

The Blicket Within: Preschoolers’ Inferences About Insides and Causes

David M. Sobel

*Department of Cognitive and Linguistic Sciences
Brown University*

Caroline M. Yoachim

*Department of Psychology and Institute for Learning and Brain Sciences
University of Washington*

Alison Gopnik

*Department of Psychology
University of California at Berkeley*

Andrew N. Meltzoff and Emily J. Blumenthal

*Department of Psychology and Institute for Learning and Brain Sciences
University of Washington*

Four experiments examined children’s inferences about the relation between objects’ internal parts and their causal properties. In Experiment 1, 4-year-olds recognized that objects with different internal parts had different causal properties, and those causal properties transferred if the internal part moved to another object. In Experiment 2, 4-year-olds made inferences from an object’s internal parts to its causal properties without being given verbal labels for objects or being shown that insides and causal properties covaried. Experiment 3 found that 4-year-olds chose an object with the same internal part over one with the same external property when asked which object had the same causal property as the target (which had both the internal part and external property). Finally, Experiment 4 demonstrated that 4-year-olds made similar inferences from causal properties to internal parts, but 3-year-olds relied more on objects’ external perceptual appearance. These results suggest that by the age of 4, children have developed an understanding of a relation between an artifact’s internal parts and its causal properties.

Classical research in cognitive development has suggested that preschool children do not understand causality (e.g., Piaget, 1929, 1930). However, contemporary studies, particularly in the framework of “naive theories,” suggest that 3- to 5-year-old children do understand the causal relations involved in everyday physics (e.g., Bullock, Gelman, & Baillargeon, 1982; Shultz, 1982), biology (e.g., Inagaki & Hatano, 1993; Kalish, 1996), and psychology (Gopnik & Wellman, 1994; Wellman, 1990). By this age, children can also make causal predictions (e.g., Shultz, 1982), generate causal explanations (e.g., Schult & Wellman, 1997), and understand counterfactual claims (e.g., Harris, German, & Mills, 1996).

Several different types of information might go into making such causal inferences. In this article, we focus on two types of information. The first is the pattern of covariation among events—events that are causally related tend to co-occur. The second involves the ways in which objects and events affect each other. Certain physical relations are more likely to be causal than others. For example, a billiard ball that moves toward and makes physical contact with a second billiard ball is seen as causing that second billiard ball to move. When such cues are interpreted as causal, they are often described in terms of *mechanisms*. In the billiard ball example, the mechanism is that an underlying force or energy is transferred when the first ball makes spatial contact with the second. In this article, we focus on a particular physical relation: between objects’ causal properties and the internal parts of objects.

There is extensive evidence that adults make causal inferences based on simple patterns of covariation (Allan, 1980; Shanks, 1995) and on more complex patterns of covariation (e.g., Cheng, 1997) in the absence of mechanistic cues. Preschoolers can also make causal inferences based on patterns of covariation (Gopnik, Sobel, Schulz, & Glymour, 2001; Schulz & Gopnik, 2004; Sobel, Tenenbaum, & Gopnik, 2004; see also Gopnik & Schulz, 2004). Ample evidence exists that adults can use physical or mechanistic cues to make causal inferences (Ahn, Kim, Lassaline, & Dennis, 2000; Michotte, 1963), and evidence suggests such information is used even if it cannot be articulated or fully specified (e.g., Rosenblit & Keil, 2002). Young children also recognize the importance of physical relations in making causal inferences (Shultz, 1982), and infants seem to register the billiard ball causal relations captured in Michottean-type displays (Cohen & Amsel, 1998; Leslie & Keeble, 1987; Oakes & Cohen, 1990; for brain correlates of Michottean causal processing in adults, see Blakemore et al., 2001).

The question we pose is whether children can integrate covariation information and mechanism information. Can young children use patterns of correlation involving objects (particularly artifacts) to make inferences about a particular physical property of those objects—that is, their internal structure? Conversely, can young children use information about the objects’ insides to make inferences about the potential patterns of covariation they will observe? To examine these questions, we introduced children to a novel causal property of objects, which led to a novel pattern of covariation between those artifacts and a new event—a machine activating. Then,

we examined the kinds of inferences children would make about the relation between this causal property and the internal structure of the objects.

Various studies have examined the relation between the causal properties of objects and those objects' category membership (often reflected by an object's label) as well as the relation between the internal physical properties of objects and category membership. Research examining causal properties and category membership has shown that young children recognized that objects with the same causal or functional properties should be placed in the same category (i.e., given the same label) even when their perceptual features were quite different. Similarly, children believed that objects with the same label would have the same causal properties, even when the two objects were perceptually dissimilar (Gopnik & Sobel, 2000; Kemler-Nelson, 1995; Kemler-Nelson, Russell, Duke, & Jones, 2000; Nazzi & Gopnik, 2000, 2003).

Category membership and inferences about internal structure also appear to be related (for a thorough review, see Gelman, 2003). Keil (1989) demonstrated that school-age children understood that an animal that has undergone external cosmetic changes (but remained the same on the inside) retained its original identity. Gelman and Wellman (1991) found that preschoolers differentiate between the insides and the outsides of an object and that children recognized that insides were important in determining to which category an object belongs. Specifically, in their second experiment, they demonstrated that 4-year-olds recognize a difference between *insides-relevant* and *insides-irrelevant* objects: Children recognized that changing the insides of a dog or an egg would fundamentally change the object's category membership, but changing the insides of a jar would not. Similarly, they found that 4-year-olds inferred that removing inside-relevant object's insides would change that object's function.

Further, investigations on induction have shown that preschoolers extend internal properties between two objects given the same verbal label, even if those objects are dissimilar in perceptual appearance. Indeed, children made this inference more often than when the two objects were given different verbal labels (i.e., assigned to different categories) and were similar in perceptual appearance (e.g., Gelman & Coley, 1990; Gelman & Markman, 1987). These data have often been taken to suggest that the link between category membership and internal properties was particularly important for biological kinds. An open question is whether young children believe that there is a specific connection between an object's causal properties and its internal structure, independent of the category knowledge presented by object labels. For instance, in the Gelman and Wellman (1991) experiments, whereas 4-year-olds recognized that changing an object's insides changed its function, children were always provided with verbal labels for the objects, and it might be argued that such labels could have facilitated children's inferences.

