Developing Explicit Understanding of Probabilistic Causation: Patterns and Variation in Young Children's Reasoning

Tina Grotzer, Leslie Duhaylongsod, and M. Shane Tutwiler
Harvard Graduate School of Education

Presented at the American Educational Research Association
(AERA) Conference, New Orleans, LA, April 9, 2011.

Developing Explicit Understanding of Probabilistic Causation:
Patterns and Variation in Young Children's Reasoning

Tina Grotzer, Leslie Duhaylongsod, and M. Shane Tutwiler
Harvard Graduate School of Education

ABSTRACT

Reflective understanding of probabilistic causality is critical for learning the science of our complex world. Microgenetic studies were carried out over the school year in kindergarten and second grade (n = 8) to assess children's assumptions about stochastic tasks from four domains: social; games; machines; and biology, and to attempt to build understanding. Sessions were analyzed to identify whether and when shifts took place and what elicited them. The results are interpreted through the lens of Siegler's "Overlapping Waves Theory" (1996) along five dimensions: path, rate, breadth, source, and variability. Most children responded deterministically across all domains with some movement towards accepting the possibility that causes can behave probabilistically. However, two children accommodated the possibility of probabilistic causal forms early and in multiple responses.

OBJECTIVES

Many everyday causal patterns do not give the appearance of being deterministic. The same causal event, pushing a button in a game or calling out to a friend, can lead to a specific outcome sometimes, but not others. Further, the patterns of contingency can vary significantly. One push may lead to an outcome sometimes, yet at others it may take three pushes for the same outcome. Sometimes these patterns fall within the boundaries of how we attend to the problem space, so we notice what the contingencies are. However, just as often, they fall beyond our attentional boundaries. Seeds take a long time to grow thus introducing time delays; the intricacies of a social interaction increase in cognitive load until it is difficult to reason well about them; plausible mechanisms become increasingly non-obvious such that detecting them invites a long reductionist investigation. Our attention shifts to other stimuli, perhaps leading us to frame the patterns in our world in more limited ways than the inherent feedback in our environment could allow for.

A significant body of research investigates how we make causal assumptions in the context of probabilistic causal events (See Gopnik & Schulz, 2007). However, little research has investigated how children develop an explicit and reflective awareness of the nature of probabilistic causation and its impact on what we know and think. This study investigated how kindergarteners and second graders reasoned about tasks with probabilistic causal features across domains given scaffolds to make the nature of the causal patterns explicit.

THEORETICAL FRAMEWORK

Probabilistic causation has long been of interest to researchers, in part, for what it can tell us about how we determine the existence of a causal relationship (e.g. Einhorn & Hogarth, 1986; Kahneman, Slovic, & Tversky, 1982) and how sum across multiple occurrences of an event and assess how other events co-vary to suggest the possibility of a causal connection. Extensive research suggests that children use co-variation data in combination with spatial and temporal contiguity (Borton, 1979; Leslie, 1982, 1984; Leslie & Keeble, 1987; Oakes, 1993; Spelke, Phillips & Woodward, 1995; Van de Walle & Spelke, 1993) and information about plausible mechanisms in assessing causality (e.g. Bullock, 1979). Earlier research suggested that children expect reliable cause-effect relationships (e.g. Bullock, 1985; Bullock, Gelman, & Baillargeon, 1982; Shultz, 1982) and that they use consistent covariation to determine whether or not a causal relationship exists. The tendency appeared age-related with the youngest children accepting less than perfect correlation, presumably due to the cognitive load of tracking perfect correlation (Shultz & Mendelson, 1975; Siegler, 1976; Siegler & Liebert, 1974).

However, recent research challenges these findings, suggesting that even young children follow Bayesian rules in summing across experiences in their causal reasoning. Gopnik and colleagues (2004) argue that young children override imperfect correlation and are able to use different patterns of probability in contiguity to make accurate causal inferences (Kushnir & Gopnik, 2007). However, Schulz and Sommerville (2006) found that preschoolers prefer deterministic over probabilistic causality, at least in the instance of a machine-like toy box mechanism. At issue appear to be two questions: 1) whether young children actually override imperfect correlation or simply don't track correlation well enough such that they make probabilistic choices; and 2) whether they make assumptions of determinism despite the structure of the causal mechanism, stochastic or not.

