Cognition 129 (2013) 494–500

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Brief article

Children's developing understanding of the relation between variable causal efficacy and mechanistic complexity

Christopher D. Erb [a,*], David W. Buchanan [b], David M. Sobel [a]

[a] Brown University, United States
[b] IBM Research Laboratories, United States

ARTICLE INFO

Article history:
Received 6 July 2012
Revised 31 July 2013
Accepted 2 August 2013

Keywords:
Causal reasoning
Causal mechanisms
Cognitive development

ABSTRACT

Two experiments investigated 3–4-year-olds' ability to infer the causal mechanisms for a pair of lights. In both experiments the exterior of the two lights appeared identical. In Experiment 1, one light displayed a stable activation pattern of a single color while the other light displayed a variable pattern of activation by cycling through a series of different colors (i.e., a more varied effect). Children were asked to judge which light had a more complex internal structure. Four-year-olds were more likely to match the light with the more variable effect with a more complex internal mechanism and the light with the more stable effect with a less complex mechanism. Three-year-olds' responses were at chance. Experiment 2 replicated this finding when the activation patterns of the two lights were described verbally but never demonstrated. Taken together, these results suggest that 4-year-olds appreciate that the variability of an object's causal efficacy is related to the complexity of its internal mechanistic structure.

1. Introduction

Imagine a child at an arcade. At one machine, she puts in a token and a gumball appears from a chute. At another machine, she puts in a token and starts to play a videogame. Identifying the efficient cause of receiving the gumball and being able to play the game is straightforward. Investigations of children's causal reasoning attempt to identify the principles underlying how children make these types of causal attributions, such as how events relate to one another in time and space (e.g., Bullock, Gelman, & Baillargeon, 1982), whether events' co-occurrence indicates a causal relation or spurious association (e.g., Gopnik, Sobel, Schulz, & Glymour, 2001), and the base rate with which events co-occur (e.g., Sobel, Tenenbaum, & Gopnik, 2004). These findings suggest that even young preschoolers would have little difficulty understanding that the tokens caused the effects in both cases.

While children might make these causal inferences, they potentially know little about the specific mechanism(s) by which the causes produce their effects. This distinction was pointed out by (among others) Gopnik et al. (2001), who suggested that very young children have knowledge of *formal* principles that underlie causal inference, like those described above, but also more *substantive* principles, such as particular kinds of mechanistic knowledge. While children might register the role of certain formal principles very early in development (e.g., Sobel & Kirkham, 2006), this latter kind of knowledge clearly develops throughout childhood. Even adults do not possess a complete understanding of causal mechanisms. They often over-attribute their understanding of how a causal system works (e.g., Rozenblit & Keil, 2002), and young children show similar effects (Mills & Keil, 2004). But just because children and adults believe they possess more mechanistic knowledge than they do, does not mean that they fail to possess *any* mechanistic knowledge. Indeed, some have argued that recognizing how causes and effects are related is more important for making causal attributions than appreciating the formal principles under-

---

* Corresponding author. Address: CLPS Department, Box 1821, Brown University, Providence, RI 02912, United States. Tel.: +1 401 863 2821.
E-mail address: Christopher_Erb@Brown.edu (C.D. Erb).

0010-0277/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cognition.2013.08.002

lying causal inference, such as correlations among events (e.g., Ahn, Kalish, Medin, & Gelman, 1995; Buchanan & Sobel, 2011; Shultz, 1982).

Our goal in the present investigation is to articulate the development of one kind of mechanism knowledge young children might have available to them when they infer the efficient cause(s) of observed effects. Returning to the arcade example, nothing of the intervening mechanism is perceptually available to the child during either of the events. But children do know that the outcomes of the two events differed greatly. In one instance, the outcome was short and simple: the gumball appeared. In the other case, the outcome was variable and extended: the game lasted for some time and (if it was at all interesting) involved presumably many more visual and auditory displays. Given this information, children might form the reasonable expectation that the mechanism underlying the latter outcome is more complex than that of the former.

