

BAYESIAN SPECIAL SECTION

Bayes nets and babies: infants' developing statistical reasoning abilities and their representation of causal knowledge

David M. Sobel¹ and Natasha Z. Kirkham²

1. Department of Cognitive and Linguistic Sciences, Brown University, USA

2. Psychology Department, Stanford University, USA

Abstract

A fundamental assumption of the causal graphical model framework is the Markov assumption, which posits that learners can discriminate between two events that are dependent because of a direct causal relation between them and two events that are independent conditional on the value of another event(s). Sobel and Kirkham (2006) demonstrated that 8-month-old infants registered conditional independence information among a sequence of events; infants responded according to the Markov assumption in such a way that was inconsistent with models that rely on simple calculations of associative strength. The present experiment extends these findings to younger infants, and demonstrates that such responses potentially develop during the second half of the first year of life. These data are discussed in terms of a developmental trajectory between associative mechanisms and causal graphical models as representations of infants' causal and statistical learning.

Introduction

Our understanding of the world depends on organizing data into a representation of causal structure. By the time children reach kindergarten, they can clearly engage in predictive inference, generate explanations of events, and reason counterfactually (e.g. Bullock, Gelman & Baillargeon, 1982; Harris, German & Mills, 1996; Schult & Wellman, 1997). Even by the end of the first year of life, infants register causal properties of physical objects, including containment, support, and contact (e.g. Hespos & Baillargeon, 2001; Leslie & Keeble, 1987; Needham & Baillargeon, 1993; Spelke, Breinlinger, Macomber & Jacobson, 1992). Very young children also register that an agent's actions are goal-directed and parse intentional actions in sequences of human action (e.g. Baldwin, Baird, Saylor & Clark, 2001; Gergeley, Nadasdy, Csibra & Biro, 1995; Meltzoff, 1995; Scaife & Bruner, 1975; Woodward, 1998).

These data suggest that infants possess a great deal of causal knowledge about objects, agents, and events. What is not definitively known is how such knowledge is represented. A growing number of psychologists have considered causal graphical models (following Pearl, 2000; Spirtes, Glymour & Scheines, 2001) as a computational

description of how children (and adults) represent causal knowledge (see e.g. Gopnik, Glymour, Sobel, Schulz, Kushnir & Danks, 2004; Tenenbaum & Griffiths, 2003). A causal graphical model defines a probability distribution over a set of variables, based on a graphical structure. Variables or events are nodes in the graph and edges indicate causal relations. A critical assumption that underlies these models is the *Markov assumption*, which states that certain conditional independence relations must exist among variables in the graph. Specifically, each variable is expected to be independent of all other variables except its effects (its descendants) conditional on its direct causes (its parents).¹ To advocate that this computational framework describes children's causal knowledge, we

¹ A specific example occurs in any causal model that contains a chain of events $A \rightarrow B \rightarrow C$. Events A and B are dependent upon each other, as are events B and C, as are events A and C. The Markov assumption states that events A and C are independent given the value of event B. Event C becomes independent of all other variables in the graph, conditional on its parent (event B). This can be seen intuitively by considering that the only way that event A has any causal influence on event C is through event B, so if you know the value of B and want to predict the value of C also knowing the value of event A provides no further information. We would refer the reader to Gopnik *et al.* (2004) for a more detailed and accessible description of the Markov assumption. Pearl (2000) provides a more in-depth description of this topic.

Address for correspondence: D. Sobel, Department of Cognitive and Linguistic Sciences, Box 1978, Brown University, Providence, RI 02912, USA; e-mail: sobel@cog.brown.edu

would want to show that children can recognize cases of conditional independence and dependence among events.

Do children (and adults) recognize the difference between cases where two events are always dependent and cases where two events are dependent, but independent conditioned on another event? Adults have little difficulty discriminating between these circumstances (e.g. Shanks, 1985; Spellman, 1996), but what about young children? Gopnik, Sobel, Schulz and Glymour (2001) examined this question using the 'blicket detector' – a machine that lights up and plays music when certain objects, controlled by the experimenter, are placed upon it. The detector presents a novel causal property of objects and can be used to examine children's causal inferences in the absence of individual differences in children's prior knowledge that might influence performance. Three- and 4-year-olds were told that 'blickets make the machine go', and were then shown two objects (A and B). Object A activated the machine by itself and object B did not. Then both objects were placed on the machine together, which activated (demonstrated twice). Critically, object B activates the detector more often than not, but it only does so dependent on the presence of object A. Therefore, B should not be judged a cause, even though it is associated with the detector's activation. Children reasoned in this manner: they judged that only object A was a blicket. Schulz and Gopnik (2004) also demonstrated that preschoolers made such inferences across different domains, not just about physical objects and machines.

