When Mr. Blicket Wants It, Children are Bayesian

David M. Sobel (Sobel@cog.brown.edu)
Department of Cognitive and Linguistic Sciences, Box 1978, Brown University
Providence, RI, 02912 USA

Sarah Munro (Sarah.Munro@ubc.ca)
Department of Psychiatry, University of British Columbia, 2255 Wesbrook Mall
Vancouver, BC, V6T 2A1 Canada

Abstract

Sobel, Tenenbaum, and Gopnik (2004) demonstrated that 4-year-olds', but not 3-year-olds' causal inferences were well-described by a Bayesian mechanism that relied on specific substantive causal knowledge. Sobel and Blumenthal (submitted) demonstrated that the particular causal knowledge needed for such inferences developed between ages 3 and 4, suggesting that if 3-year-olds had access to this knowledge, they would engage in Bayesian inference. These experiments all used a "blicket detector", which presented a novel causal property of objects. But this procedure is limited to children's understanding of physical causality. In the present paper, we first demonstrated that 3-year-olds have this requisite causal knowledge when the causal relations are presented in the psychological domain (specifically an agent's desire), and given this finding, then demonstrate that 3-year-olds' causal inferences about another's desires are Bayesian.

Bayesian Causal Inference

Researchers in cognitive development have suggested that young children possess causal knowledge across the physical (e.g., Leslie & Keeble, 1987), psychological (e.g., Gopnik & Wellman, 1994), and biological (e.g., Kalish, 1996) domains. There is much consensus on children possessing causal knowledge – even in infancy – but there is little agreement on how that causal knowledge is represented or learned.

Several researchers have recently appealed to causal graphical models as a way of representing children's causal knowledge (e.g., Gopnik, Glymour, Sobel, et al., 2004). Other researchers have appealed to Bayesian inference over a hypothesis space of these graphical models as a description of how children learn new causal relations and engage in causal inference (e.g., Sobel, Tenenbaum, & Gopnik, 2004). Under this Bayesian approach, learners use the substantive prior knowledge available to them to form a hypothesis space of causal models that potentially generated the observed data. Those hypotheses are each assigned a prior probability of being the correct hypothesis. Data are observed, and those hypotheses are updated via an application of Bayes' theorem (see Tenenbaum, Sobel, Griffiths, & Gopnik, 2006, for a more formal description of this algorithm).

What this approach suggests is that ambiguous data observed in the world can be resolved by appealing to the initial prior probabilities that hypotheses were assigned. Sobel et al. (2004) tested this prediction by introducing children to a "blicket detector": a machine that lights up and plays music when certain objects (controlled by the experimenter) are placed upon it. They used a device that presented a novel causal property of objects to children to control what prior knowledge the child bought to the learning environment. They told children that "Blickets make the machine go", and initially trained them that the base rate of blickets in the world was either rare or common. They placed 12 identical objects on the machine one at a time, and either 2 or 10 of those objects activated it (and were categorized as blickets). They then presented children with ambiguous data: a *backwards blocking* procedure in which they took two new identical objects (A and B) and placed them on the detector together, which activated, and then placed only object A on the machine alone, which also activated. Object A is clearly a blicket, but object's B's "blicketness" is uncertain. However, if children responded on the basis of the prior probability of blickets, then when blickets are rare, they should treat B as not a blicket, but when blickets are common, they should treat B as a blicket. This was exactly the pattern of performance observed in 4-year-olds, but not younger children: Three-year-olds categorized object B as a blicket across both conditions.

This procedure has been replicated and extended to other empirical paradigms that test Bayesian inference (e.g., Tenenbaum et al., 2006). At issue here is whether children are developing a mechanism that enables Bayesian inference between ages 3 and 4, or if 3-year-olds possess such a mechanism all along and lack other pieces of substantive prior knowledge.

A way of answering this question is to examine what substantive causal knowledge children must possess in order to form a hypothesis space about blickets and blicket detectors correctly. We have posited that children must have several pieces of information in order to form an accurate and representative hypothesis space. First, children must recognize that there are two categories of objects in the world: objects and detectors. Objects are placed on detectors, and detectors either activate or not, given that an object is upon it. This knowledge seems present for both 3- and 4-year-olds.

