Causal Learning Mechanisms in Very Young Children: Two-, Three-, and Four-Year-Olds Infer Causal Relations From Patterns of Variation and Covariation

Alison Gopnik, David M. Sobel, and
Laura E. Schulz
University of California, Berkeley

Clark Glymour
Carnegie-Mellon University

Three studies investigated whether young children make accurate causal inferences on the basis of patterns of variation and covariation. Children were presented with a new causal relation by means of a machine called the “blicket detector.” Some objects, but not others, made the machine light up and play music. In the first 2 experiments, children were told that “blickets make the machine go” and were then asked to identify which objects were “blickets.” Two-, 3-, and 4-year-old children were shown various patterns of variation and covariation between two different objects and the activation of the machine. All 3 age groups took this information into account in their causal judgments about which objects were blickets. In a 3rd experiment, 3- and 4-year-old children used the information when they were asked to make the machine stop. These results are related to Bayes-net causal graphical models of causal learning.

For the past 15 years, a number of cognitive developmentalists have argued that children’s cognitive development is similar to theory formation and change in science (Carey, 1985; Gopnik, 1988; Gopnik & Meltzoff, 1997; Gopnik & Wellman, 1994; Keil, 1989; Wellman, 1990). The assumption behind their work has been that there are common representations and rules—common cognitive mechanisms—that underpin both of these important types of learning. However, there has been little research specifying in detail just what those learning mechanisms might be like.

In this article we explore one such mechanism that allows children to learn about new causal relations. Causal learning is particularly important in theory formation and change. Theories specify the causal relations among objects and events. Learning a new theory involves learning about a new set of causal relations. A number of studies suggest that adults both have a great deal of causal knowledge and are adept at learning about new causal relations. In cognitive psychology, investigators have found that adults can make accurate inferences about causal relations on the basis of observations of the variation and covariation among events (see, e.g., Cheng, 1997; Shanks, 1985; Spellman, 1996; Tenenbaum, 1999). In social psychology, investigators have found that adults also make complex causal inferences when they explain the actions of others. These inferences often involve discounting

some possible causes of an action in favor of others (see, e.g., Kelley, 1973; McClure, 1998; Morris & Larrick, 1995). Are similar abilities also present in young children?

In the past, developmentalists assumed that very young children had little causal knowledge. Piaget, for example, described infants and preschoolers as “precausal” (Piaget, 1929, 1930). More recently, however, investigators have demonstrated that very young children understand important causal relations among physical objects (Bullock, Gelman, & Baillargeon, 1982; Leslie & Keeble, 1987; Oakes & Cohen, 1990; Spelke, Breinlinger, Macomber, & Jacobson, 1992), biological kinds (Gelman & Wellman, 1991; Inagaki & Hatano, 1993; Kalish, 1996), and psychological entities (Flavell, Green, & Flavell, 1995; Gopnik & Wellman, 1994; Perner, 1991). Before age 5, children seem to understand important things about how physical objects cause each other to move; how biological entities cause growth, inheritance, and illness; and how desires, emotions, and beliefs cause human actions. Young children can make appropriate causal predictions, provide causal explanations, and even make counterfactual causal claims (Harris, German, & Mills, 1996; Wellman, Hickling, & Schult, 1997).

Moreover, there appear to be systematic changes in the kinds of causal knowledge that children possess between birth and age 5 (see, e.g., Gopnik & Meltzoff, 1997). There are even some studies (Slaughter & Gopnik, 1996; Slaughter, Jaakkola, & Carey, 1999) that suggest that exposure to relevant evidence about a particular domain can accelerate the development of a causal understanding of that domain. This research suggests that children are actually learning about these causal relations from evidence.

This work has profoundly changed our view of young children’s cognitive abilities. However, simply charting children’s developing natural understanding of causal relations does not by itself tell us how that knowledge is represented or how, or even whether, it is learned. Several investigators have suggested that children’s early knowledge of folk physics, biology, and psychology might consist of elaborations of innate domain-specific explanatory prin-

Alison Gopnik, David M. Sobel, and Laura E. Schulz, Department of Psychology, University of California, Berkeley; Clark Glymour, Department of Philosophy, Carnegie-Mellon University.

We thank Josh Tennenbaum, Peter Godfrey-Smith, Tamar Kushnir, Jennifer Esterly, Andrea Rosati, and Jennie Pyers for helpful discussions; Ankush Agarwal, Valerie Castillo, Aaron Knapp, and Alex Moskovich for assistance in data collection; and all of the participants and their parents.

Correspondence concerning this article should be addressed to Alison Gopnik, Department of Psychology, University of California, Berkeley, California 94720. Electronic mail may be sent to gopnik@socrates.berkeley.edu.

ciples (Keil, 1995; Leslie & Keeble, 1987; Spelke et al., 1992). Children might not, in fact, have more general techniques for inferring new causal structure. Alternatively, children might gain much of their causal knowledge through instruction from parents or teachers, rather than through observation or inference.

In fact, some studies have suggested that children have difficulty understanding how patterns of evidence bear on new causal hypotheses (Kuhn, 1977, 1989). In these studies, children were asked to say what kinds of evidence would be needed to verify or falsify a hypothesis. Children (and sometimes adults) often made errors. Such studies suggested that young children might have difficulty actually bringing evidence to bear on causal hypotheses. However, it is also possible that children might actually draw correct causal inferences from patterns of evidence without understanding that that is what they are doing. They might have the cognitive capacity to implicitly draw appropriate causal inferences without having the meta-cognitive capacity to consciously decide which evidence will verify a causal hypothesis.

To work out the child's learning mechanisms more precisely, we need a method that will allow us to see, on-line, how children learn about a novel causal relation (similar to Siegler & Crowley's [1991] microgenetic method). Moreover, we would like to be able to control the kinds of evidence children have about this new relation and to see whether they draw genuinely causal conclusions on the basis of that evidence. And we would like to be able to test whether children can learn new causal relations that were not given innately and were not the result of explicit instruction.