There are, however, many accounts of causal learning that are consistent with these data. This procedure is similar to 'blocking', a phenomenon from the animal conditioning literature (Kamin, 1969), which various models of associative strength were designed to explain (e.g. Rescorla & Wagner, 1972). Models of human causal learning that simulate these calculations (e.g. Cramer, Weiss, Williams, Reid, Nieri & Manning-Ryan, 2002; Dickinson, 2001) easily account for these data, as do a variety of causal learning mechanisms that generate calculations of causal strength based on modifications to such associative learning models (e.g. Kruschke & Blair, 2000; Wasserman & Berglan, 1998; Van Hamme & Wasserman, 1994). Further, there are a variety of formal learning accounts that calculate an estimate of causal strength of a fixed set of causal relations given the frequency with which events co-occur (e.g. Cheng, 1997; Shanks, 1995); these accounts are also consistent with these data.

Sobel, Tenenbaum and Gopnik (2004) provided insight into whether these alternative models account for children's rapid causal learning in the blicket detector paradigm, by considering performance on a backwards blocking procedure, a phenomenon that certain associative models (e.g. Rescorla & Wagner, 1972) find problematic

(see e.g. Shanks, 1985). In this experiment, children were shown the blicket detector and two objects (A and B) that first activated the machine together. In the *one-cause* condition, children then saw that object A did not activate the detector by itself. Object A only activated the detector dependent on the presence of object B, suggesting that object B was the only possible cause. In the *backwards blocking* condition, object A did activate the detector by itself. Critically, the associative strength between object B and the detector's activation is the same between these two conditions, but our intuitions about the causal status of the objects is not. In the one-cause condition, object B should be a blicket. In the backwards blocking condition, object B's efficacy is uncertain. Preschoolers' judgments about whether each object was a blicket were consistent with these intuitions, and critically they categorized object B differently between the conditions. These data suggest that children are not making these inferences based only on recognizing the associative strength among objects and the detector.²

What about younger children? Sobel and Kirkham (2006) extended the one-cause and backwards blocking procedures to 19- and 24-month-olds. Twenty-four-month-olds' inferences were similar to older children's, but 19-month-olds responded no differently from chance (and differed from the older toddlers). However, this procedure required toddlers to put objects on the blicket detector (i.e. make a manual response). There is a long literature suggesting a dissociation between knowledge and action in toddlers across cognitive development (e.g. Hood, Cole-Davies & Dias, 2003; Kirkham, Cruess & Diamond, 2003; Munakata, 2001; Zelazo, Frye & Rapus, 1996). For example, using looking behavior as the dependent measure, several investigations suggest that infants can recognize conditional probability information in sequences of speech and visuo-temporal events (Aslin, Saffran & Newport, 1998; Kirkham, Slemmer & Johnson, 2002; Saffran, Aslin & Newport, 1996). This suggests the possibility that younger infants can recognize conditional independence information if a simpler method is used.

Sobel and Kirkham (2006) used an anticipatory eye-gaze procedure to measure whether infants made similar inferences about sequences of events that paralleled the one-cause and backwards blocking trials with the blicket detector. Eight-month-olds were shown a screen with

² Many of the contemporary learning algorithms mentioned above were designed with the backwards blocking procedure in mind. However, in follow-up studies, Sobel *et al.* (2004; see also Tenenbaum & Griffiths, 2003) proposed a model of causal reasoning based on Bayesian inference over causal graphical models, and demonstrated that 4-year-olds' inferences were more consistent with this approach than any of the alternative models mentioned above. Whether such a mechanism applies developmentally is an open question.

four square frames; two frames (A and B) appeared in the middle, while the other two frames (C and D) were on the left and right side of the screen. Events A, B, C, and D occurred inside their respective frame. Events A and B were just visual events; events C and D were both accompanied by the same appealing sound effect. Infants were first shown a sequence of events in which events A and B appeared in their respective frames together, always followed by the C event. The C event could be followed by the AB compound event or by the D event (randomly determined). Likewise, the D event could be followed by the AB compound or the C event (randomly determined). As such, the only regularity in the familiarization was that the AB events together were always followed by event C.

After infants observed four cases of this pairing, and that no other regularity existed, they were shown one of two disambiguating events. In the *backwards blocking* condition, the A event occurred by itself, followed by the C event. In the *indirect screening-off* condition (parallel to the one-cause condition above), the A event occurred by itself, followed by the D event. Infants in both conditions were then shown a test trial. They observed the B event occur by itself and the screen went blank except for the four frames. The sound effect, which had been paired with both the C and D events, was played while an eye tracker recorded where the infants looked. Their eye movements at this point give insight into what they expect to happen next.

