Running head: NAÏVE THEORIES AND AMBIGUOUS EVIDENCE

Can being scared make your tummy ache? Naive theories, ambiguous evidence and preschoolers' causal inferences

Laura E. Schulz¹, Elizabeth Baraff Bonawitz¹, and Thomas L. Griffiths²

¹Massachusetts Institute of Technology

²University of California, Berkeley

Submitted to Developmental Psychology 4/6/06

Revision submitted 11/22/06

Final revision submitted 2/8/07

This research was supported by a James H. Ferry, Jr. Fund for Innovation in Research Education Award and a McDonnell Foundation Collaborative Initiative Causal Learning grant to L.S. Thank you to Darlene Ferranti, Noah Goodman, Kate Hooppell, Adrianna Jenkins, Rebecca Saxe, Josh Tenenbaum, and Wendy Wienerman for helpful comments and suggestions. We are also very grateful to the Discovery Center at the Museum of Science, Boston and participating parents and children. Address for correspondence: L. Schulz, 46-4011, MIT Department of Brain and Cognitive Sciences, 77 Massachusetts Avenue, Cambridge, MA 02139. Phone: 617-253-7957. E-mail: lschulz@mit.edu

Abstract

Causal learning requires integrating constraints provided by domain-specific theories with domain-general statistical learning. In order to investigate the interaction between these factors, preschoolers were presented with stories pitting their existing theories against statistical evidence. Each child heard two stories in which two candidate causes co-occurred with an effect. Evidence was presented in the form: $AB \rightarrow E$, $AC \rightarrow E$, $AD \rightarrow E$, etc. In one story, all variables came from the same domain; in the other, the recurring candidate cause, A, came from a different domain (A was a psychological cause of a biological effect). After receiving this statistical evidence, children were asked to identify the cause of the effect on a new trial. Consistent with the predictions of a Bayesian model, all children were more likely to identify A as the cause within domains than across domains. While three-and-half-year-olds learned only from the within-domain evidence, four- and five-year-olds learned from the cross-domain evidence and were able to transfer their new expectations about psychosomatic causality to a novel task.

KEYWORDS: domain-general and domain-specific causal learning, Bayesian models, naïve theories, ambiguous evidence, psychosomatic causes

Can being scared make your tummy ache? Naive theories, ambiguous evidence and preschoolers' causal inferences

By the time children are five years old, they understand causal relationships in a variety of domains (Flavell, Green, & Flavell, 1995; Gelman & Wellman, 1991; Gopnik & Meltzoff, 1997; Inagaki & Hatano, 1993; Kalish, 1996; Perner, 1991; Shultz, 1982; Spelke, Breinlinger, Macomber, & Jacobson, 1992). Many researchers have suggested that children's causal knowledge can be best characterized as a set of naive theories: abstract, coherent representations of causal structure that support prediction, intervention, explanation, and counterfactual claims (Carey, 1985; Harris, German, & Mills, 1996; Hickling & Wellman, 2001; Gopnik, 1988; Gopnik & Meltzoff, 1997; Keil, 1989; Perner, 1991; Sobel, 2004; Wellman, 1990; Wellman, Hickling, & Schult, 1997). The view that children's causal representations resemble scientific theories (the *theory theory*) suggests both that patterns of evidence should affect children's causal commitments and that children's causal commitments should affect their interpretation of evidence. Indeed, this dynamic relationship between domain-appropriate causal beliefs and evidence has been taken as a defining feature of theories (e.g., Gopnik & Meltzoff, 1997).

However, despite the expectation that theory and evidence should interact, developmental psychologists have been largely divided between accounts of causal reasoning emphasizing either domain-specific causal knowledge or domain-general learning from data. Thus some researchers have suggested that children's naive theories might be generated by domain-specific modules (Leslie, 1994; Scholl & Leslie, 1999) or innate concepts in core domains (Carey & Spelke, 1994; Keil, 1995), while other

researchers have focused on children's ability to learn causal relations from statistical evidence (Gopnik et al., 2004; Gopnik, Sobel, Schulz, & Glymour, 2001; Schulz & Gopnik, 2004; Shultz & Mendelson, 1975; Siegler & Liebert, 1975; Sobel, Tenenbaum, & Gopnik, 2004; Watson & Ramey, 1972). Although some research on the development of scientific reasoning has emphasized the importance of integrating domain-specific knowledge with domain-general strategies (Barrett, Abdi, Murphy, & Gallagher, 1993; Klahr & Dunbar, 1988; Koslowski, 1996; Koslowski, Okagaki, Lorenz, & Umbach, 1989; Pazzani, 1991; Penner & Klahr, 1996; Schauble, 1990), those studies have focused primarily on adolescents and adults. Surprisingly little research has looked at how prior theories and evidence interact in young children's causal learning.

Moreover, the few studies that have directly compared preschoolers' domain-specific and domain-general causal learning have generated contradictory results. Some studies suggest that children privilege domain-specific mechanism information over domain-general evidence. Work by Shultz (1982) for instance, suggests that preschoolers will override covariation evidence to base causal judgments on the presence or absence of domain-appropriate mechanisms of transmission. In one study for instance, Shultz showed children a candle with a screen around it. He turned on a fan and then, five seconds later, turned on a second fan. While turning on the second fan, he moved the screen away from the first fan. The candle extinguished. When children were asked which fan extinguished the candle, children chose the first fan, which was in a position to transmit energy to the candle, rather than the second fan, whose activation was temporally contiguous with the effect. This was taken as evidence that children's causal judgments are more influenced by domain-specific information than domain-general

cues, like temporal contiguity. Note however, that some domain-general information (e.g., the temporal contiguity between removing the screen from the first fan and the candle extinguishing) may have reinforced the domain-specific information about mechanisms of transmission. Thus it is not clear whether children genuinely privileged the domain-specific information or whether both types of information contributed to children's judgments.

In contrast to the Shultz studies, other work suggests that preschoolers can use domain-general information to override domain-specific theories. Research suggests for instance, that four-year-olds are able to use patterns of conditional dependence and independence to learn that talking to a machine, rather than pushing a button will make the machine activate, or that a block can activate a toy, not through contact, but at a distance (Kushnir & Gopnik, 2006; Schulz & Gopnik, 2004). However, in these studies, the evidence strongly favored the implausible (domain-inappropriate) cause. Target effects never occurred spontaneously and did occur when the domain-inappropriate candidate cause was present by itself. In these contexts, children's prior knowledge seemed to have no effect on their inferences: children were able to learn causal relationships that violated domain boundaries as easily as within-domain relations. However, judgments made in the face of such unambiguous evidence may not provide either a particularly strong test of domain-general learning mechanisms or a particularly nuanced look at how theories affect the interpretation of statistical data. Thus although some studies seem to suggest the relative strength of domain-specific knowledge over domain-general learning mechanisms and others suggest the opposite, little research has closely investigated the interaction between the two.

Adding to the complexity, many researchers have suggested that the relationship between theory and evidence may be poorly understood even by older children and naïve adults (Chen & Klahr, 1999; Dunbar & Klahr, 1989; Inhelder & Piaget, 1958; Kuhn, 1989; Kuhn, Amsel, & O'Laughlin, 1988; Masnick & Klahr, 2003). For instance, some research suggests that adults interpret identical evidence differently depending on whether the data supports or conflicts with a favored theory. Thus, if two candidate causes are both independent of an effect, learners will cite instances of co-occurrence as evidence for the relationship consistent with their theories and instances of non-co-occurrence as evidence against the relationship incommensurate with their theories (Kuhn, 1989).

Critically however, such differential treatment of evidence need not be irrational: small amounts of data (e.g., seeing a vase floating in mid-air) may suffice to overturn weakly held beliefs (that the magic show was canceled) but should leave strong ones (that unsupported objects fall) intact. To the extent that children's causal judgments reflect normative interactions between naïve theories and patterns of evidence, the mixed findings across different studies are perhaps not surprising. On any given task, children's causal judgments might accord either with their prior knowledge or with the patterns of evidence, depending on the strength of children's initial theories, the strength of the evidence, and children's ability to integrate the two.

A rational answer to the question of how domain-specific theories should interact with statistical evidence can be obtained by approaching causal learning as a problem of Bayesian inference. In Bayesian inference, the learner seeks to evaluate a hypothesis, h , about the process that produced some observed data, D . The learner's a priori beliefs

about the plausibility of h are expressed in a “prior” probability distribution, $P(h)$. The learner seeks to evaluate the “posterior” probability of h , $P(h|D)$ – their beliefs about the plausibility of the hypothesis after taking into account the evidence provided by D . This can be done by applying Bayes’ rule,

$$P(h|D) = \frac{P(D|h)P(h)}{\sum_{h'} P(D|h')P(h')} \quad (1)$$

where $P(D|h)$ is the “likelihood”, indicating the probability of generating the data D if the hypothesis h were true. (The sum over all hypotheses in the denominator simply ensures that the result is a probability distribution.) The posterior distribution directly combines the evidence obtained from D , through the likelihood, with the learner’s initial beliefs about the plausibility of the hypothesis, expressed in the prior, $P(h)$. In the case of causal learning, we can imagine prior probabilities being supplied by a domain-specific theory, stipulating which causal structures are plausible (Tenenbaum, Griffiths, & Niyogi, in press; Tenenbaum & Niyogi, 2003). Thus, Bayesian inference provides a formal account of how domain-specific theories and domain-general patterns of evidence might interact to affect children’s beliefs.