Another difficulty with investigating children's causal knowledge is that there might be a strong role of prior knowledge in children's inferences. Some theorists

have suggested that the previous categorization results reflect past histories of association among stimuli or overall patterns of similarity (e.g., Rogers & McClelland, 2004). Rather than making the general assumption that objects' internal structure and causal properties are related, children may have learned a set of specific empirical generalizations about such relations. Similarly, Harris and colleagues (e.g., Harris & Koenig, 2006; Harris, Pasquini, Duke, Asscher, & Pons, 2006) suggested that children learn a great deal from "testimony" (i.e., verbal information from others), which suggests that in previous investigations children may have relied on a verbal framework provided by adults instead of spontaneously engaging in causal inference. For example, Carey (1995) argued that children in the Gelman and Wellman (1991) experiments may have simply repeated information previously told to them by adults.

Moreover, in the natural world, objects in the same category tend to have perceptual, internal, linguistic, and causal properties in common. It can be difficult to untangle the roles of these different properties. What is needed is a method for studying children's assumptions about the relation between an object's internal structure and causal properties in which (a) children do not have substantial prior knowledge that could influence their reasoning, (b) the internal structure and external perceptual appearance of objects can be pitted against each other, and (c) the amount of verbal information about the objects can be minimized.

One such method involves a device called a "blicket detector," which was designed to present a novel causal property of an object (Gopnik & Sobel, 2000). The blicket detector is a machine that lights up and plays music when certain objects (controlled by the experimenter) are placed on top of it. In the four experiments we present, we use this device to examine how preschoolers relate causal and internal properties of objects. Will children assume that objects with the same causal properties have the same *internal* physical properties and that objects with the same internal parts have the same causal properties? Or will children rely more on the perceptual similarity of objects or patterns of association to make predictions about novel causal effects and novel internal properties of objects?

In Experiment 1, we considered whether 4-year-olds understood that if an object's internal properties were moved to another object, the original object's causal properties moved as well. In some respects, this experiment paralleled Keil's (1989) investigations about transformations on artifacts. Keil (1989) found that 5 to 6-year-olds recognized that changing the insides of an artifact would change its category membership. Experiment 1 investigated 4-year-olds, looking at their inferences about objects' causal properties (as opposed to category membership).

Although Experiment 1 used a novel causal property of objects (i.e., whether they activated a novel machine), children had to have some familiarization with the objects and machine to complete the experimental procedure. This familiarization might have been sufficient to skew children's conceptions of how causal and internal properties of objects are related. Experiments 2 and 3 used no such familiarization

and examined whether 4-year-olds understood that objects with shared internal properties would have shared causal properties, even in the face of a conflict with the object's external perceptual appearance (Experiment 2) or a hidden external property (Experiment 3). As in Experiment 1, 4-year-olds were examined to parallel the majority of prior research relating categorization with causal properties of objects (e.g., Nazzi & Gopnik, 2000) and categorization with internal properties of objects (e.g., Gelman & Wellman, 1991). Finally, Experiment 4 considered whether children would make the opposite inference and recognize that objects with shared causal properties had shared internal parts. To consider this question developmentally, we examined both 3- and 4-year-olds in this experiment.

EXPERIMENT 1

Experiment 1 examined whether 4-year-olds inferred that objects with a particular internal property (in this case, artifacts with a particular part inside) shared causal properties when that internal part was transferred from one object to another. Children were shown a demonstration in which an object that contained a particular metal part caused the detector to activate. When the same object contained a different metal part inside, it had no effect on the detector. The two internal parts were then transferred to two novel objects that were externally identical to each other, and children were asked to make the machine go. We hypothesized that children could use the pattern of evidence to infer that only a specific internal property was causally efficacious, which would suggest that children recognize that there was a relation between an object having a particular internal property and its causal efficacy—its ability to activate the detector.

Method

Participants

The sample consisted of 64 four-year-olds ($M = 54.03$ months, range = 53.6–54.4 months), recruited by telephone from a university participant list. Three additional participants were recruited but were excluded from the study for failing to pass the pretest. An equal number of boys and girls participated; 51 children were White, 2 were Hispanic, 1 was Asian American, 8 were multiracial, and 2 did not provide information about race or ethnicity. No child had participated in any previous experiment in the lab.

Materials

The blicket detector used in this experiment was similar to those used in previous experiments (e.g., Gopnik & Sobel, 2000). The detector was 17.5 cm ×

12.3 cm × 8.0 cm, and the top had a slightly recessed panel of orange plastic. When an object was placed on the detector, the detector's top receded a little, suggesting that the object was affecting the detector. The detector was controlled using a remote switchbox with an on/off switch. When the switch was *on*, the detector turned on as soon as an object made contact with it and continued to light up and play music as long as the object remained in contact. When the switch was in the *off* position, the detector did not activate even if an object was placed on it, which provided a strong impression that something about the object itself caused the effect. The switchbox was located under the table so that the experimenter could surreptitiously control whether an object would activate the machine.

Three specially constructed wooden blocks and two metal inserts were used. Each block could be opened to reveal a 3.5 cm × 3.5 cm × 2.5 cm cavity. The first block was a white pentagon. The other two blocks were both blue diamonds. The first metal insert had a silver appearance and was made of stainless steel; the second insert appeared gold and was made of brass. Both inserts could be placed inside the cavity of either block. The shape of the base was the same for both inserts, but the stainless steel insert had a square top and the brass insert had a rounded top. The two inserts were identical in weight. Additional stimuli used in the pretest were three ordinary building blocks (two yellow triangles and one red square), a white ceramic knob, and a copper t-joint.

Procedure

After a warm-up period in which the child was familiarized with the experimenter, the blicket detector and two yellow triangles were placed on the table. The experimenter said, "This is my blicket machine. Blickets make it light up and play music." The experimenter placed the one triangle on the detector, and the detector lit up and played music. The experimenter verbally labeled this triangle as a blicket. The experimenter then placed the other triangle on the detector, and the detector did not activate. The child was told that this triangle was not a blicket. Each triangle was demonstrated a total of three times. This initial demonstration was the only time the child was provided with a label (i.e., *blicket*) or causal language ("blickets *make* it light up").

Children were then given a pretest to familiarize them with the blicket detector. The three pretest objects were then placed on the table: a red square block, a piece of copper pipe, and a ceramic knob. Each object was placed on the detector, one at a time. Two of the objects, selected at random, activated the detector; the third object did not. Children were then asked whether each object was a blicket. Children who correctly answered this question for all three objects were given the main task. If children did not correctly answer all of the questions, the

pretest was repeated with three new objects. Three children did not pass the pretest after two attempts and were not included in the study.