One can think of the sequences of events that infants observed to parallel the blicket detector experiments. In the *indirect screening-off* condition, children observe that two events predict a third. Then, they observe that one of those two events does not predict that third event by itself. This parallels the necessary inference in the one-cause condition of the blicket detector procedure: The A event (i.e. A object) only predicted the C event (i.e. the detector's activation) dependent on the presence of the B event (i.e. B object). As a result, children should predict that event B should be followed by event C. Infants' looking pattern was consistent with this inference: they looked longer to the C frame than the D frame in the *indirect screening-off* condition, and made the opposite response in the *backwards blocking* condition (a significant interaction between frame and condition was found). This suggests that infants can register conditional independence and dependence relations, consistent with the Markov assumption and the causal graphical model framework.

There are two open questions regarding these data. First, does this sensitivity to conditional independence information develop before 8 months of age? The goal of the present investigation is to replicate the Sobel and

Kirkham (2006) procedure on a group of younger infants to examine this question. Second, although the 8-month-old data suggest that infants' *statistical* inferences are consistent with the Markov assumption and are inconsistent with certain models of associative reasoning, the stimuli presented to infants are not necessarily *causal* in nature. The present experiment does not investigate this question, and an open question is whether infants' statistical reasoning bootstraps to a richer representation of causal structure. We will elaborate on this issue further in the discussion.

Experiment

We presented 5-month-olds with the same anticipatory eye-gaze procedure used by Sobel and Kirkham (2006). Infants observed a familiarization sequence of four events, in which two events (A and B) predicted a third event (C), but in which no other event predicted any other. After this familiarization, infants observed that event A alone either predicted event C (*backwards blocking*) or event D (*indirect screening-off*). They were then shown event B, and their eye gaze was measured, as an expectation of what event they believed would occur.

Method

Participants

Twenty-one full-term 5-month-old infants composed the final sample (11 girls, $M = 5.5$ months, range: 4.9–5.9 months, $SD = 0.24$ months). Ten infants were tested in the *indirect screening-off* condition; 11 infants were tested in the *backwards blocking* condition. Two additional infants were observed but not included in the analyses due to fussiness. Infants were recruited by letter and telephone from an established database of parents. Parents and infants received a small gift (a baby T-shirt or toy) for their participation.

Materials

The stimuli were presented on a 152 cm rear projection screen. Infants sat on their caregiver's lap approximately 180 cm away from the screen. An Applied Science Laboratories Model 504 corneal reflection eye-tracking system was used to collect eye movement data as infants were shown the stimuli. The stimuli viewed by the infants were designed using Macromedia Director and imported directly into the eye tracker from a computer. The eye tracker also fed a signal into a mini digital video recorder in the form of crosshairs superimposed on the stimulus.

