



ELSEVIER

Contents lists available at ScienceDirect

## Cognitive Development



Short report

## Rationales in children's causal learning from others' actions

David M. Sobel<sup>a,\*</sup>, Jessica A. Sommerville<sup>b</sup><sup>a</sup> Department of Cognitive and Linguistic Sciences, Box 1978, Brown University, Providence, RI 02912, United States<sup>b</sup> Department of Psychology and Institute for Learning and Brain Sciences, Campus Box 351525, University of Washington, Seattle, WA 98195, United States

## ARTICLE INFO

## Keywords:

Causal learning  
Explanations  
Rationales

## ABSTRACT

Shown commensurate actions and information by an adult, preschoolers' causal learning was influenced by the pedagogical context in which these actions occurred. Four-year-olds who were provided with a reason for an experimenter's action relevant to learning causal structure showed more accurate causal learning than children exposed to the same action and data accompanied by an inappropriate rationale or accompanied by no explanatory information. These results suggest that children's accurate causal learning is influenced by contextual factors that specify the instructional value of others' actions.

© 2008 Elsevier Inc. All rights reserved.

By the time children enter elementary school, they recognize fundamental aspects of causal structure in the physical world (Baillargeon, 2002; Spelke, Breinlinger, Macomber, & Jacobson, 1992). They understand that animate objects differ from inanimate objects in terms of various biological principles (Inagaki & Hatano, 1993; Kalish, 1996) and that an agent's behaviors are motivated by underlying mental states (Perner, 1991; Wellman, 1990). Young children also have sophisticated causal reasoning abilities. They can make predictions about future events (Bullock, Gelman, & Baillargeon, 1982), generate explanations of past events (Schult & Wellman, 1997; Wellman & Liu, 2007), and reason about counterfactuals (Harris, German, & Mills, 1996). Open questions concern how young children acquire their causal knowledge and the extent of their causal reasoning abilities.

One approach to these questions has been to investigate children's ability to recover causal structure from their observation of events, particularly events that convey conditional probability information. To use a simple example, two events may co-occur because of a causal relation between them or because of a hidden common cause (among other reasons). To decide between these possibilities, one can observe the conditional probability that these events co-occur given the presence or absence of

\* Corresponding author. Tel.: +1 401 863 3038; fax: +1 401 863 2255.  
E-mail address: Dave.Sobel@brown.edu (D.M. Sobel).

other events. If you observe that events A and B only co-occur in the presence of a third event C, it is more likely that the third event is the cause of both events A and B than is the existence of a direct causal relation between the two. Young children, and even infants, recognize this conditional probability information when they observe events (Fiser & Aslin, 2002; Gopnik, Sobel, Schulz, & Glymour, 2001; Kirkham, Slemmer, & Johnson, 2002; Saffran, Aslin, & Newport, 1996), and do so across different domains of knowledge (Schulz & Gopnik, 2004) in a manner that simple models of associative reasoning cannot easily explain (Sobel & Kirkham, 2006; Sobel, Tenenbaum, & Gopnik, 2004).

Research has shown that adults can recover sets of causal relations from conditional probability information (Lagnado & Sloman, 2004; Sobel & Kushnir, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003; Waldmann & Hagmayer, 2005). Young children also appear to have this capacity, at least under some circumstances. Schulz, Gopnik, and Glymour (2007) found that 4-year-olds used the conditional probability information inherent in another's actions to make inferences about a set of causal relations in a novel system. Children were not as accurate when they had to learn the same causal structures from just observing information that did not include equivalent conditional probability information. These findings suggest that children are engaged in the process of constructing representations of causal structure from the data they observe starting early in development.

Although it is clear that young children have sophisticated causal reasoning abilities regarding how conditional probability information relates to causal structure, an open question concerns the conditions under which they learn this information. A variety of recent developmental findings suggest that registering conditional probability information is contextually labile. For example, visual and auditory statistical learning are affected by stimulus type, attentional focus, and perceptual grouping (Baker, Olson, & Behrmann, 2005; Endress, Scholl, & Mehler, 2005; Toro, Sinnett, & Soto-Faraco, 2005; Turk-Brown, Jung, & Scholl, 2005). Similarly, inferences involving social cognition and language learning are mediated by the nature of the child's interaction with the model (Johnson, Slaughter, & Carey, 1998; Kuhl, 2007; Meltzoff, 1995). Our goal here is to further our understanding by examining whether contextual factors influence how children recover conditional probability information when learning causal structures. Considering how causal learning is gated by contextual factors can yield important information for creating ideal learning environments in both formal and informal settings.