Guided by this account, we look at how statistical evidence affects children’s causal inferences and how children’s beliefs about the plausibility of causal hypotheses affect children’s interpretation of data. Because we are interested in interactions between naïve theories and evidence, we give children evidence that is formally ambiguous: children observe events in which two candidate causes simultaneously covary with the effect. We give children a forced choice between the two causes and manipulate the

extent to which each cause is consistent with children's naive theories and with the statistical evidence.

Previous research suggests that preschoolers are able to evaluate evidence of this complexity (i.e., evidence in which candidate causes are never presented in isolation). In one study for instance, a puppet smelled a bouquet consisting of a tulip and a daisy. The puppet sneezed. The puppet then smelled a bouquet consisting of a tulip and a violet and the puppet sneezed. Children then saw that a bouquet consisting of a daisy and a violet did not make the puppet sneeze. When asked what made the puppet sneeze, children inferred that the tulip, rather than the other flowers, was the cause (Schulz & Gopnik, 2004).

Children can also evaluate ambiguous data with respect to the base rate of candidate causes. Suppose for instance, children learn that a toy will light up when particular blocks are placed on top of the toy. Children learn either that only two of ten blocks activate the toy (i.e., activating blocks are rare) or they learn that eight of ten blocks activate the toy (activating blocks are common). Children in both conditions then see two novel blocks, red and blue, placed simultaneously on the toy. The toy activates. Children subsequently see a red and yellow block placed simultaneously on the toy. Again, the toy activates. What makes the toy go – just the red block, just the blue and yellow blocks, or all three blocks? If activating objects are rare, then it is most likely that only a single block (the red one) is the cause. However, if activating blocks are common, it becomes more plausible that the blue and yellow block or all three blocks are causes. Research suggests that children's judgments about activating blocks are sensitive to such base rate information (Tenenbaum, Sobel, Griffiths, & Gopnik, in submission). Similarly,

suppose that preschoolers see a red block and a blue block together activate a toy and then see that the red block by itself activates the toy. Children know the red block is a cause but what about the blue block? If activating blocks are rare, children tend to deny that the blue block is a cause; if activating blocks are common children are more likely to think the blue block is also causally effective (Sobel, Tenenbaum, & Gopnik, 2004).

Our task is formally similar to the tasks used in these previous studies. We read children storybooks in which one cause recurs every day and the other cause is always novel (i.e., the evidence is in the form $AB \rightarrow E$; $CA \rightarrow E$; $AD \rightarrow E$... etc.). One storybook is a *Within Domain* story; all variables come from the same domain and thus all causes are a priori equally plausible. If children can engage in domain-general statistical learning from patterns of evidence, we expect that after seeing the evidence, children will infer that A is more probable than any other single cause. The other storybook is a *Cross Domains* story: the recurring candidate cause (A) is domain-inappropriate. Thus A is less plausible than the alternative given the children's naïve theories but more plausible given the pattern of evidence. By comparing children's judgments before and after seeing the data, we can evaluate the degree to which children can overcome the biases induced by their naïve theories.

Because we wanted to investigate processes that might be applicable to genuine conflicts between theories and evidence, we chose to look at a context in which preschoolers' causal beliefs are robust (and thus might affect children's interpretation of data) but distinct from adult beliefs (and thus might change with evidence). As noted, considerable research suggests that children's causal reasoning respects domain boundaries. In particular, many researchers have suggested that children respect an

ontological distinction between mental phenomena and bodily/physical phenomena (Bloom, 2004; Carey, 1985; Estes, Wellman, & Woolley, 1989; Hatano & Inagaki, 1994; Notaro, Gelman, & Zimmerman, 2001; Wellman & Estes, 1986). Indeed, some researchers have proposed that children may be innate dualists (Bloom, 2004). Thus although many adults accept the existence of psychosomatic phenomena, preschoolers typically deny that psychosomatic reactions are possible (e.g., they deny that feeling embarrassed can make you blush or that feeling frustrated can cause a headache; Notaro, Gelman & Zimmerman, 2001).

Note that this not to suggest that children deny all relations between mental states and bodily events.¹ In particular, children do understand that volitional mental states (e.g., desires and intentions) can cause intentional action (see e.g., Bartsch & Wellman, 1989; Meltzoff, 1995; Wellman, Hickling, & Schult, 1997). Indeed, in the context of voluntary action, bodily events are more typically attributed to psychological causes (“She kicked the ball because she wanted to make a goal”) than bodily causes (“She kicked the ball because she lifted her leg, extended her knee, etc.”). However, by definition, involuntary bodily events (e.g., tummyaches and headaches) are not attributed to desires or intentions. Although many adults accept that involuntary bodily states can be caused by psychological states like worrying or fear, children seem to attribute involuntarily bodily states exclusively to bodily causes (e.g., illness and injury). Thus, following Notaro et al., we investigate children’s understanding of only a subset of possible relations between psychological and bodily events: cases where a non-volitional mental state is the cause of an involuntary bodily reaction.²

We were interested in how preschool children would interpret formal patterns of evidence suggesting the presence of a psychosomatic cause in light of their strong initial belief that psychosomatic causality is improbable. Thus in our *Within Domain* task, both the candidate causes and the target effect come from the domain of physiological events. In particular, the candidate causes are bodily contact with different plants and the effect is “itchy spots”. In the *Cross Domains* task, all but one of the candidate causes are physiological events (ingestion of different foods) and the effect is also a physiological event (a tummy ache). However, the recurring cause (A) is a psychological event (feeling scared).

Bayesian inference provides us with a rational standard against which we can compare the judgments of children in this task. In this case, the hypotheses concern the relationships between the potential causes and the effect, and the data are the co-occurrences of causes and effects over the days of the story. The prior probabilities of the hypotheses are determined by children’s naïve theories, which we assume allow a greater probability for relationships within a domain than relationships that cross a domain. A detailed Bayesian model and quantitative predictions of this model for our task are described in Appendix A. This model makes two clear qualitative predictions. First, children’s naïve theories should affect their interpretation of the evidence: children should be less likely to choose A in the *Cross Domains* task than in the *Within Domain* task. Second, the evidence should affect children’s beliefs: children should be more likely to choose A after seeing the evidence than at baseline. Earlier research (Sobel, Tenenbaum, & Gopnik, 2004) suggests that older (four-year-olds) but not younger (three-year-olds) preschoolers are able to integrate knowledge about base rates and patterns of

evidence, so we test our predictions about the integration of domain-specific prior knowledge and evidence across a range of ages: three-year-olds, three-and-a-half-year-olds, and four- and five-year-olds.

Experiment 1

In Experiment 1, we read children two storybooks: a *Within Domain* story and a *Cross Domains* story. The evidence (presented in the form $AB \rightarrow E$, $CA \rightarrow E$, $AD \rightarrow E$, etc.) is formally identical in the two stories. We predict that children will be more likely to identify A as a cause when A is domain-appropriate than when it is domain-inappropriate. However, we also predict that, for both stories, children will be more likely to think A is a cause after seeing the evidence than at baseline.

Method

Participants

Eighty preschoolers were recruited from urban area preschools and the Discovery Center at a large metropolitan Science Museum. Approximately equal numbers of boys and girls participated (49% girls). While most children were from white, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented.

Children were tested in three age groups: three-year-olds (mean age: 39 months; range 36-41 months), three-and-a-half-year-olds (mean age: 45 months; range 42-48 months), and four- to five-year-olds (mean age: 60 months; range: 50– 70 months). The wider age-range was used for the oldest age group because pilot work suggested that the performance of four- and five-year-olds did not differ on this task. In each age group, 16 children were tested after seeing statistical evidence (the Evidence condition).

Additionally, 16 four- to five-year-olds and 16 three-year-olds (8 in the older group; 8 in the younger group) were tested before seeing any statistical evidence (the Baseline condition). We tested the three-year-olds as a single group at Baseline because pilot work suggested no difference in Baseline performance for the youngest two age groups.

Materials

Two storybooks were used in the experiment: a *Within Domain* book and a *Cross Domains* book. Each storybook depicted events occurring over the course of a week. Every morning (Monday-Sunday), two events and an effect occurred. One event (A) and the effect were repeated every morning; the other event varied. Each afternoon, two different events occurred and the effect failed to occur. (The afternoon events were included to eliminate the possibility that the effect was always present.) Two versions of each storybook were created to counterbalance the order of events.