This study probed the latter question. We were interested in how children chose to treat phenomenon that gave the strong appearance of being stochastic. In all cases, with considerable analysis, calculation, and cognitive load, a deterministic pattern could be uncovered. However, uncovering such a pattern would involve extensive exploration of the task constraints, for instance, determining how multiple mechanisms interacted and in what sequence, to enable an outcome. Discovering the pattern involved one or more of the following complications: 1) complex calculations beyond the capacity of most kindergartners and second graders; 2) extensive time to figure out the deterministic pattern; 3) an enclosed mechanism that obscured how it worked; and 4) a long reductionist path with many points where the information was simply not available (as in seeds that did not sprout).

Additionally, the study considered children's *growing reflective knowledge* of probabilistic causality. Bullock and colleagues (1982) argued that determinism is a fundamental, innate causal principle. Some studies have considered implicit changes in detection rates based upon level of reliability between causes and effects. Siegler and Liebert (1974) found that when the degree of covariation was varied from 100% to 50% between events and there was variation in temporal contiguity (immediate vs. 5 second delay), 8- and 9-year-olds were more sensitive to the lack of perfect covariation than the 5-year-olds who seemed to be distracted by the time delays and perhaps were less likely to notice the lack of perfect covariation.

These findings suggest that shifts are taking place during this time in how students implicitly attend to imperfect covariation. Implicit expectations about how the degree of covariation impacts causation can impact how students reason about evidence in science. Fourth graders changed their causal model for lightning when they realized that lightning does not always "strike" in a high place (Grotzer, 2003). Kalish (1998) found that 4- and 5-year-olds expected deterministic cause-effect relationships—if everyone in a classroom played with a sick child, all or none would get sick). Therefore this study looked closely at children ages five through eight when some of these shifts may be taking place to further characterize shifts that may be taking place.

A further point of discussion in the literature is whether children use separate conceptual systems (or folk theories) for biological, psychological, and physical concepts (e.g. Carey & Spelke, 1994; Keil, 1994, 1995; Leslie & Keeble, 1987; Schulz & Gopnik, 2004). This leads to the question of whether there contexts in which children prefer or accept probabilistic causality. For instance, physical causality might be more likely to invoke deterministic responses than psychological causality. A child might well realize that the same misbehavior might eventually get a response from mom, but it won't reliably get a response every time. For this reason, we interviewed students across a range of tasks aiming to increase the likelihood that children would engage in probabilistic causal reasoning and to provide insights as to whether some tasks could be used in service of building understanding of others.

METHODS

Design: Microgenetic studies were carried out with four students from each grade level in kindergarten and second grade. Across the school year, students were interviewed in depth (n = 8) on tasks with varying levels of stochastic behavior from four contexts: biological, mechanical, social, and games. The number of sessions per child varied between 4 to 6 sessions for kindergarteners and from 5 to 10 sessions for second graders. Students were interviewed at higher density during points when their experiences in the classroom and study suggested that change was likely—as called for in microgenetic analysis (e.g. Opfer & Siegler, 2004). Prior to the fifth session, unless the student generated examples from other contexts, only one context was focused on to avoid cognitive load. In later sessions, more than one context was discussed to allow for contrasting examples and mapping the analogous deep structures. Interviews proceeded from open-ended to increasingly structured, to first assess how students frame the concepts and to then assess the accessibility of concepts related to the stochastic behavior with targeted questions. Scaffolds that made use of familiar examples and compared analogous causal forms in different problem contexts through "mutual alignment" (e.g. Kurtz, Miao, & Gentner, 2001) were incorporated in the form of design studies (Brown, 1992; Collins, 1999). Sessions were videotaped for later coding and analysis.

Subjects: Students from a local charter school participated in the study. The students are from the City of Boston and Cambridge and are primarily Black and Latino with less than 1% Caucasian and 78% on "Free or Reduced Lunch."

Tasks: The tasks from four domains: games, biological, mechanical, social, and games included seed planting, hatching chicks, bubble gum machines, videotapes of brief social interactions, and a set of games. Tasks that children might be familiar with from their everyday worlds were intentionally chosen to elicit their expectations and existing knowledge. This departs from other studies where the intention is to create causes that children would not have prior expectations about that they might bring to bear on the context (e.g. Schulz & Somerville, 2006; Sobel & Buchanan, 2009). The tasks did not all fall neatly into one category or another. For instance, some of the games had mechanical aspects in addition to their game features. However, the rationale for assessing understanding across domains was to find tasks that could potentially invite probabilistic causal reasoning, not to pit domains against one another to discern fine distinctions in reasoning in different domains. Further, we chose mechanical tasks that would potentially invite concepts of probabilistic causation (i.e. a bubble gum machine that did not regularly deliver candy) instead of a machine that was characterized by highly reliable causal processes.