The particular contextual factor we considered in the present experiment was the pedagogical value of an explanation for a particular action. We examined how well the rationale behind an action related to the goal of learning the causal structure. Adult participants appear to be sensitive to this information: Structural or functional rationales related to learning causal structures promote faster and deeper encoding of information than do only a linear set of instructions (Smith & Goodman, 1984). A variety of evidence suggests that children and infants are sensitive to pedagogical cues provided by adults and use this information to guide their learning from others' actions (Gergely, Egyed, & Kiraly, 2007). For example, providing infants and children with information about the ultimate goal of an action sequence prior to viewing the sequence enhances memory for observed events, attention to those events, and the ability to learn from others' actions (Carpenter, Call, & Tomasello, 2002; Sommerville & Crane, *in press*; Sommerville & Hammond, 2007).

Building on the method developed by Schulz et al. (2007), in which 4-year-olds learned causal structure from observing conditional probability information, in the present experiment we introduced 4-year-olds to a novel causal system and showed them a set of actions that conveyed certain conditional probability information. In the first part of the procedure, the data children observed were ambiguous and specified (at least) two different causal structures. This ambiguity was pointed out to the children. Then, the experimenter performed one particular action that conveyed the conditional probability information necessary to identify which of the two alternatives was the actual causal structure. All children observed this action and the subsequent information, which specified one causal structure. Children were then asked to demonstrate their knowledge of the causal structure of the system.

The manipulation in this experiment was the pedagogical context in which that critical action was presented. Some children were told that the action was being performed for a reason related to learning the causal structure. Other children were told that the same action (with the same results) was being performed because of a personal aesthetic held by the experimenter (and thus, irrelevant to

the goal of learning causal structure). A third group of children, in a baseline condition, received no explanation for the experimenter's action; they were merely asked to watch the action. Our question was whether and how these contexts might influence children's ability to learn a set of causal relations.


One possibility is that children are not influenced by these rationales at all and merely attend to the conditional probability information inherent in the events they observe. Since this information is equivalent across the three conditions, this would suggest that preschoolers are neither affected by inappropriate information nor do they benefit from hearing an appropriate rationale. It is also possible that children benefit from hearing an appropriate rationale or are negatively affected by hearing an irrelevant rationale, or both. Appropriate rationales might focus children on the information at hand, allowing for better processing and understanding. In particular, appropriate rationales accompanying causal interventions might enable children to form hypotheses or best guesses about the underlying causal structure, enhancing attention to and encoding of the subsequent information. In contrast, irrelevant rationales might lead children to process the same information more superficially or decrease their overall attention to integrating the data they observe with the process of learning the causal structure.

## 1. Method

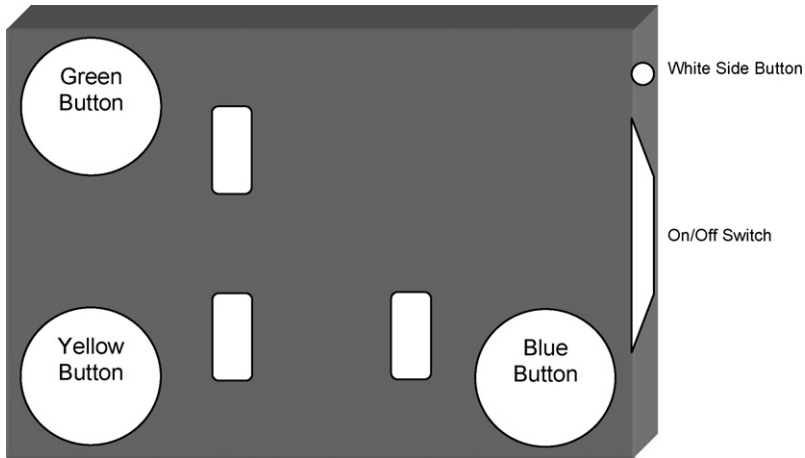
### 1.1. Participants and design

Sixty 4-year-olds (25 girls,  $M = 54.08$  months,  $S.D. = 3.46$ ) were recruited from flyers posted in local preschools and lists of hospital births. Seven additional children were tested, but excluded from the study: six because of experimental error and one refused to participate. All children were fluent English speakers. The ethnic distribution of the sample was as follows: 48 children were Caucasian, 9 children were Hispanic, 1 child was Asian, 1 child was of Middle-Eastern descent, and 1 child was of mixed descent.