This hypothesis emanates from evidence that children recognize that objects' causal properties are related to their insides, which undergoes development during the preschool years. Gelman and Wellman (1991) found that 4–5-year-olds recognized that members of the same category (i.e., objects given the same label) often shared internal structure. Children believed that category membership was a better predictor of insides than external perceptual similarity. Gelman and Wellman did not test younger children, but several follow-up studies suggest that 3-year-olds do not make similar inferences about the relation between an object's causal properties or category membership and their insides (e.g., Gottfried & Gelman, 2005; Sobel, Yoachim, Gopnik, Meltzoff, & Blumenthal, 2007). For instance, Sobel et al. (2007) showed that 4-year-olds inferred that objects with shared causal efficacy had shared insides, even when a conflicting strategy based upon the objects' visual appearance was available. Three-year-olds, in contrast, relied on external appearance when making these inferences, and not causal properties.

While 4-year-olds register that an object's insides and causal efficacy are related, the present experiments examine whether they form a more sophisticated expectation: that more variable efficacy can imply the presence of a more complex internal mechanism. In Experiment 1, we showed 3- and 4-year-olds two objects that activated given the same action, but differed in the variability of their effects. We examined whether children would match the object with the more variable effect with a more complex internal structure. In Experiment 2, the objects were not activated. Instead, the experimenter described the pattern of activation for each object verbally, thereby equating the objects' visual appearance. The question underlying both experiments is whether 3- and 4-year-olds differ in their appreciation of the relation between an effect's complexity and the underlying mechanism between it and its cause.

2. Experiment 1

Three- and 4-year-olds were presented with a matching game. Children were shown two lights that activated when pressed. One light exhibited a variable pattern of activation by repeatedly flashing a series of colors. The other light displayed a stable activation pattern by maintaining a single, solid color. Children were then presented with two pictures that were purported to match the insides of the lights. One of the pictures featured a causal mechanism with only a few internal components. The other picture featured a more complex mechanism with additional components. If children expect variability in activation to imply a more complex underlying mechanism, we would expect children to match the picture featuring the complex insides to the light featuring a variable pattern of activation. If children do not form such an expectation, we would predict matching performance to be at chance.

3. Method

3.1. Participants

Thirty-three 3-year olds (13 girls, $M = 41.81$ months, $SD = 2.73$, range 35–46 months) and thirty-two 4-year-olds (14 girls, $M = 52.66$ months, $SD = 3.72$, range 48–59 months) were included in the final sample. An additional seven children were tested but excluded from the final sample due to experimenter error ($n = 4$) or technical difficulties ($n = 3$). Participants were recruited from a local preschool, a local children's museum, and a list of hospital births.

3.2. Materials

During the familiarization phase, children were shown two transparent boxes, each $5 \times 3 \times 3''$, were presented during the familiarization phase. One box contained a black triangular wooden block and the other contained a green wooden cube of similar dimensions (see Fig. 1). These objects were paired with pictures that matched the insides of the boxes; one picture featured a black triangle and the other featured a green square.

During the test phase, children were shown three round lights that switched on or off when pressed. Two were modified such that the original bulb was replaced with a LED that could be set to a variety of activation patterns. One of the lights was set to illuminate a solid color when activated. The other light cycled through a series of seven colors at a constant rate of approximately one cycle every 2 s. We will refer to these as the *solid* and *variable* lights respectively. The third light was unmodified and

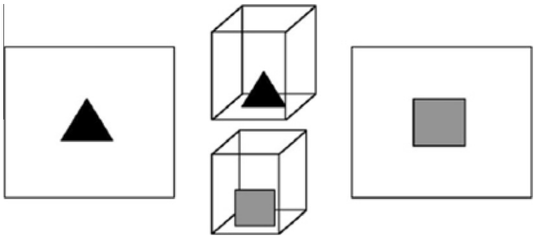

Fig. 1. A schematic representation of the materials used during the training phase as seen from the child's position.