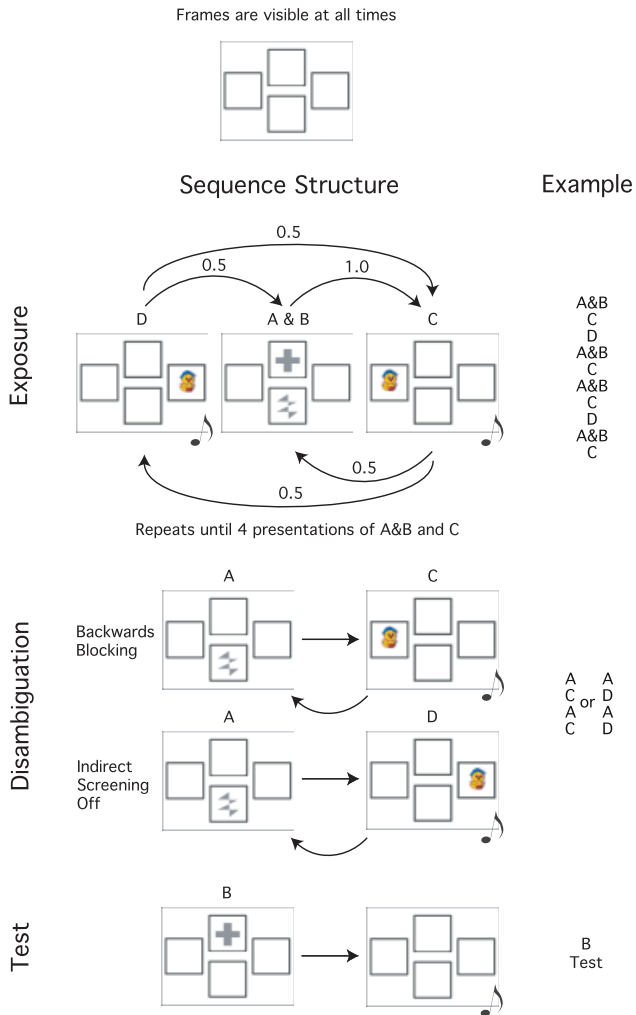


Figure 1 Schematic representation of the experimental design, including what the infant saw at test. Four frames are visible on the screen at all times. Events A, B, C, and D occur in their respective frame. During the exposure phase, infants observed that the AB compound (in the center of the screen) predicts the C event (which occurs on one side of the screen, counterbalanced), but that neither the C nor the D event (the event on the other side of the screen) predicts any event reliably. In the disambiguation phase, infants in the indirect screening-off condition observe that event A alone predicts event D; infants in the backwards blocking condition observe that event A alone predicts event C. Then in the test phase, infants observe the B event, and then the blank screen (with only the four frames visible) paired with the sound effect that has always been paired with both the C and D event. Eye gaze is measured for 8 s from the onset of the sound effect.

Stimuli were presented in four frames, which were present on the screen at all times. Each frame consisted of a box with white outlines presented on a black background. Figure 1 shows a schematic with the experimental set-up,

including the screen with no event present (i.e. what the infant saw at test). The stimuli consisted of three Quick-time movies of moving objects. One movie was of a colorful object (chosen at random from a set of four objects) that moved in time to the same repeated sound (chosen at random from a set of four sounds). This movie could appear in either of the two side frames (C and D) and lasted 8 s. The other two movies (also 8 s in duration) were of two different grey objects that moved in silence. One always appeared in the A frame, the other in the B frame. The intention was to ensure that (1) in each event visual and auditory elements were synchronous, and (2) the C and D events were highly attractive to the infant.

Procedure

Infants were given the same procedure as 8-month-olds in our previous research (Sobel & Kirkham, 2006). One experimenter controlled the eye tracker to capture the infant’s eye movements. A second experimenter controlled the presentation of stimuli to the infant. After an eye-tracker calibration check (described in detail in Sobel & Kirkham, in press), infants were shown a sequence of events on a projection screen. The initial (i.e. exposure) spatiotemporal sequence was designed such that the two central events (A and B) together always predicted the brightly colored musical side event (C). The C side was counterbalanced between subjects. A fourth event (D – the same musical event) occurred on the other side of the screen. The transitional probability between the AB events and the C event was 1.0. All other transitional probabilities among the events were 0.5. Each event lasted for 8 seconds.

After being shown this exposure sequence (which always included four AB→C pairings), infants were shown one of two disambiguating sequences. In the *indirect screening-off* condition, infants observed one of the two central events (A or B, counterbalanced) occur alone followed by the D event. In the *backwards blocking* condition, this event was followed by the C event. After two of these disambiguating sequence pairs, infants were shown a test event: the other event from the AB compound alone, followed by a measurement trial, which consisted of a blank screen (except for the four frames) and the sound that had accompanied the C and D events. Eye movements were recorded for 8 s following the onset of the measurement trial. The block of exposure, disambiguation, and test trials was then repeated. The experiment was an infant-controlled procedure, meaning that the infant could watch the blocks of trials for as long as s/he wanted (i.e. until s/he no longer was attending to the stimuli).

A critical feature of this procedure is that the predictive relation between the central A or B event shown at test

and the C event is the same across the two conditions. It is always shown to predict the C event when paired with the other central event. That other central event was then shown to predict or not predict the C event (in the *backwards blocking* and *indirect screening-off* conditions, respectively). As such, this procedure offers a good measure of whether infants respond on the basis of only calculating the associative relation among events.

Coding

Looking time data were collected by the eye-tracking system automatically. There were no human coders. Each fixation location and look duration was recorded online as the experiment was running. Areas of interest (i.e. the test trial frames) were defined by the experimenter, and the duration of looking times were outputted from those areas.

