Child Development, March/April 2006, Volume 77, Number 2, Pages 427 – 442

God Does Not Play Dice: Causal Determinism and Preschoolers' Causal Inferences

Laura E. Schulz
Massachusetts Institute of Technology

Jessica Sommerville
University of Washington

Three studies investigated children's belief in causal determinism. If children are determinists, they should infer unobserved causes whenever observed causes appear to act stochastically. In Experiment 1, 4-year-olds saw a stochastic generative cause and inferred the existence of an unobserved inhibitory cause. Children traded off inferences about the presence of unobserved inhibitory causes and the absence of unobserved generative causes. In Experiment 2, 4-year-olds used the pattern of indeterminacy to decide whether unobserved variables were generative or inhibitory. Experiment 3 suggested that children (4 years old) resist believing that direct causes can act stochastically, although they accept that events can be stochastically associated. Children's deterministic assumptions seem to support inferences not obtainable from other cues.

Many researchers have proposed that children's knowledge about the world can take the form of causal theories, in which *unobserved* causes play a central role (Carey, 1985; Gopnik, 1988; Gopnik & Meltzoff, 1997; Keil, 1989; Perner, 1991; Wellman, 1990). Children invoke unobserved mental states to explain human behavior (see, e.g., Wellman, 1990), invisible forces to explain physical events (Shultz, 1982), and invisible, internal mechanisms to explain biological events (Gelman, Coley, & Gottfried, 1994).

However, little is known about how children infer unobserved causes. Until recently, developmental psychologists have looked primarily at children's ability to infer causal structure from spatiotemporal cues (Cheng & Novick, 1992; Leslie & Keeble, 1987) and information about substantive, domain-specific mechanisms (Ahn, Gelman, Amsterlaw, Hohenstein, & Kalish, 2000; Bullock, Gelman, & Baillargeon, 1982; Carey & Spelke, 1994; Shultz, 1982; Spelke, Breinlinger, Macomber, & Jacobson, 1992). In adult cognitive psychology, by contrast, researchers have focused primarily on domain-general causal learning from the strength of association (Shanks, 1985; Shanks & Dickinson, 1987; Spellman, 1996) and patterns of covariation (Cheng, 1997, 2000) among events.

However, we can sometimes have causal knowledge even without knowing much about underlying mechanisms. If increasing serotonin levels relieve depression, we may conclude that low serotonin levels cause depression even if we do not know how. On the other hand, understanding causation seems to involve more than recognizing patterns of correlation. Lack of exercise is correlated with depression and we could imagine a plausible mechanism connecting the two (e.g., metabolic changes associated with exercise might regulate emotional arousal). However, if manipulating serotonin levels affects depression and manipulating physical activity does not, we will conclude that serotonin plays a causal role in depression and exercise does not.

Recently, psychologists, philosophers of science, and statisticians have suggested that the crucial piece missing from both mechanism and covariation accounts of causal inference is the notion of intervention (Gopnik et al., 2004; Gopnik & Schulz, in press; Pearl, 2000; Spirtes, Glymour, & Scheines, 1993; Woodward, 2003). Intuitively, if X is causally related to Y, then (all else being equal) there will be something we can do to change the value of X that will change the value of Y; that is, intervening

This research was supported by a National Science Foundation Graduate Fellowship Award (99 – 136), an American Association of University Women American Dissertation Fellowship, and a McDonnell Foundation Causal Learning Collaborative Initiative grant to Laura E. Schulz. We are very grateful to Alison Gopnik for significant contributions to this work. We are grateful also to Catherine Griffith, Carla Hudson, Anna Jenkins, Tamar Kushnir, Rebecca Saxe, and Christine Schaeffer, for their help and suggestions. Finally, we appreciate the support of the children, families, and staff at Fruit and Flower, Northeast Community Child Development Center, Helen Gordon Child Development Center, the Goddard School, and the Harold T. Jones Child Study Center.

Correspondence concerning this article should be addressed to Laura E. Schulz, MIT Department of Brain and Cognitive Sciences, 46-4011, 77 Massachusetts Avenue, Cambridge, MA 02139-4307. Electronic mail may be sent to lschulz@mit.edu.

0009-3920/2006/7702-0013

directly on X can affect Y. Critically, if intervening to change X changes Y, we can infer a causal relationship between the variables even if we do not know the underlying mechanism. Conversely, if, all else being equal, there is nothing we can do to X that will affect Y, we can conclude that there is no direct causal relationship between X and Y, even if X and Y are correlated and even if there is a plausible mechanism connecting the two.

The claim that an intervention to change X changes Y can also be represented as the claim that an intervention to change X correlates with a change in Y. However, the interventionist account differs importantly from standard covariation accounts. In general, covariation accounts of causal inference have been critiqued for their failure to distinguish between direct causal relationships ($X \rightarrow Y$) and spurious associations ($X \leftarrow U \rightarrow Y$, where U is an unobserved common cause). Although X and Y will covary under observation of both of these structures, the same correlations will not hold under interventions. If the true structure is $X \rightarrow Y$, then the value of Y will be statistically dependent on the value to which X is set by intervention; if the true structure is $X \leftarrow U \rightarrow Y$, then the value of Y will be independent of any intervention to change the value of X. Moreover, interventions can lead to inferences about the existence of unobserved causal mechanisms: if X and Y covary but "doing" X fails to change Y and "doing" Y fails to change X, then (regardless of one's prior knowledge about causal mechanisms) one should infer an unobserved common cause of X and Y (Gopnik et al., 2004; Griffiths, Baraff, & Tenenbaum, 2004; Kuhsnir, Gopnik, Schulz, & Danks, 2003; Pearl, 2000).

The interventionist account also contrasts with a mechanism-based account of causal inference. Abundant research suggests that we can and do use mechanism information to make causal claims when such information is available (Ahn, Kalish, Medin, & Gelman, 1995; Bullock et al., 1982; Koslowski & Masnick, 2002). Philosophers have also suggested that the possibility of effective intervention may always imply some underlying mechanism (see Strevens, in press; Woodward, 2003). Critically, however, we can use interventions to learn causal relations even in the absence of information about the physical relationships between events. Indeed, this is frequently the case in the history of science (as when we infer a causal relationship between serotonin and depression because experiments show that regulating serotonin relieves depression, although we do not know the mediating mechanism). Particular sequences of interventions and outcomes can lead us to infer genuinely causal links, even when the underlying mechanism remains mysterious (see Schulz, Kushnir, & Gopnik, in press, for a discussion).

Intuitions such as these underlie experimental design in science and have recently been formalized into computational accounts of causal learning (Pearl, 2000; Spirtes et al., 1993; Woodward, 2003). Although children and naïve adults lack an explicit understanding of experimental design (Chen & Klahr, 1999; Inhelder & Piaget, 1958; Kuhn, 1989; Kuhn, Amsel, & O'Laughlin, 1988; Masnick & Klahr, 2003), research suggests that even preschool children can use evidence about interventions and outcomes to learn a wide range of causal structures (Gopnik et al., 2004; Schultz, 2003; Schulz, Gopnik, & Glymour, submitted). Moreover, children can use this information to design novel effective interventions themselves, suggesting genuine causal insight (Gopnik et al., 2004; Schulz & Gopnik, 2004).