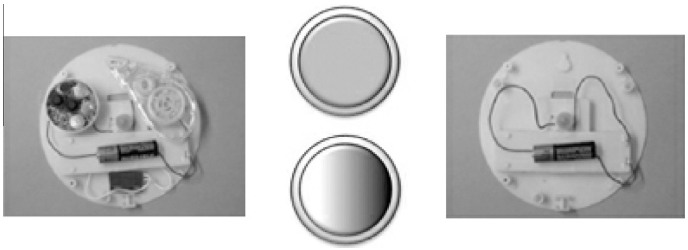

Fig. 2. The materials presented to children during the experimental phase, including the simple (right) and complex (left) mechanism pictures and illustrations of the solid (top) and variable (bottom) lights.

illuminated white using a standard bulb. Outside of their activation patterns, the lights were identical. In addition to the lights, two fictitious pictures of the lights' 'insides' were presented. In the *simple mechanism* picture, the insides consisted of a battery and a light bulb connected by wire. In the *complex mechanism* picture, three additional components were present. These additional components were selected to appear vaguely electronic. They included a piece of a voltmeter, a set of plastic gears within a transparent case, and a magnetic switch (see Fig. 2). In reality, the inside wiring of the modified lights was more complex than the pictures portrayed. However, none of the children commented on or apparently resisted the idea that these pictures were genuine.

3.3. Procedure

Children were first given a familiarization phase to acquaint them with the procedure. The experimenter showed the child two transparent boxes, highlighting that each contained a different object (see Fig. 1). The boxes were then placed one in front of the other between the experimenter and child. The experimenter then presented the two pictures corresponding to the boxes' insides, explained that the objects in the pictures matched the insides of the boxes, and placed one of the pictures on each side of the boxes. Next, the experimenter invited the child to play the matching game by matching the pictures with the insides of the toys. The experimenter then used both hands to remove the box furthest from the child and asked, "Which of these pictures goes with the insides of this toy?" After the child responded, the box on the table was replaced with the other box. Once again, the child was asked to indicate which picture matched the insides of the toy on the table. No feedback was given. After these questions, the table was cleared.

Next, in the test phase, the experimenter introduced the child to the unmodified light. The experimenter illustrated how the light worked by pressing down on it, which turned it on, and then pressed it again, which turned it off. The child was allowed to play with the light briefly. The experimenter then put away the unmodified light and explained that he had two more lights to show the child. These lights were presented to the child sequentially and were situated one in front of the other between the experimenter and the child (as shown in Fig. 2). The order of which light appeared first was counterbalanced.

The experimenter then introduced the two corresponding pictures and placed one on each side of the lights. The experimenter labeled the insides of each of the pictures. While pointing to each component in the simple mechanism picture, the experimenter said, "In this picture, there is a light and a battery on the inside." Next, the experimenter pointed to the components of the complex mechanism picture and stated, "In this picture, there is a light and a battery and a blicket and a stennet and a gazzer on the inside." The light furthest from the child was then removed from view and the experimenter asked, "Which of these pictures goes with the insides of this light?" After the child responded, the light on the table was replaced with the other light. Once again, the child was asked to indicate which of the pictures matched the insides of the toy on the table. We counterbalanced which light (i.e., the solid or variable light) the child was queried on first.

4. Results

Children were scored as passing the familiarization trial if they matched the box with the triangle inside to the triangle picture and the box with the square inside to the square picture. Sixty-two of the 65 children (95%) generated this response. When the three children who did not pass familiarization were excluded from the analysis, results did not differ from what is reported below.

During the test phase, responses that matched the simple mechanism picture to the solid light and the complex mechanism picture to the variable light were scored as correct and all other patterns of matching were scored as incorrect. The majority of children (41 of 65 or 63%) performed correctly on the test phase, which is significantly different from chance performance (50%[1]), Binomial test, $p < 0.05$, $g = 0.13$.

Fig. 3 shows the breakdown of performance by age. While more 4-year-olds responded correctly than 3-year-olds (23 of 32, or 72% vs. 18 of 33, or 55%), this difference did not reach statistical significance, $\chi^2(1, N = 65) = 2.10$, $p = 0.15$. Four-year-olds were more likely to respond above chance, Binomial Test, $p < 0.05$, $g = 0.22$, while 3-year-olds performed no different from chance levels, Binomial Test, $p = 0.73$.

