Developmental Psychology
2009, Vol. 45, No. 2, 597–603

0012-1649/09/$12.00 DOI: 10.1037/a0014727

A Self-Agency Bias in Preschoolers' Causal Inferences

Tamar Kushnir
Cornell University

Henry M. Wellman and Susan A. Gelman
University of Michigan

Preschoolers' causal learning from intentional actions—causal interventions—is subject to a *self-agency bias*. The authors propose that this bias is evidence-based, in other words, that it is responsive to causal uncertainty. In the current studies, two causes (one child controlled, one experimenter controlled) were associated with one or two effects, first independently, then simultaneously. When initial independent effects were probabilistic, and thus subsequent simultaneous actions were causally ambiguous, children showed a self-agency bias. Children showed no bias when initial effects were deterministic. Further controls established that children's self-agency bias is not a wholesale preference but rather is influenced by uncertainty in causal evidence. These results demonstrate that children's own experience of action influences their causal learning, and the findings suggest possible benefits in uncertain and ambiguous everyday learning contexts.

Keywords: causal learning, causal inference, cognitive development, intentional action, probabilistic reasoning

Consider the following examples. First, imagine you are home on the night of a thunderstorm. Just as you reach to switch off a light, the power goes out through the whole house. For an instant—before your knowledge of light switches, and of thunderstorms, enables you to make sense of the event—it feels as though your action had a surprising and unintended effect. Now imagine that you are watching a friend struggle with a key in a lock, unable to open the door. Unconvinced that the door is truly stuck, you impatiently ask for the key so that you can try it yourself. After engaging in precisely the same set of actions on the key, to the same (null) effect, you feel confident that it is time to call a locksmith.

These examples illustrate how our experience of action often results in an illusory feeling of causal efficacy. Specifically, they point to two well-studied parameters that influence our sense of agency: *contingency* (the thunderstorm example; Alloy & Abramson, 1979; Shanks & Dickinson, 1987; Thompson et al., 2004) and *perceived control* (the stuck door example; Jenkins & Ward, 1965; Langer, 1975). In both examples, when confronted with unexpected causal outcomes, we are likely to feel quite differently about why they happened if they were caused by our own actions.

The importance of our own agency can be understood within the context of intentional actions more generally. Intentionally manipulating events to produce outcomes independently of surrounding events (i.e., holding all else constant) permits stronger causal inferences than does observing covariations, as such actions can deconfound factors that ordinarily covary. There is evidence that both adults and young children use the outcomes of their own and others' actions to disambiguate correlational evidence (i.e., to infer causal directionality and control for confounding; Gopnik et al., 2004; Lagnado & Sloman, 2004; Schulz, Gopnik, & Glymour, 2007; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann & Hagmayer, 2005).

Thus, intentional actions, particularly those that have observable effects on events and objects in the world, are special types of causal events. However, intentional actions are not controlled experiments and may not always lead to predictable outcomes. In particular, actions are motivated by internal mental states and can be influenced by an actor's knowledge, skills, and abilities. While the underlying causes of others' actions are not transparent to us, we are usually convinced that we know our own internal motivations and knowledge states well and thus feel a sense of control over the outcomes of our actions (Haggard, 2005; Langer, 1975; Wegner, 2002). Therefore, our own actions may be viewed (accurately or inaccurately) as controlled, certain, unconfounded, and reliable. To be clear, we are not suggesting that one's own actions are, in fact, better than those of others—they may be even less informative, as it may be difficult to serve as an objective observer of one's own actions. Nonetheless, we propose that they may seem better for the reasons outlined above.

Our aim in this study was to explore the possibility that a self-agency bias might be particularly influential in causal learning. For this reason, we focused on 3- and 4-year-old children. Young children have no explicit training in causal inference and relatively little prior knowledge about the causal structure of the world. Yet, the preschool years are characterized by an intense interest in causal learning and causal explanation and the development of increasingly sophisticated causal beliefs (Hickling & Wellman, 2001; Gelman & Wellman, 1991; Gopnik & Meltzoff, 1997; Shultz, 1982). It is therefore important to investigate whether children at this age display a self-agency bias in their judgments of causality and, if so, under what conditions.

Tamar Kushnir, Department of Human Development, Cornell University; Henry M. Wellman and Susan A. Gelman, Department of Psychology, University of Michigan.

