Developmental Psychology
2007, Vol. 43, No. 1, 186–196

0012-1649/07/$12.00 DOI: 10.1037/0012-1649.43.1.186

Conditional Probability Versus Spatial Contiguity in Causal Learning: Preschoolers Use New Contingency Evidence to Overcome Prior Spatial Assumptions

Tamar Kushnir
University of Michigan

Alison Gopnik
University of California, Berkeley

This study examines preschoolers' causal assumptions about spatial contiguity and how these assumptions interact with new evidence in the form of conditional probabilities. Preschoolers saw a toy that activated in the presence of certain objects. Children were shown evidence for the toy's activation rule in the form of patterns of probability: The toy was more likely to activate either when objects made contact with its surface (*on* condition) or when objects were several inches above its surface (*over* condition). In Experiment 1, 61 three-year-olds saw a deterministic activation rule. In Experiments 2 and 3, 48 four-year-olds saw an activation rule that was probabilistic. In Experiment 4, 30 four-year-olds saw a screening-off pattern of activation. In all 4 experiments, children used new evidence in the form of patterns of probability to make accurate causal inferences, even in the face of conflicting prior beliefs about spatial contiguity. However, children were more likely to make correct inferences when causes were spatially contiguous, particularly when faced with ambiguous evidence.

Keywords: causal reasoning, conceptual development, probabilistic reasoning, prior knowledge, spatial contiguity

Over the past 30 years, researchers have discovered that even very young children have a great deal of causal knowledge about the physical, biological, and psychological world and that they learn more about the causal structure of the world as they grow older (Bullock, Gelman, & Baillargeon, 1982; Flavell, Green, & Flavell, 1995; Gelman & Wellman, 1991; Gopnik & Wellman, 1994; Inagaki & Hatano, 1993; Kalish, 1996; Keil, 1995; Leslie & Keeble, 1987; Oakes & Cohen, 1990; Perner, 1991; Spelke, Breinlinger, Macomber, & Jacobson, 1992; Wellman, 1990). However, we do not fully understand the mechanisms that allow this learning to take place. How do children infer causal structure from their experience?

One suggestion is that children learn from particular spatial cues that specify causal events. From early in the 1st year of life, people experience what is known as *causal perception.* In an influential series of experiments, Michotte (1962) showed that if two physical events are spatially and temporally contiguous—for example, the launching of a billiard ball when another ball collides with it—adults receive a direct impression of the causal relation without requiring repeated exposure to the stimulus. Presentation of Michotte's launching paradigm to infants reveals that they too perceive these relations as causal by at most 10 months of age (Cohen & Oakes, 1993; Leslie & Keeble, 1987; Oakes & Cohen, 1990, 1994). The studies of infants compare collision events in which the two objects are spatially contiguous with control events in which spatial contiguity is violated. The results of these studies show that the spatial cues are crucial to the perception of causal relations.

However, spatial contiguity is often absent or misleading as a causal cue. There are numerous physical events (e.g., the operation of light switches and remote controls) within the realm of our everyday experience that consistently violate spatial contiguity. Nonetheless, adults continue to regard these relations between noncontiguous events as causal. Sometimes we may know the mechanisms underlying these relations, but we make causal inferences even with little or no knowledge about mechanisms (Rozenblit & Keil, 2002). Beyond a vague sense that radiation of some sort is involved, most adults know little about the mechanism of their TV remote.

Do preschool children also understand these events as causal? Preschoolers in natural settings certainly seem to be able to act on a switch or a remote control to bring about an effect. However, it is possible that children's actions on noncontiguous causes do not arise from a belief in the causal nature of these events but rather are the result of imitation and/or reinforcement. Perhaps children initially simply imitate the actions of others on the remote or perform those actions at random and then use those actions to bring about effects through a process of reinforcement or operant conditioning. Can preschoolers use evidence to override spatial contiguity cues and infer genuine causal relations? For example,

---

Tamar Kushnir, Department of Psychology, University of Michigan; Alison Gopnik, Department of Psychology, University of California, Berkeley.

This research was supported by the James S. McDonnell Foundation Causal Learning Collaborative, National Institute of Mental Health Predoctoral Fellowship F31MH66538 to Tamar Kushnir, and National Science Foundation Grant DLS0132487 to Alison Gopnik. We thank Claire Walton and Elizabeth Seiver for assistance with data collection and Laura Schulz, Henry Wellman, and Susan Gelman for helpful discussions and comments.

Correspondence concerning this article should be addressed to Tamar Kushnir, Department of Psychology, University of Michigan, 1012 East Hall, 530 Church Street, Ann Arbor, MI 48109-1043. E-mail: tkushnir@umich.edu

can they use that information to craft new actions even when they have not performed those actions before?

There is evidence that preschoolers and possibly even infants can override perceived spatial contiguity in the physical domain if they are given information about underlying causal mechanisms (Bullock, 1985; Bullock et al., 1982; Schlottmann, 1999; Schlottmann & Surian, 1999; Shultz, 1982). In experimental situations in which contiguity information contrasts with mechanism information, knowledge of mechanism enables children to disregard immediate spatial cues. For example, Shultz (1982) trained children to understand that tuning forks could make a box ring when placed in front of its opening without actually touching it. He then showed them a situation in which one fork was touching the top of the box and another was placed in front of its opening, at which point the child had to figure out which was causing the sound. The results showed that 2–4-year-olds attributed the sound to the fork in front of the opening significantly more often than to the fork that was touching the box, showing a preference for a previously demonstrated mechanism over spatial contiguity. In another scenario, a flashlight was demonstrated shining its beam on a wall from afar and then was paired with a flashlight touching the wall but shining its beam in the opposite direction (into empty space). In this case, 3-year-olds said that the spatially contiguous flashlight was the cause of the light on the wall, but 5-year-olds were able to correctly identify the noncontiguous light as the cause.

Bullock et al. (1982) also reported that explicit instruction about causal mechanisms enabled children to explain the outcome of a sequence of events without further contiguity information. Three-, 4-, and 5-year-olds saw a rod initiate a domino effect on a sequence of blocks, which resulted in a stuffed rabbit being pushed into its bed. When the blocks were then occluded by a screen, children of all ages were able to correctly explain the outcome on the basis of the mechanism (which was now hidden from view) and also to correctly predict that relevant changes to the mechanism (e.g., the use of a shorter rod or the removal of one of the blocks) would affect the outcome.

These results show that children were able to understand a complex causal event without immediate perceptual spatial contiguity cues. However, children only did so after initially learning about a mechanism that involved spatial contact among each of its components. We might argue that these particular mechanistic explanations allow children to conceptualize events in terms of underlying spatial contiguity even if they do not directly perceive that contiguity—they give children information about spatial contiguity in another way. In Bullock et al.'s (1982) experiments, for example, it seems likely that the children represented the occluded objects as spatially contiguous, even if they did not perceive that contiguity directly. Even in Shultz's (1982) experiments, the mechanistic explanation might have given the children a conception of an invisible but spatially contiguous process.