Tasks with the following features were chosen: 1) cognitive load that was not directly related to the probabilistic causal features of the task was low (or minimized by minor modifications); 2) authenticity (that could be maintained with simple controls, for instance, in growing seeds, we would be able to adjust the number of seedlings "that grew" without impacting the authenticity); 3) manageability in the classroom; and 4) minimal competition from other forms of causal complexity or competing knowledge (for instance, in Gopnik and Schulz (2004) preschoolers are given a "biology task" wherein a stuffed monkey sneezes or not when he sticks his nose on a flower to determine whether or not he is allergic to each flower. This pits authentic knowledge that children might hold about sneezing due to allergies which typically occurs after a delay against the task goal of assessing acceptance of probabilistic causation.) Tasks with complexifying causal features were included only if these features were an inherent part of the perceived probabilistic causality such as in the non-obvious causes involved in knowing what happens to seeds under the ground. Some tasks were modified for cultural sensitivity. (For instance, a Hasbro game, "Don't Wake Daddy," was modified to "Don't Wake the Sleeping Bear" because most children were not Caucasian and were from single parent homes.) The following tasks were included:

<u>Games: Funny Bunny and "Last Bunny Standing"</u>: Funny Bunny is a commercial game created by Ravensburger for 4-8 year olds. None of the students reported seeing the game prior to the

experimental sessions. The goal of the game is to be the first one to move your rabbit along a path with two loops up the hill to the top of the big carrot. Cards tell how many steps to move. However, some cards direct player to click the carrot in the middle of the game board. When the carrot is turned, a hole opens up somewhere along the path most of the time and one's rabbit can fall through. The location of where the hole opens up gives the appearance of being stochastic in the following ways. Initially, there is no indication that the hole moves. Upon the turn of the carrot, the hole moves along the path, alternating between the top row of the path and the bottom row of the path. Periodically, no hole opens at all. The cognitive load of figuring out which hole will open involves detecting that some spaces hold the possibility of opening (nine out of 26 are "wiggly" or "soft" whereas others never open and are always safe); detecting that the holes move in a clockwise fashion around the board; that they alternate between the top and bottom rows; and that the hole also disappears at a certain point in the rotation.

"Last Bunny Standing" is a version of Funny Bunny where the child has to figure out where to put a bunny on each turn so that it will be safe when the carrot is clicked. It eliminates some of the cognitive load involved in game strategy (how many rabbits to have on the board and how the randomness of the shuffled cards interacts with outcome) and focuses directly on the goal of figuring out where the hole opens given its seemingly stochastic nature. If one turns the game over, the mechanism becomes visible and offers information that could be used to deduce a pattern. However, it involves transferring that information dynamically to the top of the board and being able to track how two moving plates under the game interact to result in whether and where a hole opens. Subjects were not stopped from turning the game over to look if they sought to look. Interestingly, while students did pick the game up to look for their rabbits, the only group of students to examine the mechanism were students in an earlier pilot study during task design.

Fig 1. Funny Bunny Game

Fig. 2. Don't Wake the Sleeping Bear

<u>Games: Don't Wake the Sleeping Bear:</u> This game was modified from a Hasbro game entitled, "Don't Wake Daddy." The game involves getting to the finish line without waking up the sleeping dad. However, when spaces are landed upon, the player must push the button on an alarm clock a given number of times and if the dad pops up, must return to start. A realistic looking stuffed teddy bear was unstuffed and placed over the sleeping dad mechanism because many of the subjects were from single parent homes and from racially diverse schools. The number of alarm clock pushes that caused the bear to pop up ranged from 6 to 20 and each of the three games that were used had a different pattern of when the bear would pop up. However, if the students were not tracking how many pushes others had entered, it could pop up on the first push (presuming five pushes occurred during other turns). There is no visible mechanism to account for what happens.

<u>Games: Uno Attack:</u> This game is a variation of the game, Uno, where the player attempts to be the first to get rid of all of his or her cards. However, it has an automated card dispenser. A player pushes the button on the dispenser and sometimes it dispenses cards (a seemingly random number of them) though most times it does not. There is no discernible regular pattern. There is no visible mechanism to account for what happens. When the dispenser is opened up to add

cards, one can see a flywheel, however, it does not work when opened up so it is not possible to test under what conditions it shoots cards or not.