This research was supported by a fellowship from the McDonnell Collaborative Initiative on Causal Learning and a National Institute of Child Health and Human Development postdoctoral fellowship to Tamar Kushnir.

Correspondence concerning this article should be addressed to Tamar Kushnir, Department of Human Development, Cornell University, Ithaca, NY 14853. E-mail: tk397@cornell.edu

One possibility is that children overvalue their own actions across the board, regardless of other evidence. Such a wholesale bias is certainly plausible for several reasons. First, developmental theorists since Piaget (1954) have speculated that children's early experience of their own agency is fundamental to their ability to understand causation. Indeed, research shows that young infants pay close attention to their own actions and contingent outcomes in both social and physical domains (Meltzoff & Moore, 1977; Watson & Ramey, 1972) and that their action experience influences their causal understanding of others' actions (Sommerville & Woodward, 2005). Moreover, preschoolers have been shown to overattribute success to their own actions that coincide with successful outcomes (Astington, 2001; Shultz & Wells, 1985) and even to remember (incorrectly) being causally responsible for the actions of others (Foley & Ratner, 1998; Sommerville & Hammond, 2007). Thus, it is conceivable that children may be particularly susceptible to the illusion that contingencies between their own actions and outcomes are causal.

A second possibility, however, is suggested by recent studies showing preschoolers' ability to make normative causal inferences based on patterns of evidence. In addition to understanding that actions can deconfound evidence, preschoolers make judgments of causal strength from observing probabilistic covariation (Kushnir & Gopnik, 2005, 2007) and conditional covariation (Gopnik, Sobel, Schulz, & Glymour, 2001). Critically, preschool children, like adults, evaluate new evidence in light of their existing knowledge; the strength of new evidence determines whether they will use it to override their prior beliefs (Kushnir & Gopnik, 2007; Schulz, Bonawitz, & Griffiths, 2007; Schulz & Gopnik, 2004). Children's reliance on self-agency may likewise be bounded in important ways by the evidence they observe. If new evidence provides strong support for causal relations and/or does not violate prior expectations, children may not be biased toward self-agency. However, when evidence is uncertain, weak, or ambiguous, children may show a self-agency bias, that is, they may rely on their own actions over the actions of other people, even when inappropriate. We aim to provide evidence for this second possibility.

Kushnir and Gopnik (2005) provided an initial demonstration that young children's causal inferences could be disproportionately affected by their own (vs. others') actions. Children were exposed to sequences of events in which one cause was overall weakly associated with an effect (33% of the time) and another was strongly associated with the effect (66% of the time). Children inferred that a weak cause that activated solely in response to their own actions was more powerful than a strong cause that activated in response to the actions of the experimenter. However, the results of that study could be explained by either of the two possibilities above. Children could have based their judgments on a wholesale preference for their own agency: choosing the weak object because it always worked for them and rejecting the alternative object because it always failed to work for them. Alternatively, their bias could have been evidence-based: When causal relations are probabilistic, children rely on evidence from their own actions more.

Put another way, in Kushnir and Gopnik's (2005) study, the effects were probabilistic overall, but with regard to the agent, they were deterministic in the sense that each cause worked (or failed to work) deterministically for each agent (experimenter, child). To find out whether preschoolers' reliance on self-agency is wholesale or evidence-based, we designed a novel method to meet two critical conditions: (a) that each cause behave either probabilistically or deterministically for each agent and (b) that we could vary the degree of uncertainty in the evidence (i.e., whether the actions are probabilistically or deterministically effective) while at the same time holding constant the associations between children's actions and their effects.

Thus, in our new method, there were two potential causes (one controlled by the child and one controlled by the experimenter) and two potential effects (the target effect and a secondary effect). The task also contained two parts. In the first part of the task, actions on two candidate causes of an effect occurred independently of each other and were either probabilistically or deterministically effective (they caused each effect some of the time, or only one effect all of the time). In the second part of the task, actions on the same two candidate causes occurred simultaneously and were always associated with both effects. Thus, when actions occurred independently, the strength of the causal relationship between each action and outcome could be estimated by their relative frequency of co-occurrence. However, when the actions occurred simultaneously, their causal efficacy was unknown; they could be spuriously associated with an outcome or causally related to it but without further evidence, it would be impossible to tell which.