The *Within Domain* storybook featured a deer (Bambi) who liked to run in different places. Sample text read: “On Monday morning, Bambi runs in the pine grove. Bambi gets excited³; Bambi runs in the cattails. Bambi has itchy spots on his legs. On Monday afternoon, Bambi runs in the cedar trees and Bambi plays on the rope swing. Bambi feels great. Bambi doesn’t have any itchy spots.” The story continued through the days of the week and ended with, “On Sunday morning, Bambi runs in the garden. Bambi gets excited; Bambi runs in the cattails. Bambi has itchy spots on his legs.”

The *Cross Domains* storybook featured a bunny who was scared of show-and-tell. Sample text read: “On Monday morning, Bunny thinks about show-and-tell. Bunny feels scared. Bunny eats some cheese. Bunny has a tummy ache. On Monday afternoon, Bunny ties her shoes and Bunny eats strawberries. Bunny feels great. Bunny doesn’t

have a tummy ache.” The story continued through the days of the week and ended with “On Sunday morning, Bunny thinks about show and tell, Bunny feels scared. Bunny eats a sandwich. Bunny has a tummy ache.”³ See Figure 1 for details and Appendix B for the full text of the stories.

Insert Figure 1 here

Procedure

Children were tested individually. The experimenter read both the *Within Domain* story and the *Cross Domains* story to every child (order of stories counterbalanced between participants). In the Baseline condition, children were read only the “Sunday” page of each story. In the Evidence condition, children were read each story in its entirety. At the end of each story, children were asked a test question. In the *Within Domain* story, the test question (in one version’s order) was: “Why does Bambi have itchy spots? Is it because of running through the garden or because of running through the cattails?” In the *Cross Domains* story, one version of the test question was “Why does Bunny have a tummy ache? Is it because of feeling scared or eating the sandwich?”

Results

The results are presented in Figure 2.

Insert Figure 2 here

An alpha level of .05 was used throughout this paper, and thus all results reported as significant are $p < .05$ or better. Preliminary analyses revealed no order effects. Because our dependent measure was the number of children making each category choice and we could not be sure the data met the normality assumptions of parametric tests, we used categorical tests (binomial and chi-square tests) throughout.

In the four- and five-year-old age group (4;0–5;6) in the *Within Domain* task, there was no significant difference in the probability with which children chose A or the alternative in the Baseline condition ($N = 16$, $p = \text{ns}$ by binomial test), but children chose A significantly more often than chance (indeed, at ceiling) in the Evidence condition ($N = 16$, by binomial test). Children were significantly more likely to choose A after seeing the evidence than at baseline ($\chi^2 (1, N = 32) = 10.67$). In the *Cross Domains* task, children had a significant preference for the domain-appropriate cause in the Baseline condition ($N = 16$, by binomial test) but did not display a statistically significant difference in their choices of A and the alternative in the Evidence condition ($N = 16$, $p = \text{ns}$ by binomial test). Children were significantly more likely to choose A after seeing the evidence than at baseline ($\chi^2 (1, N = 32) = 5.24$). Both at baseline and after seeing the evidence, children were more likely to choose A in the *Within Domain* than *Cross Domains* task (Baseline: $\chi^2 (1, N = 32) = 5.24$; Evidence: $\chi^2 (1, N = 32) = 10.67$).

In the three-and-a-half-year-old age group (3;6–4;0) in the *Within Domain* task children again did not show a statistically significant difference in their choices between A and the alternative at Baseline ($N = 16$, $p = \text{ns}$ by binomial test) and again significantly preferred A in the Evidence condition ($N = 16$, by binomial test). Children were more

likely to choose A after seeing the evidence than at baseline ($\chi^2(1, N = 32) = 7.57$).

However, in the *Cross Domains* task, children preferred the domain-appropriate cause in the Baseline condition ($N = 16$, by binomial test) and continued to prefer the domain-appropriate cause in the Evidence condition ($N = 16$, by binomial test). Children were not significantly more likely to choose A after seeing the evidence than at baseline ($\chi^2(1, N = 31) = 0, p = ns$). Both at baseline and after seeing the evidence, children were more likely to identify A as a cause in the *Within Domain* than *Cross Domains* task (Baseline: $\chi^2(1, N = 32) = 5.24$; Evidence: $\chi^2(1, N = 32) = 21.21$).

Seeing the evidence had no effect for the youngest age group (3;0-3;6). In the *Within Domain* task, children showed no statistically significant preference between A and the alternative even in the Evidence condition ($N = 16, p = ns$ by binomial test). Children were not significantly more likely to choose A after seeing the evidence than at baseline ($\chi^2(1, N = 32) = .51, p = ns$). In the *Cross Domains* task, children preferred the domain-appropriate cause in the Baseline condition ($N = 16$, by binomial test) and were not significantly more likely to choose A after seeing the evidence than at baseline ($\chi^2(1, N = 32) = .82, p = ns$). Children were more likely to choose A in the *Within Domain* than *Cross Domains* task (Baseline: $\chi^2(1, N = 32) = 5.24$; Evidence: $\chi^2(1, N = 32) = 4.57$).

The different results observed across the different age groups suggest a developmental effect, with age group interacting with the effect of statistical evidence in the *Cross Domains* task. We tested for the possibility of such an interaction using a log-linear model, predicting the frequency with which children chose A over the alternative as function of Age Group, Condition (Baseline or Evidence), and an Age Group x Condition interaction. Removal of the Age Group factor from the saturated model did not

result in a statistically significant increase in lack of fit ($\chi^2(2, N = 80) = 4.49, p = ns$), while removal of Condition or the interaction resulted in a statistically significant increase in lack of fit for both the saturated model (Condition: $\chi^2(1, N = 80) = 8.93$; Age Group x Condition: $\chi^2(2, N = 80) = 9.48$) and the model without the effect of Age Group (Condition: $\chi^2(1, N = 80) = 8.93$; Age Group x Condition: $\chi^2(2, N = 80) = 7.74$). These results indicate that the best model includes Condition and the Age Group x Condition interaction as predictors, supporting the hypothesis that the four- and five-year-olds responded differently to the cross-domain evidence than the three-year-olds.

Discussion

The results of Experiment 1 suggest both that children's domain-specific theories interact with their interpretation of evidence and that the nature of this interaction changes over the course of development. Overall, we found a graded interaction between prior knowledge and evidence of the kind predicted by our Bayesian model: all but the youngest children learned that A was a cause when A was consistent with their theories and all the children were less likely to identify A as a cause when A violated their theories. (Note that this finding rules out a simple associative explanation of children's inferences; the association between variable A and the effect was identical within and across domains.) Critically, the oldest preschoolers seemed to learn from the evidence even when the evidence conflicted with their prior beliefs. After seeing the data, four- and five-year-olds were able to entertain a causal possibility (that being scared might cause tummyaches) that they did not endorse without seeing those data.

By contrast, the younger three-year-olds (3;0 – 3;6) had a strong preference for domain-appropriate causes and apparently failed to learn from the evidence throughout.

It is not clear whether this failure is due to competence or performance deficits. Our task was quite complex and younger children might have been able to learn from the data in a simpler or more supported task. Alternatively, the youngest children might have understood the evidence but the evidence might not have overcome the children's initial inductive biases, even within domains. (That is, the youngest children might have found it relatively more difficult to believe that one type of plant might make you itch and others might not.) Further research might disambiguate these accounts.

However, for the three-and-a-half-year-olds (3;6-4;0) the discrepancy between within and cross-domain reasoning was particularly striking. Although the evidence in the tasks was formally identical, 94% of the children in this group inferred that A was the cause in the *Within Domain* task (no different than the four- and five-year-olds) while only 12% inferred that A was the cause in the *Cross Domains* task (no different than at baseline).

Why did the three-and-a-half-year-olds respond differently to the *Cross Domains* evidence than the four- and five-year-olds? There are at least three possible explanations. One possibility is that three-year-old children might have difficulty making inferences from ambiguous statistical data. If (as suggested by the failure of the younger three-year-olds to use the evidence at all) the ability of the three-and-a-half-year-olds to interpret data of this complexity is fragile, any increase in task difficulty (including a conflict with prior knowledge) might compromise children's ability to evaluate the evidence.

Alternatively, the younger children might have a stronger belief in domain boundaries than older children. The data might have been insufficient to overcome three-year-olds' initial inductive bias that psychological causes are unlikely to generate bodily effects.

Finally, the older three-year-olds and the four- and five-year-olds might not differ either with respect to their ability to evaluate evidence or their initial domain-specific theories. However, younger children might be less able than older children to update their beliefs on the basis of surprising evidence. Again, further research might disambiguate these accounts.