In earlier work, we developed such a method—"the blinket detector." The blinket detector is a machine that lights up and plays music when some objects ("blinkets"), but not others, are placed on it. Children are thus confronted with a new, nonobvious, causal relation. Something about some of the objects makes them have the causal power to light up the machine. Note that the children do not themselves act on the detector; instead, they simply observe the contingencies between the objects and the effect. In earlier studies (Gopnik & Sobel, 2000; Nazzi & Gopnik, 2000), children as young as 2 years of age quickly learned about this new causal relation and used it to name and categorize the objects. Children saw that some objects made the machine go and that other objects did not. They were then told that one of the causally efficacious objects was a blinket, and they were asked to find another blinket. Very young children extended the term "blinket" to objects with similar causal powers.

However, these experiments raised a deeper question: How do children learn about that new causal relation? What exactly is it in their experience that tells them that the blinkets are, in fact, causally related to the machine and the nonblinkets are not?

There are two broad possibilities. Children could use what we might call *substantive principles* about particular causal relations. Children could use their prior knowledge to make top-down inferences about when particular types of events are likely to be causally related to other types of events. They might assume, for example, that pushing buttons typically makes machines go. They might then apply this prior knowledge about how other machines work to the blinket detector and look for a button to push.

But children could also use what we might call *formal principles*, principles about the pattern of contingencies between the presence of the object, the activation of the machine, and other events. They could examine how the presence or absence of

particular blocks is correlated with the behavior of the machine and use this information to make more data-driven causal inferences. Such inferences would not depend on prior knowledge. Children undoubtedly use both types of information, but in the current study we focus on the latter possibility.

Screening Off

People often conclude that there are causal relations between two kinds of events by observing what Hume (1739/1978) called *constant conjunctions* between those two kinds of events. If whenever an event of kind A happens, it is closely followed by an event of kind B, one infers that As cause Bs. However, there is a notorious problem with this sort of reasoning—the association of As and Bs might be due to other events, Cs, that produce B and are associated with A. A does not cause B, but whenever C occurs, both A and B will occur together or will more probably occur together. For example, a woman may notice that when she drinks wine in the evenings, she is likely to have trouble sleeping. It could be that the wine is causing her insomnia. But suppose she usually drinks wine in the evenings when she goes to a party. The excitement of the party might be keeping her awake, independently of the effect of the wine. Drinking wine might be associated both with excitement and insomnia, but it would not cause the insomnia. Drinking wine would be associated with insomnia, and yet it would be wrong to conclude that there was a causal relation between the two.

In such cases, if A, B, and C can be manipulated, we can determine the influence of A on B by intervening to alter whether A is present or absent—without altering the presence or absence of C—and observing the resulting variation in B. In our example, the woman could try sober partying or solitary wine drinking and observe the effects of these interventions on her insomnia. But in many circumstances we may want causal information before we can carry out such informal experiments. For over a century, the standard statistical strategy for making such inferences has been to consider whether A and B are associated in the subset of all examined cases in which C is present or in the subset of all examined cases in which C is absent. If A and B are not associated in either subset, one infers that A does not cause B. In that case, A and B are commonly said to be independent conditional on C (and on the absence of C). Any causal connection between A and B is removed by holding C constant. In other words, there is no direct causal connection between A and B. Introducing a more descriptive terminology, Hans Reichenbach, the late philosopher of science, said that in such circumstances, C *screens off* A from B (Reichenbach, 1956).

Reasoning of this kind is ubiquitous in science: It is the rationale behind both techniques of experimental design and statistical methods such as partial correlation and regression. Screening-off reasoning by itself does not always lead to correct conclusions. There are various possible, if improbable, circumstances in which A does indeed influence B (and an experimental intervention would show as much) but C screens off A from B. The converse positive inference, that if A and B are not screened off by any observed C, then A causes B, is also not always correct, especially if other unobserved variables are involved.

However, while screening-off reasoning by itself does not inevitably lead to correct causal inferences, screening-off informa-

tion can be combined with additional assumptions to draw provably reliable causal conclusions. In the artificial intelligence literature on causal Bayes nets, such assumptions have been formalized and their adequacy proved (Pearl, 1988, 2000; Spirtes, Glymour, & Scheines, 1993, 2000). Moreover, work on causal Bayes nets has provided algorithms for reliably finding causal relations from data that exhibit (or fail to exhibit) screening-off relations (Glymour & Cooper, 1999). These algorithms can uncover quite complex causal structures involving many variables, and they can even generate unobserved variables. In short, there are computational methods that allow one to draw complex and normatively accurate causal inferences from data about patterns of dependent and independent probability among events.

This means that, at least in principle, considering such screening-off relations could help adults and children make accurate causal inferences. We know that there are computational procedures that can derive provably correct causal conclusions from information about dependent and independent probability. Given the evolutionary advantages of accurate causal inference, the human mind might be designed to employ such procedures, at least in part.

We can draw an analogy to the visual system. This system is designed to infer accurate information about the structure of the spatial world from the patterns of visual input that arrive at the retina. In vision science, psychologists and computer scientists have collaborated to uncover the mechanisms that allow these inferences to be made. Computer scientists can discover algorithms that will accurately recover spatial information from visual data (and implement those algorithms in computers). Psychologists can then examine whether those algorithms are instantiated in the visual system. In fact, as it turns out, there is often a great deal of convergence between the two endeavors.

Similarly, we can think of the human theory-formation system as a system designed to infer accurate information about the underlying causal structure of the world from patterns of perceptual input. Theories and theory formation let us go from the patterns of evidence we see around us to a deeper and more accurate representation of causal structure. The normative computational work means that we can investigate whether adults and children make causal inferences when correct inferences are possible and avoid making such inferences when they are not, and ideally, we can investigate the detailed procedures by which such inferences are, in fact, made in adults and children. Of course, these procedures may or may not be similar to the procedures used in computer science. However, the example of vision science holds out hope that the two enterprises might converge.

The first step in this research program is to determine whether the information that one variable does or does not screen off another pair of variables is, in fact, used in causal judgments. The cognitive studies with adults (Cheng, 1997; Cheng & Novick, 1990; Shanks & Dickinson, 1987; Spellman, 1996) suggest that adults do use screening-off information in causal inference. Similarly, the causal discounting that is studied in social psychology involves a kind of screening-off reasoning and can be represented in Bayes net terms (Pearl, 1988). Adults will make judgments about whether one event caused another using this kind of covariation information. But adults, particularly the university undergraduates who are the typical participants in these studies, have extensive experience and often explicit training in causal inference.