The demonstration and test phases of the experiment are depicted in Figure 1. At the start of the demonstration phase, the two metal inserts and the white pentagon-shaped block were placed on the table. The experimenter placed the block on the detector, which did not activate. The experimenter took the block off the detec-

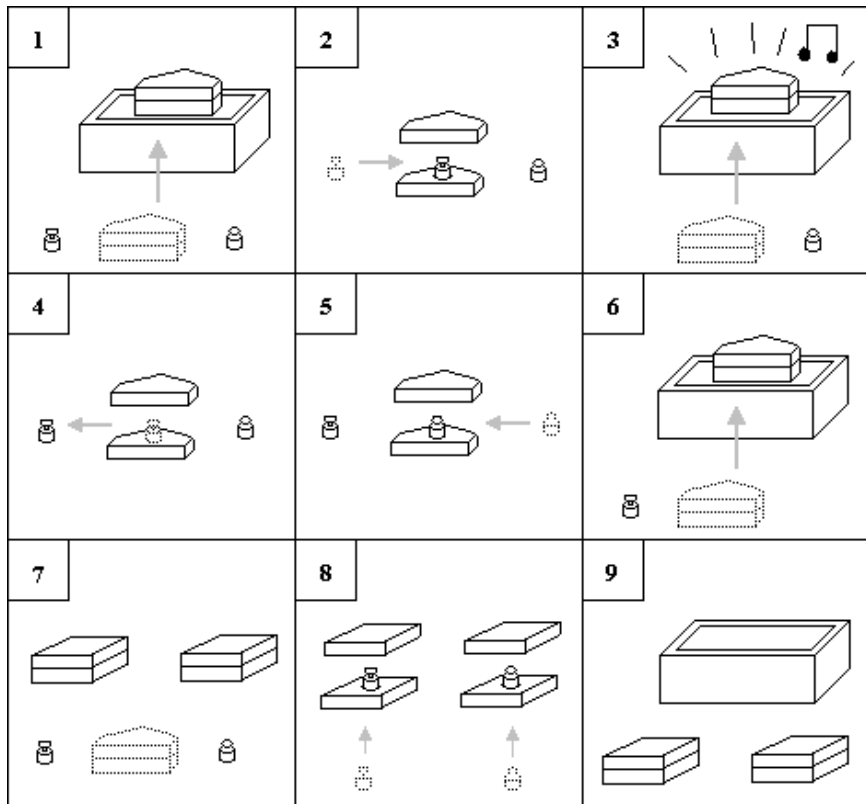


FIGURE 1. Schematic depiction of Experiment 1. The demonstration phase is shown in panels 1–6. The test phase is shown in panels 7–9. Children saw two metal inserts and one block; the block did not activate the detector (1). The block was opened, and the first insert was placed inside (2). The block was demonstrated again, and this time it activated the detector (3). Next, the first insert was removed, and the second insert was placed inside the block (4, 5). The block was placed on the detector a third time, and nothing happened (6). In the test phase, the block was then removed from view, and two externally identical blocks were placed on the table (7). The inserts were placed into the new blocks (8). The experimenter then asked the child to make the detector light up and then label objects as *blickets* (9).

tor and opened it, which revealed that the block was empty. One of the metal inserts was then put inside the block (for expository purposes, we assume that the first insert had causal efficacy, but this was counterbalanced). The block was placed on the detector again, and the detector activated. The first insert was then removed and the other insert was placed inside the block. The block was placed on the detector a third time, and the block no longer activated the detector. This pattern of events provided covariation information between a specific internal part and the machine activating.

The white pentagon-shaped block was then removed, and the two metal inserts were left on the table. The experimenter then brought out the two identical blue diamond-shaped blocks and opened them. Both blocks were empty. One metal insert was placed into each of the identical blue blocks, and then both blocks were closed.

The blocks were passed to the child, and the child was asked an *intervention* question, “Can you make my machine light up?” Note that for this measure, the child was faced with two blocks that were identical in appearance (both blue diamonds), neither of which had made the machine go in the past, and both had “insides.” The experimental question was whether children would use the block with the causally efficacious internal part based on the pattern of evidence given earlier in the demonstration. After the child’s response, the experimenter took the metal inserts out of the blocks, and asked *identification* questions: Children were asked whether each of the inserts were blickets. The question was whether children would spontaneously label the causally effective insert as a blicket. An undergraduate research assistant who was naive to the experimental hypotheses scored these two measures from videotapes. A second assistant coded the responses live. Agreement between coders was 100%.

Results and Discussion

Sixty-one percent of the children placed the block that had the causally efficacious internal part onto the detector, binomial test, one-tailed, $p = .05$.¹ Additional analyses indicated that which of the insides (silver or gold) was efficacious (counterbalanced) had no effect, $\chi^2(1, N = 64) = 0.07$, *ns*, but that there was a side preference (children selected the block with the causally efficacious insert more often when it was on the right-hand side), $\chi^2(1, N = 64) = 28.90$, $p < .001$. For the identification measure, children were asked two yes or no questions—whether each of the metal inserts was a blicket. Eighty-three percent of the children identified the causally

¹When a two-tailed test was used, this finding was not statistically significant ($p < .10$). We employed a directional test based on preliminary data, reported in Yoachim, Sobel, and Meltzoff (2005). These data showed that 4-year-olds tracked the causal power of a single internal part when it moved between two objects. Details of this experiment are available on request.

efficacious internal part as a blicket *and* the inert part as not a blicket, significantly more than expected by chance ($0.5 \times 0.5 = 0.25$), binomial test, $p < .001$. These results indicate that children were able to distinguish between two different types of insides and traced the movement of the causal agent as it was transferred from one block to another.

This experiment showed that children recognized that a specific internal property was required to activate the detector. Children made this distinction after viewing a single demonstration and could track the causal efficacy of those parts even when they were moved into new blocks (which were identical in external appearance). However, in this experiment, children were provided with covariation information relating the objects' internal properties and the detector's activation. During the initial warm-up period with the detector, they were presented with a label for objects that activated the detector (although this was not used during the demonstration phase). They were also presented with a correlation between the causally efficacious metal insert and the detector's activation. Are children responding based on this covariation information, or would inferences relating internal and causal properties of objects persist if this information was eliminated? Experiments 2 through 4 examine this question.