Overall, the results of Experiment 1 suggest that domain-general and domain-specific information interact to affect children's causal learning, consistent with the prescriptions of Bayesian inference. Moreover, older preschoolers can use statistical evidence to make inferences against their domain-specific theories even when the data is ambiguous. That is, children can infer a domain-inappropriate causal relationship even when the evidence does not formally rule out the causal relationship consistent with their initial theories.

However, Experiment 1 also suggests that children's learning is relatively conservative; children were less likely to learn from statistical evidence that conflicted with their theories than from evidence consistent with their theories. If children's learning is conservative, then the children might not generalize much beyond the task itself. That is, even those children who endorse psychosomatic causes in the *Cross Domains* story might be reluctant to endorse psychosomatic causality in general. Alternatively, the children might be more willing to accept the possibility of other psychosomatic events. In Experiment 2, we look at how exposure to the evidence in the *Cross Domains* story affects children's inferences about other psychosomatic events.

Experiment 2

In earlier research on children's understanding of the limits of psychological explanations, Schult and Wellman (1997) showed preschoolers actions that were physically and biologically possible (e.g., jumping up and down; drinking orange juice) or impossible (walking through a wall; staying awake forever). They found that preschoolers distinguished possible and impossible events in both domains; that is, children understood that you could do possible events if you wanted to, but you could not do impossible ones. In Experiment 2, we use a similar method to look at preschoolers' judgments about the possibility of physical, psychological, and psychosomatic events. We ask children to make these possibility judgments either at baseline or after reading the *Cross Domains* story used in Experiment 1. If the children interpret the storybook evidence conservatively, then children who hear the *Cross Domains* story should be no more likely than children at baseline to say that other types of psychosomatic events are possible. However, if children generalize broadly, then children who hear the *Cross Domains* story should be more likely than children at baseline to endorse psychosomatic causality.

Method

Participants

Thirty children (mean age: 58 months; range: 49-71 months) were recruited from urban area preschools. Fifty percent of the participants were girls. While most children were from white, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented. Children were randomly assigned to a

Baseline Possibility Judgments condition or an Evidence And Possibility Judgments condition.

Materials

Six pictures were used (see Appendix C). The pictures showed a physically possible event (throwing a ball in a lake and making a splash); a physically impossible event (brushing a window with a feather and breaking it); a biologically possible event (skipping rope and getting tired); a biologically impossible event (stomping on the ground and making a tomato grow); and two psychosomatic events (worrying and getting a headache; being nervous and feeling sick). In the Evidence And Possibility Judgments condition, the *Cross Domains* storybook from Experiment 1 was also used.

Procedure

Children were tested individually. In the Baseline Possibility Judgments condition, children were shown each of the six pictures in one of two fixed semi-random orders: Order 1) physically impossible; psychogenic headache; biologically impossible; biologically possible; physically possible; psychogenic sickness; Order 2) biologically possible; psychogenic sickness; physically possible; physically impossible; psychogenic headaches; biologically possible. The experimenter read the children a brief passage about the events in the picture (see Appendix C). At the end of each passage, children were asked yes or no questions about the possibility of the event. For example, for the physically impossible event, children were asked: “Can that happen? Can Tony break the window with a feather?” The two psychogenic questions were: “Can that happen? Can Leslie get a headache from worrying too much?” and “Can that happen? Can Jordan start to feel sick from being nervous and upset?” The Evidence And Possibility Judgments

condition was identical to the Baseline Possibility Judgment condition except that children were first tested on the *Cross Domains* storybook as in the Evidence condition of Experiment 1.

Results and Discussion

In the Evidence And Possibility Judgments condition, children's responses to the *Cross Domains* storybook replicated the results for this age group in Experiment 1. Sixty percent of the children chose 'being scared', not significantly different from the 50% who chose 'being scared' in the Evidence condition of Experiment 1 ($\chi^2(1, N = 31) = .125, p = ns$) and significantly more than the 12% of children who chose 'being scared' in the Baseline condition of Experiment 1 ($\chi^2(1, N = 31) = 7.63$).

In the possible/impossible picture task, one child in the Baseline Possibility Judgments condition and one child in the Evidence And Possibility Judgments condition answered, "yes" to all six questions and one child in the Baseline Possibility Judgments condition answered "no" to all six questions. To ensure that children could properly distinguish possible and impossible events, we eliminated these children from further analysis, leaving 13 children in the Baseline Possibility Judgments condition and 14 children in the Evidence And Possibility Judgments condition.

The critical question was whether children would be more likely to say that psychosomatic events were possible in the Evidence And Possibility Judgments condition than in the Baseline Possibility Judgments condition. In fact, there was no difference in children's possibility judgments between the conditions ($\chi^2(1, N = 27) = .07, p = ns$; see Table 1). In both conditions, children denied the possibility of both psychosomatic events significantly more often than expected by chance (by binomial test) and no other

patterns of responses occurred more often than chance (by binomial test). Within the Evidence And Possibility Judgments condition, children who chose ‘being scared’ in the *Cross Domains* story were no more likely than children who chose ‘food’ to say that the other psychosomatic events were possible ($\chi^2 (1, N = 14) = 2.14, p = ns$).

There were no significant differences in children’s tendency to judge the biologically impossible, physically impossible, and psychogenic events as impossible; similarly, children were equally likely to judge the physically possible and biologically possible events as possible (by McNemar’s test, $p = ns$ throughout). In both conditions, children indicated that biologically and physically possible events were possible and biologically and physically impossible events were impossible significantly more often than expected by chance (by binomial test).

The results of Experiment 2 suggest that children interpret the evidence in the *Cross Domains* story quite conservatively. Observing the evidence in the *Cross Domains* story did not affect children’s willingness to accept other causal relationships (e.g., between worrying and headaches or between being nervous and feeling sick). Why is children’s learning so constrained? One possibility is that children’s causal generalizations are affected by their understanding of the domains involved. If preschoolers think that tummyaches, headaches, and ‘feeling sick’ are distinct forms of illness, they might not readily generalize causes of tummyaches to other ailments. (Anecdotally for instance, the children seemed to identify ‘feeling sick’ primarily with vestibular upset and throwing-up; several children volunteered reminiscences on the topic in that context but never otherwise). Alternatively, children’s generalizations might have been affected by the extent to which they treated ‘being scared’, ‘worrying’, and ‘being

nervous’ as the same type of causal event; children might not have appreciated the commonality among the psychological variables. Indeed, different beliefs about the commonality among mental states or bodily states may affect even adults’ generalizations of psychosomatic causality. Adults may accept for instance, that worrying can cause tummyaches but deny that worrying can cause cancer; similarly, they may accept that anxiety can cause headaches but deny that excitement causes headaches.

Alternatively, children might have failed to generalize psychosomatic causality from the *Cross Domains* story to the possible/impossible judgment task simply because the evidence for psychosomatic causality provided by the *Cross Domains* story was relatively weak. Other plausible candidate causes (e.g., food) were always present and children saw only a total of seven trials. Given the conflict between the statistical evidence and children’s prior beliefs, such conservative learning from minimal data is rational; if children have strong prior beliefs and the evidence against these beliefs is relatively limited, children’s naïve theories should be robust to the anomalous data. The results of Experiment 2 suggest that children might engage in just this sort of authentic but conservative learning.

Experiment 3

However, given children’s failure to generalize their inferences, one might be sceptical that children genuinely learn from the statistical evidence in the first place. Although in both Experiments 1 and 2, children were significantly more likely to identify ‘being scared’ as a cause in the *Cross Domains* task after seeing the evidence than at baseline, children did not choose ‘being scared’ significantly more often than chance in either experiment (50% of children chose ‘being scared’ in Experiment 1; 60% chose

‘being scared’ in Experiment 2). Because children had a forced choice of two variables, it is not clear whether the children genuinely learned to infer psychosomatic causes from the evidence or whether the surprising evidence confused the four and five-year-olds and led them to choose at chance.

In Experiment 3, we introduce two measures to distinguish authentic learning from chance performance. First, we modify the *Cross Domains* story so that three candidate causes (one domain-inappropriate, two domain-appropriate) rather than two, covary with the effect every day. If the children are confused by the evidence and choosing at chance, they should choose ‘being scared’ 33% of the time. However, if children genuinely learn from the data, we would expect to replicate the results of Experiments 1 and 2: children should choose ‘being scared’ more often than chance and more often than either of the other variables.

Second, we ask children to extend their inferences from the forced-choice task to a free explanation task. If children genuinely learn the target psychosomatic causal relation (between being scared and tummyaches), they should be able to apply this knowledge to explain a new instance of the same target relation, even if they are unwilling to endorse psychosomatic causes in general. To assess children’s ability to engage in this near transfer of their learning, we read children a passage about a puppy who is worried about the first day of school and has a stomachache. We chose to use the free explanation task because we believed that following-up the storybook task with a possible/impossible judgment question about the same target relation (“Can that happen? Can worrying cause tummyaches?”) might lead the children to believe that we were questioning their original responses. Since preschoolers are vulnerable to changing their answers on repeated

questioning, we believed the free explanation task would be a more sensitive measure of children's understanding.

