

Research Report

Young Children Infer Causal Strength From Probabilities and Interventions

Tamar Kushnir and Alison Gopnik

University of California at Berkeley

ABSTRACT—*We examine the interaction of two cues that children use to make judgments about cause-effect relations: probabilities and interventions. Children were shown a “detector” that lit up and played music when a block was placed on its surface. We varied the probabilistic effectiveness of the block, as well as whether the experimenter or the child was performing the interventions. In Experiment 1, we found that children can use probabilistic evidence to make inferences about causal strength. However, when the results of their own interventions are in conflict with the overall frequencies, preschoolers favor the results of their own interventions. In Experiment 2, children used probabilistic evidence to infer a hidden causal mechanism. Though they again gave preference to their own interventions, they did not do so when their interventions were explicitly confounded by an alternative cause.*

By age 5, children have learned a great deal about the causal structure of the world (Bullock, Gelman, & Baillargeon, 1982; Gopnik & Wellman, 1994; Shultz, 1982). How is this learning possible? We explore the interaction of two cues that children might use to infer causal strength: the probabilistic relationships among events and the consequences of interventions.

Adults can use probability information to assess the strength of causal relations—they conclude that if X follows Y more often than X follows Z, then, other things being equal, Y is a stronger cause of X than Z (Cheng, 1997; Shanks & Dickinson, 1987; Spellman, 1996; Waldmann & Hagmayer, 2001). However, adults have extensive experience and often explicit tuition in causal inference. If children have similar causal learning abil-

ities naturally, they might play a role in the acquisition of causal knowledge.

However, studies of children’s causal reasoning have focused on deterministic rather than probabilistic causal relations (e.g., Bullock et al., 1982; Gopnik & Sobel, 2000; Gopnik, Sobel, Schulz, & Glymour, 2001; Shultz, 1982; Shultz & Mendelson, 1975). These studies have found that young children are capable of making quite sophisticated causal inferences from deterministic evidence. In a few studies (Gopnik et al., 2004; Siegler, 1976), children also inferred causal relations even when effects did not always follow causes. However, all these studies asked about causal structure (Did X cause Y?) rather than causal strength (How strongly did X cause Y?).

Interventions also play an important role in causal reasoning. They help solve one of the main problems of causal inference—the problem of confounding. When people intervene on a cause to bring about an effect, they can usually assume that the subsequent events are the result of their action, and not of other causes. When they observe the interventions of other people, this assumption is less justified. It is even less justified when they simply observe that events covary. This commonly known fact of scientific reasoning has been formalized as part of causal Bayes-net theory (Gopnik et al., 2004; Pearl, 2000; Spirtes, Glymour, & Scheines, 2001; Woodward, 2004). Adults’ performance on causal inference tasks improves when they are allowed to intervene on causes (Lagnado & Sloman, 2004; Sobel & Kushnir, in press; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003).

Studies have shown that preschool children can also use intervention information to make complex inferences about causal structure (Gopnik et al., 2004; Schulz, 2001). However, these studies involved deterministic relations and did not ask about causal strength, and children observed the interventions of other people rather than performing their own interventions.

The study we report here addresses the following questions: Do children equate frequency of co-occurrence with causal

Address correspondence to T. Kushnir, Psychology Department, University of California, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650; e-mail: tkushnir@berkeley.edu.

strength? What is the role of children's own interventions in their judgments of causal strength? How do frequency information and intervention information interact? Might children implicitly recognize that their own interventions help resolve problems of confounding?

EXPERIMENT 1

We presented children with a novel causal relation between objects and a "detector"—a toy that lit up and played music when objects were placed on it. Objects activated the toy one, two, or three out of three times. We also varied intervention information: Sometimes the children observed the experimenter's action, and sometimes the children intervened themselves.

Method

Participants

Participants in the study were nineteen 4-year-olds ($M = 4$ years 2 months, range = 3 years 10 months to 4 years 9 months) recruited from University of California, Berkeley, preschools.

Materials

The detector was a 5- × 7- × 3-in. box made of wood with a Lucite top. A hidden switch, controlled by the experimenter, could make the box's top light up and play music. The experimenter activated the switch only when an object was placed on the detector, creating the illusion that the object caused the lights and music. The objects were 22 blocks, each a different color and shape.