At the other end of the spectrum, there is evidence for a kind of primitive screening off in the animal conditioning literature. In fact, in classical conditioning, animals show changes on behavioral measures of strength of association that are similar to screening-off relations. In the phenomenon of *blocking*, an animal who is trained to respond to a light followed by a shock, and then is trained to respond to a light and tone in combination followed by the shock, will fail to have a fear response to the tone by itself (Kamin, 1969). The animal will, in a sense, screen off the tone. Presumably, the evolutionary basis for this phenomenon is that, in nature, screened-off events are unlikely to be causes of negative effects and so need not be avoided.

However, this sort of conditioning seems very different from adults' causal judgments. Classical conditioning involves only a limited number of ecologically significant events such as shock or food, but causal judgment extends to a wide variety of new events. Classical conditioning requires many trials to establish a response, whereas causal judgments can be made on the basis of just a few relevant pieces of information. Classical conditioning does not support new interventions—although animals may have a fear response when the light occurs, they do not seem able to craft an intervention that would alter the effect (say, by actively intervening to eliminate the light and thus the shock). In contrast, if people judge that one event caused another, they can try to bring about the cause in order to bring about the effect. Finally, in causal judgment, people can combine information about a new causal relation with earlier causal information to craft new and more complex interventions. For example, if one concludes that A makes B go, one may also conclude that removing A will stop B.

Will very young children use screening-off information in a way that goes beyond mere conditioning and fulfills these criteria for genuine causal judgment? If they do, that suggests that powerful capacities for causal learning are in place very early and may play an important role in the acquisition of causal knowledge. Alternatively, these capacities might themselves be the result of extensive experience or training and emerge only in relatively sophisticated adults. In that case, we would have to look to other mechanisms to explain the changes in causal knowledge in childhood. The current experiments were designed to answer the following question: In simple cases in which correct causal inference is possible from data showing patterns of dependent and independent probability, will very young children appropriately draw genuinely causal conclusions?

Experiment 1

Method

Participants. Nineteen 3-year-olds and 19 four-year-olds were recruited from two urban area preschools. Three children from each group were excluded from the experiment for failing control questions (see below), which left a sample of 16 three-year-olds ranging in age from 3 years 1 month to 3 years 10 months (mean age = 3 years 6 months) and 16 four-year-olds ranging in age from 4 years 6 months to 5 years 3 months (mean age = 4 years 9 months). Approximately equal numbers 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 same specially designed "blicket detector" box that was used by Gopnik and Sobel (2000) and Nazzi and Gopnik (2000) was used

in this experiment. The detector measured $5 \times 7 \times 3$ in. ($13 \times 18 \times 8$ cm) and was made of wood (painted gray) with a red lucite top. Two wires emerged from the detector's side. One was plugged into an electrical outlet, and the other ran to a switchbox. If the switchbox was in the "on" position, the detector would light up and play music when an object was placed on it. If the switchbox was in the "off" position, the detector would do nothing when an object was placed on it. During the experiment, this switchbox wire ran to a confederate who surreptitiously flipped the switch on to allow an object to set the machine off or flipped the switch off to ensure that an object would not set the machine off. The wire and switchbox were hidden from the children's view, and they had no suspicion of the role of the confederate, whom they never saw. The apparatus was designed so that when the switch was on, the box "turned on" as soon as the object made contact with it, and the box continued to light up and play music as long as the object continued to make contact with it. It "turned off" as soon as the object ceased to make contact with it. This design provided a strong impression that something about the object itself caused the effect.

Seventeen wooden blocks of different colors and shapes were also used. No two blocks were identical.

Procedure. All children were tested by a male experimenter with whom they were familiar. Children were brought into a private game room in their school and sat facing the experimenter at a table. The detector was on the table. Children were introduced to the blinket detector by being told that the machine was a "blinket machine" and that "blinkets make the machine go." Children were told that the experimenter "was going to put some things on the machine" and wanted the child to tell him "which things were blinkets."

Training phase. Three blocks were then placed in front of the machine. One at a time, each of the blocks was placed on the machine for approximately 3 s. Two of them (randomly determined) made the machine light up and play music; one did not. Children were then asked if each block was a blinket. If the child answered incorrectly, he or she was reminded that "blinkets make the machine go," and the demonstration and question were repeated. We then repeated this pretest with another set of three blocks to ensure that the children understood the relationship between an object's setting off the machine and its being labeled a "blinket."

Test phase. During the test phase, the children were shown three types of tasks. In the *one-cause* tasks, children were shown two blocks, A and B. The experimenter placed each block on the detector by itself, in counter-balanced order. Block A activated the machine, and Block B did not. The experimenter then placed both blocks on the machine together, simultaneously, and the machine activated. He then placed both blocks on the machine again, and again the machine activated. The experimenter then took the blocks off the machine simultaneously and placed them on the table. He then pointed to each block, in counterbalanced order, and said, "Is this one a blinket?" Children received two of these tasks, counterbalanced for the spatial location of the block that set off the machine. Different blocks were used on each trial for each child.

In the *two-cause* task, children were shown two new blocks. The experimenter placed each block on the detector by itself three times, in counterbalanced order. One of the blocks, A, activated the detector all three times. The other block, B, did not activate the detector the first time but did activate it on the following two presentations. The experimenter then pointed to each block and said, "Is this one a blinket?" Children again received two of these tasks, counterbalanced for spatial location.

If children infer causal relations from associations and the absence of causal relations from the absence of associations or the absence of conditional associations (i.e., if they pay attention to screening-off relations), they should come to different conclusions about the causal structure of the one-cause and two-cause tasks. In the one-cause task, children should conclude that A is a blinket but that B is not, because A is associated with the effect when B is absent, but B is associated with the effect only when A is also present. That is, the effect is dependent on A independent of B, but the effect is only dependent on B conditional on A. Therefore, A causes

the effect and B does not. Notice that this conclusion should be reached despite the fact that children see B positively associated with the effect twice and negatively associated with the effect only once.