Children are given the explanation task both at baseline (the Baseline Explanation condition) and after having read the revised *Cross Domains* story (the Evidence and Explanation condition). If the *Cross Domains* story does not affect children's learning, then there should be no difference between the two conditions. However, if children do learn from the evidence in the *Cross Domains* story, they should be more likely to attribute the puppy's stomachache to worrying in the Evidence And Explanation condition than the Baseline Explanation condition.

Method

Participants

Forty children (mean age: 58 months; range: 49-71 months) were recruited from urban area preschools. Fifty percent of the participants were girls. While most children were from white, middle-class backgrounds, a range of ethnicities resembling the diversity of the population was represented. Children were randomly assigned to a Baseline Explanation or an Evidence And Explanation condition.

Materials

The *Cross Domains* storybook of Experiment 1 was modified so that every morning of the week, three events (two domain-appropriate and one domain-inappropriate) and the effect occurred. The domain-inappropriate event (being scared) was repeated each day; the other two events always varied. Thus a sample test page read: "Bunny thinks about show-and-tell, Bunny feels scared. Bunny eats a sandwich. Bunny drinks apple juice. Bunny has a tummy ache." Two different versions of each book were

created so that for half the children ‘being scared’ was the last of the three events and for half the children ‘being scared’ was the first of the three events. Additionally, a novel ‘Puppy’ storybook was used. The text preceding the test question read in its entirety: “This is Puppy. Puppy is worried because next week he starts school. The first day of school makes Puppy worried. Puppy’s stomach hurts.”

Procedure

Children were tested individually. In the Baseline Explanation condition, children were read only the Puppy book. The test question was open-ended: “Why do you think Puppy’s stomach hurts?” Children were asked to offer an explanation and no prompts were given. The Evidence And Explanation condition was identical except that children were first tested on the revised *Cross Domains* book as in Experiments 1 and 2.

Results and Discussion

Children’s responses to the *Cross Domains* story in this experiment replicated the results in Experiments 1 and 2. Fifty-five percent of the children in the Evidence And Explanation condition chose ‘being scared’, comparable to the 50% of children who chose ‘being scared’ in the Evidence condition of Experiment 1 and the 60% of children who chose ‘being scared’ on the storybook task in the Evidence And Possibility Judgment condition of Experiment 2. Although in this experiment children were faced with a choice of three variables rather than two, there were no significant differences between children’s tendency to choose ‘being scared’ in this experiment and the Evidence conditions of Experiments 1 ($\chi^2 (1, N = 36) = .09, p = ns$) and 2 ($\chi^2 (1, N = 36) = .31, p = ns$). Children were significantly more likely to choose ‘being scared’ in this experiment than in the Baseline condition of Experiment 1 ($\chi^2 (1, N = 36) = 6.96$).

Within this experiment, children chose ‘being scared’ significantly above chance ($N = 20$, by binomial test) and did not choose either of the other two variables above chance ($N = 20$, $p = ns$ by binomial test). There was a trend for children to choose ‘being scared’ over any other variable ($\chi^2 (1, N = 20) = 4.0, p = .08$).

On the free explanation task, we coded children’s explanations for reference to physical/bodily variables (e.g., sickness, hunger, or injury) and psychological variables (e.g., worrying about the first day of school). Explanations fell uniquely and unambiguously into a bodily, psychological, or “I don’t know” category. In the Baseline Explanation condition, 55% of the children gave only bodily/physical explanations; 30% of the children gave only psychological explanations and 15% of the children said, “I don’t know”. By contrast, in the Evidence And Explanation condition, 20% of the children referred only to bodily/physical causes; 70% of the children referred only to psychological causes and 10% said, “I don’t know”. The bodily/physical explanations all referred to food or hunger with the exception of a single child in the Baseline Explanation condition who said “itchy stomach”. The psychological explanations all referred to being the first day of school and/or being worried, sad, scared or nervous. Children were significantly more likely to reference psychological causes in the Evidence And Explanation condition than the Baseline Explanation condition ($\chi^2 (1, N = 40) = 6.67$).

We also analyzed the data to see whether the children who chose ‘being scared’ in the *Cross Domains* story were more likely to offer psychological explanations on the puppy book than those who did not. Of the 11 children who chose ‘being scared’ in the *Cross Domains* story, 8 (73%) offered psychological explanations in the transfer task; of the 9 children who did not choose ‘being scared’ in the *Cross Domains* story, 6 (67%)

offered psychological explanations in the transfer task. There was no significant difference between these groups (Fisher's Exact Test, $p = ns$).⁴ This suggests that even those children who did not identify 'being scared' as the causal variable in the *Cross Domains* story may have learned enough from the evidence to treat worrying a relevant causal variable in the free explanation task.

The results of Experiment 3 suggest that children are not merely confused by statistical evidence that violates their prior beliefs; rather children draw accurate inferences from such evidence. Preschoolers were able to use the statistical evidence to identify a psychological variable as a likely candidate cause of a bodily effect in both a forced choice and free explanation task. These experiments suggest that four- and five-year-olds can genuinely learn novel causal relations from limited amounts of data, even when the evidence conflicts with the children's prior beliefs.

Why were children able to transfer their inferences about psychosomatic causes from the *Cross Domains* story to the free explanation task in this experiment but not from the *Cross Domains* story to the possible/impossible judgment tasks in Experiment 2? We believe the difference between the experiments can be explained in a number of ways. First, the change in stimuli across the tasks might have impaired children's transfer of information in Experiment 2; that is, children might have found it more difficult to transfer their inferences from the storybook to the picture tasks than from one storybook to another storybook. Second, the possible/impossible judgment task might have more difficulty than the free explanation task. If so, the greater difficulty of the task might have made the transfer of knowledge less likely. Finally, as hypothesized, children might have been less willing to generalize their knowledge about one psychosomatic causal relation

to psychosomatic causes in general than to extend their inferences about a single psychosomatic causal relation. As noted, given the minimal evidence for psychosomatic causality provided by the *Cross Domains* story, it is rational that children might have interpreted the data conservatively and transferred their inferences more readily in Experiment 3 than in Experiment 2.

General Discussion

Collectively, these three experiments suggest that children learn about causal relationships by taking into account both statistical evidence and constraints from their naïve theories, consistent with the predictions of Bayesian inference models. Given identical evidence, preschoolers were more likely to identify a variable as a cause when the variable was consistent with their theories than when it violated their theories. Older preschoolers were able to use ambiguous, domain-violating evidence to make inferences about psychosomatic causality that they did not make at baseline. Moreover, children were able to learn from the cross-domain evidence even though their initial domain-specific theories were not ruled out but merely rendered less probable by the data. Finally, children's learning was sufficiently robust that children who observed evidence for psychosomatic causality were more likely than children at baseline to offer psychosomatic explanations in a novel task.

The role of Bayesian inference in our analysis of children's ability to combine statistical evidence and constraints from naïve theories was intended to be similar to that of ideal observer analysis in vision (e.g., Yuille & Kersten, 2006) and rational analysis in the study of adult cognition (e.g., Anderson, 1990; Marr, 1982; Shepard, 1987). Bayesian inference provides a rational solution to the problem of updating one's beliefs in the light

of new evidence, and can thus guide us in exploring how well children solve this problem. In particular, a Bayesian model can allow us to make both qualitative predictions about how evidence and theories interact, and quantitative predictions about the conclusions that are warranted from a particular combination of observed data and constraints derived from a theory. The Bayesian model presented in the Appendix provides one such set of predictions, showing that the judgments of the four- and five-year-olds in our experiments are close to the probabilities entertained by an ideal Bayesian learner using a particular causal theory.

We do not claim that Bayesian inference is the only way to define a model that could reproduce our results. While the effect of domain on children's judgments is inconsistent with accounts of causal learning based purely on the strength of association or patterns of covariation between events (Cheng, 1997, 2000; Shanks, 1985; Shanks & Dickinson, 1987; Spellman, 1996), such accounts could predict the data reported here if augmented with initial assumptions about the strength of causal relationships within and across domains. Our intent was not to explore the mechanism by which children make these judgments, but rather whether children solve the abstract computational problem of combining theory and evidence in a way that is consistent with the prescriptions of Bayesian inference. In our studies, an appropriately augmented associative mechanism remains a possible explanation for how children could approximate the rational Bayesian solution to this problem, as it would essentially build in the two critical components of Bayesian inference: initial beliefs regarding possible causal relationships, and revision of these beliefs in light of evidence. Further experiments would be necessary to explore the adequacy of such an account, although we note that the results of several previous studies

of causal reasoning in children would seem to provide evidence against simple associative models as a general explanation for performance on this kind of task (e.g., Sobel, Tenenbaum, & Gopnik, 2004; Tenenbaum, Sobel, Griffiths, & Gopnik, in submission).