Next, children must recognize the temporal relationship between objects and the detector activating. Namely, that placing an object on the detector causes it to activate, not that the detector activating causes the experimenter to place an object on it. Similarly, children must recognize that there is some form of spatial independence among the objects. Putting any one object on the detector does not cause the experimenter to put any other object on the detector. There is also good evidence that suggests both 3- and 4-year-olds possess knowledge about temporal priority and spatial independence (e.g., Bullock, Gelman, & Baillargeon, 1982).

The final piece of knowledge children must have to reason correctly is a notion about how causes and effects are related: how these hypotheses are *parameterized*. In particular, children must recognize a deterministic (or near-deterministic) relationship between a blicket and the detector activating. Following Tenenbaum and Griffiths (2003), we will refer to this piece of knowledge as the "Activation Law": there must be something about a blicket that causes the machine to activate. Without understanding this relationship, the backwards blocking data are not ambiguous: if the detector is not deterministic, in the above procedure, it is more likely that object B is a blicket than the base rate. When A and B are on the detector together, it is possible that A, which is a blicket, failed to activate it, giving B the possible efficacy to activate the machine.

Does this concept develop between ages 3 and 4? Bullock et al. (1982) demonstrated that 3-year-olds generated few mechanistic explanations when asked to describe the connection between a marble rolling down a tube and a jack-in-the-box appearing. Mechanistic explanations were common for 4-5-year-olds. Similarly, 3-year-olds do not appear to understand the nature of causal relationships between physical events: When the marble was replaced with a sequence of lights, 3-year-olds judged that the lights would be effective at producing the effect, while 4- and 5-year-olds judged that the lights would have little effect, since there was no physical mechanism.

One way "the Activation Law" has been investigated is by examining whether children understand that there is a connection between the causal efficacy of an object (i.e., whether it activates the detector) and its internal properties (i.e., whether it has something inside it that can be construed as a mechanism for the detector's efficacy). Yoachim, Sobel, Gopnik, and Meltzoff (submitted) demonstrated that 4-year-olds used whether an object had a novel internal property (a white part inside) to make inferences about whether it would activate the blicket detector. They also demonstrated that 4-year-olds tracked a causally efficacious internal part when it was moved from one object to another. If such a part was transferred to another object that previously was shown not to activate the detector, children now selected the new object in order to make the machine go (as opposed to the original object, which previously activated the detector when it contained this internal part).

Yoachim et al. also demonstrated that 4-year-olds inferred that objects that activated the detector shared internal properties, even in light of contrary perceptual evidence. They presented children with the blicket detector and sets of three objects like those shown in Figure 1. In each set, two blocks had the same external appearance and one had a unique external appearance. Children observed that one of the identical blocks (target object) and the unique block activated the detector, while the third (externally similar to the target) did not. No information about blickets or blicket detector was provided. Each block had a dowel covering its "insides". The target object's dowel was removed to reveal it contained an internal part. The dowel was replaced, and children were asked to point to another object that also contained such an internal part. Four-year-olds chose the object that activated the detector, even though the other object was perceptually identical to the target.

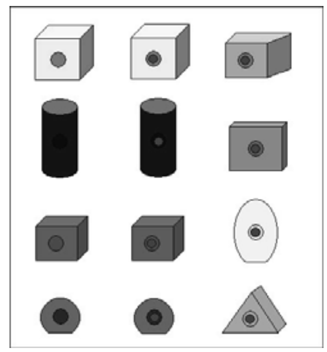

Figure 1: Objects used by Sobel and Blumenthal (submitted) and in the present Experiments.

What about younger children? Sobel and Blumenthal (submitted) replicated this procedure on a 3-year-old sample, and found that these children responded in the opposite manner. Three-year-olds chose the object with the same perceptual appearance as the target as containing the same internal property, even though it did not activate the detector. They also found that 3-year-olds did not understand that possessing an internal part made an object able to activate the detector by itself; 4-year-olds did understand this concept.

As such, it is possible that 3-year-olds's inferences can also be described by Bayesian inference, but that they lack the requisite substantive prior knowledge to form a hypothesis space correctly. To test this hypothesis, we changed the nature of the blicket detector from a physical machine to an intentional agent. There is a great deal of research in theory of mind that suggests that 3-year-olds do understand the causal structure of particular mental states, such as desire (e.g., Gopnik & Slaughter, 1991; Wellman & Woolley, 1990). In the present experiments, we first replicated the internal properties experiment, which tests whether children understand a relationship between the internal properties of objects and their causal properties (i.e., an "Activation Law") when the causal properties of the blicket detector instantiate an agent's desires or a physical machine, and then in Experiment 2 reexamined a procedure that indicated Bayesian inference, which 3-year-olds previously were shown to have difficulty.