Procedure

The experimenter introduced the detector and told the child that "sometimes things make it go and sometimes things don't make it go" and that they were going to "figure out what makes it go." To familiarize the child with the procedure, the experimenter first presented a deterministic warm-up trial, which was followed by 10 test trials. The sequence of events in each of the test trials is shown in Figure 1. Each trial involved two novel blocks (A and B) that were placed on either side of the detector (counterbalanced order). Each trial began with the experimenter placing Block A on the detector twice and Block B on the detector twice. The detector activated sometimes, but not always. In the *intervention trials*, the child was then instructed to try each object once, and the detector activated on the first or second attempt or on both attempts. The matched *observation trials* were identical to the intervention trials, but the experimenter, not the child, placed each object on the detector once more. Order of trials was counterbalanced, with intervention and observation trials alternating.

There were two types of both intervention and observation trials: 3/3 trials and 2/3 trials. On 3/3 trials, one object (Block A) set off the detector three out of three times, and one (Block B) set

off the detector one out of three times. There were four 3/3 trials: two intervention trials and two matched observation trials.

On the 2/3 trials, one object (Block A) set off the detector two out of three times, and one (Block B) set off the detector one out of three times. The 2/3 intervention trials could be set up to include a potential conflict between the child's intervention and the overall probabilities. When it was the child's turn to intervene, Block A only (*nonconflicting intervention*), both blocks (*ambiguous intervention*), or Block B only (*conflicting intervention*) activated the detector. In the matched observation trials, the child saw the same sequences of events but was not allowed to intervene. There were six 2/3 trials overall, one intervention trial and one matched observation trial of each type.

After each trial, the child was asked to "pick the best one and make it [the detector] go." In response, the child made one final intervention (which activated the detector). The child's choice was recorded as either 1 (high-frequency: Block A) or 0 (low-frequency: Block B).

Results and Discussion

Results are given in Figure 1. There were no differences in responding among the four 3/3 trials, Cochran's $Q(3) = 4.8$, n.s. Across the four 3/3 trials, children chose Block A, the high-frequency block, significantly more often than chance, as indicated by a one-sample t test with 50% (two out of four) as the comparison value ($M = 68\%$, $SE = 6.9\%$, $t(18) = 2.69$, $p < .05$).

There were, however, differences among the six types of 2/3 trials, Cochran's $Q(5) = 32.6$, $p < .001$; therefore, we analyzed the data for each type of trial separately. Children chose Block A significantly more often than chance in the observation trials matched to the nonconflicting intervention trial (84%) and the ambiguous intervention trial (95%), binomial tests, $p < .01$. They responded above chance in the observation trial matched to the conflicting intervention trial (63%), but not significantly so. The results of the 3/3 trials and the 2/3 observation trials suggest that children can use frequency as a measure of probabilistic causal strength.

Differences emerged when children were allowed to make their own interventions, particularly when the intervention information conflicted with the frequency information. In the nonconflicting intervention trial, children chose Block A significantly more often than chance (95%; binomial test, $p < .01$). There were no differences between this trial and the matched observation trial (McNemar's test, n.s.). However, in the ambiguous intervention trial, children chose Block A slightly more often than chance, but not significantly so (63%; binomial test, n.s.). This choice was significantly different from their choice in the matched observation trial (McNemar's test, $p < .05$). In the conflicting intervention trial, children picked Block A significantly less often than chance (21%; binomial test, $p < .05$). This result was also significantly different from their response to the matched observation trial (McNemar's test, $p < .001$).

