Young Children's Attribution of Causality under Uncertainty

Emma L. Johnstone (E.L.Johnstone@warwick.ac.uk)

Michelle R. Ellefson (M.Ellefson@warwick.ac.uk)

Nick Chater (N.Chater @warwick.ac.uk)
Institute for Applied Cognitive Science
Department of Psychology
University of Warwick
Coventry, CV4 7AL UK

Abstract

In previous studies, causal contingencies have been suggested to play an important role in causality judgments. However, little is known about how children might use causal contingencies to inform their judgments of causality, especially under uncertainty. In the current study, we found that young children are sensitive to both conditional and unconditional causal contingencies. Furthermore, children's performance was remarkably similar to the performance of adults in previous studies. Children's sensitivity to statistical properties is a robust finding across a number of cognitive domains and might reflect an application of a fundamental cognitive mechanism that directs cognition, more generally, rather than only causal knowledge acquisition.

Introduction

A fundamental question in human cognition is how people learn dependencies between events. These dependencies may range from the connection between light and food in an animal conditioning experiment, to the acquisition of experienced-based causal knowledge by developing children, to the creation of parsimonious scientific theories by scientists and statisticians. Questions concerning how such dependencies are learned, how far such learning is justified by normative principles, and how far different types of dependency learning, by child, adult, or animal, are analogous, and have been a central focus of cognitive science for more than a century.

Dependencies come in distinct types: contingency and causality. The contingency between Event A and Event B typically is viewed as determined by the joint probabilities of different combinations of A/$\neg$A and B/$\neg$B, and perhaps joint probabilities with other events, as we discuss later.[1] Which exact function of these joint probabilities produces the psychologically appropriate measure of contingency has been a subject of much debate. A reasonable case may be made for a wide range of measures, including functions of $P(B|A)- P(B)$, $P(B|A) - P(B|\neg A)$, $P(B|A)/ P(B|A)$, $P(B|A)/ P(B|\neg A)$, among others. Here, we shall focus on the difference between the probability of B, in the presence of A, and the probability of B, in the absence of A, $P(B|A) - P(B|\neg A)$. This measure, ΔP, is used widely in the study of dependency learning.

Causality explains how much the intervention necessary to bring about Event A leads to a change in the probability of Event B?[2] As Pearl (2000) points out, questions of causal dependency go beyond what is determined by the joint distribution. For example, the joint distribution of two perfectly correlated variables gives no clue that one causes the other. For example, a thermometer may be correlated perfectly with air temperature. However, the joint distribution of temperature and mercury levels does not capture the fact that, if the air temperature rises, the mercury would rise, whereas if the mercury rises, e.g., by warming the bulb, the air temperature remains unchanged.

Following the voluminous philosophical and technical literature (e.g., Mackie, 1974), it can be argued that causality is determined by interventions. However, in many contexts, animals, children, and scientists alike must infer causality though observation only. In classical conditioning, the animal cannot intervene to bring about the light or the food. Similarly, in astronomy, geology, and many social sciences, scientists' have limited capacity to intervene. Hume's (1739) reaction to this state of affairs was to reject the concept of causality entirely. An alternative would be to limit the power of inferences from observation to causation, subject to certain assumptions. A rich technical literature has begun to address the conditions under which such inferences can be made reasonably (Pearl, 2000; Glymour, Scheines, Spirtes, & Kelly, 1987). In this paper, we consider a very simple case. Here, background causal knowledge specifies a candidate causal structure; leaving the question of whether particular causes are facilitatory, inhibitory, or simply irrelevant.

---

[1] The joint probabilities are simply, in this case, $P(A, B)$, $P(A, \neg B)$, $P(\neg A, B)$, $P(\neg A, \neg B)$. More generally, in a set of random variables, $X_1,\ldots, X_n$, the joint probability distribution is $P(X_1,\ldots, X_n)$. As we shall see below, the contingency between X_i may X_j depend on other X_k where $i \neq j \neq k$.