As noted, some research has looked at how children might use evidence from interventions and outcomes to infer the existence of unobserved causes (Gopnik et al., 2004). However, children seem to reason about unobserved and even unobservable causes quite broadly; they are not limited to reasoning about unobserved common causes of correlated events. In this paper, we will consider the possibility that children combine inferences about interventions and outcomes with a belief in causal determinism.

In its simplest form, causal determinism is the assumption that all events have causes (see, e.g., Bullock et al., 1982; Gelman et al., 1994). If you believe that all events have causes, then you should infer unobserved causes whenever events appear to occur spontaneously. There is considerable evidence that both adults and children do this (Bullock et al., 1982; Chandler & Lalonde, 1994; Gelman et al., 1994; Luhmann & Ahn, 2003; Saxe, Tenenbaum, & Carey, in press).

However, causal determinism can entail a stronger set of commitments. In the philosophical literature, causal determinism is the assumption, not just that events have causes, but that causes deterministically produce their effects. From this perspective, the appearance of probabilistic causality is due to our ignorance of all the relevant variables. This strong version of determinism was famously articulated by the mathematician Pierre Simon-Laplace, who noted that if there were "an intelligence knowing all the forces acting in nature . . . (and) its intellect were sufficiently powerful to subject all data to analysis, to it nothing would be uncertain'' (1814/1951). This stronger kind of determinism implies that we should infer unobserved causes, not just when effects occur spontaneously, but also when effects occur stochastically.

We do not want to claim that Laplacian determinism provides an accurate picture of causal relations in the world. Chaos theory suggests that even deterministic events may be unpredictable and quantum mechanics suggests the existence of genuinely random events. Importantly, however, a belief in causal determinism need not be metaphysically accurate to be functionally adaptive. Indeed, assuming determinism may be adaptive precisely because it induces human beings to search for the existence of unobserved causal factors in indeterminate causal scenarios.

Specifically, for determinists, particular patterns of interventions and evidence will suggest the presence of unobserved variables. If X deterministically generates Y, the probability of Y, given an intervention to produce X, is 1: $p(Y \mid X) = 1$. If X and Y are not causally related, then an intervention to produce X should not change the probability of Y: $p(Y \mid X) = p(Y)$. Suppose, however, that an intervention to produce X increases the likelihood of, but does not guarantee, the occurrence of Y: $1 > p(Y \mid X) > p(Y)$. Assuming Laplacian determinism, this last pattern suggests that the causal structure should be modified. We need to add a variable Z whose presence inhibits X from generating Y, or whose absence impairs the ability of X to generate Y (see Glymour, 2001, for a complete discussion).

Importantly, assuming determinism might support inferences not only about the existence of unobserved causes but also about what *kind* of unobserved cause is present. In particular, children might be able to trade off inferences about the presence of unobserved inhibitory causes and the absence of unobserved generative causes. If for instance, children believe that a necessary unobserved generative cause of the event is sometimes absent, they should be less likely to infer the existence of an unobserved inhibitory cause (and vice versa).

This sort of trade-off, combined with other kinds of knowledge, might also allow children to decide whether unobserved causes are generative or inhibitory. Suppose that a generative cause behaves stochastically. If this is due to an unobserved inhibitory cause, then that cause should be *absent* when the observed generative cause produces an effect and *present* when the observed generative cause fails to produce the effect. However, if the unobserved cause is generative, then the unobserved cause will be *present* when the generative cause produces an effect and *absent* when the generative cause fails to produce the effect. Children might use these facts to infer whether the unobserved cause was inhibitory or generative.

A belief in causal determinism might also lead children to prefer certain causal hypotheses to others. If children are parsimonious about positing unobserved causes, then given a choice between two hypotheses, that is (1) a potential cause that deterministically produces an effect or (2) another potential cause that produces the effect stochastically because of an unobserved variable, children should prefer the former account. Critically, however, if determinism is a belief about causal relations in the world, then children's assumptions about determinism should be sensitive to the causal structure underlying the event. Children might accept that two events could be stochastically associated for an arbitrary reason (e.g., a person decided sometimes to make the events happen together and sometimes to make the events happen separately) but reject the idea that one variable could be a direct stochastic cause of another. In all these respects, the assumption of strong causal determinism could shape the way in which children learn the causal structure of the world.

Earlier research has looked at how children use patterns of covariation to make causal judgments about both generative and inhibitory observed causes (Gopnik, Sobel, Schulz, & Glymour, 2001; Schulz & Gopnik, 2004; Shultz & Mendelsohn, 1975; Siegler, 1976). However, to our knowledge, no other research has looked at how determinism and patterns of covariation and intervention affect children's inferences about unobserved causal variables. In this paper, we look at whether children are causal determinists in the domain of physical causality. Experiment 1 looks at whether children infer the existence of unobserved inhibitory causes when observed causes behave stochastically and whether children appropriately trade-off inferences about the presence of unobserved inhibitory causes and the absence of unobserved generative causes. Experiment 2 looks at whether children can use the pattern of indeterminacy to decide whether an unobserved cause is generative or inhibitory. Experiment 3 looks at whether children resist inferring direct stochastic causation.

Experiment 1

In Experiment 1, we show preschool children a generative cause of an effect and test them in one of three conditions: a deterministic causation condition and two stochastic causation conditions (unexplained and explained). In the deterministic causation condition, children see a generative cause that always produces an effect: $p(Y \mid X) = 1$, and children should not infer the existence of an unobserved cause. In the unexplained stochastic condition, children are led to believe that the presumably sufficient observed generative cause sometimes fails to produce the effect: $1 > p(Y \mid X) > p(Y)$. In this condition, children should infer the existence of an unobserved inhibitory cause. In the explained stochastic condition, children see the same evidence but are led to believe that an additional necessary unobserved generative cause might sometimes be missing. If children understand that the stochastic causation can be explained by the absence of an unobserved generative cause, then children should not infer the existence of an unobserved inhibitory cause. (Note that we call this the explained condition only because the absence of the unobserved generative cause might explain the failures of the effect—not because we actually explain this to the children; they still need to draw the appropriate inferences.)

Method

Participants

Forty-eight children ranging in age from 3 years 8 months to 5 years 5 months (mean age: 4 years 7 months) were recruited from urban preschools. An approximately equal number of boys and girls participated. Sixteen children were randomly assigned to each condition. Although most children were from White, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented.

Materials

A specially designed remote-operated light was used. The toy consisted of a light encased in a 12 cm × 17 cm × 8 cm wooden box with an orange Lucite top. When the sliding switch on the remote was put in the "on" position, the top glowed orange. When the remote was put in the "off" position, the light turned off. If the switch was only pushed part of the way, the effect failed to occur. Children were never able to see how far the switch was pushed; from the children's perspective, the switch was activated on every trial, and the toy sometimes lit and sometimes did not.

A 7 cm diameter metal ring and a 3 cm black, disc-shaped squeezable keychain flashlight were also used. The experimenter never activated the flashlight and no child identified the squeezable object as a flashlight. All children first participated in a pretest. For the pretest, a red cup, a blue cup, and a paperclip were used.

Procedure

The children were tested individually by an experimenter familiar to them. See Figure 1 for a schematic of the procedure.