Another question is whether children treated the metal inserts used in these experiments as internal parts or as separate objects that were placed in a container. We do not believe children treated these inserts as separate objects: No child attempted to remove either metal insert and place it on the machine when asked to make the machine activate. Even if children believed that the insert was a separate object, the results still parallel previous work on internal causes. For instance, Kalish (1996) found that children recognized germs as an internal cause of illnesses, but that germs were not intrinsic to the person they infected. Experiment 1 suggested that children can keep track of the causal status of this internal part (regardless of whether it is considered a separate object or a property of the larger object) and recognize that its causal properties move from one external object to another even though it is the external block that is placed on the machine.

To ensure that children specifically reasoned about internal parts of objects, in Experiments 2 through 4, we used a new set of objects for which the internal property was more clearly an intrinsic part of the object and was not removable or separable from the object. We examined whether children would treat objects with similar internal parts as having similar causal properties. If children reasoned that two objects with the same type of internal part produced the same causal property over another object that shared the same external perceptual appearance but lacked that internal part, then it would suggest that children recognized that such internal parts were connected to objects' causal properties. In Experiments 2 and 3, children were shown whether objects contained these Parts and were asked to make inferences about whether the objects would activate the

blicket detector. Children were shown whether each object activated the detector and were asked to make an inference about whether the objects contained internal parts.

EXPERIMENT 2

In Experiment 2, children were shown the internal properties of three new objects (i.e., whether they contained an internal part). They were then shown that a target object with an internal part activated the detector. Then, children were asked to select another object that would activate the detector. They could choose either an object that was externally identical to the target object, but internally dissimilar, or an object that was externally dissimilar but had a similar inside. Children were given no covariation or linguistic cues to guide their inferences, and the internal property could not be removed from the block, making it more likely to be interpreted as intrinsic to the object. Because multiple data points could be collected using this procedure, we used a smaller sample size than in Experiment 1.

Method

Participants

Twenty-four 4-year-olds ($M = 53.8$ months, range = 52.7–55.6 months) were recruited by telephone from a university participant list. An equal number of boys and girls participated; 18 children were White, 2 were multiracial, 1 was Hispanic, 1 was African American, 1 was Asian American, and 1 did not provide this information. No child had participated in any previous experiment in the lab.

Materials

This study used the same blicket detector as in Experiment 1. Twelve wooden blocks of various shapes, colors, and sizes were also used (see Figure 2). A 1.3 cm (diameter) \times 2.5 cm (depth) hole was drilled into each block. Eight of the blocks contained a large white map pin inside the hole; the remaining four blocks were empty. The map pin was placed deep enough so that the child could see it and touch it, but could not remove it from the block. Each block also had a removable dowel that covered the hole, making it impossible to see whether any individual block contained a map pin without removing the dowel. These blocks were divided into four sets of three. In each set, two blocks had identical external perceptual features. One of these blocks contained a map pin. The other did not contain a map pin. A third block differed from the other two in both color and shape but did contain a map pin.

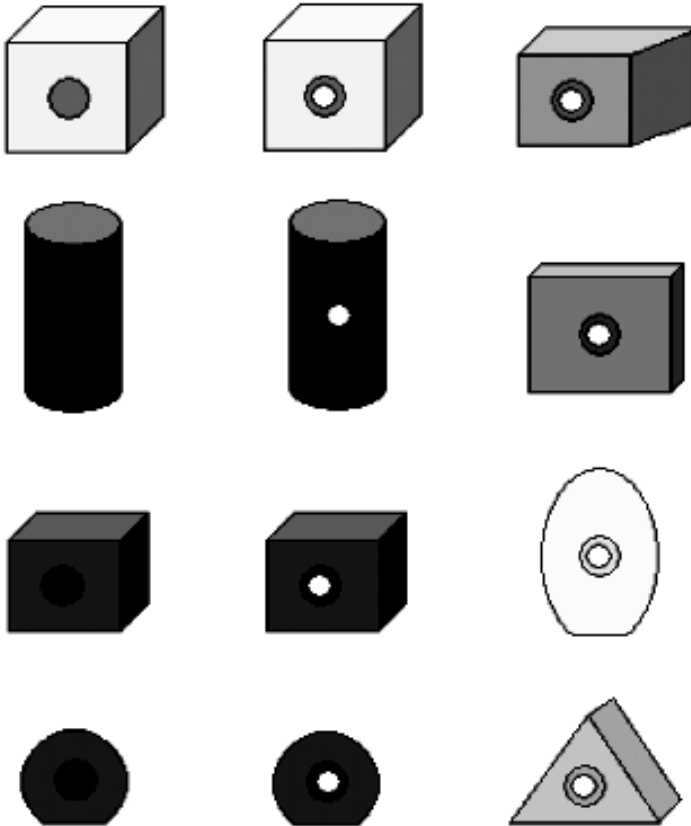


FIGURE 2. Stimuli from Experiments 2 and 4. Each row represents one object set. In each set, the object in the center is the target object (externally similar to the object on the left, internally similar to the object on the right). Similar stimuli were used in Experiment 3, but in that experiment, all three objects were the same shape, color, and size.

Four ordinary building blocks were used in a warm up: a large blue rectangular block, a red square block, a small yellow rectangular block, and a smaller green square.

Procedure

Warm up. Children were first given a warm-up activity to ensure that they would point to objects and respond to the experimenter. Four warm-up blocks were

brought out, and children were asked to point to the biggest one, then the red one, then the smallest one, and then the blue one. If children failed to answer or responded incorrectly, corrective feedback was given. Such feedback was rarely necessary (for only 6% of the questions).

Test phase. After the warm up, the blicket detector was brought out, and children were told that they were going to play a game with a very special machine. The first set of three objects was brought out (see Figure 2). The position of the three objects was counterbalanced across participants. Children were shown that each of the objects had a small dowel inserted in it. The experimenter picked up each object, removed the dowel, and pointed out an external feature of the object and whether the block had an internal property (in counterbalanced order). For instance, the experimenter might say, "This one is blue, and it has a little white thing inside," or, "This one is empty, and it is blue." One of the externally identical objects (the target object) contained a white internal part; the other did not. The third, externally distinct object contained the same internal part. This information was presented twice. We described both an internal and external property of the object to ensure that children would not pay more attention to either one of the two features. The order in which these properties were mentioned was counterbalanced across participants.

Afterward, the experimenter selected the target object (with its dowel inserted, so that its contents were not visible) and placed it on the blicket detector, which activated. The experimenter said, "Look at that, it makes the machine light up." The experimenter then asked, "Can you point to another one that makes the machine light up?" Children were not allowed to try any object on the detector, nor were they given feedback on their answers. Children were given four such trials in one of four randomly predetermined orders.