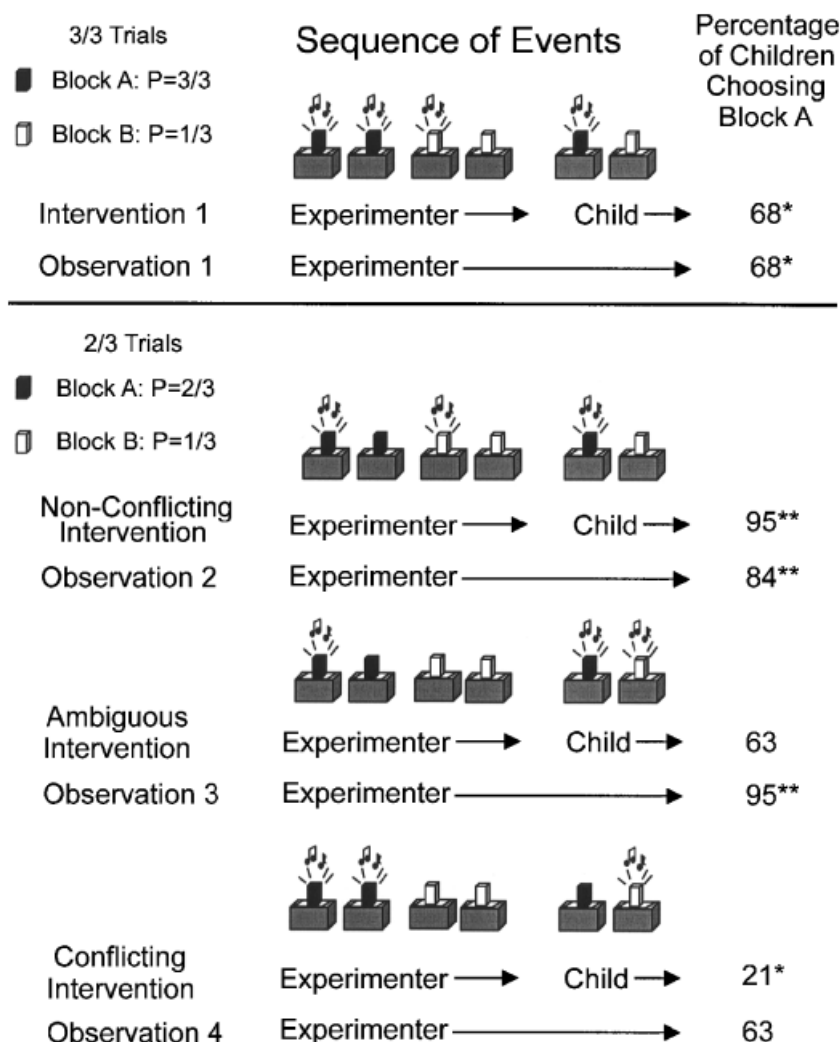


Fig. 1. Summary of the types of trials (3/3 trials, 2/3 trials, and matched observation trials) presented in Experiment 1 and the percentage of children choosing Block A on each trial. Trials are labeled by the number of times Block A activated the detector (three out of three or two out of three times). Block B always activated the detector one out of three times. On each trial, the experimenter placed the blocks on the detector the first two times, and either the child (intervention trials) or the experimenter (observation trials) placed the blocks on the detector the third time. Asterisks indicate results that differed significantly from chance responding, * $p < .05$, ** $p < .01$.

Children seemed to weigh the effects of their own intervention more heavily than the effects of another person’s intervention.

EXPERIMENT 2

Experiment 1 suggests that children are able to use probabilistic data to judge causal strength, but override these judgments when they conflict with the outcome of the children’s own interventions. There are three possible explanations for this finding. Children may weigh their own interventions more heavily than other people’s because they think their own interventions are less likely to be confounded. In fact, when the experimenter failed to make the detector activate, the children often spontaneously suggested explanations such as “you’re not pushing

hard enough” or “[it’s] on the wrong side of the toy.” When the children were allowed to intervene, they could push as hard as they wanted or put the object wherever they wanted. Thus, they could convince themselves that the intervention was free of confounding causes.

However, there are two less interesting explanations for the results on the conflicting intervention trial. First, the results could be an artifact of our test question, which required children to perform another intervention. They could simply have responded by repeating their previous successful intervention, ignoring other evidence. Second, children’s own actions may be more salient to them than the actions of other people.

In Experiment 2, we controlled for these possibilities in two ways. To rule out the first explanation, we changed the

dependent measure so that children did not activate the detector itself. Instead, they had to make an explicit judgment about the underlying causal mechanism of the detector. To rule out the second, we included a trial in which the action and outcome were identical to the action and outcome on the conflicting trial, but the children were given evidence that their own intervention on Block B was confounded. If the responses in Experiment 1 were the effect of the salience of the children's own actions, then the children would again be expected to respond by picking Block B. However, if they preferred their own interventions in Experiment 1 because they believed those interventions were not confounded, they would be expected to pick Block A. We also included a group of 6-year-old children to explore developmental differences.

Method

Participants

Participants in the study were eighteen 4-year-olds ($M = 4$ years 6 months, range = 3 years 11 months to 5 years 3 months) and eighteen 6-year-olds ($M = 6$ years 2 months, range = 5 years 7 months to 6 years 10 months). Four-year-olds were recruited from a preschool in Portland, Oregon, and 6-year-olds from a kindergarten in Lafayette, California.