However, another rather different kind of information may also enable children to override spatial contiguity cues. Substantial work in artificial intelligence and statistics has formalized—in the theory of causal Bayes nets—how normatively accurate causal conclusions can be derived from patterns of probability (Pearl, 2000; Spirtes, Glymour, & Scheines, 2001). Recent psychological research with both adults and young children has shown that people draw accurate and sometimes quite complex causal conclusions by attending to correlations between observed events and correlations between interventions and their outcomes (Cheng, 1997; Gopnik et al., 2004; Gopnik, Sobel, Schulz, & Glymour, 2001; Schulz & Gopnik, 2004, in press; Sobel & Kushnir, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann & Hagmayer, 2001). Preschool children in particular have been shown to use evidence from observations and interventions to differentiate causes from effects, learn complex causal structures, and even learn about causal relations that cross domains (Gopnik et al., 2001, 2004; Kushnir, Gopnik, Schulz, & Danks, 2003; Schulz & Gopnik, 2004, in press; Schulz, Kushnir, & Gopnik, in press; Sobel, Tenenbaum, & Gopnik, 2004).

Most of these studies involved deterministic causal relations, but there is some evidence that children can make appropriate inferences about causes that are probabilistic as well. Kushnir and Gopnik (2005) showed children objects that activated a novel toy either 66% of the time or 33% of the time. The results showed that children were able to make normatively appropriate inferences about causal strength on the basis of this probabilistic evidence.

It is important to note that these studies have demonstrated that causal learning can take place even without mechanism information or spatial causal cues. In the causal strength experiments, for example, children could use probabilities to differentiate among causal hypotheses that were equally consistent with spatial contiguity. The fact that children can learn causal relations without using spatial cues leaves open the question of how spatial cues and probabilistic learning might interact. What happens when spatial cues and probabilistic ones actually conflict? There are several possibilities. Spatial contiguity might constitute a strong, innate constraint on children's causal inferences (see, e.g., Leslie & Keeble, 1987). Although children might use probability to decide among hypotheses that respect this constraint, they might initially be unlikely to use that evidence to overturn the constraint. Alternatively, spatial contiguity might not constrain probabilistic inferences at all. The perceptual phenomena captured by Michotte (1962) and the infancy work might be quite separate from the kinds of inferences children make on the basis of probability and contingency. Indeed, there is some evidence that these two types of causal inferences are processed in different parts of the brain (Fugelsang & Dunbar, 2005; Fugelsang, Roser, Corballis, Gazzaniga, & Dunbar, 2005; Roser, Fugelsang, Dunbar, Corballis, & Gazzaniga, 2005).

Finally, prior biases about contiguity might interact with probabilistic evidence. That is, children might prefer hypotheses that respect spatial contiguity and might require more evidence to overturn those hypotheses. Eventually, however, evidence can overturn even hypotheses that are strongly supported by prior knowledge. For example, on a Bayesian view of causal inference, new evidence is weighed against the prior probability of a particular hypothesis (Tenenbaum, Griffiths, & Niyogi, in press). If a hypothesis has a high initial probability, more evidence will be required to overturn that hypothesis. However, an accumulation of new evidence may eventually overturn even an initially very likely hypothesis. Because the literature suggests that children have a strong prior bias toward causal hypotheses that involve spatial contiguity, we can explore how this bias interacts with new evidence to shape children's causal inferences.

In the following experiments, children saw a novel toy that activated in the presence of certain objects by lighting up and playing music. No information was given to the children about the underlying causal mechanism behind the activation of the toy, and children did not act on the toy. Instead, children were shown evidence for the toy's "activation rule" in the form of patterns of probability: The toy was more likely to activate either when objects made contact with its surface (*on* condition) or when objects were held several inches above its surface (*over* condition). We compared children's performance in the two conditions across three types of tasks. In Experiment 1, the activation rule was deterministic—one action (either on or over) was always effective, and one was always ineffective. To test whether they had learned the rule, we then asked children to activate the toy themselves with a new object.

In Experiments 2 and 3, the activation rule was probabilistic—both actions were effective, but one was more effective than the other. Children were asked to make judgments of causal strength and also to activate the toy with a new object. In Experiment 4, the children saw a more complex "screening-off" pattern of activation similar to that in Gopnik et al.'s (2001) study. One action by itself made the toy go, one did not, and both actions together (e.g., holding one object over the toy at the same time that a second object was placed on the toy) made the toy go. Once the toy was active in the presence of two objects, the child was asked to make it stop. The experiment was designed so that the correct response—removing only the object that could independently activate the toy—was an intervention that the children had never before seen or performed. The ability to craft novel interventions is particularly compelling evidence that children have inferred causal relations (see, e.g., Glymour, 2001; Woodward, 2003).

We hypothesized that children would be able to use new evidence in the form of patterns of probability to make accurate causal inferences, even in the face of conflicting prior beliefs about spatial contiguity. However, we also hypothesized that children's prior bias toward spatially contiguous causes might have an effect on their inferences. If children require more evidence to overturn their prior hypotheses about contiguity, then they will be more likely to make correct inferences in the *on* conditions than in the *over* conditions of the following experiments.

Experiment 1

In this experiment, we showed 3- and 4-year-old children a novel toy that could be activated to light up and play music by the presence of objects. Objects were either held a few inches above the toy or placed on its surface. In the *over* condition, the toy always activated when one object was held over it and never did so when a second object made contact with it. In the *on* condition, the activation rule was reversed—contact between an object and the toy coincided with activation, whereas lack of contact did not.

We hypothesized that, from this pattern of evidence, children would learn the rule governing the activation of the toy—either on or over—and thus be able to use any object (rather than just the previously demonstrated one) in the appropriate way to make the toy go. To test this, we asked children to intervene twice on the toy—once with the previously ineffective object, and once with a new object that differed in appearance and material from the first two.

If children can use contingency information to learn about a causal relation that violates spatial contiguity, then they should respond by activating the detector without contact in the *over* condition. However, if children have a strong prior preference for contact as a causal cue, they should be more likely to learn the causal rule and perform the correct interventions in the *on* condition than in the *over* condition.

Method

Participants. Sixty-one 3- and 4-year-olds ranging in age from 2 years 11 months to 4 years 6 months (M = 3 years 6 months, SD = 5.1 months) participated. Approximately equal numbers of boys and girls participated. Participants were recruited from a local research participation list as well as from a local Berkeley, California, preschool. The sample was predominantly middle to upper middle class and White. No detailed information on the racial/ethnic makeup of the sample was available. Forty-one children (age: M = 3 years 7 months, SD = 5.1 months) were randomly assigned to the *over* condition. Twenty were randomly assigned to the *on* condition (age: M = 3 years 6 months, SD = 5.4 months).

To examine age differences, we split each condition into two age groups—those children younger than the condition mean and those at the mean or older. In the *over* condition, 17 were in the younger group (age: M = 3 years 2 months, SD = 2.5 months) and 24 were in the older group (age: M = 3 years 11 months, SD = 2.6 months). In the *on* condition, there were 10 participants in the younger group (age: M = 3 years 2 months, SD = 1.6 months) and 10 in the older group (age: M = 3 years 10 months, SD = 4.1 months). There were no significant age differences between conditions in either age group.