The contingencies between the blocks and the effect were the same in the one-cause and the two-cause tasks. A was associated with the effect all three times, and B was associated with the effect two out of three times. In the two-cause case, however, an analogous procedure should tell the children that both A and B are blinkets, because they each independently increase the probability that the machine will be activated. There are two independent causes of the effect.

Notice also that the one-cause and two-cause tasks are similar in other ways and thus act as controls for each other. Children might use a simple strategy or bias to choose the A block rather than the B block as the blinket in the one-cause task. For example, rather than use screening-off information, they might choose the object that activated the detector more often, or reject the one that had one negative trial, or pay attention only to the first trial of each block. However, in each of those cases they should also choose the A block and reject the B block in the two-cause task.

Finally, children were given a control task to ensure that they were on-task, understood the nature of the task, and could answer the questions correctly. It was similar to the pretest. Three blocks were placed in front of the machine. The experimenter placed each block on the machine one at a time. At least one and at most two of the blocks (randomly determined) activated the machine. The experimenter then replaced the blocks on the table, pointed to each block, and asked, "Is this one a blinket?" In order to be included in the analysis, children had to label only the blocks that activated the machine as a blinket. Three children from each age group were excluded for this reason.

The five tasks were presented in a random order to the children with the constraint that the first task not be the control task.

Results and Discussion

Initial *t* tests revealed no difference in the children's performance between the two different one-cause trials or between the two different two-cause trials, so performance was averaged across the two trials of each type. Table 1 shows children's performance on both the one-cause and two-cause tasks for the block that set the machine off 100% of the time (i.e., the A block) and the block that set the machine off 66% of the time (i.e., the B block). A 2 (age: 3-year-olds vs. 4-year-olds) $\times$ 2 (contingency: 100% vs. 66%) $\times$ 2 (task: one-cause vs. two-cause) mixed analysis of variance (ANOVA) was performed with contingency and task as within-subject variables and age as a between-subjects variable. There was a main effect of age; 3-year-olds said more blocks were blinkets than did 4-year-olds, $F(1, 30) = 10.16$, $p < .005$, $MSE = 0.34$. There was a main effect of task; children were more likely to say an object was a blinket in the two-cause tasks than in

Table 1
Mean Number of Trials (Out of 2) on Which Children in Experiment 1 Said an Object Was a Blinket, as a Function of Age, Task, and Contingency

Age	One-cause task				Two-cause task			
	100%		66%		100%		66%	
	M	%	M	%	M	%	M	%
3 years ($n = 16$)	2.00	100	1.31	66	1.94	97	1.69	85
4 years ($n = 16$)	1.81	91	0.31	16	1.94	97	1.56	78

the one-cause tasks, $F(1, 30) = 27.13, p < .001, MSE = 0.21$. There was also a main effect of contingency; children were more likely to say that the 100% blocks were blickets than that the 66% blocks were, $F(1, 30) = 54.73, p < .001, MSE = 0.29$.

In addition, several interactions were significant. First, there was a significant two-way interaction between task and contingency, $F(1, 30) = 19.74, p < .001, MSE = 0.25$. Examination of Table 1 reveals that on the one-cause tasks, children in both age groups were more likely to say that the 100% A block was a blicket than that the 66% B block was: for 3-year-olds, $t(1, 15) = 3.47, p < .005$; for 4-year-olds, $t(1, 15) = 7.35, p < .001$. In contrast, on the two-cause tasks, these differences were not significant; both groups of children tended to say that both blocks were blickets. Interestingly, in these tasks, children were quite happy to say that the 66% object was a blicket in spite of the negative trial.

In addition, there were significant two-way interactions between task and age, $F(1, 30) = 10.76, p < .005, MSE = 0.21$, and between contingency and age, $F(1, 30) = 6.09, p < .05, MSE = 0.29$. These interactions are explained by the fact that the 3-year-olds were significantly more likely to say that the 66% B block was a blicket in the one-cause tasks than were the 4-year-olds, $t(1, 30) = 4.02, p < .001$. There were no other effects of age.

We also examined whether children were significantly more likely to choose the B block in the two-cause task than in the one-cause task. In fact, this was true for both age groups, although the effect was more marked in the 4-year-olds: for the 4-year-olds, $t(1, 15) = -5.00, p < .001$; for the 3-year-olds, $t(1, 15) = -1.86, p < .05$.

In short, the 4-year-olds, in particular, responded closely in accordance with the assumption that they took features that were associated and not screened off by a third feature to be causally related and did not take features that were screened off by a third feature to be causally related. Four-year-olds were very likely to say that A was a blicket and B was not in the one-cause task, and they were very likely to say that both A and B were blickets in the two-cause task. The 3-year-olds' pattern was consistent with the same assumption with one possible exception. In the one-cause task, although 3-year-olds said that the 100% A block was more likely to be a blicket than was the 66% B block, they still said that the 66% B block was a blicket a majority of the time. However, recall that the children were asked a yes/no "Is this a blicket?" question. Some of these younger children may have had a "yes bias," a tendency in general to say that objects were blickets. In fact, recall that, overall, 3-year-olds were more likely to say that objects were blickets than were 4-year-olds. This yes bias may have influenced the 3-year-olds' responding.

Experiment 2

In Experiment 2, we set out to see if we could demonstrate similar inferences in even younger children: 2-year-olds. Moreover, to eliminate the possibility of a yes bias we modified the task so that children had to answer a forced-choice question between two options rather than a yes/no question. We also made the causal nature of the detector even more salient by allowing the children themselves to place blocks on the detector during the pretest and by reminding them about the detector during the tasks.

Method

Participants. Twenty 30-month-old children were recruited from a list of babies born in the vicinity of an urban-area university. Four children were excluded for failing control questions (see below), which left a sample of 16 thirty-month-old children. Approximately equal numbers of boys and girls participated in the study. Although most children were from White, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented.

Materials. The same "blicket detector" and a similar set of 15 blocks from Experiment 1 were used in this experiment. One of those blocks was a $4 \times 4 \times 2$ cm white block. No other block was white or had that shape. The other 14 blocks were made into 7 pairs such that both blocks in each pair were either similar in shape (e.g., one $4 \times 6 \times 2$ cm rectangle and one $4 \times 4 \times 4$ cm cube) or dissimilar in shape (e.g., one 4×1.5 cm cylinder and one $4 \times 6 \times 2$ cm triangle) to the white block.