We employed a mixed design. All children were asked to learn two sets of causal relations based on watching another person act on a novel causal system. These configurations (a common cause model and a chain model) are shown in Fig. 1. Children learned these models under one of three pedagogical conditions (randomly assigned,  $n = 20$ ). In the *appropriate rationale* condition, at a critical point in the presentation of data (described below), children heard a rationale for a particular action made by the experimenter that was appropriate to learning the causal structure. In the *inappropriate rationale* condition, at that same point in the presentation, children heard a rationale for that action that was irrelevant to causal learning. Finally, in the *baseline* condition, children heard no rationale and were simply shown the action. This between-subjects manipulation was the only difference across conditions – the actions and events participants observed were identical.



**Fig. 1.** A representation of the two models children were asked to learn in the experiment. In the common cause model (a), children explicitly were asked to learn that the green light (A) activates the yellow (B) and blue (C) lights. In the chain model (b), children were asked to learn that blue (A) activated yellow (B), which in turn activated green (C), but that there was not a direct causal relation between the blue and green lights.



**Fig. 2.** A schematic of the lightbox used in the experiment. The ovals represent the colored lights (color indicated by the nearest button) that illuminated when that button was depressed.

### 1.2. Materials


A schematic of the lightbox used in this experiment is shown in Fig. 2. The box itself was 20 cm × 15 cm × 8 cm, black in color, with three colored buttons (yellow, blue, and green) and three correspondingly colored lights visible on its top. The buttons were designed such that they always activated their correspondingly colored light (e.g., pressing the blue button always turned on the blue light). After a button was pressed, the corresponding light was illuminated for as long as the button was depressed.

The box could also be programmed (controlled by the experimenter flexibly throughout the procedure and unbeknownst to the child), such that any light could cause any other light or set of lights to activate. For instance, at any point in the procedure, the box could be programmed such that the blue button turned on both the blue light and the yellow light, while the yellow button only activated the yellow light. In this case, the illusion was created that the blue light caused the yellow light to activate. This flexibility allowed us to simulate the causal structures we asked children to learn. On the side of the box, there was an on–off switch as well as a small white button, mounted on a round (3 cm) white speaker. Pressing this button emitted a cartoon-like sound effect for approximately 3 s.

Children also observed a black plastic cup, referred to as a “cover,” approximately 2.5 cm in diameter, which was large enough to fit over one of the button/light combinations. Finally, children observed a set of pictures, each denoting a causal relation between two lights. Each individual picture was drawn on a 10.2 cm × 15.3 cm white index card, mounted on an 11.5 cm × 17 cm piece of blue cardboard. Each picture depicted two lights, each of which was a different color. Each light was animated with a face and an arrow pointed from the left light to the right light. The arrow was the same color as the light on the left. Two such pictures, which constituted a set, are shown in Fig. 3.

### 1.3. Procedure

The experimenter showed the child the lightbox, and explained that the lights came on when the corresponding buttons were pressed. This was demonstrated to the child. At the outset of the experiment, the box was programmed such that no light caused any other light to activate, so each button only illuminated its corresponding light. Children were asked to turn on each light one at a time, to ensure that children knew the color names and recognized that pressing the corresponding button would activate the light of the same color. This was the only time children were allowed to act on the box themselves. Children always answered these requests correctly. The experimenter then



**Fig. 3.** An example of the two sets of pictures children observed at the test question. This example is from the common cause model in which children saw pictures of green activating yellow and blue (the set on the left) and pictures of green activating yellow and yellow activating blue (the set on the right). *Note:* In this figure, the light gray is yellow, the medium gray is green, and the dark gray is blue.

pressed each button again and reinforced the fact that each of the three buttons would always activate its correspondingly colored light.