Materials

The detector and blocks were the same as in Experiment 1. There was also an additional switch that could be plugged into the detector.

Procedure

The experimenter introduced the detector much as in Experiment 1, adding: "Things that make it go have special stuff inside. The special stuff makes it go." Before bringing out the blocks, the experimenter brought out the additional switch and plugged it into the detector. She told the child, "The switch makes the toy go when you flip it" and allowed the child to intervene on the switch and watch it activate the detector. She then unplugged the switch and set it aside.

Children were first given a warm-up with one deterministic trial and one 3/3 trial. They then saw a sequence of four 2/3 trials (counterbalanced): nonconflicting intervention, matched observation, conflicting intervention, and matched observation. The procedure is summarized in Figure 2. After each trial, the child was asked which object had "more special stuff inside." The child's choice was coded 1 (Block A) or 0 (Block B).

After the four 2/3 trials, the experimenter brought out the switch and plugged it in. She then placed two blocks on the table and proceeded with exactly the same sequence of events as in the conflicting intervention condition. The only difference was that, during the child's intervention, the experimenter flipped the switch at the exact time the child placed Block B on the

detector (see Fig. 2). The experimenter then asked the child the same test question and recorded his or her choice.

Results and Discussion

Experiment 2 replicated the strength judgments of Experiment 1 with a different and more explicit measure of causal strength. The results are given in Figure 2. There were no significant differences between the 4- and 6-year-olds on any measures, so the analysis was conducted on the combined data set. (A separate analysis of the 4-year-olds' data showed the same pattern of results as in the first study, though performance on both observation trials was significantly above chance.) A Cochran's Q test revealed significant differences among the trials, $Q(4) = 58.93$, $p < .001$. Children said that Block A had "more special stuff inside" significantly more often than chance in the observation trials (97% and 81%; binomial tests, $p < .001$). They also said this in the nonconflicting intervention trial (100%; binomial test, $p < .001$).

However, in the conflicting intervention trial, children chose Block A only 33% of the time, which was marginally below chance (binomial test, $p = .067$). This choice was also significantly different from their choice in the matched observation trial (McNemar's test, $p < .001$). This result replicated the intervention effect in Experiment 1.

Crucially, the children did not show the intervention effect in the confounded intervention trial. They chose the high-frequency object 69% of the time, which was significantly above chance (binomial test, $p < .05$) and significantly more often than in the conflicting intervention trial (McNemar's test, $p < .01$).

The results of Experiment 2 parallel those of Experiment 1; in general, children said that the object that was effective more often had more special stuff inside, but they favored the results of their own intervention. However, when children had explicit knowledge that their intervention was confounded, they did not show this preference.

GENERAL DISCUSSION

The results of this study suggest that even young children make judgments of causal strength based on covariation. They use frequency of co-occurrence to decide on the best intervention and to infer the strength of a hidden causal mechanism.

Children's own interventions also affect their judgments. When frequency information conflicted with evidence from their own interventions, children gave more weight to the intervention information. This does not seem to be simply an effect of action, because children used evidence from their own interventions to infer hidden mechanisms, and they did not give preference to their own interventions when those interventions were explicitly confounded.

These results suggest that children favor the results of their own interventions because they implicitly believe that their own