Procedure. All children were tested by a male experimenter with their mothers present. Children were brought into a game room at the University and after a brief familiarization session were introduced to the blicket detector in the same manner as in Experiment 1. The experimenter told the children that the machine was a "blicket machine" and that "blickets made the machine go." The experimenter took out the white square block and said, "This is a blicket." The blicket was then placed on the machine, which went off. The experimenter said, "See, it makes the machine go." Children were asked if they wanted to try. All children tried the blicket on the machine, which always went off. Children were told that "the blicket always makes the machine go." Then the blicket was put away and children were told that the experimenter wanted to know which of the new objects "was like the blicket."

Training phase. Children were then shown two blocks. Each was placed on the detector separately. One block made the detector go off, and the other did not. The white square block (the blicket) was then brought out and placed on the detector, which went off, and the experimenter said, "The blicket makes the machine go." Children were asked which block "was like the blicket." If the child chose the object that had set the machine off, the experimenter moved on. If not, the two test blocks were again demonstrated on the machine, and the question was asked again. We then repeated this pretest with another pair of blocks to ensure that the children understood the nature of the machine and the question.

Test phase. The one-cause and two-cause tasks were similar to those in Experiment 1. In the one-cause tasks, children saw two blocks. Each was placed on the detector by itself once. One activated the machine, and the other did not. Then both objects were placed on the machine together, and the machine activated. This was done twice. The procedure thus far was identical to the procedure in Experiment 1. Then the experimenter placed the established blicket (the white square block) on the machine, and it activated the machine. The experimenter said, "Look, the blicket makes the machine go," and then asked which of the other two objects "was like the blicket."

In the two-cause tasks, the experimenter placed each of the two blocks on the machine three times. One block activated the machine all three times. The other did not activate it the first time but did so the second and third times, again in exactly the same way as in Experiment 1. Then the established blicket (the white square block) was brought out, and it activated the machine. The experimenter said, "Look, the blicket makes the machine go," and then asked which of the other two objects "was like the blicket."

Finally, a control task similar to that in Experiment 1 was also given. Children were shown two blocks. Each was put on the machine separately. One activated it, and one did not. Children were asked which of the two "was like the blicket." Only children who categorized the block that had set off the machine with the blicket in the control trial were included in the analysis. Four children were excluded for this reason. The five tasks were presented in a random order to the children with the constraint that the first task not be the control task.

Results and Discussion

Initial *t* tests revealed no difference in the children's performance between the two different one-cause tasks or between the two different two-cause tasks, so the children's performance on those two types of trials was averaged. Table 2 shows the number of times children chose the 100% (A) and 66% (B) blocks as the one that "was like the blicket" across conditions.

In the one-cause condition, children chose the 100% A object as being "like the blicket" 78% of the time, significantly more often than they chose the 66% B object (22% of the time; binomial test, $p < .005$). In contrast, children in the two-cause condition chose the 100% A object as being "like the blicket" only 47% of the time, which was not significantly different from the 53% of the time they chose the 66% B object (binomial test, *ns*). Recall that in Experiment 1, children were equally likely to say that the A block was a blicket and that the B block was a blicket in the two-cause task in spite of the fact that B was only effective two thirds of the time. Similarly, in this experiment, in which children had to pick only one object, they were equally likely to choose the A and B blocks.

In this modified task, even these 2-year-olds behaved in accordance with the screening-off procedure. In the one-cause task, they were likely to say that the 100% block was like the blicket and the 66% block was not, but they were equally likely to choose either block in the two-cause task. Thus, introducing the forced-choice format seemed to eliminate the yes bias but replicated the basic effect of Experiment 1.

Experiment 3

The responses in Experiments 1 and 2 clearly went well beyond the kind of blocking that appears in classical conditioning. Children learned a novel fact with no immediate ecological significance, and they did so simply by observing three events. Moreover, as in the work of Gopnik and Sobel (2000) and Nazzi and Gopnik (2000), they also seemed to identify the new object's causal power by categorizing it linguistically—saying either that it was a blicket or that it was like a blicket. As we described earlier, however, genuine causal judgments should also allow children to craft appropriate interventions and to combine the new causal information with prior causal information.

In addition, in the first two experiments, we assumed that children understood our original instructions about the causal power of the blocks and were indeed using the word "blicket" to identify the blocks that made the machine go. However, it was also possible that children were simply associating the word "blicket"

with the effect and further associating the block with the effect. They need not have understood that the blicket actually made the machine go but could simply have associated the word with the effect.

In Experiment 3, we investigated whether children would combine screening-off information with prior causal knowledge in order to craft a new intervention. In the case of physical causality, and particularly in the case of machines, a likely general substantive principle is that if an event makes a machine go, the cessation of the event will make the machine stop—this principle applies to many common cases involving switches, buttons, and so forth. This is not a necessary principle, of course, but it is a plausible pragmatic assumption about this type of causal relation. If children really think that the blickets have the causal power to make the machine go, they should also infer that removing the blickets is likely to make the machine stop even if they have never seen this event. Moreover, they should intervene appropriately to bring about this effect. On the other hand, if they are merely associating the word, the object, and the effect, children should not draw this inference, nor should they be able to craft an appropriate intervention.

In Experiment 3, we modified the task so that children did not see that removing the block made the machine stop. One block, B, was placed on the machine, and nothing happened. The B block was removed, and then the other block, A, was placed on the machine, and the machine went on. After a few seconds, the original B block was replaced on the machine next to the A block, and the machine continued to stay on for an extended time. We then simply asked the children, "Can you make it stop?" If children were drawing causal conclusions from patterns of dependent probability, and combining those conclusions with their substantive causal knowledge, they should remove the A block, rather than the B block.

Method

Participants. Twelve 3-year-olds, ranging in age from 3 years 2 months to 3 years 9 months (mean age = 3 years 6 months) and 12 four-year-olds, ranging in age from 4 years 3 months to 4 years 11 months (mean age = 4 years 6 months) were recruited from two urban-area preschools. Approximately equal numbers 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 blicket detector and a set of five wooden blocks were used in this experiment. No two blocks were identical.