### 1.3.1. Familiarization phase

After this introduction, the experimenter then surreptitiously reprogrammed the box such that the yellow button would activate both the yellow and blue lights. At this point, children were shown the white side button. This button was pressed each time the experimenter wanted to change the causal structure presented by the box and thus provided a plausible mechanism by means of which the same box could be used successively to depict different causal models.

The experimenter pressed the button and it emitted a cartoon-like noise. The experimenter said, "Now, the box is different – now it is a puzzle box. In the puzzle, some lights make other lights go. Let me show you." The experimenter pressed the blue button, which resulted in only the blue light activating. The experimenter said, "See, when I press the blue button, only the blue light comes on; I don't think blue makes anything go in this puzzle." The same action and narration was repeated for the green button/light. Then, the experimenter pressed the yellow button, which resulted in yellow and blue activating. The experimenter pressed the blue button again, which resulted in only the blue light activating. All of these actions and results were narrated for the child. The experimenter then told the child that, "it looks like yellow makes blue go in this puzzle".

At this point, the experimenter introduced a method by which an individual light could be removed from the lightbox to present particular conditional probability information in the test phase of the experiment. To do this, children were shown the black cover, which could be fitted over a light. The experimenter stated that he thought yellow made blue go, and as a test, he said that he would cover the yellow light and then press the yellow button. This way, he explained, he and the child could see whether the blue light came on when the yellow light was blocked and the yellow button was pressed.

Specifically, the experimenter told children that they would see what happens when the blue light could not see the yellow light because the yellow light was covered. Schulz et al. (2007) used similar anthropomorphized language to explain their novel causal system. We believed this anthropomorphized language would be highly accessible to children. When the yellow button was pressed while that light was covered, the blue light did not activate. This action was performed a few times so that the child could observe it, and the results were narrated for the child. To make this occur in reality, unbeknownst to the child, the box was programmed at that particular point in time such that the blue light would not activate when the yellow button was pressed. Children also watched the experimenter peek under the cover while the yellow button was pressed, and told the child that yellow was coming on, but blue was not “because blue can’t see yellow.”

The cover was then removed, and the yellow button was pressed again, which resulted in both yellow and blue activating. This action was narrated for the child, and again the child was told that the experimenter thought that yellow made blue go. At this point, children were shown a picture like that in Fig. 3, of a yellow light pointing to a blue light. These pictures were similar to the ones used by Schulz et al. (2007) to test children’s understanding of causal structure. Children were also asked six individual causal structure questions, concerning the six possible causal relations among the three lights (i.e., Does the blue light make the yellow and green go? Does the green make blue and yellow go? Does yellow make blue and green go?). This got children used to responding to the experimenter. Corrective feedback was provided on these questions, if necessary.

### 1.3.2. Test phase

The experimenter then pressed the white button on the side of the box, which elicited the cartoon-like noise. Children were told that the box had changed to a different puzzle, and that the experimenter and child would learn the new puzzle together. Children were asked to learn two causal structures, a common cause model (Fig. 1a), and a chain model (Fig. 1b). The side button with the cartoon-like noise was pressed in between learning the two models, and at this point children were told the puzzle had changed.

In the common cause model, children always learned that the green light (A) activated yellow (B) and blue (C). In the chain model, children learned that blue (A) activated yellow (B), which in turn activated green (C). These color assignments were randomly determined prior to the start of the experiment with one caveat: Across the two models, we forced the (B) light in both models to be the same color (yellow). This was done in order for the experimenter’s preference in the inappropriate condition to be coherent (see below). The order in which children learned the two models was counterbalanced.

The test phase started with the experimenter pressing the (A) button for that particular model (green for the common cause model, blue for the chain model). When this occurred, all three lights illuminated. The experimenter narrated what lights illuminated, and then said, “This is pretty tough, I don’t know what to think about the lights yet.” Then, the (C) button was pressed (blue for the common cause model, green for the chain model), which resulted in only that particular light illuminating. The experimenter commented on this result by saying, “Look, only (light C) comes on. I don’t think the (C) light makes any of the other lights go. But, I don’t know about (light A).” The experimenter then specified one of two possibilities, either a common cause or a chain (e.g., in the chain model he said, “It could be that blue makes yellow go and blue makes green go. Or, it could be that blue makes yellow go, and yellow makes green go.”) The order of the two options was counterbalanced between children.<sup>1</sup>