[2] Recently, causal dependency has been characterized mathematically by Pearl (2000), by the introduction of the *do* operator. Our discussion is based on Pearl's analysis, but we avoid formal details here.

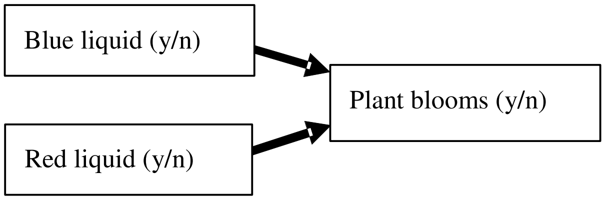

Figure 1. Presumed causal structure in Spellman's (1996) Conditions 1 to 3.

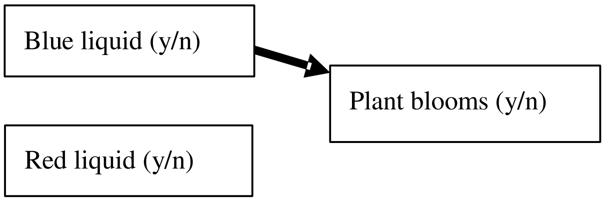

Figure 2. Presumed causal structure in Spellman's (1996) Conditions 4 to 6.

Consider the following causal scenario used by Spellman (1996; for a second related example see also Baker, Mercier, Vallée-Tourangeau, Frank, & Pan, 1993). Two fertilizers, one red and one blue, are tested to see if they help a plant to bloom. Background knowledge tells us that fertilizer is likely to influence the plant; the plant will not influence the fertilizer; and, presumably, the fertilizers will not influence each other causally. The structure of these dependencies, taken from Spellman (1996), is represented in Figure 1.

To see how causality and contingency can dissociate, consider the following cases. First, suppose that both the red and blue liquids cause the plant to be more likely to bloom and that the presence of the liquids is not independent. Specifically, when blue is used, then red is used half the time. When blue is not used, then red is used all the time. Thus, the plant may bloom often when blue is not used, because of red's presence. The numbers can be arranged so that even though blue is causally efficacious, ΔP = 0. Intuitively, this means that if we *learn* that a particular plant was given the blue liquid, then it does not alter our expectation about whether it will bloom. Although blue causes increased blooming, the absence of blue indicates the presence of red, which also causes increased blooming.[3] Nonetheless, you should *use* the blue liquid, given a chance to do so, because of its positive causal power.

Keeping the overall contingencies constant, it is possible to arrange the opposite pattern. Here, the presence of blue impedes blooming. However, when blue is present, then red is typically present, masking blue's effect and leading to ΔP = 0 (Table 1, condition 3). Alternatively, the case arises where blue has no causal impact, and ΔP = 0 (Table 1, condition 2). Thus, very different causal powers of blue may be associated with the same level of contingency.

The opposite case is also possible. Again following Spellman (1996), consider a scenario where the presence of blue reliably causes a plant to bloom, where it does not bloom otherwise. Therefore, red is causally irrelevant (Figure 2). But if red is positively correlated, uncorrelated, or negatively correlated with blue, then the ΔP for red can be positive, zero or negative, respectively.

Spellman's experimental set-up provides an efficient method for examining sensitivity to causal or contingency information during a causal judgment task. However, her discussion has a slightly different focus. She highlights the question of whether people focus on *conditional* contingencies, i.e. controlling for blue, with $\Delta P = P(\text{blooms}|\text{red, blue}) - P(\text{blooms}|\neg\text{red, blue})$ and $\Delta P = P(\text{blooms}|\text{red}, \neg\text{blue}) - P(\text{blooms}|\neg\text{red}, \neg\text{blue})$, rather than focusing on the *unconditional* conditional contingency, i.e., $\Delta P = P(\text{blooms}|\text{red}) - P(\text{blooms}|\neg\text{red})$. Here, we adopt the perspective that these conditional contingencies are indications of underlying causality only, and that understanding causality is the real objective of the cognitive system in everyday understanding of the world.[4]