Biology: Seed Planting: Inducing causal patterns in many biology concepts presents the challenge of prediction and finding out over delayed periods of time. From a research design stance, there is the additional complication of getting kids to notice the relationship between the number of seeds and how many actually sprout. Open-ended interview questions were conducted first. It became clear that students didn't notice whether all or some of their seeds sprouted. Also, in authentic biology problems, the probability of particular outcomes cannot be calculated a priori as they can when working with games. Therefore, the following task was designed to help them remember their predictions and to help us develop a priori the outcomes (acknowledging that these manipulations area departure from authenticity.) The task engaged students in prediction and we took the seeds away each week to allow us to manipulate the outcomes. Students were told that they needed to have a certain number of bean plants (to give to specific people) in a few weeks. They were then given a peat pot, soil, and seeds and invited to plant the number of seeds that they thought they should plant in order to end up with necessary plants. They engaged in the task two to three times.

Biology: Hatching Eggs: Subjects were asked to predict what the inside of an incubator might look like in 22 days after eggs were set inside it. They were told that eggs typically hatch in 21 days. They were given a drawing showing eight eggs and were given an opportunity to draw the outcome later. Afterwards, they were probed on what causes the eggs to hatch, their experiences hatching eggs, and whether they had ever seen an outcome where less than the number of eggs hatched. In one of the second grades, the teacher hatched eggs so the study took advantage of this opportunity to interview students about their expectations before and afterwards.

Social Videos: Subjects were shown two brief video clips. These included one in which a girl is calling her mom for help with her homework. The rate of calling to response varies in the following way: 1) girl calls, mom responds, 2) girl calls, calls again, calls again, mom responds; 3) girl calls, calls again, calls again, calls again, and calls again, then mom responds. Students were asked what causes the mom to come and how the versions are different from one another. A second video shows a boy pestering his sister by taking her markers and she responds. The rate of calling to response varies in the following way: 1) boy takes marker, sister responds, 2) boy takes a marker, takes another marker, takes another marker, sister responds; 3) boy takes a marker, takes another marker, takes another marker, and another, and another marker, then the sister responds.

Mechanical: Candy Dispenser: Subjects were shown a candy dispenser that you could put coins in and turn to dispense M&Ms. They were given coins and invited to make it work. The dispenser dispensed between zero and five candies with each turn with a mode of five. The actual mechanism for dispensing candies was not visible given the number of candies in the dispenser. Subjects could detect some information about the mechanism, however, because the handle was less easy to turn on some turns when it would dispense no candies and on others turning very slowly appeared to yield higher returns.

The initial sessions focused on how children reason about games. Games were chosen for the initial domain because they offer many repetitive opportunities to make predictions. Repeatedly in the course of one session, children can predict what will happen and why. Having numerous opportunities in one session cuts down on issues of cognitive load in recalling and reasoning about what happened across sessions. Games also allow for demonstrating competence on many levels. Even if the youngest students are not necessarily able to tell us what they predict will happen or, at a more advanced level, offer explanations for why, however, implicit strategies through shifts in the playing choice in the games can reveal changes in thinking. The games also offered a platform for the students to become acquainted with and comfortable with the researchers.

Following each session, we made adjustments to the game apparatus, rules, or goals structures in the game in an attempt to eliminate certain types of cognitive load and to test certain ideas that the students may hold. For instance, to eliminate some of the cognitive load and to focus on one aspect of the probabilistic causation in playing the game “Funny Bunny,” we had students focus on a version called “Last Bunny Standing” so that their goal was merely to strategize about having the last rabbit on the board, not trying to get to the end of the board. We have also manipulated the probability of certain events in the game by changing the number of rabbits and the order of the cards to see how students handled the changes.