Pretest

The test phase required the children to believe that the experimenter might deceive the confederate; hence children were given a false belief task (Wimmer & Perner, 1983) to introduce the idea that the experimenter might be deceptive and to insure that they could understand deception. The confederate put a paperclip under one of two cups. The confederate left the area and the experimenter switched the location of the paperclip. The children were asked to predict where the confederate would look for the paperclip. Two children failed the pretest and were replaced.

Training

The experimenter set the toy box and the remote control switch on the table and placed the ring on top of the toy box. The experimenter pushed the switch forward and the toy lit up. She then slid the switch back and the light extinguished. "See this switch? This switch makes my toy light up." The experimenter repeated this three times. Children thus had evidence that the switch was a generative cause of the effect.

Children were then given evidence for an observed inhibitory cause. The experimenter said, "The toy only works if this ring is on top of the toy. If I remove the ring, the switch won't work and the toy won't light up." The experimenter removed the ring and pushed the switch. The toy failed to light up (in fact, because the experimenter surreptitiously pushed the switch only part way). From the child's perspective, however, removing the ring prevented the switch from working and the toy from lighting up. The experimenter repeated this three times. The experimenter then put the ring back on top of the toy, pushed the switch (all the way) and the toy lit up. For the remainder of the experiment, the ring remained on top of the toy.

This procedure provided the children with a known way of preventing the effect. In all three

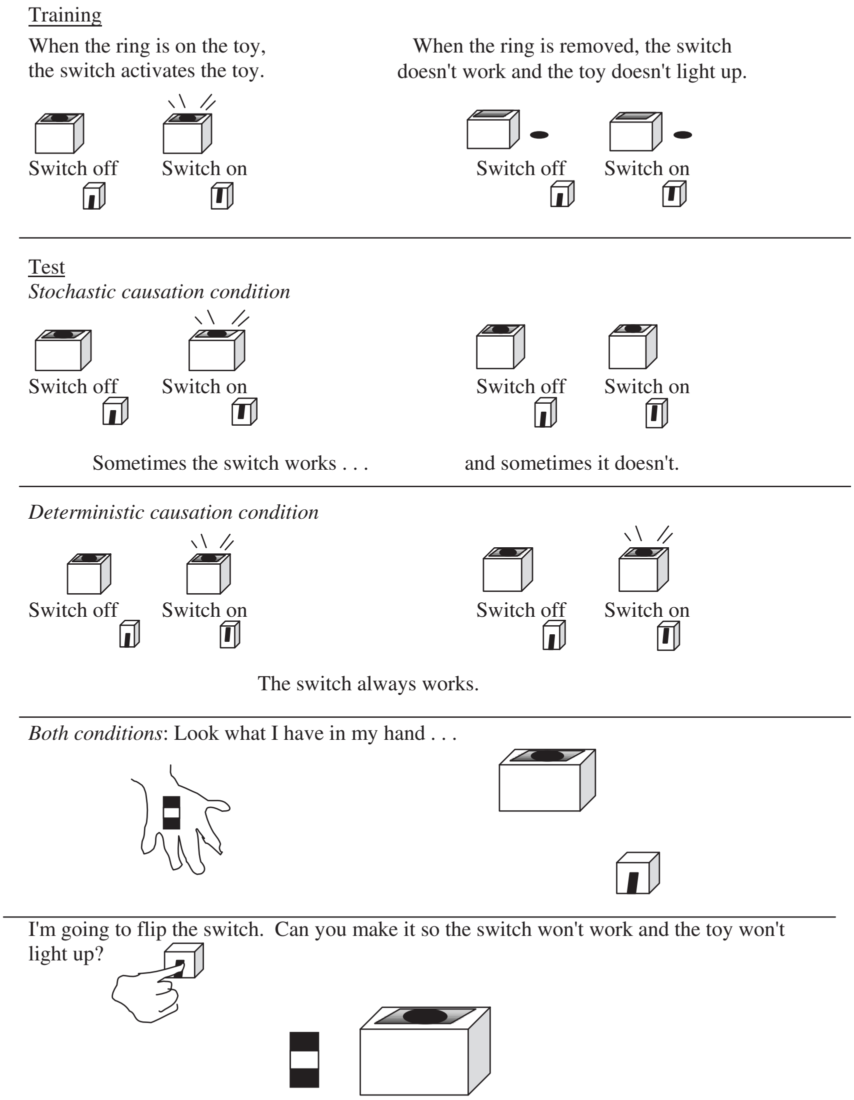

Figure 1. Schematic of the stimuli and procedure used in Experiment 1.

conditions, children had observed that removing the ring would stop the switch from working and the toy from lighting up. During the test phase, all children thus had the option of imitating the intervention that had worked in the training task.

Test Tasks

Deterministic causation condition. The experimenter gave the switch to the confederate and said, ''Now my friend Catherine is going to try to make the toy light up.'' The confederate activated the switch eight times consecutively and the toy lit up each time.

After the confederate pushed the switch the eighth time, the experimenter opened the palm of her right hand and said, ''Look what I have in my hand.'' This revealed the flashlight, which had been previously concealed. The experimenter placed the flashlight on the table, took the remote switch from the confederate, and said, ''We're going to play a game. On the count of three, I'm going to push this switch to make this toy turn on. Can you make it so the switch won't work and the toy won't turn on?''

She placed the toy with the ring on top and the flashlight within reach of the child (left/right position counterbalanced between subjects), put her own hand on the switch, and counted to three. Children did not have access to the switch but had a choice of preventing the effect by acting on the ring or by acting on the flashlight.

In this condition, the children should not infer the existence of an unobserved inhibitory cause. When asked to inhibit the effect, children should remove the ring.

Unexplained stochastic causation condition. The unexplained stochastic condition was identical to the deterministic condition, except that the switch did not always make the light go. Instead, the toy behaved in the following pattern: no effect, light, no effect, no effect, no effect, no effect, light, and no effect. From the child's perspective, the confederate pushed the switch eight times, but the toy lit up only twice.

Children had redundant evidence (intervention information and verbal instructions) that removing the ring was an inhibitory cause and would prevent the effect. However, if children are determinists and are sensitive to instances of imperfect causation, then they should infer the existence of an additional, unobserved, inhibitory cause. The flashlight concealed in the experimenter's hand might plausibly be an unobserved inhibitor of the effect. (The experimenter might have sneakily prevented the effect some of the times that the confederate attempted to generate it.) Thus, although the children never saw the flashlight do anything, they might try to inhibit the effect by intervening on the flashlight rather than on the ring.

Explained stochastic causation condition. This condition was identical to the unexplained stochastic condition, except that after the confederate pushed the button for the eighth time, the experimenter introduced a novel, necessary generative causal factor. The experimenter took the switch from the confederate and said, "You know, in order to make the toy work, you have to push the switch all the way forward. If you just push the switch part of the way, the toy won't work." For half the children, the information was introduced and the experimenter then revealed the flashlight. For half the children, the experimenter revealed the flashlight and then this information was introduced. All children received this information after all the trials were completed; hence they had no opportunity of observing, on any given trial, whether the confederate moved the switch into the correct position or not.

Note that in both stochastic conditions, children saw that the observed generative cause sometimes succeeded and sometimes failed to make the toy light up. However, in the explained condition, children were given reason to believe that there was an additional generative factor that they had not observed—namely, whether the switch was pushed all the way. This factor may have been absent on some of the trials. If the children recognize the trade-off between explaining stochastic effects in terms of the presence of an unobserved inhibitory cause and the absence of an unobserved generative cause, then children in this condition should not infer the existence of an unobserved inhibitory cause and should prevent the effect by imitating the known inhibitory intervention (removing the ring).