The results of our experiments also raise several other questions. As noted, it is not clear whether the developmental differences between three-year-olds, three-and-a-half-year-olds, and four- and five-year-olds are due to changes in children's ability to evaluate evidence, changes in children's initial inductive biases, or both. The relative influence of domain-specific and domain-general information might change dramatically over the course of development, depending on age-related commonalities and differences in children's commitment to domain-specific theories, sensitivity to evidence, and ability to integrate the two. Further research might disambiguate the role of each of these factors.

Interestingly however, the finding that three-and-half-year-olds did not use the theory-violating evidence to change their judgments is consistent with a wide variety of research suggesting that three-year-olds have particular difficulty changing their minds in the face of evidence. On paradigms as diverse as theory of mind tasks, ambiguous figure tasks, and Dimensional Change Card Sort tasks, three-year-olds' initial inferences seem to be remarkably impervious to feedback (Gopnik & Rosati, 2001; Kirkham, Cruess, & Diamond, 2003; Munakata & Yerys, 2001; Russell, Jarrold & Potel, 1994; Zelazo, Frye, & Rapus, 1996). Researchers have suggested process-level theories, such as the theory of attentional inertia (e.g., Diamond & Kirkham, 2005; Kirkham, et al., 2003) to account for many of these phenomena. In future work, it would be interesting to investigate the relationship between such process level theories and computational level accounts of the

biasing effects of prior knowledge, like the Bayesian model proposed here.

Additionally, the studies here do not tell us precisely *what* children learned. We have suggested that children can use domain-general evidence to learn at least one particular psychosomatic relationship: that being scared causes tummyaches. We also know that children did not learn to accept that psychosomatic causes were possible in general. However, children's learning might have been either more narrow or more broad than this summary suggests. Children might have learned that being scared can cause tummyaches only in the context of storybooks -- or they might have begun, but not completed, a process of fundamentally altering their understanding of domain boundaries. Additional research might establish the extent to which patterns of evidence can influence children's naive theories.

Further research might also establish the extent to which children's learning is affected by varying the amount and type of evidence children observe. In our studies, children saw candidate causes covary deterministically with an effect seven times. Differences in the quantity, quality and presentation of the data (more trials, negative evidence, evidence from interventions, probabilistic evidence, etc.) might influence both children's willingness to override domain-specific beliefs and their willingness to generalize from the data to novel events.

Note also that in this study, children saw a single consistent pattern of evidence: candidate causes were paired together and one variable was always held constant while the other variable always changed. In the real world, data are unlikely to be packaged in such a consistent manner. Causes often act stochastically, multiple variables can change simultaneously, and unobserved causes may be present. In such contexts, it may far more

difficult both to detect and to draw inferences from recurring variables. Furthermore, children might be better able to track evidence in the pedagogical context of a story than in the world at large; conversely, the inferences children make about events in a story might be particularly unlikely to generalize beyond the story itself. Further research might investigate children's ability to draw inferences from ambiguous evidence in a broader range of contexts and in cases where the presentation of the data is less tightly controlled.

Finally, in our studies, children observed the evidence in the absence of any explanation of *how* psychological events might cause bodily events. Many researchers have proposed that an understanding of causal mechanisms is fundamental to an understanding of causal relationships (Ahn, et al., 1995; Bullock, Gelman, & Baillargeon, 1982; Koslowski, 1996; Shultz, 1982) and it would be interesting to know how offering children explicit information about causal mechanisms might affect their learning. It seems probable that children might be more willing to learn from a combination of evidence and information about plausible processes of causal transmission than from evidence in isolation. Conversely, researchers have suggested that evidence about the covariation of interventions and outcomes can support inferences about causal mechanisms (Schulz, Kushnir, & Gopnik, in press; Schulz & Sommerville, 2006). It would be interesting to know whether evidence for novel observed causal relations might prompt children to posit novel mechanistic explanations.

More generally, it seems probable that children begin to learn about psychosomatic causality, not merely because they observe or are told about covariations between psychological and bodily events but because adults explicitly assert the existence (or non-

existence) of such causal relationships. We do not know how merely telling children about a causal relationship affects children's interpretation of evidence. Nor do we know how different cultural beliefs about psychosomatic causality might influence children's learning. Further research might investigate the interaction between information conveyed through cultural transmission and children's learning from evidence.

What these studies do suggest is that Bayesian inference captures a hallmark of causal learning in early childhood: conservatism with respect to prior knowledge but flexibility in the face of new evidence. Although learners might lack a metacognitive understanding of the relationship between theories and evidence, rational computations integrating new data and prior knowledge could form part of an implicit human learning mechanism, allowing the process of theory formation to be both adaptive and stable. The results of these studies suggest that even very young children can integrate prior knowledge and evidence to make normative causal inferences, giving children a powerful mechanism for developing and revising their naive theories about the world.

References

- Ahn, W. K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, 54(3), 299-352.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.
- Barrett, S. E., Abdi, H., Murphy, G. L., & Gallagher, J. M. (1993). Theory-based correlations and their role in children's concepts. *Child Development*, 64, 1595-1616.
- Bartsch, K., & Wellman, H. (1989). Young Children's Attribution of Action to Beliefs and Desires. *Child Development*, 60, (4), 946-964.
- Bloom, P. (2004). *Descartes' Baby*. New York: Basic Books.
- 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, NY: Academic Press.
- Carey, S. (1985). *Conceptual change in childhood* Cambridge, MA: MIT Press/Bradford Books.
- Carey, S., & Spelke, E. S. (1994). Domain-specific knowledge and conceptual change. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture; Based on a conference entitled "Cultural Knowledge and Domain Specificity," held in Ann Arbor, MI, Oct 13-16, 1990* (pp. 169-200). New York, NY: Cambridge University Press.
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, 70(5), 1098-1120.

- Cheng, P. W. (2000). Causality in the mind: Estimating contextual and conjunctive power. In F. C. Keil, R. A. Wilson (Eds.), *Explanation and cognition* (pp. 227-253). Cambridge, MA, US: The MIT Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367-405.
- Diamond, A., & Kirkham, N. (2005). Not quite as grown-up as we like to think. *Psychological Science*, 16(4), 291-297.
- Dunbar, K. & Klahr, D. (1989). Developmental differences in scientific discovery processes. In D. Klahr and K. Kotovsky (Eds.), *The 21st Carnegie-Mellon symposium on cognition: Complex information processing: The Impact of Herbert A. Simon*. Hillsdale, NJ: Lawrence Erlbaum.
- Estes, D., Wellman, H. M., & Woolley, J. (1989). Children's understanding of mental phenomena. In H. Reese (Ed.), *Advances in child development and behavior* (pp. 41-86). New York: Academic Press.
- Flavell, J. H., Green, F. L., & Flavell, E. R. (1995). Young children's knowledge about thinking. *Monograph of the Society for Research in Child Development*, Serial #24, Vol. 60, 1.
- Gelman, S. A., & Wellman, H. M. (1991). Insides and essence: Early understandings of the non-obvious. *Cognition* 38(3), 213-244.
- Gopnik, A. (1988). Conceptual and semantic development as theory change. *Mind and Language*, 3(3), 197-217.
- Gopnik, A., & Meltzoff, A. N. (1997). *Words, thoughts and theories*. Cambridge, MA: MIT Press.
- Gopnik, A., & Rosati, A. (2001) Duck or rabbit? Reversing ambiguous figures and

- understanding ambiguous representations. *Developmental Science*, 4, 175-183.
- Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology*, 37(5), 620-629.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111, 1-31.
- Griffiths, T. L. (2005). *Causes, coincidences, and theories*. Unpublished doctoral dissertation, Stanford University, Stanford CA.
- Harris, P. L., German, T., & Mills, P. (1996). Children's use of counterfactual thinking in causal reasoning. *Cognition*, 61(3), 233-259.
- Hatano, G., & Inagaki, K. (1994). Young children's naive theory of biology. *Cognition*, 50, 171-188.
- Hickling, A. K., & Wellman, H. M. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology*, 37(5), 668-683.
- Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind body distinction. *Child Development* 64(5), 1534-1549.
- Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence* (A. P. S. Milgram, Trans.). New York: Basic Books.
- Kalish, C. (1996). Causes and symptoms in preschoolers' conceptions of illness. *Child Development*, 67(4), 1647-1670.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