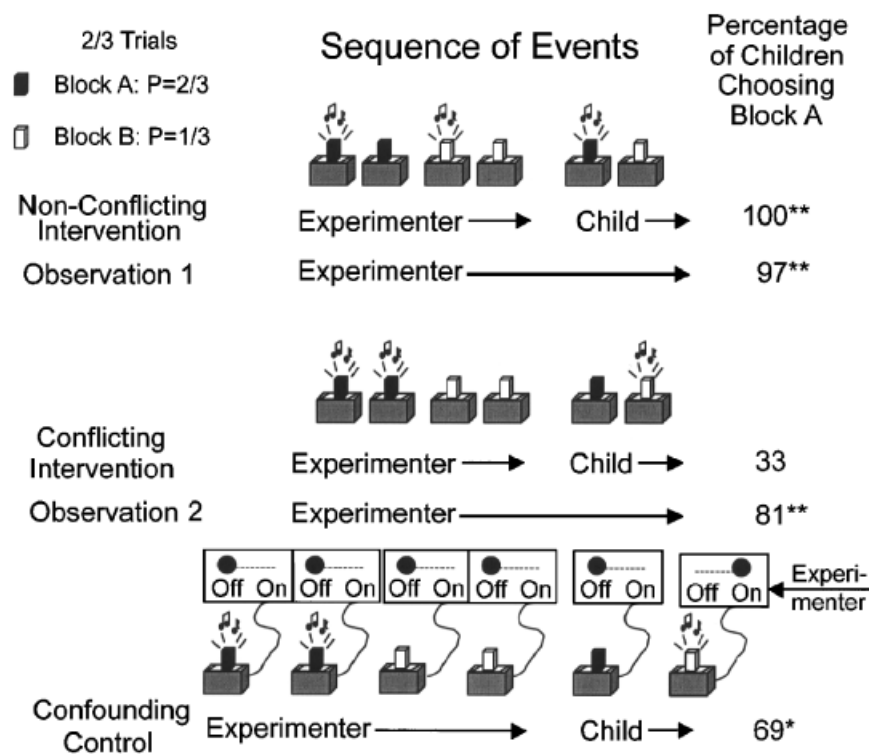


Fig. 2. Summary of the types of trials (2/3 intervention and observation trials and confounding control trial) presented in Experiment 2 and the percentage of children choosing Block A on each trial. Block A always activated the detector two out of three times, and Block B activated the detector one out of three times. On each trial, the experimenter placed the blocks on the detector the first two times, and either the child (intervention and confounding control trials) or the experimenter (observation trials) placed the blocks on the detector the third time. In the confounding control trial, an additional switch was plugged in to the detector but in the “off” position; when the child put Block B on the detector, the experimenter flipped the switch “on.” Asterisks indicate results that differed significantly from chance responding, * $p < .05$, ** $p < .01$.

interventions are less likely to be confounded than the interventions of other people. In Experiment 2, when children’s own intervention was in fact confounded by the experimenter’s parallel intervention, they no longer favored their own action. More extensive studies would be needed to confirm this hypothesis, though.

The current results do, however, clearly demonstrate that young children can use probabilistic information to infer causal strength in much the way that adults do. Moreover, children seem to differentiate their own interventions from the interventions of other people and give special weight to their own interventions. They are also able to recognize when their own interventions are confounded and do not use information from confounded interventions to infer causal strength. These causal learning abilities may be responsible, at least in part, for young children’s impressive causal knowledge.

Acknowledgments—This research was supported by National Institute of Mental Health Predoctoral Fellowship F31MH66538-01A1 to T.K. and by National Science Foundation Grant DLS0132487 to A.G. We thank Kallie Tsoukas for assistance

with data collection and Laura Schulz, Lori Markson, Carla Hudson, and Eleanor Rosch for comments.

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.

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

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

Gopnik, A., & Sobel, D.M. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, *71*, 1205–1222.

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

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

- Lagnado, D.A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 856–876.
- Pearl, J. (2000). *Causality*. New York: Oxford University Press.
- Schulz, L.E. (2001, December). *Spinning wheels and bossy ones: Children, causal structure and the calculus of intervention*. Paper presented at the Causal Inference in Humans and Machines Workshop at the Neural Information Processing Systems Conference, Vancouver, British Columbia, Canada.
- Shanks, D.R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 21, pp. 229–261). San Diego, CA: Academic Press.
- Shultz, T.R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development*, *47*(1, Serial No. 194).
- Shultz, T.R., & Mendelson, 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 inferences. *Child Development*, *47*, 1058–1063.
- Sobel, D.M., & Kushnir, T. (in press). The importance of decision-making to causal learning from interventions. *Memory & Cognition*.
- 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. (2001). *Causation, prediction, and search* (2nd ed., rev.). Cambridge, MA: MIT Press.
- Steyvers, M., Tenenbaum, J.B., Wagenmakers, E.J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453–489.
- Waldmann, M.R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition*, *82*, 27–58.
- Woodward, J. (2004). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.

(RECEIVED 9/17/04; REVISION ACCEPTED 1/12/05;
FINAL MATERIALS RECEIVED 1/24/05)

This document is a scanned copy of a printed document. No warranty is given about the accuracy of the copy. Users should refer to the original published version of the material.