–301). New York: Academic Press.

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

Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*, 189–228.

Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Smith, E. E., Shafir, E., & Osherson, D. (1993). Similarity, plausibility, and judgments of probability. *Cognition*, *49*, 67–96.

Smith, L. B. (1981). Importance of overall similarity of objects for adults' and children's classifications. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 811–824.

Spalding, T. L., & Murphy, G. L. (1996). Effects of background knowledge on category construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 525–538.

Thagard, P., & Nisbett, R. E. (1982). Variability and confirmation. *Philosophical Studies*, *42*, 379–394.

Waldmann, M. R., Holyoak, K. J., & Fratianne, A. (1995). Causal models and the acquisition of category structure. *Journal of Experimental Psychology: General*, *124*, 181–206.

Wattenmaker, W. D., Dewey, G. I., Murphy, T. D., & Medin, D. L. (1986). Linear separability and concept learning: context, relational properties, and concept naturalness. *Cognitive Psychology*, *18*, 158–194.

Wisniewski, E. J. (1995). Prior knowledge and functionally relevant features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 449–468.

Wisniewski, E. J. (2002). Concepts and categorization. In H. Pashler, & D. L. Medin (Eds.), *Steven's handbook of experimental psychology* (pp. 467–531). New York: Wiley.

Wu, M. L., & Gentner, D. (1998). Structure in category-based induction (pp. 1154–1158). *Proceedings of the 20th annual conference of the Cognitive Science Society*, Mahwah, NJ: Erlbaum.