The dominant focus of this contingency research has been with adults. However, little is known about children's sensitivity to causal contingencies. Earlier studies have indicated that children have an appropriate understanding of causal structure that is more sophisticated than described initially by Piaget (e.g., 1930), extending across both physical (e.g., Bullock, Gelman, Baillargeon, 1982; Spelke, Breinlinger, Macomber, & Jacobson, 1992) and biological domains (e.g., Gelman & Wellman, 1991; Kalish, 1996). Young children are able to make appropriate predictions of outcomes based on single causes and generate appropriate explanations regarding how single causes determine effects (e.g., Hickling & Wellman, 2001). Furthermore, studies by Gopnik and colleagues have indicated that young children can learn and understand the role of single and multiple causes, enabling them to make accurate causal predictions and effective interventions to modify outcomes.

For example, Gopnik and Sobel (2000) showed children a "blicket detector". When a "blicket" was placed on top of the blicket detector, it lit up and played a song. Blickets were various types of blocks, although not all potential blocks were blickets. The results of multiple experiments indicated that children were able to understand the causal strength of the blickets. Further studies have indicated that children could effectively activate the detector by placing a blicket on top of it and deactivate the detector by removing a blicket (e.g., Gopnik, Sobel, Schulz, & Glymour, 2001).

[3] A slight modification of these numbers can lead to blue having a negative contingency with blooming, even though it has a causally positive effect. generating Simpson's (1951) paradox, as elegantly explained in a causal framework by Pearl (2000), and attracting considerable interest in cognitive psychology (e.g., Hintzman, 1980).

[4] Contingency information, whether conditional or not, is never a completely reliable indication of causality. This is because, among other things, additional unknown causes remain possible. These unknown causes correlate with the observable causes in the ways that we have discussed, potentially leading to a disconnection between causality and contingency.

Gopnik et al. (2001) used the detector to investigate children's attribution of causality, given one versus two potential causes. In the "one-cause condition", children saw each block individually and together on the detector. Individual presentations each occurred once and the joint presentation occurred twice. The first block activated the detector, the second block did not activate the detector, and both blocks together activated the detector. Thus, the first block activated the detector on three of the three trials where it was presented and the second block activated the detector on two of the three trials where it was presented.

In the "two-cause condition", children saw each block presented individually. The first block was presented three times and activated the detector on each presentation. The second block was presented three times and activated the detector on two of the three presentations. Again, the first block activated the detector on three of the three trials where it was presented and the second block activated the detector on two of the three trials where it was presented.

The results indicated that children gave the first block a stronger causal rating than the second block in the one-cause condition. In contrast, in the two-cause condition, the children gave the two blocks similar causal ratings. These results suggest that young children are sensitive to the relative contribution of two causes, i.e., their causal contingencies.

It may be that Gopnik et al.'s (2001) causality judgments were based on the effect of individual presentation and not on the effect of dual presentations. Therefore, Sobel, Tenebaum, and Gopnik (2004) further investigated children's use of causal contingencies to make accurate interventions, i.e., activating or deactivating the detector. In both the "inference condition" and the "backwards blocking condition", children saw the detector activate after both blocks were presented together. Next, in both conditions, children saw one block presented alone. In the inference condition, the detector did not activate. In the backwards blocking condition, the detector did activate. In both conditions, children considered one block to be a blicket and one block not to be a blicket. In the inference condition the blicket was the block that was not presented alone. Conversely, in the backwards blocking condition, the blicket was the block that was presented alone. These results suggested that children might infer independent causal strength when dealing with multiple causes, indicating that children are sensitive to conditional contingencies.