Materials. The novel toy was a 5 in. × 7 in. × 3 in. (12.7 cm × 17.8 cm × 7.6 cm) box made of wood with a Lucite top. A hidden switch, controlled by the experimenter, could make the box's top light up and play music. The experimenter could control when and how the activation occurred in the presence of other objects, giving the appearance that the objects caused the activation of the toy. The objects in this experiment included four metal artifacts (silver metal rings, eye-hook screws, and other hardware), each a different shape, and one smooth rock crystal. Out of the four types of metal objects, two objects were picked at random and presented in each of the first two trials.

Procedure. Children were interviewed in a testing room at their preschool or at the Institute for Human Development at the University of California, Berkeley. Each child sat at a low table opposite the experimenter, with the toy on the table between them. The experimenter introduced the toy by saying, "This is my special toy. Sometimes things make it go and sometimes things don't make it go. We're going to figure out what makes it go." The procedure is illustrated in Figure 1.

The experimenter's action on the objects occurred in one of two ways: (a) *On* (contact) actions were ones in which the experimenter placed an object on the surface of the detector and let it sit there on its own for 4 s. (b) *Over* (no contact) actions were ones in which the experimenter held an object 4 in. (10.2 cm) above the surface of the detector for 4 s. Whenever the toy activated in the presence of an object (either on or over, depending on the condition), the activation occurred simultaneously with the action and lasted for 4 s. After 4 s, the experimenter removed the object—placing it back in its location on the table—and stopped the activation of the toy simultaneously with its removal.

The experimenter placed two objects, A and B, on the table in front of the toy. She then placed each object on (or over) the toy one at a time, with the order and side of presentation counterbalanced across participants. In the *over* condition, the experimenter held Object A over the toy, causing it to activate. The experimenter then placed Object B on the toy, and the toy did not activate. The experimenter repeated each of these actions a second time. In the *on* condition, the experimenter placed Object A on the toy, causing it to activate. The experimenter then held Object B over the toy, and the toy did not activate. The experimenter repeated each of these actions a second time.

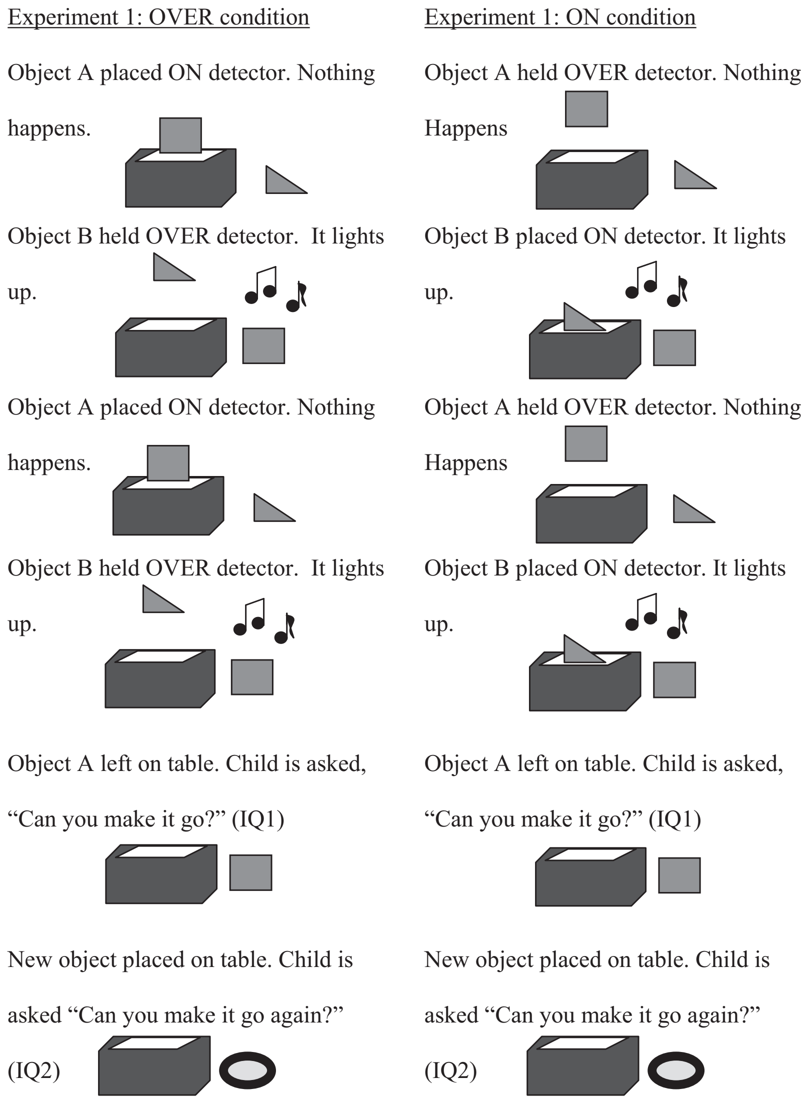

Experiment 1: OVER condition

Object A placed ON detector. Nothing happens.

Object B held OVER detector. It lights up.

Object A placed ON detector. Nothing happens.

Object B held OVER detector. It lights up.

Object A left on table. Child is asked, “Can you make it go?” (IQ1)

New object placed on table. Child is asked “Can you make it go again?” (IQ2)

Experiment 1: ON condition

Object A held OVER detector. Nothing Happens

Object B placed ON detector. It lights up.

Object A held OVER detector. Nothing Happens

Object B placed ON detector. It lights up.

Object A left on table. Child is asked, “Can you make it go?” (IQ1)

New object placed on table. Child is asked “Can you make it go again?” (IQ2)

Figure 1. Procedure of Experiment 1. IQ = Intervention Question.

The children were then asked two intervention questions. Intervention Question 1 (IQ1) was as follows: The experimenter removed Object A from sight, leaving Object B (the one that did not activate the toy the first two times) on the table. She asked the child, “Can you make it go?” Note that this question required the child to perform a different action on the object than the one the child had observed the experimenter perform on that object. For Intervention Question 2 (IQ2), the experimenter placed a new kind of object, a rock crystal, on the table in front of the child. She then asked the child, “Can you make it go with this one?”

Responses to each intervention question were coded as either *on* or *over* on the basis of the way the child activated the detector with the new object. *On* responses were ones in which the child made contact between the object and the surface of the toy. *Over* responses were ones in which the child held the object any distance directly above its surface. Twenty-five percent of the responses were coded by a researcher who was blind to condition, with an equal number of responses randomly sampled from each condition. Agreement was 97% (29 of 30 responses).

Results and Discussion

The percentages of each type of response to IQ1 and IQ2 by condition and age are summarized in Figure 2. In the *on* condition, 100% of the children performed the appropriate action (placing the object on the toy), regardless of age or intervention question (binomial tests, $p < .01$). In the *over* condition, however, there were differences between the age groups. In the older group, a significant majority of the children made the correct *over* response (19 out of 24, or 79%) for each question (binomial tests, $p < .01$). In the younger group, the children responded at chance: Only 10 out of 17 (59%) held the object over the toy for IQ1 and 12 out of 17 (71%) for IQ2 (binomial test, *ns*).