Results and Discussion

Children were coded as choosing the unobserved cause if they picked up the flashlight, aimed it at the toy, and either activated it or attempted to activate it by pushing on its surface. Children were coded as choosing the observed inhibitory cause if they removed the ring.

Alpha was set at .05, and thus all results reported as significant are $p < .05$ or better. In the deterministic causation condition, only 2 of the 16 children (12.5%) chose the flashlight (the rest chose the ring). In the otherwise similar unexplained stochastic causation condition, however, 15 of the 16 children (94%) intervened on the unobserved inhibitory cause (the flashlight) and only 1 child (6%) intervened on the observed inhibitory cause (the ring). In the explained stochastic causation condition, only 4 of the 16 children (25%) intervened on the flashlight (the rest chose the ring). Children were significantly more likely to choose the unobserved inhibitory cause in the unexplained stochastic causation condition than in either the deterministic causation condition, $\chi^2(1, n = 32) = 21.21$, or the explained stochastic condition, $\chi^2(1, N = 32) = 15.68$.

Within the unexplained stochastic condition, children were significantly more likely to choose the unobserved inhibitory cause than the observed inhibitory cause, $\chi^2(1, n = 16) = 12.25$. By contrast, within both other conditions, children were significantly more likely to choose the observed inhibitory cause than the unobserved inhibitory cause, $\chi^2(1, n = 16) = 4.00$ and $\chi^2(1, n = 16) = 9.00$, respectively. Indeed, when the possible absence of the generative cause could explain the effect, children were no more likely to infer an unobserved cause, given stochastic causation than deterministic causation, $\chi^2(1, n = 32) = 0.82$.

These results are consistent with the idea that children are causal determinists and can use patterns of interventions and evidence to infer the existence of unobserved causes. When children believed that the generative intervention always occurred but the effect occurred stochastically, children inferred the existence of an unobserved inhibitory cause and created a novel, appropriate intervention on a previously unobserved variable. However, when children believed that the stochastic causation might be explained by the absence of a generative cause, children did not infer the presence of an unobserved inhibitory cause. Nor did they posit unobserved causes when the effects occurred deterministically.

Many factors presumably contributed to making the flashlight a plausible, potential cause. Concealing the flashlight in the experimenter's hand allowed for the possibility that the experimenter might intervene on it. The flashlight itself had affordances (as a manipulable object with a squeezable depression) that let it to be treated as a potential button. Finally, the false belief pretest introduced the possibility that the experimenter liked to play tricks and thus that an object concealed in her hand might be causally relevant. In the absence of any of these factors, children might have been less likely to infer that the flashlight was an unobserved cause of the effect. Further research might look at the factors that determine children's willingness to treat particular variables as causally relevant.

Critically, however, these features were held constant across the conditions. Thus, although many factors might explain why children who inferred the existence of an unobserved cause chose the flashlight (rather than any other object in the room), these factors do not explain why children looked for unobserved causes in the stochastic condition but not in the other two conditions. Indeed, the results of this experiment suggest that children are relatively parsimonious about inferring the existence of unobserved causes. Children did not look for an unobserved inhibitory cause when the effects occurred deterministically or when the absence of a generative cause might explain the stochasticity. Children only inferred the existence of an unobserved inhibitory cause when they observed otherwise unexplained stochastic effects.

Experiment 2

In Experiment 1, children were only asked to infer whether or not there was an inhibitory unobserved cause. In Experiment 2, we look at whether children can use particular patterns of indeterminacy to infer whether unobserved causes are generative or inhibitory. If an observed generative cause sometimes fails to produce an effect and children believe that a potential unobserved cause could not have been activated (e.g., was not available for intervention) on failed trials but could have been activated on successful trials, then children should infer that the unobserved cause is generative. By contrast, if children believe that the potential unobserved cause could not have been activated on successful trials but could have been activated on failed trials, they should infer that it must be inhibitory.

Method

Participants

Thirty-two children ranging in age from 4 years 1 month to 5 years 7 months (mean age: 4 years 7 months) were recruited from urban preschools. One child was unable to complete the training task and was replaced. Children were randomly assigned to an effect first condition and an effect last condition. An approximately equal number of boys and girls participated. Although most children were from White, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented.

Materials

The metal ring was not used in this experiment. Otherwise, the same materials used in Experiments 1 and 2 were used in this study. Additionally, a monkey puppet, six switches, and a remote control airplane were used in a training procedure. Three of the switches were toggle switches, and were identical except for color (green, white, and red). The remaining three switches were uniquely colored and shaped (a blue panel switch, a gold push-button switch, and a cream-colored toggle switch). None of the switches were functional, although some of them appeared to be, as described in the procedure below. In the training task, the airplane was always activated or deactivated surreptitiously, using the remote control.

Procedure

Pretest

The procedure was identical to pretest in Experiment 1, except that a monkey puppet took the place of the human confederate. All children passed the pretest.

Labeling Training Task

Because the test task involved deciding whether a candidate cause was generative, inhibitory, or not really a cause at all, we trained the children to identify switches as "starters," "blockers," or "do-nothings." The experimenter set the toy airplane on the table and concealed the remote in her hand. She then set the green switch on the table.

The experimenter flipped the green switch and simultaneously (surreptitiously) triggered the remote so that the plane began to make a whirring noise. She said, "Look, this is a starter switch. See, it made the toy turn on."

The experimenter removed the starter switch and brought out the white switch. She flipped the white switch back and forth. Nothing happened. She said, "See this is a do-nothing switch. This switch doesn't do anything." The children were allowed to play with the do-nothing switch.

Finally, the experimenter took out the red switch and gave the child the starter switch. She asked the child to flip the starter switch on and the toy started whirring. The experimenter then flipped the red switch on. The toy stopped whirring. The experimenter said, "See this is a blocker switch. The blocker blocks the starter switch and makes it so the toy won't go." The child was encouraged to flip the starter on again. Nothing happened. The experimenter said, "See, the blocker is blocking it." The experimenter then flipped the red switch off. The child flipped the green switch on and the toy started whirring.

The experimenter repeated the entire demonstration and then brought out all three switches. She pointed to each switch in turn and asked the child to identify each switch and explain its function. Children were corrected if necessary.

The experimenter removed the old switches and brought out three new switches. She said, "Now I want you to figure out what these switches do." She repeated the above procedure, except that she did not label the switches or describe their effects. All the children were able to identify the three new switches as a starter, a do-nothing, and a blocker.

The training procedure differed slightly between the two conditions. In the effect last condition the procedure was exactly as described above. The children learned the labels easily; hence, to shorten the procedure, the children in the effect first condition watched the experimenter manipulate the switches (as above) but never flipped the switches themselves.

Test Tasks

Effect first condition. The experimenter set the toy light and the switch on the table and concealed the flashlight in her right hand. She brought the monkey puppet back out and said, "monkey is going to try to turn this light on." The monkey flipped the switch four times, and all four times the toy lit up. The experimenter opened her right hand and said, "Look what I have in my hand." She set the flashlight down on the table and said, "monkey, try again." The monkey flipped the switch four times and the toy failed to light up. The experimenter pointed to the flashlight on the table and said, "What do you think this does? Do you think this is a starter, a blocker or a do-nothing?" (order of the choices counterbalanced between children).