These and other studies have been able to address a number of questions regarding the nature of children's sensitivity to causal relations with single and multiple cues and how children use this knowledge to accurately predict and modify outcomes. However, there are a number of questions that remain unanswered. For example, how do children deal with uncertainty? Gopnik et al. (2001) began to address the influence of children's causal ratings when one of the blocks produced different effects. However, little is known about how children will respond when both blocks produce different effects on different trials. Using a fusion of Gopnik and colleagues' blicket detector and Spellman's (1996) conditions, we can assess how children might deal with uncertainty in multiple causes. Using this paradigm, we can begin to explore children's relative sensitivity to conditional and unconditional contingencies. To have a complete comparison of conditional and unconditional contingencies, participants need exposure to each block alone, the blocks together, and no blocks. This creates a situation where children are exposed to apparently unexplainable activation effects in the absence of causal cues, as well as uncertain outcomes of regular causal cues over a series of trials; allowing for an exploration of the consistency of children's causal predictions given additional experience.

Based on the consistency of the causal judgments of adults and young children, as found in the work by Gopnik and colleagues (e.g., Gopnik, Glymour, Sobel, Schulz, Kushnir, & Danks, 2004), we expected that young children in this experiment would perform similarly to the adults of Spellman (1996). When considering two potential causes, we expected that young children's causal predictions and causal ratings of the strength of these causes would be sensitive to changes in the relative strength of the conditional and unconditional probabilities of each of these causes. Second, we expected that children's causal predictions would be dependent on the certainty of the outcome. Specifically, when the outcome is consistent, children's causal predictions would reflect the actual outcome, whereas when the outcome is inconsistent, children will make more positive predictions based on children's tendency to perseverate to the more salient positive predictions.

Method

Participants

One-hundred-fifty student volunteers were recruited from local middle-class primary schools. The mean age was 5.72 years (SD = 0.88, range = 4.16 to 7.40). There were 78 females and 72 males. Twenty-five children were randomly assigned to each condition, with the range and distribution of children's ages and sex similar across the six conditions.

Materials

The stimuli were computerized versions of Gopnik and colleagues' blicket detector. The detector was used in this study to diminish the impact of any prior knowledge of causal structures. In this version, the detector was a gray box. The detector activated when certain blocks were placed on top of it. There were two different color blocks in each condition (Block A and Block B).

Design and Procedures

The experiment was conducted with individual children in a quiet space outside the students' normal classroom. The experiment lasted about 10 minutes and consisted of one block of 40 trials, presented randomly on a laptop computer.

Table 1. The trial presentations and contingencies for conditions 1 to 6.

	Condition					
	1	*2*	*3*	*4*	*5*	*6*
Trial Presentation[1]						
A , B	5/5	8/10[2]	10/15	15/15	10/10	5/5
A, ¬B	10/15	8/10[2]	5/5	0/5	0/10	0/15
¬A, B	5/15	2/10[2]	0/5	5/5	10/10	15/15
¬A, ¬B	0/5	2/10[2]	5/15	0/15	0/10	0/5
Contingencies						
A Unconditional	.5	.6	.5	.5	0	-.5
A Conditional	.67	.6	.67	0	0	0
B Unconditional	0	0	0	1	1	1
B Conditional	.33	0	-.33	1	1	1

Note: [1] Number of trials where machine was activated over number of presentations
[2] This is a slight modification of Spellman's (1996) trial presentations to accommodate a lesser number of overall trials.

The distribution of the trials for each condition was a slight modification of Spellman (1996) and is listed in Table 1. Spellman presented a total of 80 trials, which allowed for an exact consistency of conditional probabilities across Conditions 1 to 3. Here, the number of trials was decreased to 40 to prevent fatigue. This required the proportion of trials within each cell for Condition 2 to be modified slightly.

At the beginning of the experiment, children were shown the detector and told, 'This is a blicket machine'. Next they were shown an array of blocks and told 'Blickets make the blicket machine light up and play a song. Some of these may be blickets.' During each trial, children were asked to predict whether or not the detector would activate, based on the visual presentation: Block A alone, Block B alone, Block A and B, and no blocks. The position of the block was consistent across trials to prevent the children from using position as a cue to whether or not the detector activated. Block A was always on the left side and Block B was always on the right side. After each prediction, children would see whether or not the detector activated.