Experiment 1

The goal of this first experimenter is to replicate the Sobel and Blumenthal (submitted) procedure on a group of 3-year-olds introducing the detector either as a psychological agent or as a physical machine. To introduce the detector as an agent, we relied on a procedure used by Johnson, Slaughter and Carey (1998), in which the experimenter and detector had a contingent relationship – the experimenter talked to the detector as an agent, and the detector "responded" by activating, which was interpreted for the child as language by the experimenter.

Methods

Participants Thirty-two 3-year-olds (20 girls, M = 41.12 months, range 37-47 months) were recruited from flyers posted at preschools in an urban area.

Materials The blicket detector was 8" x 6" x 3", made of black plastic with a white top, which was wired to a set of LCDs (see Figure 2). The detector worked via a remote-controlled switch (hidden from the child) with two functions. One button on the remote control activated the detector – as long as it was pressed, the detector would light up and play music. This enabled the experimenter to activate the detector without the child knowing that the experimenter caused the activation. Another set of buttons enabled the detector: When the detector was "enabled" any object placed on top of it would cause it to activate. This caused a strong illusion of causality.

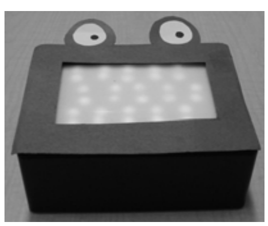

Figure 2: Mr. Blicket: An active blicket detector in the agency condition used in Experiment 1 and 2.

Procedure Children were assigned to one of the two conditions (randomly determined). In the *agency* condition, the experimenter introduced the detector as "Mr. Blicket" and then conducted a brief dialogue in which the experimenter talked to the detector, and Mr. Blicket responded to questions and comments by activating spontaneously (modeled after Johnson et al., 1998). The experimenter also encouraged the child to interact with Mr. Blicket by saying hello, in which case Mr. Blicket activated in response. The children were then told that they were going to play a game in which Mr. Blicket would tell them whether he liked a set of objects.

In the *machine* condition, the box was presented to the child as a "Blicket Machine". Instead of the dialogue used during the agency condition, the experimenter spoke only to the child. During this time, the detector activated spontaneously an equivalent number of times as in the agency condition. No explanation for activation was provided and activation was not contingent on the experimenter's speech.

The trials in both conditions were equivalent except for the language used. In each of four trials, three blocks from the sets shown in Figure 1 were placed in front of the detector. Each block was individually placed on it. The detector activated for the two items containing red map pins and was disabled for the remaining block. In the agency condition, the experimenter said "Look, Mr. Blicket likes this one" when the detector activated or "Look, Mr. Blicket does not like this one" when it did not. In the machine condition, the child was told either "Look, the machine is going" or "Look, the machine is not going."

The dowel on the block that activated the detector and belonged to the perceptually identical pair was removed to reveal that it contained a red internal part. In both conditions, the experimenter said, "Look, this one is made of red stuff" and asked the child to point to "another one made of red stuff." Children were given four such trials.

Results and Discussion

Preliminary analyses demonstrated that children did not differ in responses across the four trials, nor did the order in which the trials were presented affect performance. Overall, children chose the externally distinct object (the other object that activated the machine) on 70% of the trials in the agency condition, significantly more than in the machine condition (where they made this responses on the 41% of the trials), $t(30) = 3.24$, $p = .003$.

This finding was supplemented by a nonparametric analysis. We categorized children's response patterns into three groups: Children were categorized as *causal* if they chose the externally distinct object on 3 or more trials. Children were categorized as *perceptual* if they chose the externally distinct object on 1 or fewer trials. Children were categorized as being *neutral* if they made exactly 2 perceptual and 2 causal responses. The distribution of responses differed between the two groups, $\chi^2(2, N = 32) = 6.54$, $p = .03$. Eleven out of sixteen children in the agency condition were categorized as causal (69%), compared with only four children in the machine condition (25%). The 69% of children who were categorized as making causal responses in the agency condition was significantly greater than what would be expected by chance (31.25%), Binomial test, $p < .001$. This was not the case in the machine condition; the number of children categorized as causal did not differ from chance.