Consistency between IQ1 and IQ2 was very high—57 out of 61 (93%) children responded the same way to each question. A McNemar’s test confirmed that there were no significant

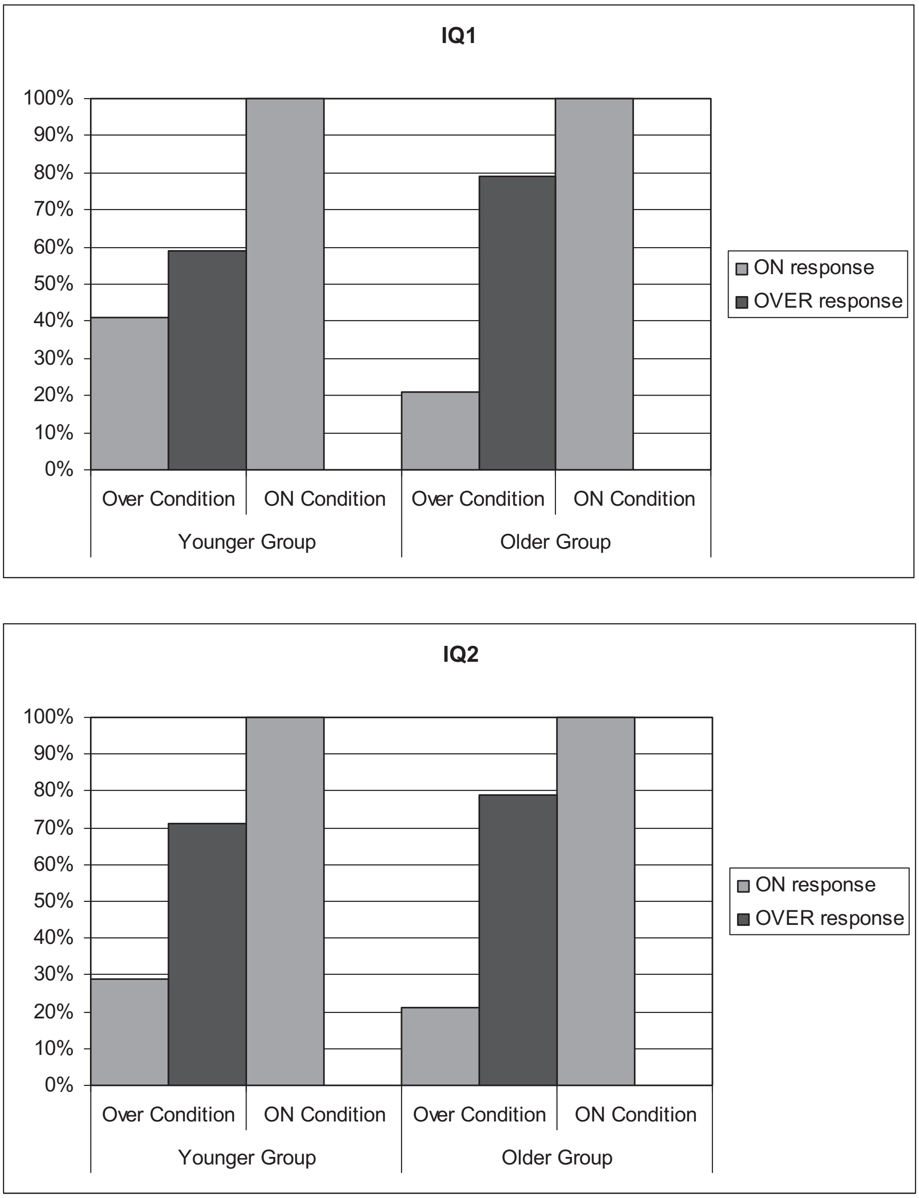

Figure 2. Percentage of *on* and *over* responses to Intervention Question 1 (IQ1) and Intervention Question 2 (IQ2) in Experiment 1 by age and condition.

changes in responding between IQ1 and IQ2 for either age group. This indicates that the 3- and 4-year-olds in the *on* condition and the 4-year-olds in the *over* condition not only were able to perform the appropriate action on the first attempt but also were able to generalize the appropriate causal action to a different type of object.

The results show that the older children in the *over* condition were able to learn the causal rule governing the activation of the toy despite the fact that the causal relation was completely new, that children had no information about the causal mechanism, and that children had never acted on the machine themselves. We also compared correct responses to IQ1 and IQ2 between the *on* and *over* conditions to determine whether children did better when the new inference was congruent with spatial contiguity. There were no differences between the distributions of correct responses for the older children (Fischer's exact test, *ns*). This might, however, have been due to the fact that children were at ceiling in the *on* condition.

The pattern for the younger children was different. The difference in correct responses between the *on* and *over* conditions was significant for IQ1 (Fischer's exact test, $p < .05$) but not for IQ2 (Fischer's exact test, *ns*). This may indicate that the younger children were sometimes able to overturn their preference for contiguity but had more difficulty overriding this prior preference for contact in the face of new evidence.

Experiment 2

How robust is children's ability to learn from new patterns of contingency? One of the features of Experiment 1 is that children were given evidence that one action was deterministically effective and that the other action was deterministically ineffective. However, in real-world scenarios, we often have only imperfect evidence of cause–effect relationships. It is easy to think of examples in which several interventions may be successful in achieving a desired outcome but some of those interventions will work better than others. In fact, outside of the limited context of children's mechanical toys, probabilistic data are much more common. Indeed, models of causal reasoning in adults (e.g., Cheng, 1997; Shanks & Dickinson, 1987) explicitly propose a consistent relation between probability and causal strength. Will children continue to prefer evidence from patterns of probability to spatial contiguity cues when all of the data are probabilistic? How does probabilistic reasoning interact with children's prior knowledge of spatial contiguity?

Given the poorer performance of the younger children in the deterministic study, we focused on 4-year-old children in the subsequent study. These children would have no difficulty learning the activation rule in the deterministic task. Children were presented with a probabilistic scenario in which one action (on or over, depending on condition) worked more often (66% of the time) and the other action worked less often (33% of the time). Children thus had evidence that both types of action could be effective, but one action was more effective. In addition, we asked the children to intervene on the toy before the trials began. This allowed us to get a more explicit measure of how strong their prior preference for contact would be given no previous exposure to the toy. We predicted that, when asked to make the toy go with no prior exposure, children would attempt to make contact between the object and the toy.

We then examined children's causal beliefs about the objects in two ways—by asking them to make causal strength judgments, and by asking them to craft novel interventions. We predicted that if children can use patterns of probability to learn the activation rule, (a) they should equate causal strength with probabilistic strength rather than with type of action and (b) when asked to make the toy go after exposure to the probabilities, they should use the action that worked more frequently to do so.

Method

Participants. Thirty-two children ranging in age from 3 years 11 months to 5 years 2 months (M = 4 years 7 months, SD = 4.6 months) participated in the study. Approximately equal numbers of boys and girls participated. Participants were recruited from local preschools in Berkeley, California, and Ann Arbor, Michigan. The sample was predominantly middle to upper middle class and White. No detailed information on the racial/ethnic makeup of the sample was available.

Materials. The novel toy was the same as in Experiment 1. The objects were various colorful blocks of different sizes and shapes.