At the end of the experiment, children were asked to indicate whether or not each block, presented randomly, was a blicket. The seven possible options, given a score on a 7-point blicket rating scale were:

1 - *Always NOT a blicket.*
2 - *Mostly NOT a blicket.*
3 - *Sometimes NOT a blicket.*
4 - *Equally both a blicket and not a blicket.*
5 - *Sometimes a blicket.*
6 - *Mostly a blicket.*
7 - *Always a blicket.*

This rating scale was used instead of Spellman's (1996) scale because a range of –100 to 100 would have been too difficult for children. However, we used terms similar to Spellman's in determining these options.

Results

Analyses

The data were screened to ensure a normal distribution and to identify outliers. The mean proportion of positive predictions and blicket ratings were subjected to mixed model repeated measures ANOVAs (Block Type × Condition). Condition was manipulated between participants and Block Type was manipulated within-subjects. Planned comparisons were conducted for the hypothesized main effects and interactions described earlier, with Bonferroni correction procedures applied to control for Type I error. Unless otherwise stated, all main effects, interactions, and planned comparisons were tested using an α of .05.

Causal Predictions

Overall, the mean proportion of 'Yes' predictions supported our hypotheses. There was a significant main effect of Block Type, F (3,144) = 242.96, *MSE* = .356, a main effect of Condition, F (5,144) = 4.38, and a Block Type × Condition interaction, F (15,144) = 23.00 (see Figure 3). Further analyses indicated that the main effect of Condition was a result of children responding 'Yes' more often to Conditions 1 to 3 than to Conditions 4 to 6, F (1,149) = 17.82. Overall, children responded 'Yes' most often when both blocks were presented, followed by Block A alone, F (1,144) = 37.85, then Block B alone, F (1,144) = 47.49, then no blocks, F (1,144) = 381.93. The significant interaction indicated that when both blocks were presented, children responded 'Yes' more often in Conditions 1 to 3 than in Conditions 4 to 6.

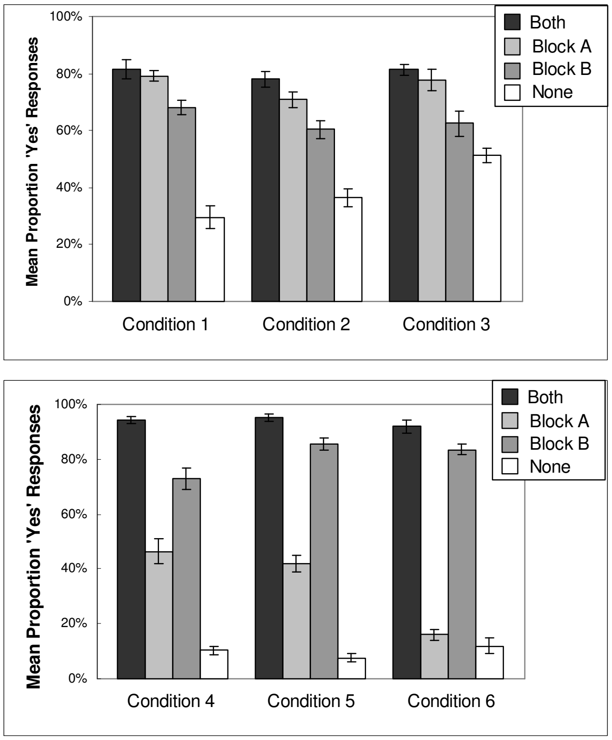

Figure 3. The mean proportion of 'Yes' responses.