---

[1] Technically, children had to respond to two 2AFC questions, and thus a liberal interpretation of chance could be 25% correct. However, given that the responses were mutually exclusive, and only one child (a 3-year-old) picked the same picture twice, we thought that 50% was the most conservative interpretation of chance performance.

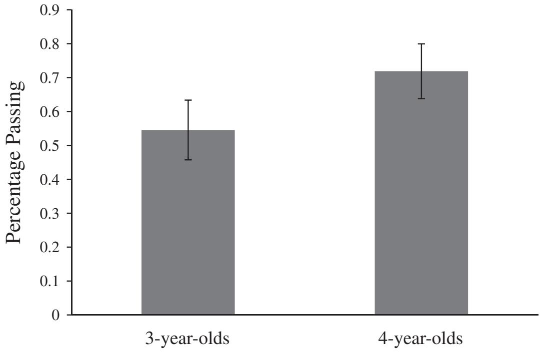

Fig. 3. The percentage of 3- and 4-year-olds passing the experimental phase.

5. Discussion

Four-year-olds matched pictures portraying different levels of mechanistic complexity (low or high) to novel causal objects displaying different levels of variability in activation (stable or variable, respectively). Three-year-olds' performance did not differ from chance. Consistent with previous studies demonstrating that older preschoolers recognize the relation between objects' causal efficacy and their insides, 4-year-olds in this study recognized that the nature of that efficacy indicated something about the complexity of each object's internal structure.

A difficulty with this conclusion, however, is that children might have based their responses not on an inference about the nature of the causal relation and mechanism, but on a simpler association between the visual complexity of the effect and the perceptual complexity of the picture. Experiment 2 controlled for the visual complexity of the lights by substituting the demonstration phase in which each light was activated for a verbal description of each light's activation. If 4-year-olds were simply matching based on visual complexity, they would now respond at chance levels. In contrast, if 4-year-olds inferred that the complexity of the causal relation and not the external perceptual appearance mattered, they would continue to match the variable light with the more complex internal mechanism.

6. Experiment 2

Experiment 2 replicated Experiment 1 with an important modification: The activation patterns of the two test lights were never demonstrated. Instead, the experimenter verbally described the activation pattern of each light. This allowed both of the lights to remain perceptually identical throughout the procedure.

7. Method

7.1. Participants

Twenty-five 3-year olds (14 girls, $M = 42.93$ months, $SD = 2.82$, range 37–47 months) and twenty-five 4-year-olds (12 girls, $M = 53.82$ months, $SD = 2.92$, range 48–59 months) were tested. No children were excluded from the final sample. Participants were recruited from a local preschool, a local children's museum, and a list of hospital births.

7.2. Materials and procedure

The materials used were identical to those of Experiment 1. The familiarization phase of the procedure was the same as Experiment 1. Children matched the triangle and square with the pictures of the triangle and square. There were four differences in the test phase of the procedure. First, neither of the modified lights was activated. Rather, the experimenter verbally described the activation pattern of each light. For the solid light the experimenter said, "When I push down on this one, it lights up and stays the same color." For the variable light the experimenter said, "When I push down on this one, it lights up and changes colors."

Second, before introducing the simple and complex mechanism pictures, the experimenter explained how the pictures were generated in order to make the relations between the pictures and the light more transparent ("Earlier, I took the tops off of each of these toys and took a picture of what they look like on the inside. Let's take a look at the pictures."). Third, in Experiment 1, the simple mechanism picture was always explained before the complex mechanism picture. In Experiment 2, this order was randomized. Finally, the child was reminded about the activation pattern of each item before selecting a picture (e.g., "When I push down on this one, it lights up and stays the same color. Which of these pictures goes with the insides of this light?").