Results

The dependent variable was amount of time spent looking inside the C and D frames during the test trial. Looking time within the frames, and not just to the absolute side of the screen, was used as a conservative measure of infants' expectations. Because of this conservative measure, looking times were relatively short. Since infant eye position signals are often extremely noisy, much noisier than adult signals, average looking times (rather than fixation counts, normally used in adult eye-tracking experiments) were used as the dependent variable (e.g. Richardson & Kirkham, 2004). If the infant attended to less than 50% of the familiarization events (each event lasting 8 seconds) the test trial from that block was not included in data analysis. Sixty percent of the recorded test trials were included in the analyses. This reflects the nature of our infant-controlled paradigm: We continued to run the experiment until we were sure that the infant was no longer attending, as a consequence of which the last few test trials were always dropped. Attending was defined as fixating on the frame that was currently playing. Infants successfully completed between one and six blocks of trials (median = 3).

Figure 2 shows the mean looking time to the C and D trials across conditions, by experiment (the current experiment and data from Sobel & Kirkham, 2006). A 2 (Gender) \times 2 (Condition: *backwards blocking* vs. *indirect screening-off*) \times 2 (Absolute location of the C event: left vs. right) \times 2 (Frame: C vs. D) ANOVA revealed a significant interaction between condition and frame, $F(1, 13) = 4.72, p = .05$. Post-hoc comparison showed that infants in the *backwards blocking* condition looked

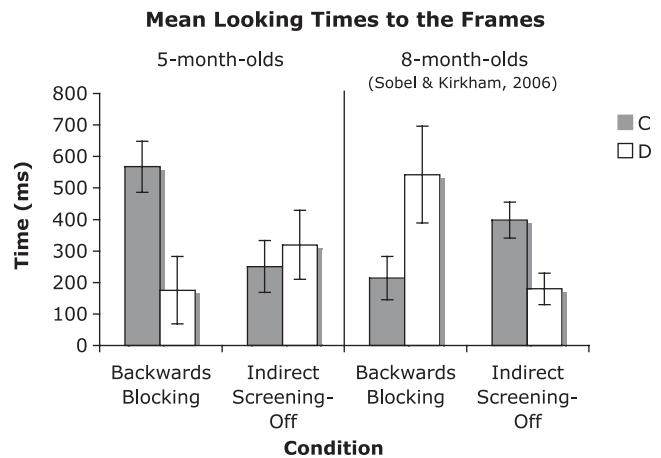


Figure 2 Mean looking times (ms) to frames C and D during the test trials, by condition, and by experiment (5-month-olds from the current experiment, and 8-month-olds from Sobel & Kirkham, 2006).

significantly longer at the C frame ($M = 565$ ms, $SE = 81$ ms) than at the D frame ($M = 173$ ms, $SE = 107$ ms) during the test trials, Tukey's HSD, $p < .05$, whereas infants in the *indirect screening-off* condition showed no significant looking time difference, looking equally at the C frame ($M = 248$ ms; $SE = 82$ ms) and the D frame ($M = 317$ ms; $SE = 109$ ms) during the test trials. In addition, infants looked longer at the C frame in the *backwards blocking* condition than in the *indirect screening-off* condition, Tukey's HSD, $p < .05$.

In Sobel and Kirkham (2006) there was also a significant interaction between condition and frame. Eight-month-olds in the *backwards blocking* condition looked longer at the D frame than at the C frame during the test trials, whereas infants in the *indirect screening-off* condition looked longer at the C frame than at the D frame during the test trials. To consider whether these infants differed from the 5-month-olds, we performed a 2 (Age: 5 months vs. 8 months) \times 2 (Condition: *backwards blocking* vs. *indirect screening-off*) \times 2 (Frame: C vs. D) Analysis of Variance looking across the two experiments. No main effects or two-way interactions were significant, but this analysis revealed a significant three-way interaction between age, condition, and frame, $F(1, 42) = 12.81, p = .001$. This interaction is explained by the developmental difference in how the stimuli were interpreted. Younger infants stared longer at the C frame in backwards blocking, while they showed no discernable difference in indirect screening-off. In contrast, the older infants stared longer at the C frame in the indirect screening-off condition, and tended to stare longer at the D frame in the backwards blocking condition.

Discussion

Previous investigations (Sobel & Kirkham, 2006) demonstrated that 8-month-olds can engage in statistical inference consistent with recognizing conditional independence relations in such a way that is problematic for certain models of associative reasoning (e.g. Rescorla & Wagner, 1972). Five-month-olds did not respond in the same manner as 8-month-olds: they looked longer at the C frame in the backwards blocking condition and showed no looking time difference in the indirect screening-off condition.

A possible interpretation of these data is that younger infants lack the ability to register conditional independence information, and instead responded by calculating the associative strength among the events. However, this possibility on its own seems unlikely, since it would predict a different response pattern than what was actually observed. Specifically, 5-month-olds did not look equally at the C location across the conditions: They looked significantly longer at the C frame in the *backwards blocking* condition than in the *indirect screening-off* condition.