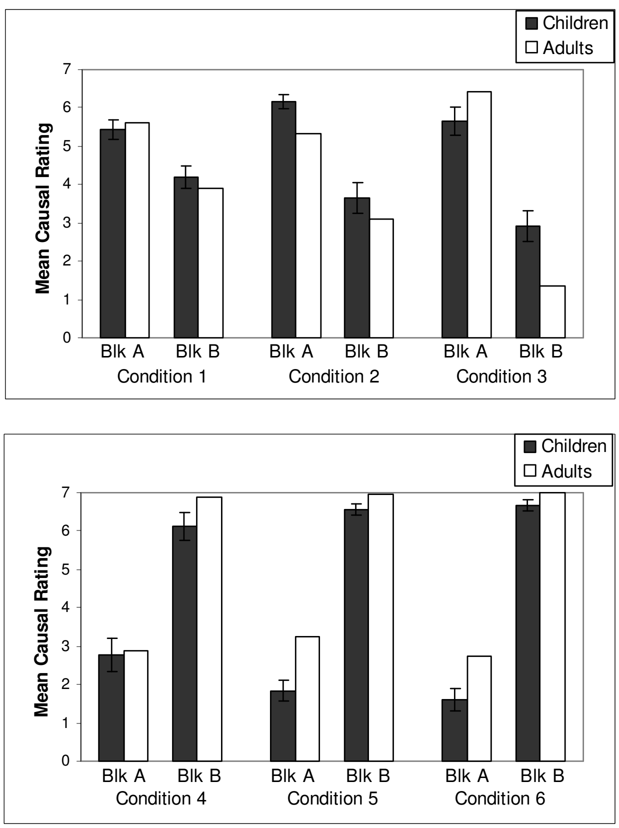

Figure 4. The mean causal ratings for Block A and Block B, as compared to the adults in Spellman (1996).

Children responded 'Yes' more often to Block A than to Block B in Conditions 1 to 3. Conversely, they responded 'Yes' more often to Block B than to Block A in Conditions 4 to 6. When no blocks were presented, children responded 'Yes' more often in Conditions 1 to 3 than in Conditions 4 to 6.

Causal Ratings

For the most part, the causal ratings supported our hypotheses. There was a significant main effect of Block Type, $F(1,288) = 36.63$, $MSE = 1.55$, and Block Type × Condition interaction, $F(5,288) = 69.84$ (see Figure 4). Overall, children gave Block B a higher causal rating than Block A. Planned comparisons indicated that the causal ratings of Block B were larger in Condition 1 than in Condition 3, with no significant differences in the causal ratings of Block A across Conditions 1 to 3. The causal ratings of Block A were larger in Condition 4 than in Condition 5 and Condition 6, with no significant differences in the causal ratings of Block B across Conditions 4 to 6. We transformed the results of Spellman (1996) to fit our 7-point rating scale. The results of the young children in our experiment are remarkably similar to the Spellman's results, with the exception of Block A across Conditions 4 to 6. Here, children's causal ratings decrease with decreased unconditional contingencies, whereas Spellman's adults showed consistent causal ratings despite decreases in unconditional contingencies.

Discussion

Previous research suggests that adults use only conditional contingencies (e.g., Spellman, 1996). Here, we found that children used both conditional and unconditional contingencies to inform their attributions of causal strength.

The results confirmed our hypotheses about children's use of conditional contingencies. When faced with uncertainty, as in Conditions 1 to 3, children made more positive predictions than during consistent outcomes, as in Conditions 4 to 6. Specifically, children's proportion of 'Yes' responses in Conditions 4 to 6 reflected their accurate learning of the outcomes. The greater proportion of 'Yes' responses in the uncertain Conditions 1 to 3 appeared to indicate that children might be optimizing, predicting 'Yes' with a greater frequency. Across the conditions, the proportion of 'Yes' responses was higher when both blocks were on the detector than when either block was presented individually, even under uncertainty. This may suggest an additive effect of presenting both blocks over individual blocks, with children predicting that the detector *must* activate if both blocks were present. However, additional studies designed to specifically address this issue must be conducted before we can make conclusions about additive effects.

An obvious difference between our results and Spellman's (1996) results occur in Conditions 4 to 6. Here, the children appear to be sensitive to changes in unconditional contingencies, as evidenced by the decrease in their causal ratings of Block B across Conditions 4 to 6. Specifically, when conditional contingencies vary alongside consistent unconditional contingencies, then causality judgments were sensitive to changes in these conditional contingencies. Conversely, when unconditional contingencies vary alongside consistent conditional contingencies, then causality judgments were sensitive to the changes in these unconditional contingencies. It may be the case that children benefit from attention to unconditional contingencies in a manner that adults do not. When judging causal strength, unconditional contingency information may be irrelevant to adults, but used by young children. Regardless of how children treat the causal efficacy of unconditional contingencies, it is clear from the results here, that young children are sensitive to and utilize both types of contingency in their causal judgments, even under uncertainty. This provides support for the notion that children are able to simultaneously consider multiple cues when learning about the causal strength of various potential causes.