Results and Discussion

The child's task was to pick the other object that would make the machine light up. For each trial, children were given a score of 1 if they selected the internally similar object and zero if they selected the externally similar object. Preliminary analyses revealed no effect of the order in which the object sets were presented, Kruskal–Wallis test, $\chi^2(3) = 1.30$, *ns*, and children were not more likely to select the internally similar object as a function of trial number, Cochran's $Q(3) = 1.53$, *ns*. Therefore, the data from the four trials were combined to make an overall total score that ranged from zero to 4. These data are shown in Table 1, as well as what pattern of performance would be expected by chance responding.

Children chose the internally similar object on approximately 66% of the trials ($M = 2.63$, $SD = 1.21$). The data were first compared against chance responding. Because two of the expected values were below 5, we collapsed the data into three

TABLE 1
Experiments 2, 3 and 4: Number of Children Making
Internal/Causal Responses

	<i>Number of Internal/Causal Responses (out of 4)</i>					<i>Mean</i>	<i>SD</i>
	<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>		
Experiment 2							
4-year-olds (<i>N</i> = 24)	1	4	5	7	7	2.63	1.21
Expected by chance	1.5	6	9	6	1.5		
Experiment 3							
4-year-olds (<i>N</i> = 16)	0	1	4	11	0	2.62	0.62
Expected by chance	1	4	6	4	1		
Experiment 4							
Causal group							
3-year-olds (<i>N</i> = 16)	7	4	1	2	2	1.25	1.48
4-year-olds (<i>N</i> = 16)	2	0	3	4	7	2.88	1.36
Expected by Chance	1	4	6	4	1		
Associative Group							
3-year-olds (<i>N</i> = 16)	10	3	1	1	1	0.75	1.24
4-year-olds (<i>N</i> = 16)	10	2	3	1	0	0.69	1.01
Expected by chance	1	4	6	4	1		

response patterns: internal responses (selected the internally similar object on 3–4 trials), neutral responses (selected internally and externally similar objects equally often), and external responses (selected the internally similar object on 0–1 trials). The distribution of selections differed significantly from chance, $\chi^2(2, N = 24) = 8.24, p = .016$. Most important, significantly more children (58.3%) fell into the internal response category (selected the internally similar object on three or four of the trials) than would be expected by chance (31.25%), binomial test, $p = .005$.

Children inferred that objects that shared a similar internal part (a little white thing inside) also shared a novel causal property (activating the detector). Children relied on the presence of this internal part, rather than the object's external perceptual appearance, to predict the object's efficacy in activating the detector. Children made this inference without any training or prior exposure to the detector. Similarly, children did not receive feedback, and there were no differences in performance in earlier versus later trials, indicating that children did not simply learn the relation between internal parts and causal properties over the course of the experiment.

These data suggest that 4-year-olds presented with these stimuli believe that objects' internal properties are more reliable than their external appearance when making an inference about those objects' causal properties. However, the procedure used here might have called more attention to the internal property than the

perceptual similarity between the objects. Telling children about the internal part also revealed a nonobvious property of the object: Each object's external perceptual appearance was always visible, but the internal parts were not because the dowels were present. That the internal parts were hidden and then revealed might have made them particularly salient. In Experiment 3, both an internal and an external property were revealed during the course of the trial, so they should be equally salient.

EXPERIMENT 3

In Experiment 3, children were shown three objects of similar shape, color, and size. One object had both a hidden internal part and hidden external property (situated on its rear face). One object had the internal part but not the external property, and one had the external property but not the internal part. After pointing out the internal and external features of each object, we showed children that the object with both the internal and external property activated the blicket detector and asked them to pick out another object that did so. Would children choose objects with a shared internal part or a shared external feature?

Methods

Participants

The sample consisted of 16 four-year-olds (9 girls, $M = 53.1$ months; range = 49.0–62.0 months) recruited by telephone from a list of hospital births. Eleven children were White, and there was one child each from the following other groups: African American, Asian American, Middle Eastern, Hispanic, and Pacific Islands. No child had participated in any previous experiment in the lab.

Materials

A blicket detector similar to the one used in previous experiments was used. This detector was gray with a red top, but all other functions were the same. Four new sets of three blocks were constructed. In each set, the three blocks were identical in size, color, and shape, and each had a hole and dowel as the blocks in Experiment 2. One block in each set contained a white map pin and also had a white circular sticker (1.9 cm in diameter) attached to it, facing away from the child, so that it could be placed on the table without the child seeing it. Another block possessed a sticker in a similar location but had no map pin inside. The third block had no sticker but did contain a map pin. The same blocks in the warm up in Experiment 2 were also used.

Procedure

Children received the same warm-up procedure as in Experiment 2. Corrective feedback was rarely necessary (for only 3% of the questions).

The first set of three blocks was brought out and placed on the table such that children could not see whether each had a sticker on it. Children were told that they were going to play a game with these toys and, specifically, that “some of the toys have things inside them, and some do not; and some of these toys have stickers on them and some do not.”

Children were then shown whether each block had a white part inside, a sticker on its back, or both. Specifically, the dowel was removed, revealing the insides of the object, and then the object was flipped over, revealing whether it had a sticker on its back. The experimenter narrated these properties, for example by saying, “This one has a white thing inside and a sticker on its back.” The order of these statements was constant for the trials, but counterbalanced across participants (i.e., half were told about the internal property first, half about the sticker first). After children observed and were told about the internal and external properties of each object, the experimenter brought out the blinket detector, and children were told that some of the toys made the machine go, and some did not. On each trial, children were shown that the object with both the internal and external property activated the machine. Children were asked to point to another object that would also make the machine go. At this point in the procedure, all three objects had dowels on them, concealing their insides, and the objects were placed on the table such that children could not see whether any had a sticker on it.

Children received four trials of this procedure. The blinket detector was present on the table for the second through fourth trial. The order of the trials was presented in one of four quasi-random orders, counterbalanced across participants. The spatial location of the three blocks was randomly determined for each trial.

Results and Discussion

For each trial, children were given a score of 1 if they selected the internally similar object and zero if they selected the object with the shared external property. Preliminary analyses revealed no effect of the order in which the object sets were presented, Kruskal–Wallis test, $\chi^2(3) = 5.63$, *ns*, nor were children affected by whether they were shown the sticker first or the internal property first, Mann–Whitney $U = 29.50$, $z = -0.32$, *ns*. Children were also not more likely to select the internally similar object differently among the four trials, Cochran’s $Q(3) = 1.98$, *ns*. The data were combined to make a total score that ranged from zero to 4. These data are shown in Table 1.