The combination of independent and simultaneous actions allowed us to distinguish an evidence-based self-agency bias from a wholesale one. If our hypothesis is correct, and children's self-agency bias is bounded by the evidence they receive, then it should influence their causal judgments about simultaneous actions that follow probabilistic effects (and thus are causally ambiguous) but not those that follow deterministic effects (and thus are clearly spurious). The critical test of our hypothesis was thus to contrast these two scenarios. Experiment 1 also included two additional control conditions, both designed to ensure that, in general, children make normative inferences about probabilistic evidence. Under an evidence-based hypothesis, probabilistic evidence may lead children to prefer their own actions more, not to disregard frequency information entirely. Furthermore, we expected that the self-agency bias is not simply due to a tendency to appeal to more salient causes (such as one's own actions) in ambiguous situations. Experiment 2 controlled for this possibility.

Experiment 1

Method

Participants

Eighty 3- and 4-year-olds (40 female; $M = 4$ years, 0 months; $SD = 6.8$ months) from a Midwestern university town participated. The sample was predominantly middle- to upper middle-class and White, reflecting the composition of this community at large. Twenty children were randomly assigned to each condition, with the constraint that the ages be roughly equal across conditions.

Materials

The novel device was a white plastic cylinder (height = 4.25 in. [10.8 cm]; diameter = 2.5 in. [6.4 cm]). The bottom half was opaque, and the top half was clear. Two identical buttons were attached to wires, which were plugged into the device, and a control box was hidden under the table. The control box contained two sliding switches that determined the effects of the buttons. One sliding switch enabled/disabled a 3-s bell sound, and one enabled/disabled a light on the clear top of the device to glow red for 3 s. Thus, the experimenter could determine whether the button presses caused a sound, a light, or both effects together (synchronously for 3 s).

Procedure

Each child was individually seated opposite the experimenter, with the device in between them on a table and one button on either side. The experimenter began with, "This is my special toy. It has two buttons. Let's figure out what the buttons do." In order to keep the simultaneous button presses synchronous, all button presses (in all conditions) were prompted by the experimenter saying, "One, two, three, go!"

The four conditions are described below, and depicted in Figure 1. In each condition, Button A refers to the button that independently produced Effect A more often. Button B refers to the button that independently produced Effect B more often. The button that the child intervened on (in all conditions except the observed effects condition) was Button A. Effect B was designated to be the target effect.

The inclusion of an alternate effect ensured that all actions were associated with a positive outcome of one kind or another. This method—different from previous studies in which effects were either present or absent—further ensured that any sense children had that they had caused the target effect did not arise from the simple expectation that all actions they perform must be efficacious.

Focal condition. The focal condition was hypothesized to elicit the self-agency bias. The initial independent effects of the