Note that the children had never observed any intervention on the flashlight. However, if children are causal determinists, then they should infer the existence of an unobserved cause when the switch behaves stochastically. As the flashlight could have been activated when the effect occurred (i.e., although it was concealed in the experimenter's hand) but could not have been activated when the effect failed (i.e., although it was sitting on the table), the children might infer that the flashlight was an unobserved generative cause of the effect (i.e., a "starter").

Effect last condition. This condition was identical, except that the first four times the monkey flipped the switch, the toy failed to light up. After the experimenter set the button down on the table and said "monkey, try again," the monkey flipped the switch four times and all four times the toy lit up.

Again, the children never observed any intervention on the flashlight. However, as the flashlight could have been activated when the effect failed to occur but could not have been activated when the effect did occur, the children might infer that the flashlight was an unobserved inhibitory cause of the effect (i.e., a "blocker").

Results and Discussion

The pattern of indeterminacy (effect first vs. effect last) and the possibility of intervention affected the children's decisions. Children were significantly more likely to identify the unobserved cause as a starter in the effect first condition than in the effect last condition, $\chi^2(1, N = 32) = 10.49$, and significantly more likely to identify the unobserved cause as a blocker in the effect last condition than in the effect first condition, $\chi^2(1, N = 32) = 8.00$.

In the effect first condition, 11 of the 16 children (69%) identified the unobserved cause (the flashlight) as a starter or spontaneously reported that the flashlight "made the toy turn on." Four children (25%) identified the unobserved cause as a blocker or spontaneously reported that the flashlight "stopped the toy." One child (6%) said that the flashlight did not do anything. By contrast, in the effect last condition, 12 of the 16 children (75%) identified the unobserved cause as a blocker or spontaneously reported that the button "stopped the toy." Two children (12.5%) identified the unobserved cause as a starter and 2 children identified the unobserved cause as a do-nothing switch.

Within the effect first condition, there was a trend for children to identify the unobserved cause as a starter more often than as a blocker, $\chi^2(1, n = 15) = 3.27$, $p = .07$, and children were significantly more likely to identify the unobserved cause as a starter than as a do-nothing switch, $\chi^2(1, n = 12) = 8.33$. Within the effect last condition, children were significantly more likely to identify the unobserved cause as a blocker than as a starter, $\chi^2(1, n = 14) = 7.14$, or as a do-nothing switch, $\chi^2(1, n = 14) = 7.14$.

Note that in both conditions, children had the option of treating the switch as a sufficient but stochastic cause of the effect. That is, they could have simply accepted that the switch sometimes worked and sometimes failed. If children accepted that the switch behaved stochasically, they could have labeled the flashlight as a do-nothing switch (particularly as they had never seen the flashlight do anything and they had been trained to label some switches as do-nothings). Instead, consistent with a belief in determinism, children seemed to resist accepting that the switch might act stochastically. Children were able to combine information about the pattern of indeterminacy of the switch with information about whether the flashlight could have been activated on particular trials to infer whether the flashlight was a starter, a blocker, or a do-nothing.

Again, although many features of the flashlight (its presence in the experimenter's hand, its affordances, etc.) may have made it a plausible candidate for children's causal inferences, these do not explain children's different inferences in the two conditions. The only actual difference in the two conditions was the pattern of indeterminacy. Consistent with a belief in determinism, the pattern of indeterminacy seemed to dictate the particular judgments children made.

Experiment 3

If determinism is a belief about causal relations in the world, then children's inferences should be sensitive to the causal structure underlying events. Suppose, for instance, that children believe that there is a direct causal link between a purple button and a light so that intervening to depress the button specifically makes the light turn red. If the light does not turn red, but instead turns yellow, then children should make one of three inferences: (A) that the purple button was not depressed, (B) that the purple button was depressed and an unobserved inhibitory cause was present, or (C) that the purple button was depressed and a necessary unobserved generative cause was absent. If, as suggested by Experiment 1, children are parsimonious about inferring the presence of unobserved causes, they should prefer inference A.

However, suppose instead that there is no direct causal link between the button and the light. Instead, a timer set to different intervals (deterministically) triggers each event: whenever the timer triggers the purple button, the purple button depresses and whenever the timer triggers the light, the light turns red. As the two events are the result of a cause that can (because of the particular nature of a timing device) be set to create arbitrary associations, the events might sometimes co-occur and sometimes fail to co-occur. In this case, children should not infer from the fact that the light did not turn red, that the purple button was not depressed. If children are causal determinists, they should accept the idea that two events can be stochastically associated (one event can occur with or without the other) as long as there is no direct causal link between them.

In Experiment 3, we looked at whether children's inferences about determinism are affected by the causal structure underlying the events. Specifically, we looked at whether children would resist inferring stochastic relationships between variables when one variable was a direct cause of another.

Method

Participants

Sixty-four children ranging in age from 3 years 11 months to 5 years 3 months (mean age: 4 years 6 months) were recruited from urban preschools. Children were randomly assigned to a direct cause novel pairing, a direct cause unexpected, a no direct cause novel pairing, or a no direct cause unexpected condition. An approximately equal number of boys and girls participated. Although most children were from White, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented.

Materials

A black cardboard box (30 cm × 15 cm × 10 cm) was used in this experiment. One side of the box had a wax-paper cutout of a moon; the other had a cutout of a flower. Two colored lights (green and purple) were on top of the box and two lights (red and yellow) were hidden inside the box. A blue cellophane filter was also used. The lights were encased in plastic rings and wired so that one concealed switch turned on the green and yellow lights (simultaneously); another concealed switch turned on the purple and red lights (simultaneously). The lights underneath the box were positioned so that the wax-paper cutout glowed (red or yellow) when the appropriate switch was flipped. The red and yellow filters on top of the lights could be removed, allowing the lights to glow white (when no filter was on top) or blue (when the cellophane filter was used) instead of red or yellow. A black cardboard screen was also used. See Figure 2 for a schematic diagram of the stimuli and procedure.

Procedure

Direct Cause Novel Pairing Condition

The experimenter introduced the toy to the children, showing them both sides of the box (the moon side and the flower side) and both buttons (encased in the plastic rings, the purple and green lights resembled colored buttons). Half the participants saw the moon first; half saw the flower first.