The current study addresses the sensitivity of children's causal predictions and causal ratings to relative alterations in conditional and unconditional contingencies under uncertainty. However, we have yet to address whether holding these contingencies equal and altering the relative frequency of presentations results in performance changes. Goedert and Spellman (in press) found that adults are sensitive to frequency alterations. However, this sensitivity appears to be dependent on the framing of the task.

The current sample of children ranged from 4- to 7-year-olds. Children within this age range may approach this task differently. Therefore, further studies should address whether there are quantitative differences in the causal predictions and causal ratings of children of different ages.

In summary, the results here suggest that children use statistical dependencies like causal contingencies to inform their judgments of causality, to determine the relative strength of multiple causes, and to discern uncertainty. Children have remarkable sensitivity to statistical dependencies during language acquisition (e.g., Thiessen & Saffran, 2003) learning complex structures (e.g., Ellefson, Young, Christiansen, & Espy, 2005) and memory (e.g., Kirkham, Slemmer, Johnson, 2002); and they effectively integrate statistical information from multiple cues (e.g., Christiansen & Monaghan, in press). If attending to statistical properties is a fundamental cognitive skill, then it is entirely reasonable that children would find statistical dependencies useful in determining the causal efficacy of potential causes.

References

Baker, A.G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *JEP:LMC*, *19*, 414-432.

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.

Christiansen, M.H. & Monaghan, P. (*in press*). Discovering verbs through multiple-cue integration. In K. Hirsh-Pasek & R.M. Golinkoff (Eds.), *Action meets words: How children learn verbs*. New York: Oxford University Press.

Ellefson, M.R., Young, M.E., Christiansen, M.H., & Espy, K.A. (2005). *The Utility of Structural Cues to Facilitate Student Learning of Complex Information*. Manuscript submitted for publication.

Gelman, S.A., & Wellman, H.M. (1991). Insides and essence: Early understandings of the non-obvious. *Cognition*, *38*, 213-234.

Glymour, C., Scheines, R., Spirtes, P., & Kelly, K. (1987). *Discovering causal structure: Artificial intelligence, philosophy of science, and statistical modeling*. San Diego, CA: Academic Press.

Goedert, K.M., & Spellman, B.A. (*in press*). Non-normative discounting: There is more to cue interaction effects. *Learning and Behavior.*

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

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. *Psych. Rev.*, *111*, 3-32.

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. *Dev. Psych.*, *37*, 620-629.

Hickling, A.K., & Wellman, H.M. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Dev. Psych.*, *37*, 668-683.

Hintzman, D. L. (1980). Simpson's paradox and the analysis of memory retrieval. *Psych. Rev.*, 87, 398-410.

Hume, D. (1739). *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*.

Kalish, C.W. (1996). Preschoolers' understanding of germs as invisible mechanisms. *Cog. Dev.*, *11*, 83-106.

Kirkham, N.Z, Slemmer, J.A., & Johnson, S.P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*, B35-B42

Mackie, J. (1974) *The Cement of the Universe*. Oxford, UK: Clarendon Press.

Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. New York: Cambridge University Press.

Piaget, J. (1930). *The child's conception of physical causality*. New York: Harcourt, Brace.

Simpson, E.H. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society (Series B)*, *13*, 238-241.

Sobel, D.M., Tenenbaum, J.B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking in Bayesian reasoning in preschoolers. *Cognitive Science*, *28*, 303-333.

Spellman, B.A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psych. Science*, *7*, 337-342.

Thiessen, E.D., & Saffran, J.R. (2003). When cues collide: Use of stress and statistical cues to word boundaries by 7- to 9-month-old infants. *Dev. Psych.*, *39*, 706-716