8. Results

Responses were scored in the same manner as Experiment 1. In the familiarization phase, all 50 of the children matched the boxes and pictures appropriately. The majority of children (32 of 50, or 64%) performed correctly on the test phase, however this level of performance was only marginally above chance (50%), Binomial test, $p = 0.065$, $g = 0.14$.

Fig. 4 shows the proportion of 3- and 4-year-olds that responded correctly in the test phase. Once again, more 4-year-olds responded correctly than 3-year-olds (18 of 25 or 72% vs. 14 of 25 or 56%), but this difference was not statistically significant $\chi^2(1,\ N = 50) = 1.39$, $p = 0.38$. Four-year-olds performed significantly better than chance (determined in the same manner as Experiment 1), Binomial test, $p < 0.05$, $g = 0.22$, whereas 3-year-olds did not, Binomial Test, $p = 0.69$.

9. Discussion

The results of Experiment 2 are similar to those of Experiment 1. Only 4-year-olds reliably matched the object with the complex causal effect to the complex internal structure. These results provide evidence that 4-year-olds'

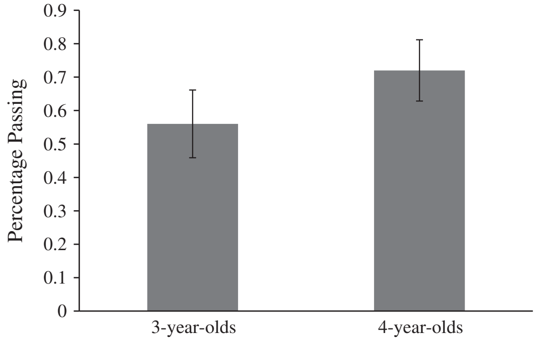

Fig. 4. The percentage of 3- and 4-year-olds passing the experimental phase.

performance in Experiment 1 was not due to a tendency to match the pictures and lights on the basis of observed perceptual complexity alone. Rather, 4-year-olds appear to have inferred that greater variation in activation corresponded to greater mechanistic complexity.

10. General discussion

The present study demonstrates that 4-year-olds can form inferences about the complexity of an object's internal mechanism based on the variability of the object's effects. This research illustrates how children go beyond simply identifying causes and effects in their environment to form deeper inferences about the causal systems at work around them (see also Schulz, Goodman, Tenenbaum, & Jenkins, 2008). The developmental difference is consistent with previous research suggesting that children do not reliably understand of the relation between an object's nonobvious internal properties and its causal efficacy until around the fourth birthday (e.g., Sobel et al., 2007).

The current results demonstrate that children know something about the link between mechanisms and their effects by the age of 4, but that younger children might not understand this connection. An important question in the causal inference literature is how children make causal attributions from the interaction of formal learning principles, like information about statistical regularity, and more substantive principles, like knowledge of specific causal mechanisms. In many investigations suggesting that young children have sophisticated causal reasoning capacities, what is usually measured is the formal principle independent of a causal mechanism, or with very little mechanism knowledge present (e.g., Gopnik et al., 2001; Schulz & Gopnik, 2004). Given these data, we do not suggest that 3-year-olds lack causal inference capacities. Rather, what we are suggesting is that an aspect of their substantive mechanistic knowledge is developing. But what specific knowledge is indicated by the difference that we observed?

One possibility is that the 4-year-olds have specific knowledge about how machines or electronics work that the 3-year-olds simply lacked. This possibility, however, seems unlikely. For instance, Cook and Sobel (2011) showed that children's developing knowledge of whether machines with different functions are or are not possible extends well beyond the age of 4.

An alternative possibility, which we support, is that during the preschool years children are developing knowledge of a more general sort; namely, an appreciation that causal relations involve underlying mechanisms. This possibility is supported by work by Buchanan and Sobel (2011). They showed that 4-year-olds reason about novel causal relations as if there are underlying mechanisms connecting the efficacy of the cause to that of the effect. Three-year-olds, in contrast, could only do so under very specific conditions. Taken with the present data, this suggests that by the age of 4, children have developed the understanding that causes affect their outcomes through a set of mechanisms and the more complicated the effects, the more complicated those mechanisms are likely to be.