Experiment 1 revealed that 3-year-olds were much more likely to relate an artifact's internal properties (i.e., whether it contained an internal part, which constituted what it was "made of") with its causal properties when those causal properties represented an agent's desires as opposed to the workings of a physical device. This suggests that changing the way the blicket detector's activation is interpreted –

from a physical device to indicating the desires of an agent – allows 3-year-olds to recognize the relationship between an object's internal properties and its causal power. This potentially provides them with a conception akin to the "Activation Law", which might allow them to engage in Bayesian inference. Experiment 2 tests this particular possibility.

Experiment 2

In Experiment 2, we use the procedure in Experiment 1 to introduce the detector as an agent, and then parallel the procedure used to investigate Bayesian inference in a new group of 3-year-olds (following Sobel et al., 2004). We trained children that Mr. Blicket liked very few things, and then presented them with ambiguous evidence as to what he liked: that he liked two objects together, and then one of those two objects alone. If they interpret these data in a manner described by Bayesian inference, then they should be inclined to say that Mr. Blicket will not like the other object. We then perform a similar procedure to Experiment 1 to examine whether there is a relationship between children's inclination to response based on a Bayesian algorithm and their understanding of the relationship between an object's causal and internal properties.

Methods

Participants Twenty-five 3-year-olds (11 girls, M = 42.90 months, range 37-48 months) were recruited from flyers posted at preschools in an urban area.

Materials The same blicket detector and three sets of blocks in Experiment 1 were used. In addition, the same eighteen identical 3.5 x 7.7cm blue cylindrical blocks and three boxes used by Sobel et al. (2004) were used here.

Procedure All children first received a measure of Bayesian inference, which paralleled the procedure of Sobel et al. (2004), except that the blicket detector was introduced as an agent, and its activation was described as indicating Mr. Blicket's desires. Children were introduced to Mr. Blicket in the same manner as Experiment 1. They were then shown a box of identical blue blocks. Twelve blocks were taken out of the box one at a time, and children were shown that Mr. Blicket only liked 2 out of the twelve blocks. Children were asked to categorize each object as being one that Mr. Blicket liked or did not like. Corrective feedback was given if the child erred, but this was rarely necessary (less than 9% of the time). The results of this categorization were present on the table (objects were categorized into a box of things Mr. Blicket liked and a box of things Mr. Blicket did not like). Children were also shown that Mr. Blicket would activate (revealing his preference) if at least one object that he liked was on him. The training procedure was exactly parallel to the one used by Sobel et al. (2004, Experiment 3), and we would refer the reader to that paper for more details.

Children were then given the critical *test* trial. Two new objects (A and B), identical to the first twelve, were taken from the original box, and placed on Mr. Blicket, who activated. No causal language was used to describe this demonstration. Then one of those two objects (A) activated Mr. Blicket on its own. Children were asked to categorize these two objects (one at a time). Next, children were given a *baseline* condition: two new identical objects (X and Y) were taken from the original box. They were placed on Mr. Blicket together, who activated, and children were asked to categorize each. Finally, in a *control* trial, two new identical objects were taken from the original box. One activated Mr. Blicket, and the other did not, and children had to categorize them. If children failed to succeed on this trial, they were not included in the procedure. Three children were replaced for this reason.

After children participated in an unrelated experiment, they were given an Internal Properties measure, which was identical to the measure given to them in Experiment 1. There was one exception. Because the blocks used in the Bayesian measure were very similar to one set, only three trials were given here. The goal of this measure was to see whether children's tendency to engage in Bayesian inference related to their connecting an object's internal and causal properties (a way of examining whether children understood the "Activation Law").

Results and Discussion

On the test trial, all children responded that Mr. Blicket liked object A; only 44% of the children (11 out of 25) claimed that Mr. Blicket liked object B. On the baseline trial, children claimed that Mr. Blicket liked objects X and Y no differently from each other, and 80% of the time, significantly more often than their treatment of object B in the test trial, Wilcoxon Signed Ranks Test, z = -3.29, p = .001. This suggests that children retrospectively reevaluated whether Mr. Blicket liked object B consistent with the base rate of such objects being rare.

Performance on the Bayesian Measure was also compared with performance on the similar procedure used by Sobel et al. (2004). There, sixteen 3-year-olds were introduced to a "Blicket Machine" and were told that "Blickets make the machine go". Children were trained that blickets were equally rare, in the same manner as they were trained here that the things Mr. Blicket liked were rare, given then same three test trials. For each, they were asked whether each object was a blicket. In the test trial of that experiment, all children claimed that object A was a blicket, and 13 of the 16 children (81%) also claimed that object B was a blicket. This represents a greater percentage than the 44% of the children who did so here, $\chi^2(1, N = 41) = 5.58$, p = .018. Since only three children in the previous experiment claimed that object B was not a blicket (i.e., one part of the Chi-Squared table had fewer than 5 data points), these data were also subjected to a Fisher's Exact test, which also revealed a significant difference, p = .019. This suggests that when 3-year-olds were presented with the same inference as the workings of a machine (as opposed to Mr. Blicket), they fail to make similar inferences consistent with a Bayesian algorithm.