Procedure. All children were tested by a female experimenter with whom they were familiar. Children were brought into a quiet hallway in their school and sat facing the experimenter at a table. The blicket detector was on the table. Children were introduced to the blicket detector by being told, "Some blocks make this machine go, and some blocks don't." Children were asked if they could help "figure out which blocks make the machine go." However, to ensure that the child had no opportunity to see that removing the blocks made the machine stop, we did not actually demonstrate the blicket detector for the child, and the word "blicket" was not used.

One-cause task. During the one-cause task, the children saw two blocks, A and B. The experimenter placed Block B on the blicket detector, and nothing happened. Then she removed Block B. Block A was then placed on the blicket detector, and it activated the machine. With Block A still on the detector, the experimenter placed Block B back on the machine,

Table 2
Mean Number of Blocks (Out of 2) That 30-Month-Old Children Chose as Being "Like the Blicket" in Experiment 2, as a Function of Task and Contingency

Age	One-cause task				Two-cause task			
	100%		66%		100%		66%	
	<i>n</i>	%	<i>n</i>	%	<i>n</i>	%	<i>n</i>	%
30 months	1.56	78	0.44	22	0.94	49	1.06	53

with its placement to the left or the right of Block A counterbalanced. The machine remained on. Then the experimenter asked the child, "Can you make it stop?"

If children were using screening-off assumptions to draw causal conclusions and combining those assumptions with their prior knowledge of physical causality, they should remove Block A rather than Block B. As in the earlier one-cause task, children saw that A was associated with the effect independent of B but that B was only associated with the effect dependent on A. That should lead them to conclude that A, but not B, caused the effect and that they should remove A, but not B, to stop the effect.

Two-cause task. Children were also given a similar two-cause task. In this task, the experimenter placed Block B on the blivet detector, and the machine activated. Then she removed Block B and placed Block A on the detector. Again the machine activated. With Block A still on the detector, the experimenter placed Block B back on the machine, with its placement to the left or the right of Block A counterbalanced, just as in the one-cause task. Then the experimenter asked the child, "Can you make it stop?"

This time, if children understood the causal properties of the blocks and everyday principles of physical causality, they should have removed both Blocks A and B from the detector. In this case, children saw that A and B were both independently associated with the effect. Children should conclude that both A and B cause the effect and that both blocks will have to be removed to stop the effect. Note also that although in this case the children see that removing Block B will make the detector stop, this is not actually the correct response. Children have to remove both blocks, a response they have never seen. In addition, as in Experiment 1, including the one-cause and two-cause tasks helps eliminate the possibility of simple strategies or biases such as choosing the block that was placed on the detector first.

Control task. Finally, children were given an additional control task to ensure that they were on task and that they had indeed made the prior pragmatic assumption that taking the block off would make the machine stop. In this task, a single block was placed on the detector, and the machine turned on. Children were asked, "Can you make it stop?"

Children were always given the one-cause task first. The other two tasks were then presented in random order. We did this to ensure that the child's insight into how to stop the machine in this task was not derived from prior associations. At the time of the one-cause task, children had never actually seen that taking a block off made the machine stop, just as in the two-cause task they had never actually seen that taking both blocks off made the machine stop. This also meant that children could not simply imitate the experimenter's action but had to craft a genuinely new intervention.

Results and Discussion

Initial analysis of performance on both the one-cause and two-cause tasks revealed no difference between the 3-year-olds and the 4-year-olds (Fisher's exact test, *ns*). Therefore the children's responses were analyzed as a single group. All of the children successfully completed the control task, indicating that they did indeed know the general pragmatic rule that if a block made the machine go, removing it would make the machine stop. Children in the other two tasks made four types of responses. They either initially took off the A block, initially took off the B block, took off both blocks together, or made no response. Table 3 shows the distribution of responses from the 24 children on the one-cause and two-cause tasks.

Children were more likely to reach for Block A in the one-cause condition than in the two-cause condition, McNemar's χ^2 (1, $N = 24$) = 11.53, $p < .001$. Children were also more likely to reach for both blocks in the two-cause condition than in the one-cause condition, McNemar's χ^2 (1, $N = 24$) = 7.11, $p < .01$.

Table 3

Number of Children per Type of Response on One-Cause and Two-Cause Tasks in Experiment 3

Response	One-cause task		Two-cause task	
	<i>n</i>	%	<i>n</i>	%
Remove Block A	18	75	3	13
Remove Block B	3	13	7	29
Remove both A and B	3	13	12	50
No response	0	0	2	8

Even when children in the two-cause condition picked only one block, their initial choice was random. In the two-cause condition, there were no significant differences between the performances of children initially choosing the A and B blocks (13% vs. 29%; binomial test, *ns*). In the one-cause condition, in contrast, the children were significantly more likely to pick Block A, the causal block, than Block B, the noncausal block (75% vs. 13%; binomial test, $p < .001$).

It might be objected that children in the two-cause condition were as likely to remove one object (or the other) initially as they were to remove both objects simultaneously. When children removed both objects in the two-cause task, they clearly demonstrated that they thought both blocks were causes. Note, however, that some children may have intended to remove the blocks one after another. This would also be an appropriate causal response. As a matter of fact, all of the children in the two-cause condition who initially removed one block did rapidly remove the other block after their initial choice, which might indicate that they did think both objects were causes. Children in the one-cause condition did not do this. However, there is a problem with this interpretation of the data. In the two-cause task, the machine continued to remain on after the first block was removed, whereas it did not in the one-cause task. This means that we cannot be sure whether this second serial response did indeed reflect a decision to remove both blocks serially in the first place or was simply a reaction to the fact that the initial response was ineffective. Thus, these serial responses do not indicate either that the children thought that both blocks were causes or that they thought that only one block was the cause. Instead, these responses could be consistent with either a two-cause or a one-cause interpretation. However, the single-block responses and the simultaneous responses do discriminate between these two cases, and they were distributed in a way that accords with the screening-off hypothesis.