Children chose the internally similar object on approximately 66% of the trials ($M = 2.62$, $SD = 0.62$). As in Experiment 2, the data were compared against chance

responding, and as before we collapsed the data into three patterns of response: internal responses (selected the internally similar object on three or four trials), neutral responses (selected internally and externally similar objects equally often), and external responses (selected the internally similar object on zero or one trials). The distribution of selections differed significantly from chance, $\chi^2(2, N = 16) = 11.07, p = .004$. Furthermore, significantly more children ($11/16 = 69\%$) fell into the internal response category than would be expected by chance (31.25%), binomial test, $p = .002$.

Children inferred that objects with shared internal properties were more likely to produce a novel causal property (i.e., activate the blicket detector) than objects that shared similar external features. Children made this inference with no training or prior exposure to the blicket detector, and because children did not appear to be changing their responses over trials, they were not simply learning the relation between the internal parts and causal property of the objects during the course of the experiment.

EXPERIMENT 4

In Experiment 4, we explored inferences in the other direction: from causal properties to internal properties. We also wanted to ensure that children were focusing on the *causal* properties of these objects and not just associations between specific object properties and the machine's activation. In the previous experiments, children might not have recognized that the interaction between the objects and detector was causal—that the object caused the detector to activate. It is possible that they only thought that the objects were associated with the detector's activation and responded based on this association. Experiment 4 included a control group in which objects were associated with the detector's activation but did not cause the activation. We examined whether children would make inferences about internal properties when objects shared causal properties but not when objects were only associated with the detector's activation.

Moreover, Experiments 1 through 3 demonstrated that 4-year-olds make inferences from an object's internal properties to its causal properties, but what about younger children? Several investigations have suggested children younger than 4 use an object's functional properties as a guide to categorization (e.g., Kemler-Nelson, 1995; Kemler-Nelson et al., 2000). However, in these cases, the artifacts' functions were directly observable. Kelemen (2001, 2004) demonstrated that 4- to 5-year-olds, but not 3-year-olds, relied on an artifact's intended function (an unobservable property) as a guide to categorization, as opposed to the artifact's actual function (see also Matan & Carey, 2001), which suggests that children's ability to make inferences about unobservable causal properties might lag behind inferences about what is directly observable.

Method

Participants

The participants were 32 three-year-olds (16 girls, $M = 42.50$ months, range = 36.0–47.0 months) and 32 four-year-olds (17 girls, $M = 56.0$; range = 48.0–63.0 months). Children were recruited from two preschools and flyers posted in local preschools. Seven other children were tested but not included: 2 were excluded because of experimental error, 1 was a nonnative speaker of English, and 4 had participated in a different experiment with the blicket detector prior to the experimental session. Fifty-two children were White, 3 were Hispanic, and 7 were Asian, and 2 were African American. No child in the final sample had participated in any previous experiment in the lab.

Materials

The same blicket detector as in Experiment 3 and the same blocks as in Experiment 2 were used here.

Procedure

Children were given the same warm up as in Experiments 2 through 3. Corrective feedback was rarely necessary (7% of the time). Children who required corrective feedback were included in the final sample. An equal number of children in each age group were randomly assigned to one of two groups. In the *causal* group, a set of three objects was brought out: the target object, the externally similar object, and the internally similar object. Children were not shown the internal properties of the objects. The objects were placed on the machine one at a time. The target object and the internally similar (but externally distinct) object activated the detector; the third did not. Children saw this demonstration twice. No causal language or object labels were used to describe these events; children were only told to look.

A second group was given an *association control* similar to the one used by Gopnik and Sobel (2000). This group was treated identically to the causal group, but rather than placing each object on the detector, the experimenter held each object approximately 6 in. above the detector with one hand. The experimenter's other hand was positioned on top of the detector. When objects that would have activated the detector for the causal group (i.e., the objects with the map pins inside) were held over the detector, the experimenter pressed down on the top of the detector (causing the top to depress slightly, just like an object being placed on it would), activating it. Children in the association group observed the same association between the objects and the detector as children in the causal group; however, the experimenter's hand was a clear alternative cause, which indicated that the objects themselves did not cause the blicket detector to light up but were merely associated with that effect.

For both groups, the experimenter then said, referring to the objects, "Look, they have little doors on them. Let's open one up." The experimenter then selected the target object and opened it to reveal that it had a white pin inside it. The experimenter said, "Oh, look, it has a little white thing inside of it. Can you point to another one with a white thing inside it?" Children were then given the opportunity to respond. The experimental question was whether they chose the object that had the same external features as the target or the one that had the same causal properties (made the machine light up). Children were not allowed to open up the other objects, nor were they given feedback on their answer to this question. Children were given four such trials in one of four quasi-random orders, counterbalanced across participants. The spatial location of the three objects was randomly determined for each trial.

Results and Discussion

For each trial, children were given a score of 1 if they selected the other object paired with the machine's activation and a score of zero if they selected the externally similar object. Preliminary analyses revealed no effect of the order in which the object sets were presented, Kruskal-Wallis test, $\chi^2(3, N = 64) = 2.78, ns$, and children were not more likely to make a causal response on any individual trial, Cochran's $Q(3) = 1.80, ns$. Therefore, the data were combined to form an overall score that ranged from zero to 4. These data are shown in Table 1.

Three-year-olds chose the internally similar object on approximately 31% of the trials in the causal condition ($M = 1.25, SD = 1.48$), and 19% of the trials in the association condition ($M = 0.75, SD = 1.24$). Four-year-olds chose the internally similar object 72% of the time in the causal condition ($M = 2.88, SD = 1.36$) and 17% of the time in the association condition ($M = 0.69, SD = 1.01$). These data were analyzed by a 2 (Age Group) \times 2 (Condition) analysis of variance, which revealed main effects of age, $F(1, 60) = 5.91, p = .018$, and condition, $F(1, 60) = 17.48, p < .001$. A significant interaction between age and condition was also found, $F(1, 60) = 6.89, p = .011$. This interaction resulted from 4-year-olds in the causal group scoring higher (i.e., producing more causal responses) than 3-year-olds in the causal group, $t(30) = -3.23, p < .01$ with a Scheffe correction, with no such age difference in the association group. Similarly, 4-year-olds scored higher in the causal group than the association group, $t(30) = -5.16, p < .01$ with a Scheffe correction, and 3-year-olds showed no significant difference between the causal and association conditions. We also supplemented this with a nonparametric analysis of the data. For the 4-year-olds, there was a significant difference between the causal and associative groups, Mann-Whitney $U = 31.50, z = -3.78, p < .001$. The 3-year-olds did not show such a difference, Mann-Whitney $U = 101.50, z = -1.09, ns$.