Procedure. Children were interviewed in a testing room at their preschool. Each child sat at a low table opposite the experimenter, with the toy on the table between them. The experimenter first introduced the toy by saying, "This is my special toy. Sometimes things make it go and sometimes things don't make it go." With no further instruction, she handed the child an object (chosen at random) and said, "Can you make it go?" The children's responses were divided into four categories: (a) contact with the surface of the toy, (b) contact with some other part of the toy, (c) holding the object over the surface of the toy, (d) some other action with the object (e.g., spinning or tapping the object on the table), or (e) no attempt. All actions that fell into Categories a and c (which were either *on* or *over* responses) resulted in successful activation of the toy.

The experiment then proceeded with the four probability trials. On each trial, the experimenter put two new objects (chosen at random) on the table. She then demonstrated each object on the toy three times. The actions (on and over) were performed in the same way as in Experiment 1. One object was held over the toy three times. The other was placed on the toy three times. In the *over* condition, the *over* object made the toy go two out of three times, and the *on* object made the toy go one out of three times. In the *on* condition, the probabilities were reversed. Order of presentation and starting side were counterbalanced across participants. Each of the four trials had a different sequence of activations to ensure that children were not just responding on the basis of the first or last outcome for each object. The order of presentation of these four trials was counterbalanced across participants via a Latin square design.

After each of the four trials the children were asked, "Which one works best?" Their choice of object was recorded. The experimenter did not allow the children to intervene with the objects between trials, to preserve the overall probabilities of successful activation. At the end of the four trials, the child was given a new object, and the experimenter said, "Can you make the toy go with this one?" The child's intervention was categorized as *on* or *over* by the same criteria used in Experiment 1.

Twenty-five percent of the initial and final interventions were coded by a researcher who was blind to condition, with an equal number of responses randomly sampled from each condition. Agreement was 100%.

Results and Discussion

Children's initial interventions showed an overwhelming prior preference for contact as a method of activating the toy, $\chi^2(4, N = 32) = 42.38$, $p < .001$. Twenty-two children who attempted an action made contact between the object and the toy—21 made contact between the object and the surface of the toy, and 1 made contact between the object and the side of the toy. Only 3 children made a response with no contact (held the object over the toy). An additional 4 children just played with the object on the table (tapping or rolling it around), and the remaining 3 children did not want to attempt any action.

Despite their prior preference for contact, children in both the *on* and the *over* conditions were equally able to make causal strength judgments on the basis of probability rather than type of action. Children in the *on* condition were correct on an average of 3.4 out of 4 trials—significantly more often than chance, $t(15) = 7.65$, $p < .001$. Children in the *over* condition were correct on an average of 3.6 out of 4 trials ($SD = 0.62$)—also significantly more often than chance, $t(15) = 10.50$, $p < .001$. In addition, the means did not differ significantly between the conditions, $t(30) = 1.05$, *ns.* If children had genuinely inferred a difference in the causal strength of the two objects, they should have used this difference to determine the most effective intervention. Table 1 shows the distribution of children's initial and final interventions by condition.

In the *on* condition, 14 out of 16 children made the *on* response, 1 out of 16 made the *over* response, and 1 refused to intervene. A comparison of initial to final *on* responses (within subjects) showed no significant difference (McNemar's test, *ns*). In the *over* condition, 6 out of 16 children made the *over* response, and 10 out of 16 made the *on* response. A comparison of initial with final *over*

Table 1
The Distribution of Children's Initial and Final Interventions by Condition in Experiment 2

	On condition		Over condition	
Intervention response	Initial intervention	Final intervention	Initial intervention	Final intervention
On	9	14	12	10
Over	3	1	0	6
Other	4	1	4	0

responses showed a significant difference (McNemar's test, $p < .05$). This was due to 4 of the children switching from an initial *on* response to a final *over* response and 2 of the children switching from another response to a final *over* response. There were also significantly more *over* responses in the *over* condition than in the *on* condition (Fischer's exact test, $p < .05$).

These results indicate that a significant number of children in the *over* condition revised their intervention in response to new evidence that the *over* action was probabilistically more effective. Note also that in the *on* condition, which was identical to the *over* condition in every way except the pattern of evidence, there was no change between the pretest and posttest performance. This suggests that the change in performance in the *over* condition was due to the pattern of evidence and not to some other factor, such as regression to the mean or fatigue.

However, unlike in Experiment 1, the number of children who chose the *over* response was not a significant majority, and there were also significant differences between the proportions of correct to incorrect responses between conditions (14:15 vs. 6:10; Fischer's exact test, $p < .01$). Overall, children seemed less willing to override their prior hypothesis about contiguity in this probabilistic task than in the earlier, deterministic one.

Although children in this experiment were shown probabilistic relations between actions and outcomes, the pattern of final intervention responses cannot be explained by simple probability matching. First, the distribution of the children's posttest intervention responses is significantly different from the distribution of observed probabilities (2 of 3 vs. 1 of 3) for both the *on* and the *over* conditions: *on* condition, $\chi^2(1, N = 14) = 4.23, p < .05$; *over* condition, $\chi^2(1, N = 16) = 6.30, p < .05$. In fact, the most frequent response in both conditions was the *on* response, which indicates that children's responses depended partially on their prior causal beliefs, not just on observed contingencies (as would be the case if they were simply probability matching).

Experiment 3

The results of Experiment 2 show that observing probabilistic data can affect children's causal strength judgments and their own new interventions, even in the face of an initial preference for contact as a causal cue. The experiment also suggests that children take into account the overall pattern of the data, rather than relying on deterministic information to make causal inferences. However, because one object activated the toy twice and one only once, it could still be the case that children were only paying attention to one piece of the pattern of data, namely the frequency of positive associations between the object and the activation of the toy. To strengthen the claim that children were relying on the pattern of probability, we need to show that they used information about both positive and negative instances. Experiment 3 was designed to test whether children truly take probabilistic evidence into account by controlling the absolute frequency of positive instances while keeping the overall probabilities the same as in Experiment 2.

Moreover, it is possible that the pattern of events in Experiment 2, particularly the lower performance relative to Experiment 1, was due to the fact that children were explicitly asked to make a causal judgment at the start of the task and that they did so in a way that assumed spatial contact. This might have strongly activated or even created a contact hypothesis, which then conflicted with the pattern of evidence. To eliminate this possibility, Experiment 3 contained the same probabilistic evidence as in the *over* condition of Experiment 2 but no initial intervention.

Method

Participants. Sixteen children ranging in age from 3 years 11 months to 5 years 1 month (M = 4 years 6 months, SD = 5.1 months) participated in the study. Approximately equal numbers of boys and girls participated. Participants were recruited from a local research participation pool and from a local Berkeley, California, preschool. The sample was predominantly middle to upper middle class and White. No detailed information on the racial/ethnic makeup of the sample was available.

Materials. The novel toy and objects were the same as in Experiment 2.