To summarize, in the one-cause condition, both the 3- and 4-year-olds screened off and selectively removed the causal block. In the two-cause condition, the most frequent response was to remove both blocks simultaneously. Moreover, in the one-cause condition, children preferentially chose the causal over the noncausal block, whereas in the two-cause condition, the children who did choose only one block initially chose at random. These results are consistent with the hypothesis that preschool children reason in a genuinely causal way and use formal principles to reach causal conclusions.

General Discussion

The results of these experiments suggest that children as young as 2 years of age will indeed draw causal conclusions from patterns

of dependent and independent probability. Several things are significant about these results. Children had never seen or heard about this causal relation before the experiment itself, yet they swiftly learned about it after only a few exposures. In addition, there were no obvious perceptual clues in the blocks themselves that indicated which blocks would set off the machine. Children imputed an underlying nonobvious, causal power to the blocks. Moreover, the children initially simply observed the patterns of contingency between the blocks and the machine, and yet they generalized from that information in two ways. First, in Experiments 1 and 2, they used the information appropriately to say which blocks were blickets or were similar to a blicket. Second, in Experiment 3, they combined that information with prior causal information to appropriately predict the effects of a new action and to produce that action themselves. Thus it appears that children used the screening-off information to make a genuinely causal judgment.

Earlier we mentioned that children could use both more top-down substantive causal learning mechanisms, and more data-driven formal learning mechanisms. Although we have focused on data-driven formal learning in this article, we do not wish to imply that these procedures are more important than the more substantive principles that also underlie causal inference. In fact, in Experiment 3, we showed that children combined both kinds of information in their causal reasoning. We think that both types of causal reasoning are complementary and interact in useful ways in development. Innate substantive causal schemes—innate theories, in effect—may be important in initially telling children how to divide the world up into candidate causal variables and relations and in pointing out which variables to consider (see Gopnik & Meltzoff, 1997). However, data-driven formal causal learning mechanisms provide children with a powerful method of learning about new causal relations and modifying the causal schemes with which they start.

These formal causal learning mechanisms are an interesting kind of halfway point between domain-general and domain-specific mechanisms of cognitive development. Unlike the usual domain-specific mechanisms, causal inference procedures can be applied to input from many domains. Causal learning need not be limited to information about people or plants or objects. Causal learning need not be restricted to postulating particular types of causes for particular effects. We might explain a human action in terms of some combination or interaction among physical, psychological, and biological causes. Both children and scientists may postulate genuinely new types of causal entities and mechanisms to explain the data. For example, the theory-of-mind literature suggests that children postulate a new type of causal entity, a mental representation, to explain certain psychological phenomena (Gopnik & Wellman, 1994; Perner, 1991).

However, causal learning of the kind we are describing is more constrained than traditional domain-general learning mechanisms, such as logical inference on the one hand or associations or connections on the other. Causal inferences are not themselves necessarily deductive; they depend on contingent assumptions about how evidence and causal structure are related. Causal inferences also go beyond mere associations. The process of causal learning we have described is quite different from the entirely domain-general process of simply capturing or matching regularities, even high-level regularities, in the input, as classic associa-

tionist accounts or contemporary connectionist and dynamic systems accounts typically do.

In fact, there are many domains of cognitive development—such as the acquisition of syntactic or phonological knowledge, mathematical knowledge, or musical knowledge—that do not seem to involve the recovery of causal structure. Causal learning mechanisms might not apply to these areas. Interestingly, these are also not domains where “the theory theory” (Gopnik & Wellman, 1994) seems to be naturally applicable.

Another important point is that we think these learning procedures may well be unconscious; 3-year-olds and even adults would be hard-pressed to formulate explicitly the assumptions they use in causal inference. Indeed, we think it is very unlikely that children could consciously predict what sorts of evidence are necessary to draw causal conclusions, and children might, in fact, make errors if they were asked to do so, as they did in Kuhn’s (1989) “control of variables” studies. Instead, it appears that children simply do draw the right causal conclusions when they are presented with appropriate patterns of evidence, without any conscious access to the mechanisms that allow them to do this.

There are also several important limitations of this study. First, it did not explore many important facets of children’s uses of dependencies in data to make causal inferences, such as the sensitivity of their judgments to varying strengths of probabilistic dependency or to sample size. Interestingly, the children in this study appeared willing to accept that an object could produce an effect probabilistically—they said that the block that made the detector light up two out of three times was indeed a blicket. However, we do not know how children use this probabilistic information in their judgments.

Second, it was possible that the children paid attention only to the blocks that were placed on the detector by themselves and simply ignored the cases in which both blocks were placed on the detector. This procedure would lead to results similar to those in the present study. However, another blicket study showed that this is not the case (Sobel, 2001). In that study, we presented the children with the following sequence of events. The experimenter placed two blocks on the detector together, and the detector lit up. Then the experimenter placed the B block on the detector by itself, and the detector did not light up. Then we asked the children whether each block was a blicket. Three- and 4-year-olds in this experiment reported that the A block was a blicket and that the B block was not. This was in spite of the fact that they had only seen the A block in conjunction with the B block.

More fundamentally, there are at least two different general formal models of how screening-off relations can be used to make causal inferences. First, there is the normative account of the relationship between causality and probabilistic dependence that is given in the literature on causal Bayes nets (Pearl, 1988, 2000; Spirtes et al., 2000), which we have already described. Cheng’s causal power model is also a special case of a causal Bayes-net model (Cheng, 1997, 2000). However, there are also models of causal inference in adults—in particular, Shanks’s (1985) delta rule model—that are more similar formally to the Rescorla-Wagner learning model (Rescorla & Wagner, 1972) and to the learning procedures used in connectionist systems (Glymour, *in press*; Shanks, 1985). These latter models, unlike the Bayes-net models, do not produce normatively accurate causal judgments.

However, they might, in fact, be the models implicit in human causal judgment.

In some cases, these two models make different sets of predictions, and some preliminary work suggests that adult causal judgments accord better with the Bayes-net models (Glymour, in press; Tenenbaum, 1999). However, in simple cases like the ones in the current experiments, they yield similar predictions. Our experiments do not test which of these accounts better explains the causal judgments of young children.