The experimenter turned the box so that one side (e.g., the flower) faced the child and the child could not see the shape on the other side. She said, ‘‘I’m going to push the purple button and let’s see what happens.’’ She pushed the top of the purple light and simultaneously (surreptitiously) flipped the switch. The ‘‘button’’ turned purple and the flower turned red. Pilot work with adults suggested that this provided a strong illusion of causality; it looked as if the experimenter had pushed a purple button that

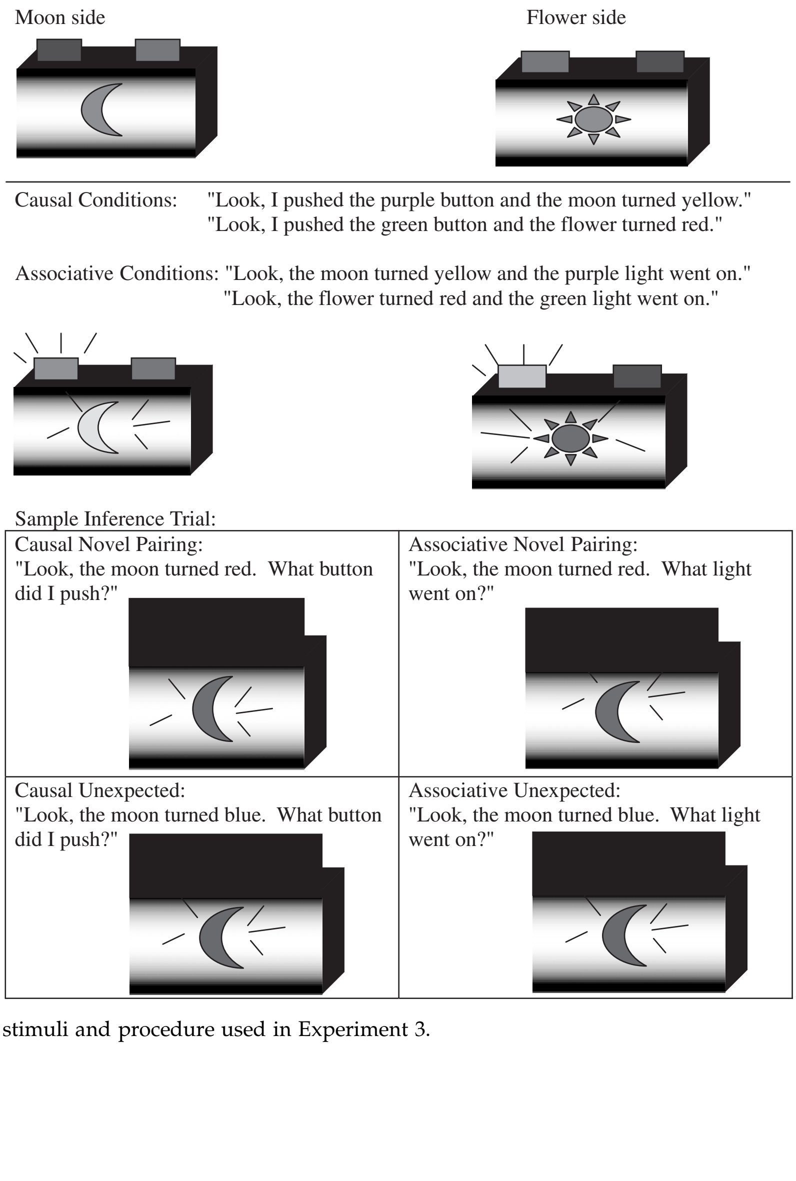

Figure 2. Schematic of the stimuli and procedure used in Experiment 3.

simultaneously lit up and caused the flower to turn red. The experimenter then turned the box to the moon side, and the children learned that pushing the green button made the green button light up and the moon turn yellow. However, children never observed the effect of the purple button on the moon or of the green button on the flower, as that side of the toy was facing away from them. The order of presentation (moon vs. flower) was counterbalanced between participants and children saw the entire procedure three times.

The children were then given two inference tasks (order counterbalanced between participants) and asked to make an inference about each of the two novel effects (also counterbalanced). In one inference task, the experimenter placed the black cardboard screen over the top two lights so that the "buttons" were no longer visible. She told the children, "I'm going to hide the buttons now." The children saw a novel effect: either the moon turned red or the flower turned yellow. The experimenter asked, "Which button did I push?"

The other inference task was similar, except that children were asked to make an intervention. The screen was not used. The experimenter turned the toy so that one of the shapes faced the child and asked the child to produce a novel effect: "Can you make the flower turn yellow?" or "Can you make the moon turn red?"

Which button would the children choose? The children had seen that the purple button turned the flower red and the green button turned the moon yellow. But they had never seen any button turn the flower yellow or the moon red, and they had never seen the relation between the purple button and the moon or the green button and the flower. Thus the children had no direct causal evidence to answer the test questions.

However, for each choice, one button suggested stochastic causation (we will call this the stochastic choice), whereas the other did not (the alternative). The children knew, for instance, that the green button caused the moon to turn yellow. If the children inferred that the green button also caused the moon to turn red, they would have to infer that the green button behaved stochastically: sometimes turning the moon yellow and sometimes not (i.e., when turning the moon red). The children had no information about what the purple button did to the moon; they had only seen it turn the flower red. However, as the shapes were on different sides of the box, the purple button could have also turned the moon red (even though the children had not seen it). If children resist inferring stochastic causation, they should consistently prefer the alternative to the stochastic choice.

Direct Cause Unexpected Condition

In the previous condition, children might have inferred that the purple button turned things red in general (and the green button turned them yellow in general) and applied this logic without being determinists. However, if children are genuinely determinists, they should resist inferring stochastic causation and prefer an alternative (potentially deterministic) cause even when it would mean attributing a novel effect to the alternative cause. To look at this, we tested children in a direct cause unexpected condition. This condition was identical to the novel pairing condition, except that immediately before the inference tasks, the experimenter surreptitiously changed the colored filters so that the moon would turn blue and the flower would turn white. As in the novel pairing condition, one button for each event suggested stochastic causation whereas the other did not. When, for instance, the moon turned blue, the children in the direct cause condition could infer that the green button behaved stochastically (sometimes turning the moon yellow and sometimes turning the moon blue). Alternatively, the children could infer that the purple button turned the moon blue, although they had no information about what the purple button did to the moon and no information that the purple button could turn anything blue. However, if the children are determinists, they should nonetheless resist the stochastic hypothesis and again prefer the alternative, choosing the purple button rather than the green one when the moon turns blue and the green button rather than the purple when the flower turns white. That is, they should think that the purple button simultaneously turns the flower red and the moon blue and that the green button simultaneously turns the moon yellow and the flower white.

No Direct Cause Novel Pairing Condition

The no direct cause conditions were perceptually quite similar to the direct cause conditions, but we introduced several changes so that the children could not infer a direct causal link between the events. First, the top (purple and green) lights were identified as lights rather than buttons. Second, children were given a cover story attributing the activation of the lights to a preset timer: "The lights are on a timer. Sometimes the lights go on and off. Let's see what happens." Third, the experimenter pointed to the lights rather than pushed them. These modifications were intended to prevent children from inferring that one light caused another and to lead them to infer instead that the lights were (deterministically) triggered by the timer but arbitrarily associated with one another. The lights, of course, were not really on a timer; the mechanism and the timing of events were the same across all conditions.

On one inference trial, children saw the moon turn red; on the other, they saw the flower turn yellow (order counterbalanced between participants; because the lights were presumably caused by the timer, children were not given an intervention task in this condition). In each case, the experimenter pointed toward the screen and asked, ''Which light do you think went on?''

In this condition, the lights were presumably triggered by a timer and were thus causally unrelated. However, the associations between the lights were identical to those in the casual condition. The green light was consistently associated with the moon and the color yellow, and the purple light with both the flower and the color red. Given a red moon, children might think the green light went on (and accept the stochastic association with color) or that the purple light went on (and accept the stochastic association with shape). A similar story applies to the yellow flower. We expected that on both trials, children would choose between the lights at chance.