The internal properties measure was scored in the same manner as in Experiment 1, except that there were only three trials. Overall, children made causal responses on 39% of the trials (an average score of 1.16 out of a possible 3, *SD* = 0.91). In contrast with performance on the agency condition in Experiment 1, the overall pattern of response was no different from chance, $t(23) = -1.69$, *ns*.

When these data were contrasted with performance on the Bayesian measure, a different pattern emerged. Thirteen of these 24 children[1] claimed that Mr. Blicket did not like object B (consistent with a Bayesian description). These children made causal responses on 56% of the internal property trials (*M* = 1.69 out of a possible 3, *SD* = 0.95). In contrast, the 11 children who claimed that Mr. Blicket liked object B (inconsistent with a Bayesian description) only made causal responses on 18% of the trials (*M* = 0.55 out of a possible 3, *SD* = 0.52). These percentages were significantly different from each other, Mann-Whitney $U = 23.00$, $z = -2.94$, $p = .004$. This suggests that if children resolved ambiguous data in a manner that was consistent with a Bayesian description, they scored higher on the internal properties measure; they were more likely to recognize that there was a relationship between the detector's activation and an internal property of the objects.

A concern here is that the overall level of response on the internal property question was slightly lower than in Experiment 1. One possible explanation for this difference is experimental fatigue: in the present experiment, children might have paid less attention to the procedure, given the length of the experiment. This could have resulted in more of a reliance on external perceptual similarity (e.g., the working of a "dumb attentional mechanism", e.g., Smith, Jones, & Landau, 1996). Further, in the Bayesian inference procedure, all the objects were perceptually identical, which might have similarly primed a perceptual response. Such factors do not seem to solely motivate performance. The critical finding is that children whose responses were best described by a Bayesian mechanism were more likely to override perceptual similarity in the internal properties procedure and recognize that objects that activate the detector share internal properties.

General Discussion

Several researchers have suggested that children's causal inference can best be described by a rational mechanism that is consistent with Bayesian inference (e.g., Sobel et al., 2004; Tenenbaum & Griffiths, 2003). On this approach, children use the substantive causal knowledge they possess to form a hypothesis space about the data they observe. They assign a set of prior probabilities that each of those hypotheses actually generated the data they observe in the world. Those priors are then updated via an application of Bayes' theorem given observed data, and causal inference is made by considering the resulting posterior probabilities. In previous research, and the present Experiment 2, the data the child observes is ambiguous, and what differs is the number of blickets posited by each model. As such, the base rate of blickets disambiguates the observed data. When the causal relationship between object and detector represented a physical relationship, 3-year-olds' inferences were not consistent with a Bayesian description. In the present Experiment 2, when the causal relationship between object and detector are interpreted as reflecting an intentional agent's desires, children's inferences were more consistent with this Bayesian description.

Why is this the case? The present procedure offers a clear explanation: Specifically, 3-year-olds recognize a particular piece of prior substantive causal knowledge about agent's desires, but not about physical events: that there is a relationship between that desire and an internal property of the object that is responsible for the machine's activation at a deterministic (or near deterministic) level. If Mr. Blicket wants something, it will have a particular property that (children presume) is responsible for his desire[2].

A concern that runs through both experiments is that the agency manipulation might not actually tap a different set of causal knowledge, but rather make the procedure more interesting to children. This might allow them to succeed more easily on the measures presented here. We do not believe that interest alone motivates the difference between conditions here. Rather, we think that these two conditions (Agency vs. Machine) tap different domain-specific causal knowledge structures. Further, we hypothesize that children's understanding of the activation law would be limited to it being phrased as a mental state easily understood by 3-year-olds, such as desire. If Mr. Blicket's activation revealed another mental state such as surprise, which is not well-understood by 3-year-olds, the present findings would not replicate. Further, the hypothesis that attention only motivates these results does not necessarily predict that children's more rational inferences about base rates would be related to their knowledge of internal structure, as found in Experiment 2.