Another possibility is that the 5-month-olds categorized all of the events in the center of the screen (the AB compound, as well as either the A or B event occurring alone) as the same event. If this were the case, associative models could easily explain these data. In the backwards blocking procedure, the disambiguation event (A alone followed by C) only reinforced the predictive nature of the central event. As such, infants treated event B (still parsed as the central event) as more predictive of C than D. In the indirect screening-off procedure, if infants treat the AB event, as well as the A and B events occurring alone, as the same event, then the data they have observed is less predictive overall. This event was followed by event C four times, and then was followed by event D twice. This indicates that this event has less predictive power in this condition, resulting in more equivalent looks to the two locations.

It is certainly possible that the simultaneous occurrence of the two central events (A and B) was problematic for the young infants, and they registered it and the individual presentation of A and B as a single event. Infants' visuo-temporal reasoning abilities do develop over the course of the first year of life (see e.g. Johnson, 2001). Young infants reliably attend to contour and motion in the visual field (Slater, 1995), but might lack the ability to treat two events that occur simultaneously as separate. This could indicate that the present procedure was not sensitive enough to make definitive conclusions about these young infants' reasoning abilities. Associative models could account for these data, but if A and B were not considered separate events, these data are also potentially consistent with the hypothesis that even these young

infants possess a reasoning mechanism that registers conditional independence relations; that is, even these younger infants might respond consistently with the Markov assumption.

There is some indirect support for this possibility: Evidence from Younger and colleagues (Younger, 1990; Younger & Cohen, 1983, 1986; Younger & Gottlieb, 1988) suggests that for many types of correlated events, infants younger than 7 months have some difficulties integrating certain information about correlations among certain object features into representations of object categories. Several computational models focus on the complexity of the stimuli as critical for determining whether infants can recognize the statistical regularity (e.g. Westermann & Mareschal, 2004; see also Younger, Hollich & Furrer, 2004). The stimuli used in the present investigation are complex: They require the infants to integrate motion, sound, and colorful figures together to form a representation of an individual event, and then to integrate those events together in a temporal sequence to form a predictive relation. We suggest that if these stimuli were simplified, 5-month-olds might generate similar responses to the older infants.

These alternatives are not mutually exclusive. It is possible that infants' visuo-temporal reasoning abilities develop between 5 and 8 months, enabling them to interpret the A and B events separately in our experiment, and it is possible that infants' statistical reasoning abilities develop at the same time. Since research in statistical learning has suggested that even 2-month-olds can recognize transitional probabilities among visual events (Kirkham *et al.*, 2002), it is also possible that this early-developing sensitivity to associations between events can bootstrap later and more complex reasoning about conditional independence (and causal relations).

How might we elaborate on this possibility? Clearly, the methods used to investigate these issues are of utmost importance; for instance, experiments must take into account infants' working memory capacity as well as how infants actually perceive the presented stimuli. More generally, recall that an open question with this line of work was whether infants' reasoning about these visual sequences reflected their *causal* knowledge. Do infants perceive the visuo-temporal sequences in this experiment as events that cause one another? One way to consider this question is to examine a similar procedure using real-world events, such as placing objects on a blivet detector. We are currently examining 5- and 8-month-olds using a violation of expectation procedure, in which we can show infants live action scenes of objects producing causal effects, which reflects the kinds of conditional independence relations in this experiment. This might allow us to judge whether infants' reasoning is causal or only statistical.

But another assumption behind causal graphical models is that vertices between nodes specify a form of causal mechanism that supports interventions (Pearl, 2000; Woodward, 2003). As such, it would be interesting to consider whether infants' own interventions reflect an understanding of the causal structure they observe. Such a method is not novel in infant testing: Studies have used infant intervention to measure their understanding of changes in the perceptual environment (e.g. Eimas, Siqueland, Jusczyk & Vigorito, 1971). In the classic high amplitude sucking paradigm (Siqueland & DeLucia, 1969), if the infant desires to see or hear that particular stimulus she will change her sucking behavior accordingly. An important question is whether infants would generate interventions to produce the *predictors* of the desired stimuli.

Consider the present exposure and disambiguation sequences. In the present procedure, suppose that only the C event was interesting to observe (i.e. the D event was just another silent grayscale event). If infants were trained to produce both the central A and B events separately by looking to particular spatial locations prior to observing the sequence, what intervention would they generate after observing either the indirect screening-off or backwards blocking sequence (assuming event A was the disambiguation stimuli)? Would they generate the B event with the same frequency? We hypothesize that infants would be more likely to generate the B event in the *indirect screening-off* condition than in the *backwards blocking* condition. This would suggest that infants register that their own interventions can generate new events. This is an empirical question, one which we are also currently investigating.