Procedure. Children were interviewed in a testing room at their preschool. Each child sat at a low table opposite the experimenter, with the toy on the table between them. The experimenter first introduced the toy by saying, "This is my special toy. Sometimes things make it go and sometimes things don't make it go." We did not include an initial intervention in this experiment to keep the probabilities exact (two thirds and one third) over the whole task and to test whether prior knowledge might have an effect even without an explicit causal judgment. As in Experiment 2, the experiment consisted of four trials. On each trial, the experimenter put two new objects (chosen at random) on the table. She then demonstrated each object on the toy. One object was held over the toy and activated it two out of three times. The other was placed on the toy and activated it two out of six times. Thus, the absolute frequency of successful activations was the same for both objects, whereas the probability of successful activation was two thirds for the *over* object and one third for the *on* object, just as in Experiment 2. Order of presentation and starting side were counterbalanced. Each of the four trials had a different sequence of activations. The order of the trials was counterbalanced across participants via a Latin square design. After each of the four trials, the child was asked, "Which one works best?" His or her choice of object was recorded. At the end of four trials, the child was given a new object and asked to "make the toy go." These responses were coded according to the same criteria used in Experiment 1. Twenty-five percent of the final interventions were coded by a researcher who was blind to the experimental set-up and hypothesis. Agreement was 100%.

Results and Discussion

As in Experiment 2, the results of the four trials again showed that children made causal strength judgments on the basis of probabilistic evidence rather than contact. Because there were no order effects, averaging over the four trials revealed that children chose the two thirds effective *over* object as best, on average, 2.75 out of 4 times (68.8% of the time), which is significantly more often than chance, $t(15) = 3.50, p < .01$.

A comparison of the average causal strength judgments reveals that they were significantly lower in Experiment 3 than in the *over* condition of Experiment 2. This could be due to the fact that some children were unable to use relative frequency as a measure of causal strength and instead used the strategy of comparing positive activations only (two on vs. two over). However, this task was also more than twice as long as the task in Experiment 2 because of the amount of negative evidence presented. Therefore, we considered another possible interpretation—that children were simply behaving erratically toward the end of the task because of fatigue. Indeed, there was a downward trend in correct responses from Trial 1 through Trial 4 (decreasing steadily from 14 to 9) Because of this, we compared correct responses to Trial 1 in this experiment with correct responses in the *over* condition in Experiment 2 and found no significant differences (15 out of 16 correct vs. 14 out of 16 correct) by a Fischer's exact test. This indicates that children were capable of using relative frequency as a measure of probability and that the differences in average causal strength judgments were most likely caused by fatigue.

Because children in this experiment were not given an initial intervention to perform, the results of the final intervention were compared with the results of the initial intervention in Experiment 2. Children in the present experiment intervened with no contact at the end of the four trials more often than children did at the beginning of Experiment 2 (7 out of 16 vs. 3 out of 32; Fischer's exact test, $p < .01$). Similarly, children in this experiment made more *over* responses than did children in the *on* condition in Experiment 2 (7 out of 16 vs. 1 out of 16; Fischer's exact test, $p < .05$). However, there was no difference in performance on the posttest in this experiment and in the *over* condition of Experiment 2 (6 of 16 vs. 7 of 16; Fischer's exact test, *ns*).

This experiment confirms the fact that children took into account the overall pattern of probability to make causal inferences. They relied neither on substantive knowledge about spatial relations nor on simple counts of positive associations between interventions and outcomes. As in the previous experiment, however, there is evidence that the spatial contiguity bias did influence children's judgments.

Experiment 4

These results suggest that 4-year-old children can override spatial contiguity on the basis of simple contingency information whether that information is probabilistic or deterministic. Can children also use more complex patterns of evidence to override contiguity assumptions? Earlier studies suggest that children can use a more complex pattern of conditional dependency to make causal inferences (Gopnik et al., 2001, 2004; Sobel et al., 2004).

In these experiments, children were shown associations between both objects and the activation of the detector. However, the B object only activated the detector in the presence of the A object, whereas the A object activated the detector without the presence of the B object. Children reported that the A object was causal, used that object to activate the machine, and removed that object to make the machine stop. According to normative theories of causal inference, the effects of A should "screen off" the effects of B (Reichenbach, 1956). A phenomenon called *blocking* suggests that adults use this information appropriately in causal inference, and there may even be parallel phenomena in animals' performance in classical conditioning experiments (Shanks & Dickinson, 1987). Will children also use a more complex pattern of evidence to override prior knowledge about spatial contiguity?

This screening-off design also allows us to control for a potential problem in Experiment 1. In the earlier experiment, children could not have solved the problem by simple imitation because they always saw both *over* and *on* actions. Moreover, they generalized the effective action to new objects. However, it is possible that children simply ignored the objects and only focused on imitating the actions. In that case, the fact that a particular type of action was followed by the effect might have made that action more salient to the children and thus might have made them more likely to imitate that action, even without a causal inference. The screening-off design allowed us to test whether children would use the contingency evidence to design an entirely new action they had not seen before at all—namely an intervention to make the machine stop.

Method

Participants. Thirty children ranging in age from 3 years 6 months to 4 years 8 months (M = 4 years 2 months, SD = 3.65 months) participated in the study. Fifteen were assigned to the *on* condition and 15 to the *over* condition. Approximately equal numbers of boys and girls participated. Participants were recruited from local Ann Arbor, Michigan, preschools. The sample was predominantly middle to upper middle class and White. No detailed information on the racial/ethnic makeup of the sample was available.

Materials. The novel toy and objects were the same as in Experiments 2 and 3.

Procedure. The procedure (summarized in Figure 3) was identical to that in the previous experiments, with the following exceptions. Rather than seeing each block placed over or on the detector separately, children were shown the following sequence of events. In the *over* condition, one block was placed on the detector, and the detector did not activate. The block was taken away. Then a second block was held over the detector, and the detector activated. With the second block still over the detector and the detector still activating, the experimenter placed the first block back on the detector for 4 s. Thus, the children saw one block on the detector and a second block simultaneously held over the detector. The experimenter then asked the child, "Can you make it stop?" Children had never seen any action stop the detector. If children inferred that the hovering-over block was causally effective and assumed that removing an effective block would stop the detector, they should have taken the hovering-over block away from the detector.

The *on* condition was identical except that the block was first placed over the detector and was ineffective. Then the experimenter removed that block, placed a block on the detector, which activated, held the first block over the detector for 4 s, and then said, "Can you make it stop?" In this condition, children should have removed the block that was on the detector.

Children's responses were coded as having removed the *on* or the *over* object only if the object was removed from the vicinity of the toy's surface. That is, responses in which an object was taken but not removed (e.g., taking the *over* object and putting it on the toy) were coded in a separate category. Twenty percent of the responses (including a representative sample of the different types) were coded by another researcher who was blind to the research hypothesis and experimental condition. Agreement was 100%.

In the *over* condition, the correct response required that the children take the block from the experimenter. Pilot testing suggested that some children were reluctant to do this, so we included a warm-up task at the start of both conditions in which children were encouraged to take the blocks. The experimenter held up two blocks and said, "Can you take the red block?