The experimenter then said, “Let’s take the cover and put it over the yellow light (light B) and press the (A) button. At this point, children were assigned to one of three conditions (between subjects). In the *appropriate rationale* condition, children were told, “We can cover (light B, which was always yellow) so we can see what happens when we press the (A) button without the yellow light.” This way, children received a rationale for this action related to learning the causal structure. Children in the *inappropriate rationale* condition were told, “We can cover (light B, which was always yellow) because I don’t like

<sup>1</sup> Note that these two causal structures are not the only models consistent with the observed data (e.g., a model in which light A causes B and C and light B causes C is also consistent), but the subsequent data specifically discerned between the common cause and chain models.

it that much, and we can press the (A) button because I think the (A) light is pretty.” Thus, in this condition, the rationale children received for this action was focused on the experimenter’s personal aesthetic, and irrelevant to learning the causal structure. Because it made sense for the experimenter’s preferences not to change during the course of the procedure, we initially assigned one color at random (yellow) to always be the light that is covered across the two models. Finally, in the *baseline* condition, the action was performed without a rationale.

At this point, across all three conditions, children were reminded about the nature of the cover, “Remember, when we cover the yellow light, yellow can’t see (light A) and (light A) can’t see yellow, and yellow can’t see (light C) and (light C) can’t see yellow. The (A) button was then pressed several times. In the common cause model, lights A and C (green and blue respectively) illuminated. In the chain model only the A light (blue) illuminated. These results were generated several times, and the experimenter made sure that the child observed these interventions. Children were then told to look at two sets of pictures (see Fig. 3), which contrasted the appropriate common cause or chain model from the other alternative the experimenter mentioned. The content of these pictures was described, and children were asked to pick the pictures they thought were right. The spatial location of the picture choices was counterbalanced.

## 2. Results

Overall, children required corrective feedback on an average of 0.32 of six questions during the familiarization phase (about 5% of the time), and accuracy did not differ across the three conditions. This suggested that they understood the nature of the task and how the lightbox and cover operated. Children were assigned a score of 1 if they chose the appropriate set of pictures for each causal structure, and a score of 0 otherwise. Preliminary analyses suggested that gender, model order, and spatial location of questions did not affect these scores (Chi-squared tests, all *p*-values *ns*). Overall, and in each condition separately, scores between the common cause and chain models did not differ (across the three conditions, children responded correctly on the common cause model 58% of the time as opposed to 48% of the time on the chain model, McNemar tests, all *p*-values *ns*), so these data were combined into an aggregate score between 0 and 2.

Children in the appropriate rationale condition received an overall score of 1.40 (of 2; S.E. = 0.11), compared with a score of 0.80 (S.E. = 0.14) in the inappropriate rationale condition and 1.00 in the baseline condition (S.E. = 0.13). These scores were subjected to an analysis of variance, which resulted in a main effect of condition,  $F(2, 57) = 5.91, p = .005$ , partial  $\eta^2 = .17$ . Simple effect analysis revealed that performance in the appropriate rationale condition was significantly different from that in the baseline condition ( $p = .05$  with a Dunnett correction), but that performance in the inappropriate rationale was not significantly different from that in the baseline condition ( $p > .10$ ). Children were also more likely to respond correctly in the appropriate than inappropriate rationale condition ( $p = .004$  with a Scheffé correction).

These findings were supplemented by nonparametric analyses. Table 1 presents the distribution of children’s responses according to response category. Overall, the distribution of responses differed across conditions,  $\chi^2(4, N = 60) = 10.98, p = .03$ , Cramer’s  $\phi = .428$ . Specifically, the distribution of responses differed between the appropriate and inappropriate rationale conditions,  $\chi^2(2, N = 40) = 9.60, p = .008$ , Cramer’s  $\phi = .490$ , and the distribution of responses between the appropriate rationale and baseline conditions were marginally different from one another,  $\chi^2(2, N = 40) = 5.43, p = .06$ , Cramer’s  $\phi = .368$ . There was no significant difference between responses in the inappropriate

**Table 1**  
Number of children in each response category.

	Neither correct	Only common cause correct	Only chain correct	Both correct
Condition				
Appropriate rationale	0	6	6	8
Inappropriate rationale	6	9	3	2
Baseline	3	7	7	3

rationale and baseline conditions,  $\chi^2(2, N=40)=1.35$ , *ns*. Examining performance in each condition separately, a Chi-squared goodness-of-fit test revealed that the distribution of responses in the appropriate rationale condition was significantly different from chance,  $\chi^2(2, N=20)=7.20$ ,  $p < .01$ . The distributions of responses in the baseline and inappropriate rationale conditions were not significantly different from chance,  $\chi^2(2, N=20)=3.20$  and  $2.40$ , both *ns*.