We would also like to clarify another concern: More contemporary associative models, like the ones designed with the backwards blocking phenomenon in mind, can explain the response patterns generated by the 8-month-olds. One question is whether they do so rationally: Infants appear capable of making inferences based on a very small sample of data. These models often require large amounts of data to make meaningful estimations. More critically, Sobel *et al.* (2004) found that 4-year-olds' causal inferences involving theblicket detector were best explained by Bayesian inference over causal graphical models (see also Tenenbaum & Griffiths, 2003). This work considered all of these contemporary associative models as contrast cases, and found that 4-year-olds' inferences were inconsistent with these accounts. Sobel and Munro (2006) modified the procedure slightly and found that 3-year-olds' inferences were also consistent with this Bayesian description. In both cases, children made inferences about the backwards blocking procedure based on the base rate of blickets: when blickets were rare, the ambiguous B object was categorized as not a blicket; when blickets were common, this object

was a blicket. We suspect that a similar manipulation could be applied to the visuo-temporal sequence learning paradigm we have used here to examine young infants' reasoning. We are also currently exploring this possibility.

To conclude, although the nature of the observed developmental difference remains uncertain, what we do know is that infants' inferences about conditional independence relations appear to develop between the ages of 5 and 8 months. It is plausible that such development reflects infants' increased visuo-spatial abilities, or their attentional and memory capacities. It is also plausible that such development reflects a more general ability to integrate statistical information together to recognize conditional independence relations, in which infants move from recognizing associations to making inferences that are more causal in nature. Interestingly, researchers have also suggested that infants develop abilities to interpret Michotte-like collision stimuli in more causal ways around these ages (e.g. Leslie & Keeble, 1987; Oakes & Cohen, 1990). Whether infants' developing understanding of the causal nature of these displays is related to their understanding of conditional independence and its role in representing causal knowledge is purely speculative, but seems important for future research.

Finally, we would like to draw attention to a point that we have always found compelling about the causal graphical modeling framework – that it has suggested this line of experiments, and that we are not modeling already-collected data. Considering whether this framework offers a computational description of children's causal knowledge and reasoning abilities suggests this meaningful question about cognitive development: Does causal knowledge emerge from sensitivity to statistical regularity? We believe the answer is yes. How infants use the statistical regularities among events to form a representation of causal structure is an open question.

Acknowledgements

This work was funded by NSF (DLS 0518161 to DMS). We would like to thank Jonathan Slemmer for assistance in stimuli design and programming, Lisa Smythe for help in data collection, and Denis Mareschal and Daniel Richardson for helpful discussions about this manuscript.

References

- Aslin, R.N., Saffran, J.R., & Newport, E.L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, *9*, 321–324.

