Published in final edited form as:
Dev Psychol. 2009 March ; 45(2): 597–603. doi:10.1037/a0014727.

A self-agency bias in preschoolers' causal inferences

Tamar Kushnir,
Cornell University

Henry M. Wellman, and
University of Michigan

Susan A. Gelman
University of Michigan

Abstract

Preschoolers' causal learning from intentional actions – causal interventions – is subject to a *self-agency bias*. We propose that this bias is evidence-based; 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 establish 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 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 if 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 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. They point to two well-studied parameters which 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 observing covariations, as such actions can deconfound factors that ordinarily co-vary. There is evidence that both adults and young children use the outcomes of their own and others' actions to disambiguate correlational evidence (that is, 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, in particular 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 paper is to explore the possibility that a self-agency bias might be particularly influential in causal learning. For this reason, we focus 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 (Shultz, 1982; Hickling & Wellman, 2001; Gelman & Wellman, 1991; Gopnik & Melzoff, 1997). 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.

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 of all, 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 over-attribute success to their own actions based on observing successful outcomes (Shultz & Wells, 1985; Astington, 2001) and even to remember (incorrectly) being causally responsible for the actions of others (Foley & Ratner, 1998; Sommerville & Hammond, in press). 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 based on 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, in press; 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 – 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 a 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 which activated solely in response to their own actions was more powerful than a strong cause which 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. Or, 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 & 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 distinguish between the two types of self-agency bias, we designed a novel method to meet two critical conditions. First, that each cause behave either probabilistically or deterministically *for each agent*. Second, 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. So, 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 is unknown – they may 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 allows 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 is 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 expect 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 controls for this possibility.

Experiment 1

Methods

Participants—Eighty 3- and 4-year-olds (40 Female; Mean = 4;0, SD = 6.8 months) from a Midwestern university town participated. The sample was predominantly middle- to upper-middle-class and white, reflecting the make up 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″, diameter = 2.5″). 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 3s bell sound and one enabled/disabled a light on the clear top of the device to glow red for 3s. Thus, the experimenter could determine whether the button presses caused a sound, a light, or both effects together (in sync for 3s).

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 in sync, 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 caused the target effect did not arise from the simple expectation that all actions they perform must be efficacious.

Focal Condition: This condition was hypothesized to elicit the self-agency bias. The initial independent effects of the 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. So, 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). Then the child and experimenter pressed their buttons simultaneously four times, each time the 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 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 propose that children will 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: This condition provides 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), though the number of associations was exactly the same as in the Focal Condition. So, the child pressed button A three times. The sequence of effects was AAA. Then the experimenter pressed button B three times. The sequence of effects was BBB. Then the child and experimenter pressed their buttons simultaneously *five* times, each time the 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 this condition the independent interventions on the buttons were identical to 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 this condition the pattern of independent and dependent probabilities was identical to those 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 button B four times (resulting in the sequence BBAB), then 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 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/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 provide evidence 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

NIH-PA Author Manuscript

NIH-PA Author Manuscript

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

Moreover, this difference was not exclusively due to children 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 based on 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 Focal Condition, Fisher's exact test, $p < .001$. Thus children can, in general, make accurate causal inferences based on 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 towards causes that are more salient in ambiguous situations. The contrast conditions provide some evidence that speaks against this interpretation. Children's accurate responses in the Deterministic Effects contrast show that they do not simply 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 based on probabilistic evidence when salience is equated (same agent presses both buttons). In Experiment 2, we explore this possibility further by asking whether children can make accurate inferences about the stronger cause when the *weaker* cause has a more salient agent – a puppet that they find more interesting and prefer to play with.

Experiment 2

Methods

Participants—Fourteen 3- and 4-year-olds (5 Females; Mean = 4;2, 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 Experiment 1, except that the experimenter introduced two puppets to the child and said they were going to press the buttons. One puppet (high-salience agent) was a life-like 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 as was the child's button in the Focal Condition of Experiment 1. “Monkey's” button was strongly associated with the target effect, exactly as was 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 validated that indeed “Monkey” was the preferred (more salient) causal agent: 11/14 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) though 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 Monkey, no child who answered incorrectly (picking 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 – a tendency to weigh their own actions as more effective than the actions of others in ambiguous situations. The results suggest that this bias neither based on a wholesale preference for one's own actions, nor on an inability to evaluate probabilistic causal evidence, nor is it based on 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 due to 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 is may often be difficult for children to determine accurately which action causes a particular outcome. A self-agency bias which 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.