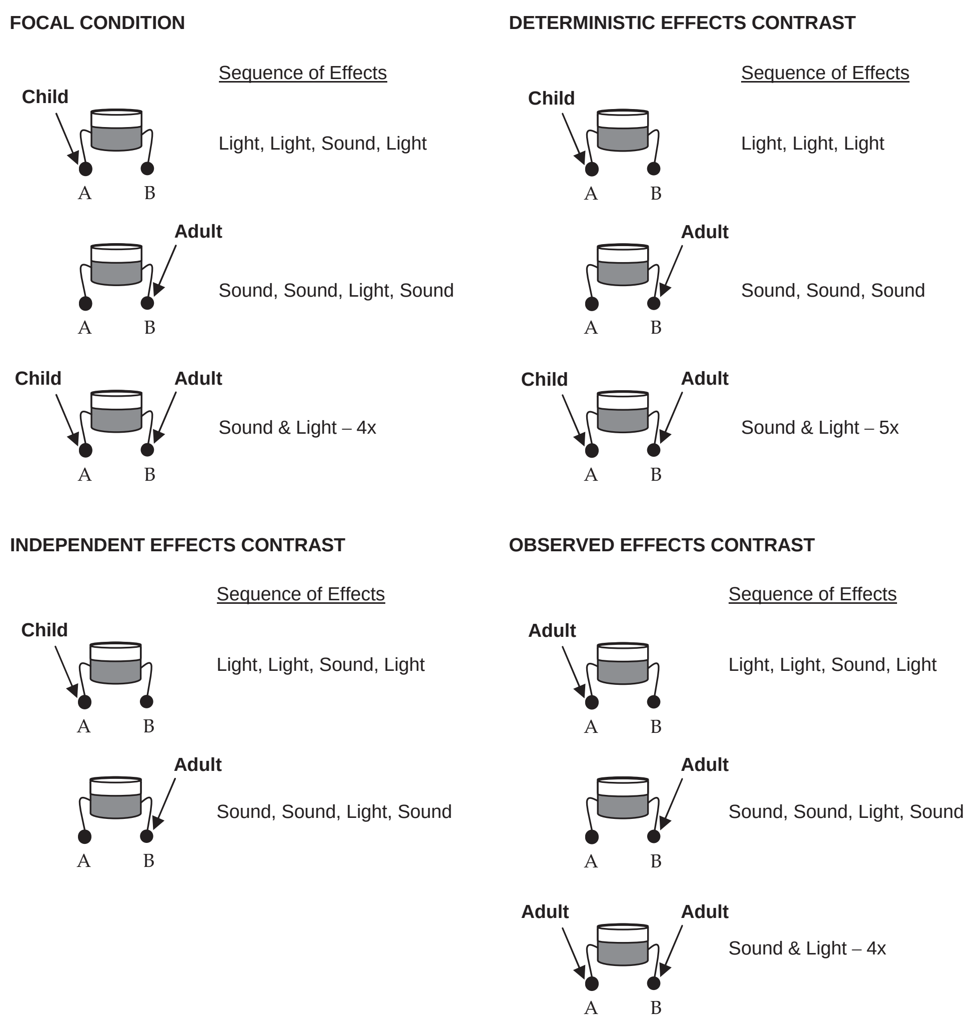

Figure 1. An illustration of a sequence of events in each of the four conditions of Experiment 1. Interventions (button presses) are indicated by an arrow and are labeled (child, adult). In each example shown, the child intervenes first, the child's button (Button A) is the one on the left, and the target effect is the sound. In the actual procedure, the first intervener (child, adult), the side of the child's button (left, right), and the target effect (sound, light) were counterbalanced.

buttons were probabilistic. The child's button was weakly associated with the target effect and strongly associated with the alternate effect. The experimenter's button was strongly associated with the target effect and weakly associated with the alternate effect. Simultaneous actions were then associated with both effects. Therefore, the child pressed Button A four times. The sequence of effects was AABA (e.g., light, light, sound, light). The experimenter pressed Button B four times. The sequence of effects was BBAB (e.g., sound, sound, light, sound). The child and the experimenter then pressed their buttons simultaneously four times, each time producing the combined effect (sound and light).

The most straightforward interpretation of these events assumes that the independent actions are representative of the general behavior of the buttons: that Button A causes Effect A about 75% of the time and that Button B causes Effect B about 75% of the time. However, the probabilistic nature of the effect makes the simultaneous actions causally ambiguous. Thus, we proposed that children would be biased to treat their actions on Button A as causing Effect B during some (or all) of the simultaneous actions.

Contrast 1: Deterministic effects. The first contrast condition provided the critical contrast to the focal condition above. In this condition, all independent effects of the buttons were deterministic. Thus, all associations between the child's action and the target effect were spurious (occurred in the presence of a known, deterministic cause—the experimenter's action on the other button), although the number of associations was exactly the same as in the focal condition. Therefore, the child pressed Button A three times. The sequence of effects was AAA. The experimenter then pressed Button B three times. The sequence of effects was BBB. Next, the child and the experimenter pressed their buttons simultaneously five times, each time producing the combined A and B effect.

Under an evidence-based hypothesis, children should weigh the effects of their own actions in light of all of the supporting evidence, including the actions of the experimenter. Thus, the simultaneous actions should have a clear interpretation due to the deterministic independent evidence: The child's action on Button A is causing Effect A, whereas the experimenter's simultaneous action on Button B is causing Effect B.

Contrast 2: Independent effects. In the second contrast condition, the independent interventions on the buttons were identical to those in the target condition, but there were no simultaneous interventions. Thus, the child pressed Button A four times (resulting in the sequence AABA); then the experimenter pressed Button B four times (resulting in the sequence BBAB), and nothing else was done or shown.