No Direct Cause Unexpected Condition

This condition was identical to the novel pairing condition, except that the experimenter replaced the filters so that the shapes would shine blue or white. Again, the lights are causally unrelated, but in this condition the associations are not equivalent. When the moon turns blue, the green light is associated with the moon but not with the color blue. However, the purple light is associated with neither the moon nor the color blue. Similarly, when the flower turns white, the purple light is associated with the flower, whereas the green light is associated with neither the flower nor the color white. In direct contrast to the direct cause unexpected condition, we predicted that children in the no direct cause unexpected condition would prefer the stochastic inference to the alternative.

Results and Discussion

There were no significant differences between the inference and intervention task in either direct cause condition or between the two inference trials in either no direct cause condition; hence for all conditions, we will report the data across both trials. Children's responses are shown in Figure 3.

All the predictions were confirmed. The children were more likely to make the stochastic inference across trials in the no direct cause novel pairing condition than in the direct cause novel pairing condition, $\chi^2(1, N = 32) = 5.00$; in the no direct cause

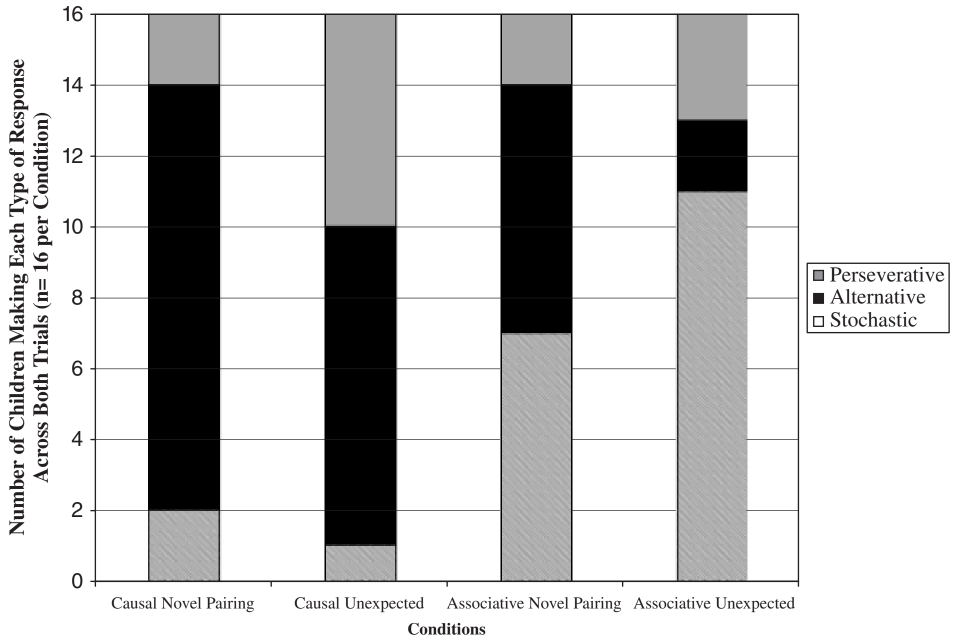

Figure 3. Children's responses in Experiment 3.

unexpected condition than in the direct cause novel pairing condition, $\chi^2(1, N = 32) = 10.49$; in the no direct cause novel pairing condition than in the direct cause unexpected condition, $\chi^2(1, N = 32) = 7.31$; and in the no direct cause unexpected condition than in the direct cause unexpected condition, $\chi^2(1, N = 32) = 13.33$.

Direct Cause Novel Pairing Condition

Within the direct cause novel pairing condition, 12 of the 16 children (75%) made the alternative inference on both trials. Only 2 children (12%) made the stochastic inference across both trials. Two children responded perseveratively (preferring the green button both times or the purple button both times). As predicted, across trials children were significantly more likely to make the alternative inference than expected by chance and significantly more likely to make the alternative inference than the stochastic inference. Children were also more likely to make the alternative inference than to perseverate on a button (all results by binomial test).

Direct Cause Unexpected Condition

In the direct cause unexpected condition, 9 of the 16 children (56%) made the alternative inference on both trials. Only 1 child (6%) made the stochastic inference across both trials. Six children (38%) chose perseveratively. As predicted, across trials children were significantly more likely to make the alternative inference than expected by chance and significantly more likely to make the alternative than the stochastic inference. The children were not more likely to make the alternative inference than to perseverate on a button (all results by binomial test).

A relatively large proportion of children made the perseverative response in this condition—choosing the same button for both trials. This condition introduced two novel effects to the children: the blue moon and the white heart. These effects might have distracted some of the children and exacerbated their tendency to perseverate. However, as predicted, children were as likely to make the alternative inference across trials (vs. anything else) in the direct cause novel pairing condition as in the direct cause unexpected condition, $\chi^2(1, N = 32) = 1.25$, $p = ns$, and as likely to resist the stochastic inference, $\chi^2(1, N = 32) = 1.25$, $p = ns$.

No Direct Cause Novel Pairing Condition

As predicted, children in the no direct cause novel pairing condition chose between the buttons at chance. Across both trials, 7 children (44%) consistently chose the button associated with color; seven consistently chose the button associated with shape; they were exactly as likely to make the stochastic inference as the alternative inference. Two children (12%) performed perseveratively.

No Direct Cause Unexpected Condition

In the no direct cause unexpected condition, 11 of the 16 children (69%) made the stochastic inference across both trials. Two children (12%) made the alternative inference across both trials and 3 children (19%) perseverated on a single light. As predicted, children were significantly more likely to make the stochastic inference across trials than expected by chance and significantly more likely to make the stochastic inference than to make the alternative inference (all results by binomial test). There was a trend for children to make the stochastic inference more often than to perserverate on a light ($p = .057$ by binomial test).

These findings suggest that children's belief in determinism helps constrain their inferences about the cause of novel events. A number of features (human intervention, mechanism knowledge, etc.) presumably influenced children's tendency to perceive a direct causal link between the events. However, critically, when children *do* think that one variable is a direct cause of another, they assume that the relationship is deterministic. Children resist making causal attributions that imply stochastic causation and look instead for alternative causal accounts. The results also suggest that children's belief in determinism is sensitive to the causal structure underlying the events. When children do not think there is a direct causal link between variables, children accept that the events might be stochastically associated.

General Discussion

These studies suggest that young children make fundamental assumptions about the causal structure of the world. At least in these mechanical cases, children seem to believe that physical causes produce their effects deterministically. This assumption allows children to use probabilities to learn about unobserved causes. It also allows children to trade off inferences about the presence of unobserved inhibitory causes and the absence of unobserved generative causes. Combined with other kinds of causal assumptions, it allows children to make inferences about whether an unobserved cause is generative or inhibitory. Moreover, the assumption of causal determinism seems to constrain children's inferences about the cause of novel effects: given an alternative, children seek to avoid inferring probabilistic causation. Finally, children's assumptions about determinism reflect other beliefs about causal structure. Although they resist the idea that direct causes are stochastic, they are willing to make stochastic inferences when there is no direct causal link between two events. In all these respects, the assumption of causal determinism seems to play an important role in shaping children's causal learning. In particular, the assumption of determinism allows children to make sophisticated inferences about unobserved events. Such inferences might allow children to discover new causal structures and might support changes in children's intuitive theories.