Acknowledgments

This research was supported by a fellowship from the McDonnell Collaborative Initiative on Causal Learning and an NICHD post-doctoral fellowship to T.K.

References

Astington, JW. The paradox of intention: assessing children's metarepresentational understanding. In: Malle, BF.; Moses, LJ.; Baldwin, DA., editors. Intentions and intentionality: Foundations of social cognition. Cambridge, MA: MIT Press; 2001. p. 85-103.

Alloy L, Abramson L. Judgment of contingency in depressed and nondepressed students: Sadder but wiser? Journal of Experimental Psychology: General. 1979; 108:441–485. [PubMed: 528910]

Hickling A, Wellman H. The emergence of children's causal explanations and theories: Evidence from everyday conversation. Developmental Psychology. 2001; 37:668–683. [PubMed: 11552762]

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

Gelman S, Wellman H. Insides and essence: Early understandings of the non-obvious. Cognition. 1991; 38:213–214. [PubMed: 2060270]

Gilbert DT, Malone PS. The correspondence bias. Psychological Bulletin. 1995; 117:21–38. [PubMed: 7870861]

Gopnik A, Glymour C, Sobel DM, Schulz LE, Kushnir T, Danks D. A theory of causal learning in children: Causal maps and Bayes nets. Psychological Review. 2004; 111:1–30.

Gopnik, A.; Meltzoff, A. Words, thoughts, and theories. Cambridge, MA, US: The MIT Press; 1997.

Gopnik A, Sobel DM, Schulz L, Glymour C. Causal learning mechanisms in very young children: Two, three, and four-year-olds infer causal relations from patterns of variation and covariation. Developmental Psychology. 2001; 37:620–629. [PubMed: 11552758]

Haggard P. Conscious intention and motor cognition. Trends in Cognitive Sciences. 2005; 9:290–295. [PubMed: 15925808]

Jenkins HH, Ward WC. Judgement of contingency between responses and outcomes. Psychological Monographs. 1965; 79

Koenig M, Harris P. Preschoolers mistrust ignorant and inaccurate speakers. Child Development. 2005; 76:1261–1277. [PubMed: 16274439]

Kushnir T, Gopnik A. Children infer causal strength from probabilities and interventions. Psychological Science. 2005; 16:678–683. [PubMed: 16137252]

Kushnir T, Gopnik A. Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. Developmental Psychology. 2007; 44:186–196. [PubMed: 17201518]

Kushnir T, Wellman HM, Gelman SA. The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. Cognition. 2008; 107:1084–1092. [PubMed: 18039543]

Lagnado DA, Sloman S. The advantage of timely intervention. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2004; 30:856–876.

Langer E. The illusion of control. Journal of Personality and Social Psychology. 1975; 32:311–328.

Meltzoff AN, Moore MK. Imitation of facial and manual gestures by human neonates. Science. 1977; 198:75–78. [PubMed: 17741897]

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

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

Ross, L. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In: Berkowitz, L., editor. Advances in experimental social psychology. Vol. 10. New York: Academic Press; 1977. p. 173-220.

Schulz LE, Bonawitz EB, Griffiths T. Can being scared make your tummyache? Naive theories, ambiguous evidence and preschoolers' causal inferences. Developmental Psychology. in press.

Schulz LE, Gopnik A. Causal learning across domains. Developmental Psychology. 2004; 40:162–176. [PubMed: 14979758]

Schulz LE, Gopnik A, Glymour C. Preschool children learn about causal structure from conditional interventions. Developmental Science. 2007; 10

Shanks, DR.; Dickinson, A. Associative accounts of causality judgment. In: Bower, GH., editor. The psychology of learning and motivation: Advances in research and theory. Vol. 21. San Diego, CA: Academic Press; 1987. p. 229-261.

Shultz TR. Rules of causal attribution. Monographs of the Society for Research in Child Development. 1982; 47(Serial No. 194)

Shultz T, Wells D. Judging the intentionality of action-outcomes. Developmental Psychology. 1985; 21:83–89.

Sommerville JA, Hammond AJ. Treating another's actions as one's own: Children's memory of and learning from joint activity. Developmental Psychology. in press.

Sommerville J, Woodward A. Pulling out the intentional structure of action: the relation between action processing and action production in infancy. Cognition. 2005; 95:1–30. [PubMed: 15629472]

Sobel DM, Tenenbaum JB, Gopnik A. Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. Cognitive Science. 2004; 28:303–333.