This account is also consistent with work by Bullock et al. (1982), who found that 4–5-year-olds, but not 3-year-olds, rejected events that possess irrelevant mechanisms as potential causes for an outcome. Children are developing a more sophisticated understanding of what constitutes a causal mechanism between the ages of 3–4. Similarly, in Shultz's (1982) experiments on generative transmission, 5-year-olds, but not 3-year-olds, recognized that correlational information only indicated a causal relation when the child could posit a mechanism linking cause to outcome. These experiments are indicative of several other studies that point out differences between 3-year-olds and 4- or 5-year-olds' depth of understanding about how causal inferences are mediated by mechanistic knowledge (e.g., Das Gupta & Bryant, 1989; Nazzi & Gopnik, 2000; Sobel et al., 2007).

Critically, there is nothing about the child being 3 that prevents them from possessing this kind of mechanism knowledge. Sobel and Munro (2009) presented 3-year-olds with an ambiguous causal attribution problem that required children to apply substantive mechanistic knowledge (specifically about whether causal relations acted deterministically or stochastically). Over various conditions that were designed to support more deterministic or stochastic responses from these children, they found that if 3-year-olds believed the causal relations were deterministic, they resolved the data differently than if they believed the relations were not.

Moreover, there are cases in which even younger children potentially understand something about causal mechanisms. Madole and Cohen (1995) showed that 18-month-olds, but not 14-month-olds would limit the kind of statistical regularity they learned to cases that were mechanistically plausible (specifically through contact relations). Although in this case the mechanism was underspecified, these results suggest that even in infancy, children are integrating substantive causal principles with more formal principles for learning.

That said, an important limitation with the present work has to do with the linguistic nature of our tasks. In both experiments, the "correct" response was to match the light with the simpler effect with an internal structure that had a simpler linguistic description (i.e., shorter, no novel labels) and to match the light with the more complex effect with an internal structure that was described with more complicated language. It is possible that children's responses did not reflect their knowledge of causal mechanism. Four-year-olds could have generated their responses because they appreciated the relational mapping between a more complex effect and a more complex linguistic statement. The results of Experiment 2, however, do speak against this possibility, as children never actually observe a more complex visual stimulus, and thus would have to construct a more abstract relational mapping than in Experiment 1, yet show no difference in responses. The training given to children in both experiments also suggests that they are making a literal mapping between the objects and their insides, as opposed to an abstract, relational one.

It is also possible that children in the two experiments may have relied on a strategy akin to frequency matching by assuming that variable effects are more rare and, therefore, more likely to correspond to comparatively rare objects with unfamiliar linguistic markers (like the novel labels children heard). However, this sort of frequency matching account seems insufficient to explain a particular aspect of the observed data, namely the developmental difference. Three-year-olds can understand analogical mappings in a variety of domains, including space (e.g., Lowenstein & Gentner, 2005) and causality (Bullock et al., 1982; Goswami & Brown, 1990). Although none of these investigations test this concern directly, they all suggest that 3-year-olds would potentially have the capacity to recognize this relational mapping. In order to explain the observed developmental difference, the frequency matching account requires additional claims regarding how this kind of strategy is acquired. In light of the previously reviewed literature on children's developing knowledge of causation and mechanism, we feel that the most supported explanation of this developmental difference is that 4-year-olds are developing an appreciation of the relation between variable causal efficacy and mechanistic complexity. That said, to address these concerns fully, a subsequent investigation could equate the language used to describe the simple and complex insides (both in terms of length and familiarity). We would expect similar results to what we have reported here.

Demonstrating that preschoolers can use information regarding variability to infer the complexity of an object's underlying mechanism has potential implications for our understanding of how children represent causal information. For instance, causal representations can be understood as capturing purely statistical (Cheng, 1997) or counterfactual (Lewis, 1973) information. Other theories have posited that causal representations include links between causes and effects in the form of force or energy transmission (Shultz, 1982; Wolff, 2007). These theories differ in the degree to which they feature a role for mechanism information. The present study suggests that mechanisms become increasingly relevant to children's causal reasoning during the preschool years, supporting the conclusion that children's representations of causal structure are mechanism-laden. We would argue that the present data supports various hypotheses that posit children's causal inferences integrate mechanistic knowledge, which is developing throughout childhood with more formal principles that govern causal inference, which might be available to the child very early (e.g., Gopnik, Glymour, Sobel, Schulz, Kushnir, & Danks, 2004; Sobel & Munro, 2009).