### 3. Discussion

Although children observed the experimenter perform the same actions with the same conditional probability information, their ability to recover a causal structure varied as a function of the experimenter's rationale for generating the action that produced the information. Not all observations of action equally afforded accurate causal learning. Specifically, children appeared sensitive to the positive instructional value of the experimenter's appropriate rationale. Only in this condition was accuracy above chance, while receiving an inappropriate rationale did not appear to impair learning beyond baseline.

One question raised by these data is whether children are learning an overall causal structure or are just recognizing the presence or absence of a particular causal relation. In deciding between a common cause and chain model, children might simply consider whether a link from the A light to the C light is present. If this is the case, then the information specifies a common cause and if not it specifies a chain. Thus, children might be learning causal structures as a coherent whole or as a set of individual relations. We speculate that the former is more likely, simply because if children were only paying attention to whether a causal relation between lights A and C was present, they might be less sensitive to the contextual information across the rationales, since in all conditions they directly observe the experimenter press the button that causes the A light to illuminate while only looking at the A and C lights.

Whether children are learning individual causal relations or an overall causal structure, there appears to be an effect of the rationale children hear on their ability to respond correctly to the test question. Why does an appropriate rationale benefit causal learning compared to baseline, while an inappropriate rationale results in levels of accuracy no different from baseline?

#### 3.1. *Why do appropriate rationales benefit causal learning?*

One possible explanation for the pattern of results we observed is based on children's ability to focus their attention. The inappropriate rationale was about specific lights, and potentially focused the child on whether an individual event occurred (i.e., the occurrence of the "pretty" light). In contrast, the appropriate rationale was about whether a relation between two lights existed, which might have focused the child's attention on what other lights occurred in addition to the light that corresponded to the button that was pressed. This account might not be completely consistent with the present data, however. The baseline condition did not focus children on individual events, but revealed similar performance to that of the inappropriate rationale condition. That said, it is possible that explicitly focusing children on learning the causal relations positively affected accuracy of responses.

A second possibility is that hearing the appropriate rationale enhanced children's overall engagement with the procedure, not just whether they attended to one light or another. Scaffolding what children observed with a rationale for why they were going to see it might have led children to process the conditional probability information more deeply and thus be more likely to draw an accurate causal conclusion.

A related possibility is that the appropriate causal rationale may have encouraged children to form a hypothesis or best guess about the causal outcome, which in turn may have enhanced their attention to or memory of that outcome. A variety of evidence suggests that learning environments that facilitate prospective processing of actions and outcomes promote greater learning of those actions and outcomes than learning environments that do not do so (Ratner, Foley, & Gimpert, 2002; Sommerville & Hammond, 2007). These studies demonstrated that collaborative contexts, in which a child and an adult share a goal and the child has access to the adult's upcoming action plans, enhanced children's ability to encode and recall their partner's actions. An intriguing question for future research is whether

explicitly manipulating factors that influence prospective processing, such as the degree of collaboration between the child and experimenter, might affect children's causal learning from another person's interventions.

An account that emphasizes the role of appropriate rationales in engagement and/or prospective processing nicely explains a facet of the present data that is, on the surface, incongruent from one previous study of children's causal learning. In the present experiment, children responded at chance on the baseline condition. In contrast, Schulz et al. (2007) found that 4-year-olds recovered causal structures (specifically about how gears interacted on a box) reliably above chance when they observed another person act without necessarily hearing a specific rationale for any action (similar to the baseline condition). There were several differences between the two procedures which might have affected performance. In the present experiment, children had to learn about the causal structure among three events instead of two (suggesting, incidentally, that they are not learning individual causal relations, but rather overall causal structures). Further, in the present experiment the mechanism among the lights is not directly perceivable, while the mechanism linking the gears in Schulz et al.'s procedure was directly perceivable, since they must interlock. These differences might have made our task slightly more difficult. But in addition, in Schulz et al.'s (2007) procedure, children were given a great deal of information as to the goal of the experimenter's actions during the procedure. This information might have encouraged children to attend to those actions in more detail, similar to the present appropriate rationale condition. This suggests that if children were given less structure in the demonstration, or rationales that were unrelated to causal learning, in Schulz et al.'s experiment, performance might drop to chance levels.