Steyvers M, Tenenbaum JB, Wagenmakers EJ, Blum B. Inferring causal networks from observations and interventions. Cognitive Science. 2003; 27:453–489.

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

Watson, JS.; Ramey, CT. Reactions to response-contingent stimulation in early infancy. In: Oates, J.; Sheldon, S., editors. Cognitive Development in Infancy. Hillsdale, NJ: Erlbaum; 1987. p. 77-85.

Wegner, D. The illusion of conscious will. Cambridge, MA: MIT Press; 2002.

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

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

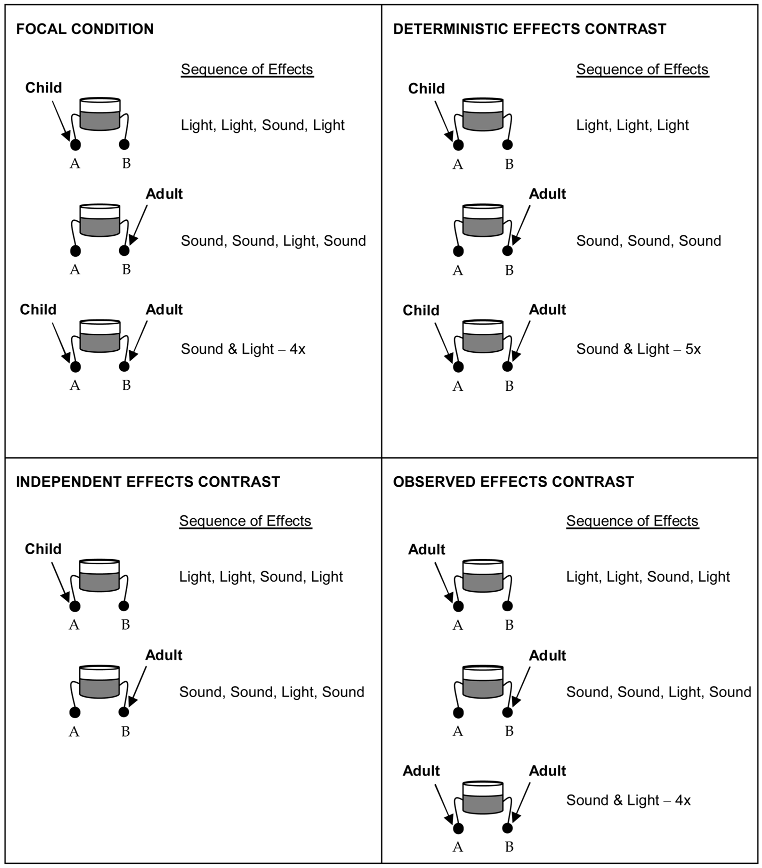

Figure 1.

An illustration of 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.

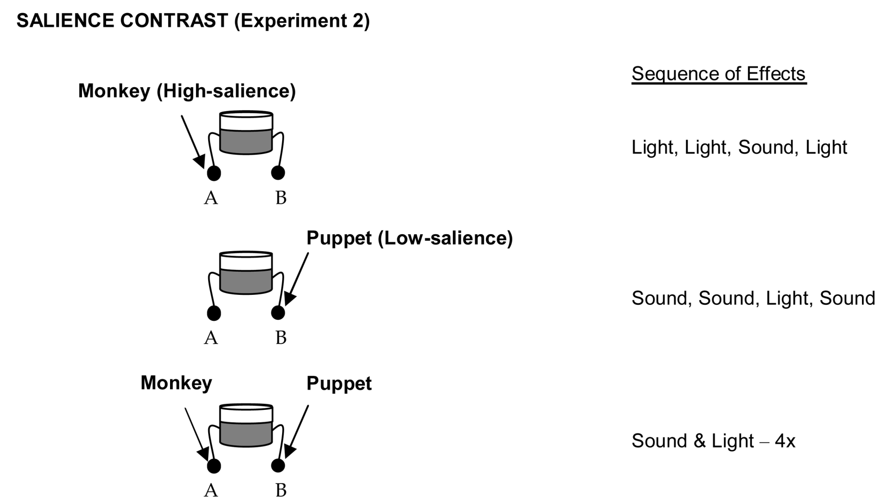

Figure 2.
Methods and results of Experiment 2.

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%) *

Fisher's exact test comparisons with Focal Condition:

* $p < .05$

** $p < .01$

*** $p < .001$