Similarly, within the causal Bayes-net framework itself, there are many alternative detailed learning algorithms that allow for inferences from screening-off data to causal relations and that might explain our experimental results. Some of these procedures build causal hypotheses from identified patterns of probabilistic independence and dependence. Other methods rely more on Bayesian inference: They assign a prior probability to causal relations and, using various possible heuristics, compute an approximate posterior probability given the data (see, e.g., Glymour & Cooper, 1999). Again, the current experiments do not discriminate among these possibilities.

What we have shown is that even very young children can and do infer new causal relations from information about dependent and independent probability. Developmentalists have already demonstrated that very young children learn a great deal about the causal structure of the world and do so with remarkable accuracy and speed. Computationalists have shown that learning mechanisms can derive normatively accurate causal inferences from information about dependent and independent probability. Such learning mechanisms might well play an important role in the acquisition of causal knowledge. Future investigations should help us to uncover the nature of these mechanisms in more detail.

References

- 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.
- 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. Wilson (Eds.), *Explanation and cognition* (pp. 227–253). Cambridge, MA: MIT Press.
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, 58, 545–567.
- 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, Serial No. 243).
- Gelman, S. A., & Wellman, H. M. (1991). Insides and essence: Early understandings of the non-obvious. *Cognition*, 38, 213–244.
- Glymour, C. (in press). *The mind's arrows: Bayes nets and graphical causal models in psychology*. Cambridge, MA: MIT Press.
- Glymour, C., & Cooper, G. F. (1999). *Computation, causation, and discovery*. Cambridge, MA: MIT Press/AAAI Press.
- Gopnik, A. (1988). Conceptual and semantic development as theory change. *Mind and Language*, 3, 163–179.
- Gopnik, A., & Meltzoff, A. (1997). *Words, thoughts and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Sobel, D. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, 71, 1205–1222.
- 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.
- Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, 61, 233–259.
- Hume, D. (1978). *A treatise of human nature*. Oxford, England: Oxford University Press. (Original work published 1739)
- 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.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 234–267). New York: Clarendon Press/Oxford University Press.
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, 28, 107–128.
- Kuhn, D. (1977). Conditional reasoning in children. *Developmental Psychology*, 13, 342–353.
- Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, 96, 674–689.
- Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, 25, 265–288.
- McClure, J. (1998). Discounting causes of behavior: Are two reasons better than one? *Journal of Personality and Social Psychology*, 74, 7–20.
- Morris, M., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review*, 102, 331–335.
- Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science*, 3, 389–396.
- Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, 5, 193–207.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Mateo, CA: Morgan Kaufmann.
- Pearl, J. (2000). *Causality*. New York: Oxford University Press.
- Perner, J. (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press.
- Piaget, J. (1929). *The child's conception of the world*. London: Routledge & Kegan Paul.
- Piaget, J. (1930). *The child's conception of physical causality*. London: Routledge & Kegan Paul.
- Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, 37B, 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.
- Siegler, R. S., & Crowley, K. (1991). The microgenetic method: A direct

- means for studying cognitive development. *American Psychologist*, 46, 606–620.
- Slaughter, V., & Gopnik, A. (1996). Conceptual coherence in the child's theory of mind: Training children to understand belief. *Child Development*, 67, 2967–2988.
- Slaughter, V., Jaakkola, R., & Carey, S. (1999). Constructing a coherent theory: Children's biological understanding of life and death. In M. Siegal & C. Peterson (Eds.), *Children's understanding of biology and health* (pp. 71–96). Cambridge MA: Cambridge University Press.
- Sobel, D. M. (2001, April). *Children's ability to make inductions about unobserved causal properties of objects based on indirect evidence*. Poster presented at the 2001 Biennial Meeting of the Society for Research in Child Development, Minneapolis, MN.
- 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, 337–342.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search* (Springer Lecture Notes in Statistics). New York: Springer-Verlag.
- Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction and search* (Springer Lecture Notes in Statistics, 2nd ed.). Cambridge, MA: MIT Press.
- Tenenbaum, J. B. (1999). *A Bayesian framework for concept learning*. Unpublished doctoral dissertation, Massachusetts Institute of Technology.
- Wellman, H. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., Hickling, A., & Schult, C. A. (1997). Young children's explanations: Psychological, physical, and biological reasoning. In H. Wellman & K. Inagaki (Eds.), *New directions for child development: Children's theories* (pp. 7–25). San Francisco, CA: Jossey-Bass.

Received December 4, 2000

Revision received March 14, 2001

Accepted March 29, 2001 ■

New Editors Appointed, 2003–2008

The Publications and Communications Board of the American Psychological Association announces the appointment of five new editors for 6-year terms beginning in 2003.

As of January 1, 2002, manuscripts should be directed as follows:

- For the *Journal of Applied Psychology*, submit manuscripts to **Sheldon Zedeck, PhD**, Department of Psychology, University of California, Berkeley, CA 94720-1650.
- For the *Journal of Educational Psychology*, submit manuscripts to **Karen R. Harris, EdD**, Department of Special Education, Benjamin Building, University of Maryland, College Park, MD 20742.
- For the *Journal of Consulting and Clinical Psychology*, submit manuscripts to **Lizette Peterson, PhD**, Department of Psychological Sciences, 210 McAlester Hall, University of Missouri—Columbia, Columbia, MO 65211.
- For the *Journal of Personality and Social Psychology: Interpersonal Relations and Group Processes*, submit manuscripts to **John F. Dovidio, PhD**, Department of Psychology, Colgate University, Hamilton, NY 13346.
- For *Psychological Bulletin*, submit manuscripts to **Harris M. Cooper, PhD**, Department of Psychological Sciences, 210 McAlester Hall, University of Missouri—Columbia, Columbia, MO 65211.

Manuscript submission patterns make the precise date of completion of the 2002 volumes uncertain. Current editors, Kevin R. Murphy, PhD, Michael Pressley, PhD, Philip C. Kendall, PhD, Chester A. Insko, PhD, and Nancy Eisenberg, PhD, respectively, will receive and consider manuscripts through December 31, 2001. Should 2002 volumes be completed before that date, manuscripts will be redirected to the new editors for consideration in 2003 volumes.