- Baldwin, D.A., Baird, J.A., Saylor, M.M., & Clark, A.M. (2001). Infants parse dynamic action. *Child Development*, **72**, 708–717.
- 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**, 367–405.
- Cramer, R.E., Weiss, R.F., Williams, R., Reid, S., Nieri, L., & Manning-Ryan, B. (2002). Human agency and associative learning: Pavlovian principles govern social process in causal relationship detection. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology*, **55B**, 241–266.
- Dickinson, A. (2001). Causal learning: association versus computation. *Current Directions in Psychological Science*, **10**, 127–132.
- Eimas, P.D., Siqueland, E.R., Jusczyk, P.W., & Vigorito, J. (1971). Speech perception in infants. *Science*, **171**, 303–306.
- Gergeley, G., Nadasdy, Z., Csibra, G., & Biro, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, **56**, 165–193.
- Gopnik, A., Glymour, C., Sobel, D.M., Schulz, L.E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological Review*, **111**, 1–30.
- Gopnik, A., 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.
- Griffiths, T.L., & Sobel, D.M. (2006). Learning about probabilistic and deterministic causal systems. Manuscript submitted for publication, Brown University.
- Harris, P.L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, **61**, 233–259.
- Hespos, S.J., & Baillargeon, R. (2001). Infants' knowledge about occlusion and containment events: a surprising discrepancy. *Psychological Science*, **12**, 141–147.
- Hood, B., Cole-Davies, V., & Dias, M. (2003). Looking and search measures of object knowledge in preschool children. *Developmental Psychology*, **39**, 61–70.
- Johnson, S.P. (2001). Visual development in human infants: binding features, surfaces, and objects. *Visual Cognition*, **8**, 565–578.
- Kamin, L.J. (1969). Predictability, surprise, attention, and conditioning. In B.A. Campbell, & R.M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.
- Kirkham, N.Z., Cruess, L.M., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science*, **6**, 449–467.
- Kirkham, N.Z., Slemmer, J.A., & Johnson, S.P. (2002). Visual statistical learning in infancy. *Cognition*, **83**, B35–B42.
- Kruschke, J.K., & Blair, N.J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin and Review*, **7**, 636–645.
- Leslie, A.M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, **25**, 265–288.
- Meltzoff, A.N. (1995). Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, **31**, 838–850.
- Munakata, Y. (2001). Graded representations in behavioral dissociations. *Trends in Cognitive Sciences*, **5**, 309–315.
- Needham, A., & Baillargeon, R. (1993). Intuitions about support in 4.5-month-old infants. *Cognition*, **47**, 121–148.
- Oakes, L.M., & Cohen, L.B. (1990). Infant perception of a causal event. *Cognitive Development*, **5**, 193–207.
- Pearl, J. (2000). *Causality*. New York: Oxford University Press.
- Rescorla, R.A., & Wagner, A.R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In A.H. Black & W.F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–99). New York: Appleton-Century-Crofts.
- Richardson, D.C., & Kirkham, N.Z. (2004). Multi-modal events and moving locations: eye movements of adults and 6-month-olds reveal dynamic spatial indexing. *Journal of Experimental Psychology: General*, **133**, 46–62.
- Saffran, J.R., Aslin, R.N., & Newport, E.L. (1996). Statistical learning by 8-month-old infants. *Science*, **274**, 1926–1928.
- Scaife, M., & Bruner, J.S. (1975). The capacity for joint visual attention in the infant. *Nature*, **253**, 265–266.
- 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.
- Shanks, D.R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology*, **37B**, 1–21.
- Shanks, D.R. (1995). Is human learning rational? *Quarterly Journal of Experimental Psychology A*, **48**, 257–279.
- Siqueland, E., & DeLucia, C. (1969). Visual reinforcement of non-nutritive sucking in human infants. *Science*, **165**, 1144–1146.
- Slater, A. (1995). Visual perception and memory at birth. In C. Rovee-Collier & L.P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 9, pp. 107–162). Norwood, NJ: Ablex.
- 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., & Munro, S. (2006). When Mr Blicket wants it, children are Bayesian. Proceedings of the 2006 Meeting of the Cognitive Science Society, Vancouver, BC.
- 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.
- Spelke, E.S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, **99**, 605–632.
- Spellman, B.A. (1996). Acting as intuitive scientists: contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, **7**, 337–342.
- Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction, and search* (Springer lecture notes in statistics, 2nd edition, revised). Cambridge, MA: MIT Press.

- Tenenbaum, J.B., & Griffiths, T.L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems 15* (pp. 35–42). Cambridge, MA: MIT Press.
- Tenenbaum, J.B., Sobel, D.M., Griffiths, T.L., & Gopnik, A. (submitted). Bayesian inference in causal learning from ambiguous data: Evidence from adults and children. Manuscript submitted for publication, MIT.
- Van Hamme, L.J., & Wasserman, E.A. (1994). Cue competition in causality judgments: the role of nonpresentation of compound stimulus elements. *Learning and Motivation*, **25**, 127–151.
- Wasserman, E.A., & Berglan, L.R. (1998). Backward blocking and recovery from overshadowing in human causal judgment: the role of within-compound associations. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology*, **51**, 121–138.
- Westermann, G., & Mareschal, D. (2004). From parts to wholes: mechanisms of development in infant visual object processing. *Infancy*, **5**, 131–151.
- Woodward, A.L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, **69**, 1–34.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.
- Younger, B.A. (1990). Infants' detection of correlations among feature categories. *Child Development*, **61**, 614–620.
- Younger, B.A., & Cohen, L.B. (1983). Infant perception of correlations among attributes. *Child Development*, **54**, 858–869.
- Younger, B.A., & Cohen, L.B. (1986). Developmental change in infants' perception of correlations among attributes. *Child Development*, **57**, 803–815.
- Younger, B.A., & Gottlieb, S. (1988). Development of categorization skills: changes in the nature or structure of infant form categories? *Developmental Psychology*, **24**, 611–619.
- Younger, B.A., Hollich, G., & Furrer, S.D. (2004). An emerging consensus: Younger and Cohen revisited. *Infancy*, **5**, 209–216.
- Zelazo, P.D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development*, **11**, 37–63.