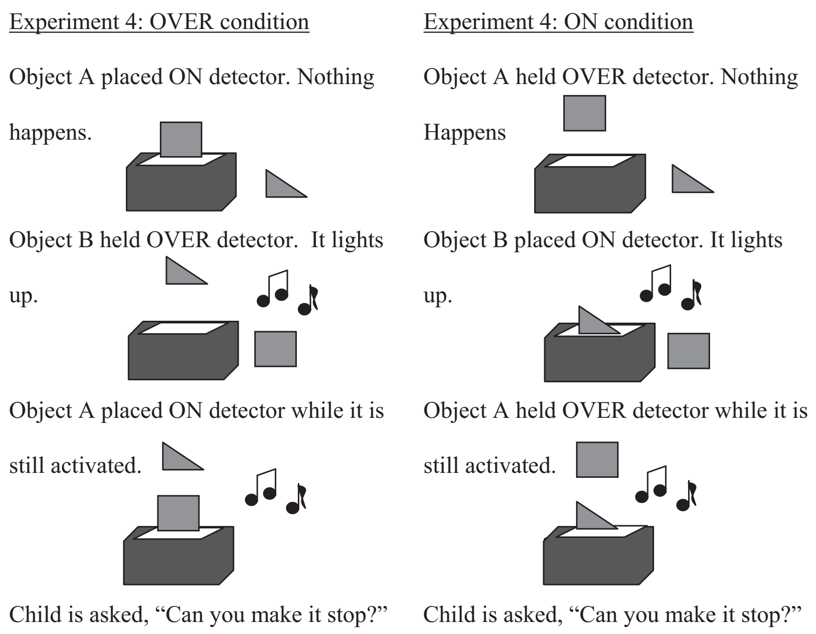

Experiment 4: OVER condition

Object A placed ON detector. Nothing happens.

Object B held OVER detector. It lights up.

Object A placed ON detector while it is still activated.

Child is asked, “Can you make it stop?”

Experiment 4: ON condition

Object A held OVER detector. Nothing Happens

Object B placed ON detector. It lights up.

Object A held OVER detector while it is still activated.

Child is asked, “Can you make it stop?”

Figure 3. Procedure of Experiment 4.

Can you take the blue block?” and then repeated this with two new blocks. All the children took the blocks in the warm-up task.

Results and Discussion

Children in this experiment performed a variety of actions in response to the “make it stop” question, which is not surprising given that they had never seen anything stop the toy before. A majority—19 out of 30—of the responses, however, involved removing the blocks from the vicinity of the toy. Eleven children removed the object on the toy, and 8 took the object held by the experimenter over the toy and moved it away from the toy. There was also a range of inappropriate responses, with no dominant alternative. These included taking the hovering-over object and placing it on the toy (6 children), pressing the object that was on the toy more firmly on the toy (2 children), picking up the object on the toy and holding it over the toy (1 child), placing a finger an inch over the toy (1 child), and instructing the experimenter to drop the *over* block (1 child). It was unclear whether these other responses involved alternative hypotheses about stopping (e.g., that the same action would both start and stop the toy) or were experimental actions designed to figure out how the toy worked.

Among the appropriate responses, however, there was a clear difference between conditions. In the *over* condition, 8 children took the hovering block away from the toy. Only 1 child took the *on* block off the toy (binomial test, $p < .05$). In the *on* condition, this pattern was reversed: Eleven children removed the block that was on the toy, and no child removed the hovering block (binomial test, $p = .001$). Children were significantly more likely to take the hovering object away in the *over* condition than in the *on* condition (Fischer’s exact test, $p < .0001$). Similarly, they were more likely to take the *on* block off in the *on* condition than in the *over* condition (Fischer’s exact test, $p < .01$). As in Experiment 1, however, although children performed slightly better in the *on* condition, there was no significant difference in overall correct performance between the two conditions.

This pattern of responses suggests that the similar responses in Experiment 1 were not simply the result of differential imitation. In the current experiment, children had to make an entirely novel response that they had never seen before. Moreover, children in the *over* condition overrode spatial contiguity even when they were presented with a more complex pattern of contingencies.

General Discussion

Taken together, these studies demonstrate that children can quickly revise previous beliefs about spatial contiguity and infer noncontiguous causal relations in light of new contingency evidence. They can do so regardless of whether the evidence is deterministic or probabilistic, simple or complex. In Experiment 1, a large majority of children were able to learn the no contact rule on the basis of a small amount of evidence and to use it to craft interventions with novel objects. In Experiments 2 and 3, children clearly demonstrated an initial preference for spatially contiguous causes but used probabilistic evidence to override that preference and to make judgments of causal strength. A significant number of children in Experiments 2 and 3, although fewer than in Experiment 1, also used what they had learned to revise their initial preference for contact when intervening on new objects. In Experiment 4, children again were able to override assumptions about spatial contiguity on the basis of more complex patterns of conditional dependence, and they used this information to craft entirely novel interventions they had never observed. These results add to the growing body of evidence that young children can make quite sophisticated causal inferences from patterns of probability. Young children’s causal learning is neither restricted by rigid, domain-specific constraints nor contingent on their understanding of underlying mechanisms. Rather, it seems to be flexible, domain general, and normative.

In addition, however, these results suggest that there may be interactions between prior knowledge and new evidence. New evidence may override prior biases, but those biases may also modify the influence of the new evidence. The pattern of interaction between prior knowledge and new evidence is not simple, however. In Experiments 1 and 4, which involved deterministic inferences with unambiguous patterns of evidence, 3.5- to 4-year-old children were slightly more likely to make causal inferences that accorded with prior knowledge than to make inferences that did not, but these differences were not statistically significant. However, the youngest children, under 3.5, were significantly more likely to be correct in the *on* condition than in the *over* condition.

Children in Experiments 2 and 3, which involved probabilistic judgments, showed a more complex pattern. There was no difference in children's judgments of causal strength in the two conditions, but children were significantly less likely to intervene correctly in the *over* condition than in the *on* condition. This suggests that in the probabilistic case, both prior knowledge and evidence had a substantial effect on children's judgments.

At least the 4-year-olds' pattern, however, is consistent both with principles of Bayesian inference and with other recent results in the literature. When causal relations are deterministic, a single piece of evidence may be enough to rule out a hypothesis, no matter how likely it may be on the basis of prior knowledge. In Experiments 1 and 4, one ineffective attempt was enough to rule out the hypothesis that putting a block on the detector makes it go, even if that hypothesis seemed very likely at first. However, when the evidence is consistent in principle with more than one hypothesis but supports one hypothesis more strongly than another, prior knowledge may have a stronger effect. This is true when the causal relations are probabilistic, particularly with small amounts of evidence. When the child sees that placing the block on the toy makes it go only one of three times, he or she cannot be certain about the effect of the block on the next attempt—perhaps the evidence the child has seen is unrepresentative. The new evidence may be enough to make the initially less likely hypothesis more likely but not enough to overturn the initial hypothesis altogether.

In Schulz and Gopnik (2004), children showed a prior preference for within-domain causal relations (e.g., physical causes leading to physical rather than psychological effects) over between-domains relations. However, with deterministic and unambiguous evidence for between-domains relations, they could entirely override that preference, and there was no evidence for an effect of their prior beliefs on their later inferences. More recently, however, Bonawitz, Griffiths, and Schulz (2006) found that when the evidence was ambiguous, children did initially prefer within-domain hypotheses but eventually overturned that preference as more evidence accumulated. Similarly, Sobel et al. (2004) found that prior knowledge about the probability that objects activated the detector had no effect on an unambiguous inference but did have systematic effects when the evidence was ambiguous. The intervention results in Experiments 2 and 3 may suggest a similar pattern.