We next considered the data against chance responding. As in Experiment 2, we collapsed the data into three categories: causal responders (selected the other object paired with the machine's activation on three or four trials), neutral responders (selected causally and externally similar objects equally often), and external responders (selected the externally similar object on three or four trials). The distribution of 4-year-olds' selections differed significantly from chance in both the causal group, $\chi^2(2, N = 16) = 10.50, p = .005$, and the associative group, $\chi^2(2, N = 16) = 14.50, p = .001$, but, importantly, did so in opposite directions. Four-year-olds were causal responders in the causal condition (69% of the time) and external responders in the associative condition, (75% of the time) significantly more often than would be expected by chance responding (31.25%), binomial tests, both p values $< .005$. The distribution of 3-year-olds' selections also differed significantly from chance in both the causal condition, $\chi^2(2, N = 16) = 18.38, p < .005$, and the associative condition, $\chi^2(2, N = 16) = 27.78, p < .001$. In both cases, 3-year-olds were external responders more often than would be expected by chance (69% and 81% of the time respectively, compared with 31.25%), binomial tests, both p values $< .005$.

The data suggest that 4-year-olds used information about an object's causal properties to predict whether it contained an internal part, even though this information conflicted with the external perceptual appearance of the objects. Younger children did not show this pattern and, instead, relied more on the external perceptual features of the objects to make inferences about internal parts. As in Experiments 2 and 3, 4-year-olds made this inference without any familiarization with the machine or the objects. Moreover, the control group showed that when the objects were only associated with the machine's activation, 4-year-olds relied more on the external perceptual features of the objects to make inferences, which suggests that by age 4, children recognize that there is a link between the object's causal and internal properties.

GENERAL DISCUSSION

This study suggests that by the age of 4, children can use causal information derived from patterns of covariation to make inferences about the internal structure of objects, and vice versa. These inferences were made even when they were in conflict with the external perceptual similarity of the objects. Instead of choosing an object that was the same shape, size, and color as the target, children chose an object that was internally (Experiment 2) or causally (Experiment 4) consistent. Children also do not simply rely on their specific prior knowledge about objects and machines. In Experiments 2 through 4, children made such inferences based on their first exposure to the blicket detector—a completely new device.

We are certainly not claiming that external perceptual appearance has no bearing on inferences about an object's internal properties. In Experiment 4, when 3- and 4-year-olds were only given associative (and not causal) information, children relied on the external perceptual appearance of the objects to make inferences about their internal structure. This finding is consistent with the hypothesis that, in general, objects that share external perceptual features also share linguistic category membership and nonobvious features, such as internal properties. However, when this information is in conflict, 4-year-olds appear to rely more on causal properties than external perceptual information when making inferences about objects' internal properties. Experiments 2 and 3 also suggest that the inferences can run in the other direction as well: By age 4, children relied more on internal properties than external perceptual features when making inferences about an object's causal properties. These data are consistent with the hypothesis that by age 4, children's causal knowledge integrates covariation and mechanism information (Cheng, 1997; Gopnik, Glymour, Sobel, Schulz, Kushnir, & Danks, 2004; Woodward, 2003).

How can we explain the developmental change between ages 3 and 4 that was observed in Experiment 4? One possibility is to consider children's past history of exposure to objects (e.g., Rogers & McClelland, 2004). Because children usually see that objects with similar external perceptual properties have similar causal properties and similar internal properties, they might initially associate all three types of information. Learning that causal and internal properties may still be related, even in light of competing perceptual appearances, simply takes more time, because there are few examples of such conflict in the environment.

Another possibility is that children initially believe that category membership, an object's perceptual appearance, internal structure, and causal properties are all related, but their understanding of this relation is not fully developed. Younger children may understand that category membership and internal properties are connected, especially in the biological domain (Gelman, 2003; Gelman & Wellman, 1991) and equally appreciate a link between categories and causal properties (Gopnik & Sobel, 2000; Kemler-Nelson et al., 1995, 2000) but not yet appreciate that there is relation between object's causal properties and internal properties. To build the link between objects' causal properties and their internal structure, children might rely on their earlier belief that an object's causal properties are more important than its perceptual appearance in determining its linguistic category membership. This understanding would suggest that providing linguistic labels for the objects—a key to category membership—might affect performance. Specifically, in Experiment 4, if objects that activated the detector were all labeled *blickets*, then the younger children might connect objects' causal properties with their internal structure more easily. Similarly, if objects with the same perceptual appearance were all given the label *blicket*, perhaps 4-year-olds' responses would be more like the younger children.

There might be still other reasons for the developmental difference. Unlike most previous cases in which children inferred that the internal properties of biological kinds were responsible for their causal properties, these experiments involved novel physical artifacts. One possibility is that young children's knowledge of the relation between causal and internal properties in the physical domain does not develop until age 4. Indeed, Gottfried and Gelman (2005) demonstrated that although 4-year-olds make appropriate inferences about the internal parts of biological and physical entities (namely hearts and gears respectively), younger children had more difficulty with these inferences, especially in the physical domain.

Finally, the procedures used across all these experiments were quite difficult. For instance, in Experiment 4, children had to remember the causal property of each object—a nonobvious property—and use that information instead of the objects' more obvious perceptual information to base a response, which shows a strong demonstration that 4-year-olds understand the relation between objects' causal properties and insides. But, one could imagine the developmental difference observed in Experiment 4 resulted from these task demands, and if such demands were reduced, 3-year-olds might respond in a similar manner to the 4-year-olds here. This question remains for future investigation.

Another open question is whether children, at least by age 4, recognize that in these experiments an object's internal parts can be responsible for its causal properties—that is, those insides act as a *mechanism* for the object's causal effects. Some evidence exists that this is the case: In Experiment 1, 4-year-olds could keep track of a causally efficacious and inefficacious internal part, which suggests that they could use covariation information to identify a specific mechanism (a single type of metal insert) rather than a more general mechanism (the presence of any metal insert) as being responsible for the activation of the blicket detector. Being able to make this type of distinction is critical to understanding phenomena in the real world. For instance, just from looking at a battery, we cannot tell whether it is charged or dead. However, we know that charged batteries will operate electronics, whereas dead batteries will not, and this is important for real-world actions (e.g., if your remote control fails to operate your TV, the first thing many people do is replace the batteries) as well as a scientific understanding of electricity. Further, the results of Experiment 4 suggest that children (at least 4-year-olds) relate internal parts and causal efficacy only when the object appears to cause an event, as opposed simply to being associated with its occurrence.