To conclude, the present findings address one part of describing the process by which children engage in causal learning. We do not wish to minimize the importance of conditional probability information in causal learning. We also do not wish to argue that understanding the instructional value of another person is more important than observing conditional probability data when children engage in causal learning. Rather, we conclude that young children can appreciate conditional probability information, but do so in light of contextual factors that affect how inferences are made based on this information. These findings add to a growing literature concerning the ways in which children's learning can be gated by various contextual factors and may provide important information for structuring and creating optimal learning environments.

## Acknowledgments

This work was supported by NSF (DLS-0518161). We would like to thank all of the parents and children who participated in this research. We would also like to thank Esra Aksu, Emily Blumenthal, Sheridan Brett, Claire Cook, Emily Hopkins, Cesalie Stepney, and Kristen Sylvester who helped with data collection and coding, and Philip Parker who helped with stimulus design.

## References

- Baillargeon, R. (2002). The acquisition of physical knowledge in infancy: A summary in eight lessons. In U. Goswami (Ed.), *Blackwell handbook of childhood cognitive development* (pp. 47–83). Malden, MA: Blackwell Publishing.
- Baker, C. I., Olson, C. R., & Behrmann, M. (2005). Role of attention and perceptual grouping in visual statistical learning. *Psychological Science*, *15*, 460–466.
- 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.
- Carpenter, M., Call, J., & Tomasello, M. (2002). Understanding “prior intentions” enables two-year-olds to imitatively learn a complex task. *Child Development*, *73*, 1431–1441.
- Endress, A. D., Scholl, B. J., & Mehler, J. (2005). The role of salience in the extraction of algebraic rules. *Journal of Experimental Psychology: General*, *134*, 406–419.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences of The United States of America*, *99*, 15822–15826.
- Gergely, G., Egyed, K., & Kiraly, I. (2007). On pedagogy. *Developmental Science*, *10*, 139–146.
- 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 co-variation. *Developmental Psychology*, *37*, 620–629.
- Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, *61*, 233–259.
- Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind body distinction. *Child Development*, *64*, 1534–1549.

- Johnson, S., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? The elicitation of gaze-following in 12-month-olds. *Developmental Science*, 1, 233–238.
- Kalish, C. W. (1996). Preschoolers' understanding of germs as invisible mechanisms. *Cognitive Development*, 11, 83–106.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy. *Cognition*, 83, B35–B42.
- Kuhl, P. K. (2007). Is speech learning 'gated' by the social brain? *Developmental Science*, 10, 110–120.
- Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 856–876.
- Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31, 838–850.
- Perner, J. (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press.
- Ratner, H. H., Foley, M. A., & Gimpert, N. (2002). The role of collaborative planning in children's source-monitoring errors and learning. *Journal of Experimental Child Psychology*, 81, 44–73.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
- Schult, C. A., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition*, 62, 291–324.
- 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 causal structure from conditional independence. *Developmental Science*, 10, 322–332.
- Smith, E. E., & Goodman, L. (1984). Understanding instructions: The role of an explanatory scheme. *Cognition and Instruction*, 1, 359–396.
- Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology*, 42, 1103–1115.
- Sobel, D. M., & Kushnir, T. (2006). The importance of decision demands in causal learning from interventions. *Memory & Cognition*, 34, 411–419.
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28, 303–333.
- Sommerville, J. A., & Crane, C. C. (in press). Ten-month-old infants use prior information to identify an actor's goal. *Developmental Science*.
- 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.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, 99, 605–632.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27, 453–489.
- Toro, J. M., Sinnott, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, 97, B25–B34.
- Turk-Brown, N. B., Jung, J., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, 134, 552–564.
- Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 216–227.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., & Liu, D. (2007). Causal reasoning as informed by the early development of explanations. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation*. New York, NY: Oxford University Press, pp. 261–279.