- Keil, F. C. (1995). The growth of causal understandings of natural kinds. In D. P. D. Sperber (Ed.), *Causal cognition: A multidisciplinary debate. Symposia of the Fyssen Foundation; Fyssen Symposium, 6th Jan 1993, Pavillon Henri IV, St-Germain-en-Laye, France*. New York, NY: Clarendon Press/Oxford University Press.
- Kirkham, N., Cruess, L., & Diamond, A. (2003) Helping children apply their knowledge to their behavior on a dimension switching task. *Developmental Science*, 6(5), 449-476.
- Klahr, D. & Dunbar, K. (1988). Dual Space search during Scientific Reasoning. *Cognitive Science*, 12, 1-48.
- Koslowski, B. (1996). *Theory and Evidence: The Development of Scientific Reasoning*. Cambridge, MA: MIT Press.
- Koslowski, B., Okagaki, L., Lorenz, C., & Umbach, D. (1989). When covariation isn't enough: The role of causal mechanism, sampling method and sample size in causal reasoning. *Child Development*, 60, 1316-1327.
- Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, 96(4), 674-689.
- Kuhn, D., Amsel, E., & O'Laughlin, M. (1988). *The development of scientific thinking skills*. Orlando, FL: Academic Press.
- Kushnir & Gopnik (2006). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*.
- Leslie, A. M. (1994). ToMM, ToBy, and Agency: Core architecture and domain specificity. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture; Based on a conference entitled "Cultural Knowledge and Domain*

- Specificity," held in Ann Arbor, MI, Oct 13-16, 1990* (pp. 119-148). New York, NY: Cambridge University Press.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman.
- Masnick, A. M., & Klahr, D. (2003). Error Matters: An initial exploration of elementary school children's understanding of experimental error. *Journal of Cognition and Development*, 4(1), 67-98.
- 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., & Yerys, B.E. (2001) All together now: when dissociations between knowledge and action disappear. *Psychological Science*, 12, 335-337.
- Notaro, P. C., Gelman, S. A., & Zimmerman, M. A. (2001). Biases in reasoning about the consequences of psychogenic bodily reactions: Domain boundaries in cognitive development. *Merrill-Palmer Quarterly*, 48, 427-449.
- Pazzani, M.J. (1991A Computational Theory of Learning Causal Relationships. *Cognitive Science*, 15(3): 401-424 (1991).
- Penner, D. E., & Klahr, D. (1996). The interaction of domain-specific knowledge and domain-general discovery strategies: A study with sinking objects. *Child Development*, 67, 2709-2727.
- Perner, J. (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press.
- Russell, J., Jarrold, C., & Potel, D. (1994) What makes strategic deception difficult – the deception or the strategy? *British Journal of Developmental Psychology*, 12, 301-314.
- Schauble, L. (1990). Belief Revision in Children: The Role of Prior Knowledge and Strategies for Generating Evidence. *Journal of Experimental Child Psychology*, 49,

31-57.

Scholl, B.J., & Leslie, A.M. (1999). Modularity, development and ‘theory of mind’. *Mind & Language*, 14, 131–153.

Schultz, 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(2), 162-176.

Schulz, L. E., Kushnir, T., & Gopnik, A. (in press). Learning from doing: Interventions and causal inference. In A. Gopnik & L. E. Schulz (Eds.), *Causal Learning; Psychology, Philosophy and Computation*. New York: Oxford University Press.

Schulz & Sommerville, (2006). God does not play dice: Causal determinism and preschoolers’ causal inferences. *Child Development*, 77, 2, 427-442.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 37(1), 1-21.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 21 (pp. 229-261). San Diego, CA, US: Academic Press, Inc.

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317–1323.

Shultz, T. R. & Mendelson, R. (1975). The use of covariation as a principle of causal analysis. *Child Development*, 46, 394-399.

Shultz, T. R. (1982). Rules of causal attribution. *Monographs of the Society for Research in*

- Child Development*, 47(1), 1-51.
- Siegler, R. S., & Liebert, R. M. (1975). Acquisition of formal scientific reasoning by 10- and 13-year olds: Designing a factorial experiment. *Developmental Psychology*, 10, 401-402.
- Sobel, D. M. (2004). Exploring the coherence of young children's explanatory abilities: Evidence from generating counterfactuals. *British Journal of Developmental Psychology*, 22, 37-58.
- Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28(3), 303-333.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. (1992). Origins of Knowledge. *Psychological Review*, 99(4), 605-632.
- Spellman, B. A. (1996). Acting as intuitive scientists: Contingency judgments are made while controlling for alternative potential causes. *Psychological Science*, 7(6), 337-342.
- Tenenbaum, J.B., Griffiths, T. L., and Niyogi, S. (in press). Intuitive theories as grammars for causal inference. In Gopnik, A., & Schulz, L. (Eds.), *Causal learning: Psychology, philosophy, and computation*. Oxford University Press.
- Tenenbaum, J. B. and Niyogi, S. (2003). Learning Causal Laws. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*.
- Tenenbaum, J. B., Sobel, D. M., Griffiths, T. L., and Gopnik, A. (in submission). Bayesian inference in causal learning from ambiguous data: Evidence from adults and children.
- Watson, J. S., & Ramey, C. T. (1972). Reactions to response-contingent stimulation in early infancy. *Merrill-Palmer Quarterly*, 18(3), 219-227.

- Wellman, H. M., & Estes, D. (1986). Early understanding of mental entities: a reexamination of childhood realism. *Child Development*, 57, 910-923.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., Hickling, A. K., & Schult, C. A. (1997). Young children's psychological, physical, and biological explanations. In H. M. Wellman & K. Inagaki (Eds.), *The emergence of core domains of thought: Children's reasoning about physical, psychological, and biological phenomena. New directions for child development No. 75* (pp. 7-25). San Francisco, CA: Jossey-Bass/Pfeiffer.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, 10, 301-308.
- Zelazo, P.D., Frye, X., & Rapus, T. (1996) An age-related dissociation between knowing rules and using them. *Cognitive Development*, 11, 37-63.

Appendix A

Bayesian inference provides a rational account of how children should go about combining theory and evidence. It can also be used to make precise quantitative predictions as to whether conclusions are justified by the observation of data, given a set of assumptions about the constraints provided by naïve theories. In Experiment 1, children are asked “Why does [character] have [symptom]? Is it because of [A] or because of [B]?” We model the probability that children choose explanation A as

$$\frac{P(\text{Explanation A} \mid D)}{P(\text{Explanation A} \mid D) + P(\text{Explanation B} \mid D)} \quad (2)$$

This directly contrasts the two possible explanations given the data observed (a similar expression applies in Experiment 2, contrasting the three possible explanations). The probability of each candidate explanation being selected given the data is computed by summing over all possible causal models that are consistent with the explanation. This is formalized as:

$$P(\text{Explanation A} \mid D) = \sum_{h \in H} P(\text{Explanation A} \mid h)P(h \mid D) \quad (3)$$

where h is a hypothesis as to the underlying causal structure, and H is the space of all hypotheses.

We represent hypotheses using causal graphical models (Pearl, 2000; Spirtes, Glymour, & Scheines, 1993), where nodes correspond to variables, arrows from cause to effect represent relationships between variables, and a set of conditional probability distributions captures the probability that each variable takes on a particular value given the values of its causes. We assume that the probability of a cause being selected as an

explanation given a particular causal structure h in is $1/k$, where k is the size of the set of candidate causes that are present and possess a causal relationship with the effect in h , and where the proposed explanation is a member of this set. The probability of a particular causal structure given the data is obtained via Bayes' rule (Equation 1), using a prior $P(h)$ and likelihood $P(D|h)$ derived from a causal theory.

As proposed by Tenenbaum and Niyogi (2003), Griffiths (2005), and Tenenbaum, Griffiths, and Niyogi (in press), we model the framework theory that guides children's inferences as a simple scheme for generating causal graphical models. In this scheme, we allow for different domains. Causal variables have relationships with effect variables; causes are likely to have relationships to effects within their domain, however, there is also a small probability that a cause from one domain can lead to an effect in another domain.

The prior probability associated with each model is simply its probability of being generated by the theory. The process of generating a causal graphical model from this theory breaks down into three steps. First, we identify the nodes (causes and effects) in the model. In our case, the nodes simply correspond to the set of causes and effects that appear in the story. Second, we generate the causal relationships between these nodes. If cause and effects are within-domain, then the probability a relationship exists is relatively high and given by p . If the link between two variables crosses domains, then a relationship is unlikely, and is given a lower probability, q . With n causes, there are 2^n possible causal models. Assuming that each relationship is generated independently, we can evaluate the prior probability of each of these models by multiplying the probabilities of the existence or non-existence of the causal relationships involved. The particular

values of the probabilities p and q depend on the child's theory. Such theories might change with age and experience; that is, younger children might think cross-domain events are more or less probable than older children. We assume that children think the probability of cross-domain events is low (but not extremely low) by setting $q = .1$, and by setting a higher within-domain probability $p = .4$.