Again, an evidence-based bias predicts that children's inference in the focal Condition is based in part on their interpretation of the simultaneous evidence. Thus, weak independent evidence of children's own efficacy (1 out of 4 successes) should not be sufficient to override strong evidence in favor of the experimenter's actions (3 out of 4 successes), and children should choose Button B as the cause of Effect B.

Contrast 3: Observed effects. In the third contrast condition, the pattern of independent and dependent probabilities was identical to the pattern in the focal condition, but all of the actions were generated by another person (no self-agency). The experimenter pressed Button A four times (resulting in the sequence AABA), then pressed Button B four times (resulting in the sequence BBAB), and finally pressed both buttons simultaneously four times (resulting in the combined effect four times).

Without self-agency to interfere, children should reason about the probabilistic evidence normatively. Thus, they should infer that the independent evidence is representative of the true causal strength of each button—that Button A is the stronger cause of Effect A and that Button B is the stronger cause of Effect B.

At the end, each child was asked which button produced Effect B (e.g., "Which button makes the sound?"), and his or her first choice was recorded. In all conditions, starting side (right or left) and Button A's effect (sound or light) were counterbalanced. In the first three conditions in which the child was allowed to intervene, the starting intervener (child or experimenter) was counterbalanced. In the observed effects contrast, the effect mentioned in the question ("Which button makes the [sound/light]?") was counterbalanced.

Results and Discussion

The results provided support for an evidence-based self-agency bias that affects children's causal inferences. The proportion of children choosing Button B as the cause of Effect B in each condition is shown in Table 1. In the focal condition, only 6 out of 20 children said that Button B—the experimenter's button, and the button that was more often independently associated with Effect B—was the cause of Effect B. Instead, 14 out of 20 children claimed that their own button (Button A) caused Effect B. In contrast, in the deterministic effects contrast, a significant majority of children (17 out of 20) said that Button B caused Effect B (binomial test, $p < .01$). Thus, as hypothesized, children were significantly more likely to choose Button A (their own button) as the cause of Effect B when the independent effects of the buttons

Table 1
The Proportion of Children in Experiment 1 Choosing Button B in Response to the Question "Which Button Makes [Effect B]?"

Condition	Agency	Probabilistic effects	Simultaneous interventions	Proportion of children choosing Button B
Focal condition	Yes	Yes	Yes	6/20 (30%)
Deterministic effects	Yes	No	Yes	17/20 (85%)**
Observed effects	No	Yes	Yes	18/20 (90%)***
Independent effects	Yes	Yes	No	14/20 (70%)*

Note. The significance tests represented Fisher's exact test comparisons with focal condition.
* $p < .05$. ** $p < .01$. *** $p < .001$.

were probabilistic (70% of children) rather than deterministic (15% of children) (Fisher's exact test, $p < .01$). Thus, the children did not base their judgments of self-agency on associative evidence alone.

Moreover, this difference was not exclusively due to children's producing the target outcome 1 out of 4 times independently of the experimenter's actions. In the independent effects contrast, 14 out of 20 children (70%) chose Button B as the cause, significantly more than in the focal condition (Fisher's exact test, $p < .05$). This contrast shows that children considered the strength of the probabilistic evidence against their own agency; they did not disregard frequency information on the basis of being able to produce Effect B once.

Finally, a significant majority of children (18 out of 20) who merely observed the experimenter perform the same pattern of actions as in the focal condition chose Button B as the cause of Effect B (binomial test, $p < .01$). This was also significantly different from the results in the focal condition (Fisher's exact test, $p < .001$). Thus, children can, in general, make accurate causal inferences on the basis of combinations of independent probabilistic actions and simultaneous actions.

Overall, the contrasts show that children's self-agency bias affected their interpretation of the simultaneous actions in the focal condition. Strikingly, this occurred even though children knew that their button had another strong effect, Effect A. This suggests that children's bias toward self-agency is quite strong, yet evidence-based, occurring only under conditions of uncertainty.

An alternative interpretation of the findings, however, is that the salience of children's own actions makes them easier to attend to or encode and that children gravitate toward causes that are more salient in ambiguous situations. The contrast conditions provide some evidence against this interpretation. Children's accurate responses in the deterministic effects contrast show that they do not disregard another's actions simply because their own actions may be more salient. The observed effects contrast shows that children can make accurate causal inferences on the basis of probabilistic evidence when salience is equated (same agent presses both buttons). In Experiment 2, we explored this possibility further by asking whether children can make accurate inferences about the stronger cause when the weaker cause has a more salient agent, specifically, a puppet that the children find more interesting and prefer to play with.