However, these studies also raise a number of questions. It is not clear how the degree of indeterminacy might affect children's inferences. In Experiments 1–3, the observed cause failed to produce the effect more often than it succeeded. Would children still make inferences about unobserved causes if the observed cause succeeded more often than it failed? Would even a single failure of the observed generative cause suffice for children to infer unobserved variables?

Moreover, even if children believe that unobserved causes always *exist* in indeterminate scenarios, it is not clear that they (or we) always attend to or search for such causes. In our first two experiments, we presented children with a clear and plausible candidate for the unobserved cause—the flashlight—and children acted on that cause. As noted, numerous factors (its salience, its affordances, its presence in the experimenter's hand) made the flashlight an attractive candidate cause. However in the real world, candidate causes might be difficult to identify and a variety of factors—ranging from children's temperaments, to the consequences at stake, to estimates about the number of likely unobserved variables, to the frequency of indeterminacy—might influence children's willingness to search for unobserved causes.

Additionally, although we have discussed causal determinism quite generally, our experiments looked only at children's inferences about physical causal events. Intuitively, it seems possible that children might be more willing to accept indeterminacy in psychological cases than in physical ones. Thus children might be less likely to infer the existence of unobserved variables to account for the stochastic efficacy of psychological causes (e.g., a smile sometimes but not always being reciprocated). On the other hand, in great part, what it means to understand an event *as* psychological is to understand that it can be explained with reference to unobservable variables (beliefs, desires, feelings, etc.). Thus children might not explicitly infer the existence of unobserved variables to account for psychological indeterminacy simply because unobserved psychological causes are always assumed to exist. Future research might examine domain differences in causal determinism.

Finally, a belief in strong causal determinism is not the only route by which children could infer the existence of unobserved causes. As we noted earlier, children might use spatiotemporal information or mechanism knowledge to make such inferences. There is also evidence that children infer unobserved causes whenever an event appears to occur spontaneously, implicating a weaker form of causal determinism (Bullock et al., 1982). Moreover, as noted, patterns of interventions and outcomes may indicate unobserved causes even without assuming causal determinism at all. The causal Bayes net formalism, for instance, provides a different account of how combinations of interventions and evidence could lead to the introduction of unobserved common causes (see, e.g., Gopnik et al., 2004; Pearl, 2000).

Overall, however, our findings suggest that a belief in causal determinism gives children a powerful and systematic basis for inferring the existence of unobserved causes. When children observe causal indeterminacy, they appropriately infer the existence of unobserved causes, but given deterministic causation, they parsimoniously refrain from such inferences. Children's fundamental assumptions about causal determinism appear to both enable and constrain their search for unobserved causes. In turn, this kind of inference may play an important role in the development of children's everyday theories.

References

Ahn, W. K., Gelman, S. A., Amsterlaw, J. A., Hohenstein, J., & Kalish, C. W. (2000). Causal status effect in children's categorization. *Cognition*, *76*(2), B35–B43.

Ahn, W., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *57*, 299–352.

Ahn, W., Kalish, C. W., Medin, D. L., Gelman, S. A., Luhmann, C., Atran, S., et al. (2001). Why essences are essential in the psychology of concepts. *Cognition*, *82*, 59–69.

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 Press.

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

Carey, S., & Spelke, E. S. (1994). Domain-specific knowledge and conceptual change. In S. A. G. L. A. Hirschfeld (Ed.), *Mapping the mind: Domain specificity in cognition and culture;* Based on a conference entitled "Cultural Knowledge and Domain Specificity," held in Ann Arbor, MI, Oct 13–16, 1990 (pp. 169–200). New York: Cambridge University Press.

Chandler, M. J., & Lalonde, C. E. (1994). Surprising, miraculous, and magical turns of events. *British Journal of Developmental Psychology, 12*, 83–95.

Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, *70*(5), 1098–1120.

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

Cheng, P. W. (2000). Causality in the mind: Estimating contextual and conjunctive power. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 227–253). Cambridge, MA: MIT Press.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review, 99*, 365–382.

Gelman, S. A., Coley, J. D., & Gottfried, G. M. (1994). Essentialist beliefs in children: The acquisition of concepts and theories. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 341–365). New York: Cambridge University Press.

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

Gopnik, A. (1988). Conceptual and semantic development as theory change. *Mind and Language*, *3*, 163–179.

Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*, 1–31.

Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts and theories*. Cambridge, MA: MIT Press.

Gopnik, A., & Schulz, L. E. (Eds.) (in press). *Causal learning; Psychology, philosophy and computation*. New York: Oxford University Press.

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.

Griffiths, T. L., Baraff, E. R., & Tenenbaum, J. B. (2004). *Using physical theories to infer hidden causal structure*. Proceedings of the 26th annual conference of the Cognitive Science Society.

Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence* (A. P. S. Milgram, Trans.). New York: Basic Books.

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

Koslowski, B., & Masnick, A. (2002). Causal reasoning. In U. Goswami (Ed.) *Blackwell handbook of childhood cognitive development* (pp. 257–281). Malden, MA: Blackwell Publishers.

Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review, 96*(4), 674–689.

Kuhn, D., Amsel, E., & O'Laughlin, M. (1988). *The development of scientific thinking skills*. Orlando, FL: Academic Press.

Kushnir, T., Gopnik, A., Schulz, L. E., & Danks, D. (2003). Inferring hidden causes. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th annual meeting of the Cognitive Science Society* (pp. 699–703). Mahwah, NJ: Lawrence Earlbaum Associates.

Laplace, P.-S. (1951). *A philosophical essay on probabilities* (Truscott, F. W. & Emory, F. L. Trans.). New York: Dover. (Original work published 1814.)

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

Luhmann, C. C., & Ahn, W. (2003). Evaluating the causal role of unobserved variables. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th annual conference of the Cognitive Science Society* (pp. 734–739). Mahwah, NJ: Lawrence Earlbaum Associates.

Masnick, A. M., & Klahr, D. (2003). Error matters: An initial exploration of elementary school children's understanding of experimental error. *Journal of Cognition and Development*, *4*(1), 67–98.

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

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

Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science*, *16*(12), 995–1001.

Schulz, L. E. (2003, April). *The play's the thing: Interventions and causal inference*. Paper presented at the Biennial meeting of the Society for Research in Child Development, Tampa, FL.

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40*(2), 162–176.

Schulz, L. E., Gopnik, A., & Glymour, C. (2005). *Preschoolers learn causal structure from conditional interventions*. Manuscript submitted for publication.

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. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology: Comparative & Psychological Psychology, 37*(1), 1–21.

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, Inc.

Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, *47*, 1–51.

Shultz, T. R., & Mendelsohn, R. (1975). The use of covariation as a principle of causal analysis. *Child Development*, *46*, 394–399.

Siegler, R. S. (1976). The effects of simple necessity and sufficiency relationships on children's causal inference. *Child Development*, *47*, 1058–1063.

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

Spellman, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, *7*(6), 337–342.

Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search (Springer lecture notes in statistics)*. New York: Springer-Verlag.

Strevens, M. (in press). Why represent causal relations? In A. Gopnik & L. E. Schulz (Eds.), *Causal learning; psychology, philosophy and computation*. New York: Oxford University Press.

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

Wimmer, H., & Perner, J. (1983). Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, *13*, 103–128.

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