Finally, we specify the conditional probability of the effect given the causes present in the causal model. This allows us to evaluate the probability of a specific model, h , generating the data observed on the m th day, $P(d_m|h)$. These data consist of the values taken on by all variables on that day – the presence or absence of the causes and effects. We assume that the probability of each cause being present or absent is constant across all of the causal models, and the only difference is in the probability they assign to the occurrence of the effect on that day. We then take the conditional probability of the effect given the set of causes to be 1 if any cause which influences the effect is present, and ϵ otherwise, corresponding to a noisy-OR parameterization (Pearl, 1988) where each cause has a strength of 1 and the background has a strength of ϵ . We assumed that the probability of an effect in the absence of any causes was low, with $\epsilon = .001$. The probability of the full set of data, D , accumulated over the course of the story is given by

$$P(D|h) = \prod_m P(d_m | h) \quad (4)$$

where the data observed on each day are assumed to be generated independently.

As can be seen comparing the results predicted by the Bayesian model in Figure 4 with the four-year-olds' responses in Experiments 1 and 2, our model accurately predicted the responses of the oldest children, with a Pearson correlation coefficient of $r(9) = .85$. The model gives correct relative weights to the variables at baseline in both

the *Within Domain* and *Cross Domains* conditions. Critically, the model predicted the increased A responses after evidence in all conditions, while still capturing the more subtle graded interaction between theory and evidence.

Insert Figure 3 here

Appendix B

Text of *Within Domain* storybook used in Experiment 1

Title: Bambi's Adventures

This is Bambi. Bambi likes to prance and run in lots of different places. Running is fun for Bambi. On Monday morning Bambi runs in the pine grove. Bambi gets excited. Bambi runs in the cattails. Bambi has itchy spots on his legs. On Monday afternoon Bambi runs in the cedar trees and Bambi plays on the rope swing. Bambi feels great! Bambi doesn't have any itchy spots. On Tuesday morning Bambi gets excited. Bambi runs in the cattails. Bambi runs in the grass. Bambi has itchy spots on his legs. On Tuesday afternoon Bambi reads a book and Bambi runs through the rock bed. Bambi feels great! Bambi doesn't have any itchy spots. On Wednesday morning Bambi runs in the marsh. Bambi gets excited. Bambi runs in the cattails. Bambi has itchy spots on his legs. On Wednesday afternoon Bambi runs through the apple orchard and Bambi plays with his toy truck. Bambi feels great! Bambi doesn't have any itchy spots. On Thursday morning Bambi gets excited. Bambi runs in the cattails. Bambi runs in the leaves. Bambi has itchy spots on his legs. On Thursday afternoon Bambi plays jump rope and Bambi runs in the sand. Bambi feels great! Bambi doesn't have any itchy spots. On Friday morning Bambi runs in the bushes. Bambi gets excited. Bambi runs in the cattails. Bambi has itchy spots on his legs. On Friday afternoon Bambi runs through the playground and Bambi roller skates. Bambi feels great! Bambi doesn't have any itchy spots. On Saturday morning Bambi gets excited. Bambi runs in the cattails. Bambi runs in the grass. Bambi has itchy spots on his legs. On Saturday afternoon Bambi gets his hair brushed and Bambi runs through the blueberry patch. Bambi feels great! Bambi doesn't have any itchy spots. On Sunday morning Bambi runs through the garden. Bambi gets excited. Bambi runs in the cattails. Bambi has itchy spots on his legs. The next day Bambi's spots were all gone. Have fun Bambi! The End.

Text of Cross Domain storybook used in Experiment 1

Title: Bunny's Big Week

This is Bunny. Bunny is scared because next week she has to give show-and-tell. Show-and-tell makes Bunny scared. On Monday morning Bunny thinks about show-and-tell. Bunny feels scared. Bunny eats some cheese. Bunny has a tummyache. On Monday afternoon Bunny ties her shoes and Bunny eats strawberries. Bunny feels great! Bunny doesn't have a tummyache. On Tuesday morning Bunny eats a popsicle. Bunny thinks about show-and-tell. Bunny feels scared. Bunny has a tummyache. On Tuesday afternoon Bunny eats some toast and Bunny takes a bath. Bunny feels great! Bunny doesn't have a tummyache. On Wednesday morning Bunny thinks about show-and-tell. Bunny feels scared. Bunny eats French fries. Bunny has a tummy ache. On Wednesday afternoon Bunny plays bingo and Bunny eats pasta. Bunny feels great! Bunny doesn't have a tummyache. On Thursday morning Bunny eats a muffin. Bunny thinks about show-and-tell. Bunny feels scared. Bunny has a tummyache. On Thursday afternoon Bunny eats some yogurt and Bunny brushes her teeth. Bunny feels great! Bunny doesn't have a tummyache. On Friday morning Bunny thinks about show-and-tell. Bunny feels scared. Bunny eats some soup. Bunny has a tummyache. On Friday afternoon Bunny plays on the monkey bars and Bunny eats a banana. Bunny feels great! Bunny doesn't have a tummyache. On Saturday morning Bunny eats a carrot. Bunny thinks about show-and-tell. Bunny feels scared. Bunny has a tummyache. On Saturday afternoon Bunny eats some tofu and Bunny builds a snowman. Bunny feels great! Bunny doesn't have a tummyache. On Sunday morning Bunny thinks about show-and-tell. Bunny feels scared. Bunny has a tummyache. The next day Bunny gave show-and-tell. She did very well and everyone clapped! Hurray for Bunny! The End

Appendix C

Text and pictures used in the *Far Transfer* study of Experiment 2

Physically possible: This is Alex. Alex is at the playground. Alex is throwing a ball. Alex is throwing the ball near a lake. Alex' friend tells Alex that if he throws the ball into the water it will make a big splash. Can that happen? Can Alex make a splash by throwing the ball into the water?

Physically impossible: This is Tony. Tony is at the park. Tony picks up a feather. Tony brushes the feather against a car window. Tony's friend tells Tony that if he keeps brushing the feather on the window, the window will break. Can that happen? Can Tony break the window with a feather?

Biologically possible: This is Erin. Erin is in her backyard. Erin is jumping rope. Erin has been jumping rope for a long time. Erin's friend tells Erin that if she keeps jumping rope s/he will get very tired. Can that happen? Can Erin get tired from jumping rope for a long time?

Biologically impossible: This is Mel. Mel is in the garden. Mel is playing with the soil. Mel pats the soil with her hand. Mel's friend tells her that if s/he pats the soil a lot, the soil will sprout a tomato. Can that happen? Can Mel make a tomato grow by patting the soil?

Psychogenic headache: This is Leslie. Leslie is on the school bus. It's Leslie's first day of school today. Leslie's is worried about the first day of school. Leslie worries and worries. Leslie's friend tells her that if she keeps worrying she'll get a headache. Can that happen? Can Leslie get a headache from worrying too much?

Psychogenic sickness: This is Jordan. Jordan is in his bedroom. Jordan is upset and nervous because he has to stay with a babysitter. Jordan feels very upset about the babysitter. Jordan's friend tells Jordan that if he keeps being upset and nervous he will start to feel sick. Can that happen? Can Jordan start to feel sick from being nervous and upset?

Footnotes

¹ We are grateful to an anonymous reviewer for stressing the importance of this distinction.

² Note that this experiment compares only children's different judgments about the mind and the body. It does not require children to have a fully elaborated naïve biology (see e.g., Carey, 1985).

³ In the *Cross Domains* story, the sentence "Bunny thinks about show-and-tell" precedes the target variable ("Bunny feels scared") in order to delimit the onset of the psychological cause. In order to match the evidence in the two books precisely, we included the sentence "Bambi gets excited" before the target cause: "Bambi runs in the cattails". It is possible that the inclusion of the lead-in sentence aided children in tracking the evidence in both conditions. However, the inclusion of the lead-in sentence did not affect children's causal inferences at Baseline, nor can it account for differences between the *Within* and *Cross Domains* story.

⁴ A Fisher's Exact Test was used in this case because the expected counts in the cells of the contingency table were less than four.

Table 1. Children's possible/impossible judgments in Experiment 2

	Number of children who said that the event could happen (all others said the event could not happen).	
	Baseline Possibility Judgments (n = 13)	Evidence and Possibility Judgments (n = 14)
Biologically Possible	12 (92)	11 (79)
Biologically Impossible	2 (15)	2 (21)
Physically Possible	13 (100)	14 (100)
Physically Impossible	0 (0)	0 (0)
Psychogenic (headaches)	4 (31)	5 (36)
Psychogenic (sickness)	3 (23)	5 (36)

Note: Percentages in parentheses.

Figure Captions

Figure 1. Sample pages from the storybook used in Experiments 1 and 2.

Figure 2. Children's responses to the storybook task in Experiment 1. The vertical axis shows the number of children selecting the different responses.

Figure 3. Predictions of the Bayesian inference model compared to four and five-year-olds' responses to the storybook task in Experiments 1, 2, and 3, when q (cross-domain probability) = .1 and p (within-domain probability) = .4.

(A) Within Domain

(B) Cross Domains