To conclude, the current study attempts to gain traction on understanding the development of a complex form of causal inference. For this reason, the task presented to children was intentionally simplistic. While the results of Experiments 1 and 2 are supportive of a developing appreciation of the relation between variability and complexity, it should be noted that this relation is likely not as straightforward as the present study may suggest. For example, there may be instances in which we would expect a more complex mechanism to result in more stable effects, as in systems that feature multiple failsafe mechanisms. Additionally, our operationalization of mechanistic complexity was based solely on the number of components purportedly featured within a causal system. However, complexity can take many other forms. For instance, a system with few components may support greater functional complexity than a system with a greater number of components based upon, for example, the arrangement or functional connectivity of the components. Further research is needed to explore such possibilities more fully, and integrate them with descriptions of the means by which children engage in causal inference.

Acknowledgements

This research was supported by NSF 1223777 to DMS. A portion of this research was part of the second author's doctoral dissertation. We would like to thank all of the parents and children who participated in this research. We would also like to thank Hannah Boettcher, Brianna Doherty, Julia Franckh, Kathryn Green, Naomi Heilweil, Dannielle John Whiley, Natalie Kleeman, Rebecca Specking, Katherine Williams, Andrea Wister, and Rachel Zolno for their help with data collection and analysis.

References

Ahn, W. K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition, 54*(3), 299–352.
Buchanan, D. W., & Sobel, D. M. (2011). Mechanism-based causal reasoning in young children. *Child Development, 82*(6), 2053–2066.
Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York: Academic Press.
Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*(2), 367–405.
Cook, C., & Sobel, D. M. (2011). Children's beliefs about the fantasy/reality status of hypothesized machines. *Developmental Science, 14*(1), 1–8.
Das Gupta, P., & Bryant, P. E. (1989). Young children's causal inferences. *Child Development, 60*(5), 1138–1146.
Gelman, S. A., & Wellman, H. M. (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), 3.
Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology, 37*(5), 620.

Goswami, U., & Brown, A. L. (1990). Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition, 35*(1), 69–95.

Gottfried, G. M., & Gelman, S. A. (2005). Developing domain-specific causal-explanatory frameworks: The role of insides and immanence. *Cognitive Development, 20*, 137–158.

Lewis, D. (1973). Causation. *The Journal of Philosophy, 70*(17), 556–567.

Lowenstein, J., & Gentner, D. (2005). Relational language and the development of relational mapping. *Cognitive Psychology, 50*(4), 315–351.

Madole, K. L., & Cohen, L. B. (1995). The role of object parts in infants' attention to form-function correlations. *Developmental Psychology, 31*(4), 637.

Mills, C. M., & Keil, F. C. (2004). Knowing the limits of one's understanding: The development of an awareness of an illusion of explanatory depth. *Journal of Experimental Child Psychology, 87*(1), 1–32.

Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science, 3*, 389–396.

Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26*(5), 521–562.

Schulz, L. E., Goodman, N. D., Tenenbaum, J. B., & Jenkins, A. C. (2008). Going beyond the evidence: Abstract laws and preschoolers' responses to anomalous data. *Cognition, 109*(2), 211–223.

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40*(2), 162.

Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47*(1), 1–51.

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42*(6), 1103.

Sobel, D. M., & Munro, S. A. (2009). Domain generality and specificity in children's causal inferences about ambiguous data. *Developmental Psychology, 45*, 511–524.

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*(3), 303–333.

Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N., & Blumenthal, E. J. (2007). The blicket within: Preschoolers' inferences about insides and causes. *Journal of Cognition and Development, 8*(2), 159–182.

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology, 136*(1), 82–111.