Experiment 2

Method

Participants

Fourteen 3- and 4-year-olds (5 female; $M = 4$ years, 2 months; $SD = 9$ months) from a Midwestern university town participated. The sample was predominantly middle to upper middle-class and White.

Procedure

The procedure was identical to that of Experiment 1, except that the experimenter introduced two puppets to the child and said that the puppets were going to press the buttons. One puppet (high-salience agent) was a lifelike monkey (introduced as "Monkey"), and the other (low-salience agent) was made of two pieces of sewn-together brown felt in a gingerbread man shape (but without facial features, introduced as "Puppet").

Figure 2 shows the sequence of events. "Puppet's" button was weakly associated with the target effect and strongly associated with the alternate effect, exactly the case for the child's button in the focal condition of Experiment 1. "Monkey's" button was strongly associated with the target effect, exactly the case for the experimenter's button in the focal condition. The nature of the target effect (sound or light), the side of the toy of each puppet, and the starting puppet were all counterbalanced.

Children were then asked two questions. The first was the same question from Experiment 1 ("Which button made [Effect B]?"). The experimenter then asked "Did you like my puppets today? Which puppet would you like to play with?" This question vali-

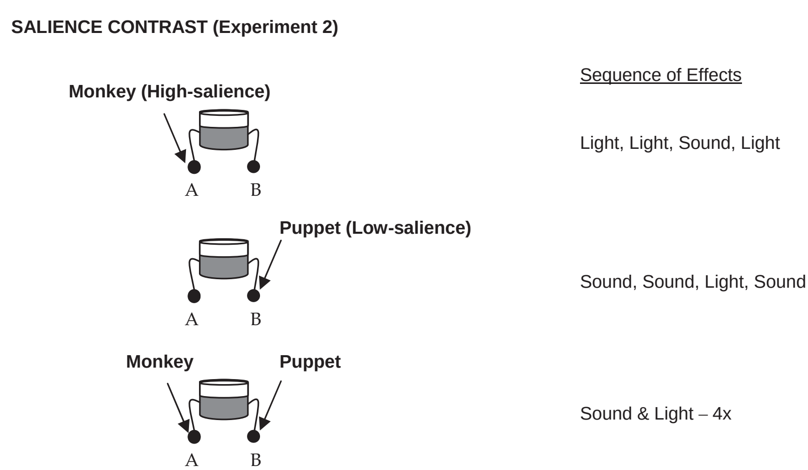

Figure 2. An illustration of sequence of events in Experiment 2. Interventions (button presses) are indicated by an arrow and are labeled (Monkey, Puppet). In this example shown, Monkey intervenes first, on Button A (the left button), and the target effect is the sound. In the actual procedure, the first intervener (Monkey or Puppet), the side of each button (left, right), and the target effect (sound, light) were counterbalanced.

dated that indeed "Monkey" was the preferred (more salient) causal agent: 11 of the 14 children preferred to play with the monkey puppet (binomial test, $p < .05$, one-tailed).

Results and Discussion

The results demonstrate that children do not simply choose the most salient cause—in this case, a preferred agent—in ambiguous situations. A significant majority of children (11/14) correctly stated that Button B was the cause of Effect B (binomial test, $p < .05$, one-tailed), although it was operated by the less salient agent. Moreover, this response pattern is significantly different than the responses in the focal condition of Experiment 1 (Fisher's exact test, $p < .05$). Importantly, children's causal responses were correct despite their play preferences: Whereas 8 children who answered the causal question correctly then said they preferred to play with the monkey puppet, no child who answered incorrectly (selecting the monkey's button) preferred to play with Puppet (McNemar's test, $p < .001$).

General Discussion

Taken together, these results support an evidence-based self-agency bias in preschool children, that is, a tendency to weigh their own actions as more effective than the actions of others in ambiguous situations. The results suggest that this bias is not based on a wholesale preference for one's own actions, an inability to evaluate probabilistic causal evidence, or the salience of certain actions produced by certain agents.