Overall, these results suggest that children recognize a link between objects' causal properties and their internal structure. This study involves a particular class of objects: mechanical artifacts. We do not know whether insides are equally important for other types of objects or for causal relations beyond the domain of physical mechanical causation. These results also do not demonstrate that children spontaneously seek out internal causes as explanations of objects' causal properties. Such explanation data would support the hypothesis that children represent

their causal knowledge more abstract ways than simple association (e.g., Wellman & Liu, in press). However, what we do show is that by the age of 4, children appear to use covariation and some mechanism information together to recover causal relations among objects and events. Some questions about the nature of children's mechanistic understanding remain, which are directions for future research, but by the age of 4, children appear to connect causes and insides.

ACKNOWLEDGMENTS

David M. Sobel, Department of Cognitive and Linguistic Sciences, Brown University. Caroline M. Yoachim, Andrew N. Meltzoff, and Emily J. Blumenthal, Department of Psychology and Institute for Learning and Brain Sciences, University of Washington. Alison Gopnik, Department of Psychology, University of California at Berkeley.

First authorship on this paper is shared between David M. Sobel and Caroline M. Yoachim. This work was supported by NSF (DLS-0518161 to David M. Sobel and DLS-0132487 to Alison Gopnik), NIH (HD-22514 to Andrew N. Meltzoff), and the John S. McDonnell Foundation Causal Learning Collaborative.

We would like to thank Laura Case, William Fill, Denise Hill, Amy Hoff, Emily Hopkins, Adrienne Meyer, Kathryn Noe, Philip Parker, Phuong Pham, Ben To, Lena Tsui, and Sara Yerry for their assistance in constructing stimuli, data collection, and data analysis.

Portions of this paper were presented at the biennial meeting of the Cognitive Development Society (October, 2003) and the biennial meeting of the Society for Research on Child Development (March, 2005).

REFERENCES

- 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.
- Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, *15*, 147–149.
- Blakemore, S.-J., Fonlupt, P., Pachot-Clouard, M., Darmon, C., Boyer, P., Meltzoff, A. N., et al. (2001). How the brain perceives causality: An event-related fMRI study. *NeuroReport*, *12*, 3741–3746.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York: Academic.
- Carey, S. (1995). On the origin of causal understanding. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 268–308). New York: Oxford/Clarendon.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant Behavior & Development*, *21*, 713–731.

- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford: Oxford University Press.
- Gelman, S. A., & Coley, J. D. (1990). The importance of knowing a dodo is a bird: Categories and inferences in 2-year-old children. *Developmental Psychology, 26*, 796–804.
- Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development, 58*, 1532–1541.
- Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings of the non-obvious. *Cognition, 38*, 213–244.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*, 3–32.
- Gopnik, A., & Schulz, L. E. (2004). Mechanisms of theoryformation in young children. *Trends in Cognitive Science, 8*, 371–377.
- Gopnik, A., & Schulz, L. E. (in press). *Causal learning: Psychology, philosophy, and computation*. New York: Oxford University Press.
- Gopnik, A., & Sobel, D. M. (2000). Detectingblickets: How young children use information about novel causal powers in categorization and induction. *Child Development, 71*, 1205–1222.
- Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology, 37*, 620–629.
- Gopnik, A., & Wellman, H. M. (1994). The theory theory. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 257–293). Cambridge: Cambridge University Press.
- Gottfried, G.M., & Gelman, S.A. (2005). Developing domain-specific causal-explanatory frameworks: The role of insides and immanence. *Cognitive Development, 20*, 137–158.
- Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition, 61*, 233–259.
- Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child Development, 77*, 505–524.
- Harris, P. L., Pasquini, E. S., Duke, S., Asscher, J. J., & Pons, F. (2006). Germs and angels: The role of testimony in young children's ontology. *Developmental Science, 9*, 76–96.
- Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind–body distinction. *Child Development, 64*, 1534–1549.
- Kalish, C. W. (1996). Preschoolers' understanding of germs as invisible mechanisms. *Cognitive Development, 11*, 83–106.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Kelemen, D. (2001, April). *Intention in children's understanding of artifact function*. Paper presented at the biennial meeting of the Society for Research in Child Development, Minneapolis, MN.
- Kelemen, D. (2004). Are children “intuitive theists”? Reasoning about purpose and design in nature. *Psychological Science, 15*, 295–301.
- Kemler-Nelson, D. G. (1995). Principle-based inferences in young children's categorization: Revisiting the impact of function on the naming of artifacts. *Cognitive Development, 10*, 347–380.
- Kemler-Nelson, D. G., Russell, R., Duke, N., & Jones, K. (2000). Two-year-olds will name artifacts by their functions. *Child Development, 71*, 1271–1288.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition, 25*, 265–288.
- Matan, A., & Carey, S. (2001). Developmental changes within the core of artifact concepts. *Cognition, 78*, 1–26.
- Michotte, A. E. (1963). *The perception of causality* (T. R. Miles & E. Miles, Trans.). London: Methuen.
- Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science, 3*, 389–396.

- Nazzi, T., & Gopnik, A. (2003). Sorting and acting with objects in early childhood: An exploration of the use of causal cues. *Cognitive Development, 18*, 299–317.
- Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development, 5*, 193–207.
- Piaget, J. (1929). *The child's conception of the world* (J. Tomlinson & A. Tomlinson, Trans.). London: K. Paul, Trench, Trubner.
- Piaget, J. (1930). *The child's conception of physical causality* (M. Gabain, Trans.). London: Lund Humphries.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rozenblit, L. R., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26*, 521–562.
- Schult, C. A., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition, 62*, 291–324.
- Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40*, 162–176.
- Shanks, D. R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 48*, 257–279.
- Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47*(1, Serial No. 194).
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science, 28*, 303–333.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., & Liu, D. (in press). Causal reasoning as informed by the early development of explanations. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy and computation*. New York: Oxford University Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.
- Yoachim, C. M., Sobel, D. M., & Meltzoff, A. N. (April, 2005). *Causes, insides, and categories*. Paper presented at the 2005 biennial meeting of the Society for Research in Child Development, Atlanta, GA.