A more general question is where does this prior causal knowledge come from? We would offer the following speculations. One possibility is that children acquire information from recognizing regularity among events. Many researchers have demonstrated that infants recognize statistical regularity in language and visual events (e.g., Kirkham, Slemmer, & Johnson, 2002; Saffron, Aslin, & Newport, 1996). Sobel and Kirkham (in press) found that 8-

[1] Note that one child was excluded because of experimental error.

[2] These data do not definitively show that 3-year-olds interpret the internal property as being responsible for the causal property. Using a separate procedure, Sobel and Blumenthal (submitted) did demonstrate that 4-year-olds recognize this relationship in the physical domain (using an ordinary blicket detector), and that 3-year-olds do not. This procedure can be replicated using the agency manipulation, and we are currently working on this endeavor to definitely answer this question.

month-olds can also recognize conditional independence and dependence relations in visual sequences. This is consistent with infants representing statistical knowledge using the same kind of computational framework – causal graphical models – that we suggest underlies children's representation of their causal knowledge.

Recognizing statistical regularity is only one way that prior knowledge might be acquired. Algorithms with similar Bayesian architecture may allow learners to acquire new pieces of prior knowledge, which can subsequently be applied to novel inferences. Kemp, Perfors, and Tenenbaum (2004) suggested that a Bayesian inferential system could make meaningful inferences in biological and political datasets: critically that the causal structures across these domains tended to have different shapes. Tenenbaum, Griffiths, and Niyogi (in press) suggested that Bayesian inference can be considered at various levels of analysis: Inferences about causal structure can be made from the data, but also inferences about how such causal structures are formed. However, these findings were only with adult data; a clear extension of this work is to see if applications of this algorithm generalize onto children's inferences.

The present experiments suggest that changing the nature of causal information from that of a physical device to that of the desires of a psychological agent allow 3-year-olds access to a specific piece of causal information: that objects that share causal power tend to share a particular internal structure. This knowledge may allow 3-year-olds to recognize that there is a deterministic relationship between those objects and the agent's desires, which in turn helps them to reason in a rational manner, consistent with the predictions of a Bayesian inferential system. We hypothesize that what develops is not a learning mechanism consistent with this Bayesian approach. Rather, even young children can make rational inferences, consistent with a Bayesian account when they possess appropriate domain-specific substantive causal knowledge.

Acknowledgments

This work was supported by NSF (DLS-0518161 to DMS). We would like to thank all of the parents and children who participated in this research. We would also like to thank Emily Blumenthal, Emily Hopkins, and Philip Parker who helped with data collection and coding.

References

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.

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., & Slaughter, V. (1991). Young children's understanding of changes in their mental states. *Child Development*, *62*, 98-110.

Gopnik, A., & Wellman, H. M. (1994). The theory theory. In L. Hirschfield & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 257-293). New York: Cambridge University Press.

Johnson, S., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? The elicitation of gaze-following in 12-month-olds. *Developmental Science*, *1*, 233-238.

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

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2004). Learning domain structures. *Proceedings of the 2004 meeting of the Cognitive Science Society, Chicago, IL.*

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy. *Cognition*, *83*, B35-B42.

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*, 265-288.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926–1928.

Smith, L. B., Jones, S. S., & Landau, B. (1996). Naming in young children: a dumb attentional mechanism? *Cognition*, *60*, 143-171.

Sobel, D. M., & Blumenthal, E. J. (2005). *Preschoolers' developing knowledge of causal and internal properties of artifacts.* Manuscript submitted for publication, Brown University.

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

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

Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal inference. *Proceedings of the 14th Annual Conference on the Advances in Neural Information Processing Systems*. Vancouver, BC.

Tenenbaum, J. B., Sobel, D. M., Griffiths, T. L., & Niyogi, S. (in press). Intuitive Theories as Grammars for Causal Inference. In A. Gopnik & L. E. Schulz (Eds.), *Causal Learning: Psychology, Philosophy, and Computation.* Cambridge: MIT Press.

Tenenbaum, J. B., Sobel, D. M., Griffiths, T. L., & Gopnik, A. (2006). *Bayesian inference in causal learning from ambiguous data: Evidence from adults and children*. Manuscript in preparation, MIT.

Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, *35*, 245-275.

Yoachim, C. M., Sobel, D. M., Gopnik, A., & Meltzoff, A. N. (2005). *The blicket within: Preschoolers' inferences about insides and hidden causes.* Manuscript submitted for publication. Brown University.