The current results are also consistent with earlier data in suggesting that there may be developmental changes in the ways that children integrate prior knowledge and new evidence. Bonawitz et al. (2006) found that young 3-year-olds were less likely to override prior knowledge in the light of new evidence than were older children, and a similar pattern seemed to emerge for the young 3-year-olds in our study. It is unclear, however, whether this is a reflection of the fact that the prior biases are simply stronger for younger children or whether it reflects a more general developmental difference in techniques for causal inference. Further research on younger children seems necessary. Moreover, there is other evidence that even infants are sensitive to patterns of probability and contingency in certain circumstances (e.g., Saffran, 2001; Sobel & Kirkham, in press). It would be interesting to see whether they use this information in causal inference.

Another question for further research involves exactly how children integrate probabilistic and spatial evidence. In Experiments 2 and 3, there was simply a contrast between more probable and less probable effects and between spatial contact and lack of contact. We do not know whether children would be sensitive to more fine-grained differences in probability or in spatial configuration, and this also remains a topic for future research.

These results also raise another important question about the interaction of causal knowledge and new evidence. This concerns the relationship between evidence from patterns of probability and children's inferences about causal mechanisms. A dominant view in developmental psychology is that, from infancy, children understand causal relations in terms of notions of force and generative transmission (Leslie & Keeble, 1987; Shultz, 1982)—that is, in terms of underlying causal mechanisms. Does evidence in the form of patterns of probability also lead children to make inferences about the existence of causal mechanisms? (See Bullock, 1985, for some evidence that suggests 5-year-olds may do this.) Yoachim, Sobel, and Meltzoff (2005) found that children inferred that when two objects both activated a toy, they were also likely to have similar internal structure. Did children in our study similarly infer that the no contact rule was due to some hidden transmission of force? The interaction of new evidence (in the form of patterns of probability) and underlying assumptions about mechanism (e.g., generative transmission) may be one important way children learn about new hidden causes and is an important line of inquiry for future research.

What does emerge from the present study is that children are able to use small amounts of evidence about contingency and probability to overturn earlier biases about spatial contiguity. However, our work also suggests that causal learning is both conservative and flexible. New probabilistic evidence does not completely overturn previous hypotheses, but it does make children more likely to consider alternatives. This kind of learning is consistent with the idea that children's learning is analogous to theory formation in science (Gopnik & Wellman, 1994). Theories are resistant to change but eventually can be overturned. Bayesian causal learning may provide a mechanism for this sort of change.

References

Bonawitz, E. B., Griffiths, T. L., & Schulz, L. (2006). *Modeling cross-domain causal learning in preschoolers as Bayesian inference*. Retrieved October 30, 2006, from http://www.cogsci.rpi.edu/CSJarchive/Proceedings/2006/docs/p89.pdf

Bullock, M. (1985). Causal reasoning and developmental change over the preschool years. *Human Development, 28,* 169–191.

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York: Academic Press.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104,* 367–405.

Cohen, L. B., & Oakes, L. M. (1993). How infants perceive a simple causal event. *Developmental Psychology, 29,* 421–433.

Flavell, J. H., Green, F. L., & Flavell, E. R. (1995). Young children's knowledge about thinking. *Monographs of the Society for Research in Child Development, 60*(1, Series No. 243).

Fugelsang, J. A. & Dunbar, K. N. (2005). Brain-based mechanisms underlying complex causal thinking. *Neuropsychologia, 43,* 1204–1213.

Fugelsang, J. A., Roser, M. E., Corballis, P. M., Gazzaniga, M. S., & Dunbar, K. N. (2005). Brain mechanisms underlying perceptual causality. *Cognitive Brain Research, 24,* 41–47.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essence: Early understandings of the non-obvious. *Cognition, 38,* 213–244.

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

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111,* 1–30.

Gopnik, A., Sobel, D. M., Schulz, L., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of dependent and independent probability. *Developmental Psychology, 37,* 620–629.

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

Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind–body distinction. *Child Development, 64,* 1534–1549.

Kalish, C. (1996). Causes and symptoms in preschoolers' conceptions of illness. *Child Development, 67,* 1647–1670.

Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 234–267). Oxford, England: Clarendon Press.

Kushnir, T., & Gopnik, A., (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science, 16,* 678–683.

Kushnir, T., Gopnik, A., Schulz, L. E., & Danks, D. (2003). *Inferring hidden causes.* Retrieved October 30, 2006, from http://www.ccm.ua.edu/cogsci/prof141.html

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition, 25,* 265–288.

Michotte, A. E. (1962). *Causalite, permanence et realite phenomenales: Etudes de psychologie experimentale* [Phenomenal causality, permanence and reality]. Louvain, Belgium: Publications universitaires.

Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development, 5,* 193–207.

Oakes, L. M., & Cohen, L. B. (1994). Infant causal perception. In C. Rovee-Collier & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 9, pp. 1–54). Norwood, NJ: Ablex.

Pearl, J. (2000). *Causality.* New York: Oxford University Press.

Perner, J. (1991). *Understanding the representational mind.* Cambridge, MA: MIT Press.

Reichenbach, H. (1956). *The direction of time.* Berkeley: University of California Press.

Roser, M. E., Fugelsang, J. A., Dunbar, K. N., Corballis, P. M., & Gazzaniga, M. S. (2005). Dissociating processes supporting causal perception and causal inference in the brain. *Neuropsychology, 19,* 591–602.

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

Saffran, J. (2001). Words in a sea of sounds: The output of infant statistical learning. *Cognition, 81,* 149–169.

Schlottmann, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. *Developmental Psychology, 35,* 303–317.

Schlottmann, A., & Surian, L. (1999). Do 9-month-olds perceive causation-at-a-distance? *Perception, 28,* 1105–1113.

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40,* 162–176.

Schulz, L. E., & Gopnik, A. (in press). Preschool children learn about causal structure from conditional interventions. *Developmental Science.*

Schulz, L. E., Kushnir, T., & Gopnik, A. (in press). Learning from doing: Interventions and causal inference. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation.* New York: Oxford University Press.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 21, pp. 229–261). San Diego, CA: Academic Press.

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. & Kirkham, N. Z. (in press). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology.*

Sobel, D. M. & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition, 34,* 411–419.

Sobel, D., Tenenbaum, J., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science, 28,* 303–333.

Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review, 99,* 605–632.

Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction, and search: Springer lecture notes in statistics* (2nd ed., rev.). Cambridge, MA: MIT Press.

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

Tenenbaum, J. B., Griffiths, T. L., & Niyogi, S. (in press). In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation.* New York: Oxford University Press.

Waldmann, M. R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition, 82,* 27–58.

Wellman, H. M. (1990). *The child's theory of mind.* Cambridge, MA: MIT Press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation.* New York: Oxford University Press.

Yoachim, C., Sobel, D. M., & Meltzoff, A. (2005, April). *Causes, insides and categories.* Paper presented at the meeting of the Society for Research in Child Development, Atlanta, GA.

Received June 27, 2005
Revision received April 20, 2006
Accepted May 18, 2006 ■