This finding has implications for everyday causal learning. Young children's inexperience makes them especially likely to encounter causal actions that are only probabilistically effective as a result of mitigating circumstances that they cannot perceive and/or do not fully understand. Moreover, young children's explorations of new or difficult situations are often facilitated by others (Rogoff, 1990). In collaborative contexts, children's actions and the actions of others can occur simultaneously, by virtue of being motivated by common intentions and goals. Thus, it may often be difficult for children to determine accurately which action causes a particular outcome. A self-agency bias that is evidence-based, and thus specific to these sorts of uncertain situations, may very well offer children a learning advantage, by helping them resolve confounding quickly, without generally compromising their ability to make accurate inferences.

Children learn much about the causal structure of the world from intentional actions precisely because such actions are often effective controls for confounding. However, preschoolers have accumulated enough social knowledge to potentially understand the pitfalls of relying on others' actions for causal inference. Perhaps children make use of their knowledge of intentional action to remain skeptical about the causal actions of others, in particular, when evidence is uncertain. They might implicitly assume that their own actions are more trustworthy—less likely to be confounded—than the actions of other people. It is not known for certain whether children (or adults for that matter) make such an assumption about the "trustworthiness" of their own actions. One way to explore this question would be to vary the expertise of the other actor. Children appropriately rely on evidence from knowledgeable sources in causal learning (Kushnir, Wellman, & Gelman, 2008). Varying another person's expertise should influence their trustworthiness as a source of evidence, and therefore, such information should interact with a self-agency bias in a systematic way.

The current investigation demonstrates one way in which children's own experience of action influences causal learning. It also contributes to the growing evidence that children are affected in important ways by the psychological and social context in which learning takes place.

References

Alloy, L., & Abramson, L. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology: General, 108,* 441–485.

Astington, J. W. (2001). The paradox of intention: Assessing children's metarepresentational understanding. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition* (pp. 85–103). Cambridge, MA: MIT Press.

Foley, M., & Ratner, H. (1998). Children's recoding memory for collaboration: A way of learning from others. *Cognitive Development, 13,* 91–108.

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

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

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

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

Haggard, P. (2005). Conscious intention and motor cognition. *Trends in Cognitive Sciences, 9,* 290–295.

Hickling, A., & Wellman, H. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology, 37,* 668–683.

Jenkins, H. H., & Ward, W. C. (1965). Judgement of contingency between responses and outcomes. *Psychological Monographs, 79.*

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

Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology, 44,* 186–196.

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition, 107,* 1084–1092.

Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 856–876.

Langer, E. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311–328.

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science, 198,* 75–78.

Piaget, J. (1954). *The construction of reality in the child.* New York: Basic Books.

Rogoff, B. (1990). *Apprenticeship in thinking: Cognitive development in social context.* New York: Oxford University Press.

Schulz, L. E., Bonawitz, E. B., & Griffiths, T. (2007). Can being scared make your tummy ache? Naive theories, ambiguous evidence and preschoolers/ causal inferences. *Developmental Psychology, 43,* 1124–1139.

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

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science, 10,* 322–332.

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* (Serial No. 194).

Shultz, T. R., & Wells, D. (1985). Judging the intentionality of action-outcomes. *Developmental Psychology, 21,* 83–89.

Sommerville, J. A., & Hammond, A. J. (2007). Treating another's actions as one's own: Children's memory of and learning from joint activity. *Developmental Psychology, 43,* 1003–1018.

Sommerville, J., & Woodward, A. (2005). Pulling out the intentional structure of action: The relation between action processing and action production in infancy. *Cognition, 95,* 1–30.

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

Thompson, S., Kyle, D., Osgood, A., Quist, R., Phillips, D., & McClure, M. (2004). Illusory control and motives for control: The role of connection and intentionality. *Motivation and Emotion, 28,* 315–330.

Waldmann, M., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 216–227.

Watson, J. S., & Ramey, C. T. (1972). Reactions to response-contingent stimulation in early infancy. In J. Oates & S. Sheldon (Eds.), *Cognitive development in infancy* (pp. 77–85). Hillsdale, NJ: Erlbaum.

Wegner, D. (2002). *The illusion of conscious will.* Cambridge, MA: MIT Press.

Received July 28, 2007
Revision received October 21, 2008
Accepted October 27, 2008 ■