Each task was analyzed in terms of the way that it presented as a probabilistic causal problem space. Some tasks appear probabilistic because our perception is imperfect or limited and we are unlikely to take in all the information. For instance, a game that requires students to hold information about the nature of the causes in their heads might appear probabilistic even if from the stance of a perfectly unlimited reasoner (such as a powerful computer), it is deterministic. This would also be the case with information that is gained over time. For instance, the game of Funny Bunny holds three linked causes that lead deterministically to an outcome. However, but unless one can figure out how the patterns work together, the game gives the appearance of a probabilistic causal mechanism. First of all, the large carrot in the middle of the board turns and when it is clicked (in response to a game card). It moves (unseen) a mechanism under the game that has holes in particular places. Secondly, there are only some spots on the board where the holes can open up. These spots are “soft’ or wiggle when you push on them. Thirdly, the mechanism has two plates that move and this causes the opening to shift from the bottom path towards the carrot to the top path towards the carrot. Finally, to further the appearance of probabilistic causation, there is a place on the turning mechanism below where no hole will open up when the carrot is turned making the relationship between turning the carrot and producing a hole probabilistic. So what causes a rabbit to fall? The carrot has to be turned. The rabbit has to be on a soft spot. A spot has to open. And the rabbit has to be on the particular spot that opens. The interaction between these factors makes it appear stochastic.

Other tasks might be viewed as deterministic if you could follow a long reductive trail and have answers about all the variables along the way. For instance, when hatching eggs, even if you control for temperature, moisture, type of care, and the physical features of the eggs, the eggs will typically not all hatch. If you could know detailed specifics of each egg, perhaps you could know why it didn’t hatch, but this information is unavailable. This is also the case with mechanisms that appear to function randomly. If you had detailed information about tension on the spring, you might be able to tell what turn will cause it to snap, but the information is not available. It involves thinking about what the knower knows versus what is knowable (and with how much effort and with what outcomes.)

DATA SOURCES

Sessions were intensively analyzed to identify when shifts appeared to be taking place and why. The data sources include audiotapes and transcripts from each session; the series of students’ predictions for each task (including mathematical analyses of moves that support probabilistic or deterministic assumptions) and students’ reflections upon previous sessions in later ones. Our analysis included using ATLAS.TI to consider 1) transcripts and videos of students responses using etic, top-down questions to discern what students perceive about the nature of the causality; and 2) Emergent analysis of patterns in students’ reasoning where we identified themes in the students’ reasoning. Independent coders coded the interpretive aspects and their levels of agreement were assessed with refinements made until there was at least 85% agreement. Emergent codes independently generated by two of more coders are reported here. Narratives were developed of how the students changed in their explanations of the task, accompanied by detailed task analyses and consideration of leverage points that may be useful in teaching concepts that embed the causal concepts.

The analysis was conducted using the five criteria set forth by Siegler in his 1996 "overlapping waves theory" for microgenetic analysis to reflect how children's understanding changes along five dimensions: path, rate, breadth, source, and variability. The narratives consider how children's understanding develops along five dimensions: path, rate, breadth, source, and variability. Path describes how children sequenced their behaviors to get to the change. Rate describes how quickly and with what supports the child moved from the realization of the new concept to the consistent application of that concept. What kinds of time and experience were needed to make the shift? Was it abrupt, gradual, and what kinds of investments may have led to that shift? The breadth of the change refers to how narrowly or broadly the child gained the concept. Did it apply just to the particular game, to other games with the same inherent features, or to other contexts? The source addresses what the child did to make the change. What changes in behavior appear to have led to the change? What realizations seemed to have helped the child to see the new feature? The variability refers to the difference between students on the dimensions above. Are there similarities in the patterns of change that would suggest instructional approaches beyond this set of students? Are there idiosyncratic patterns that are important for teachers to be aware of?

RESULTS

The following findings emerged across the subjects and grade levels with multiple forms of support as discussed below:

1. The children began the study offering highly deterministic interpretations of the tasks. Six of the eight children held a predominantly deterministic stance across the domains tested. This was expected for machines that are designed to work, but not for games, biology, and social tasks. However, most children treated even these with an expectation of reliable one to one outcomes. For instance, after playing Funny Bunny for two sessions and having many opportunities to witness the stochastic patterns, when asked how the game worked, Jordan (K, male) replied:

"You pick a card and then if you get one and you count like one two three and if you get a card and if you get the bunny on the circle, you turn the carrot and then the rabbit will fall." (ln. 10)

He predicts that he can determine where the hole will open next. He focuses on a soft spot predicts it will open next. He turns out to be correct. When asked how he knew that one would open, he replies, "because I am smart" and pushes down on other soft spots. He continues to study the game. After repeatedly guessing wrong, Jordan appears to be guessing where he places the pieces and when playing Last Bunny Standing, remarks, "The game can tell where they are supposed to be. The game knew."

Similarly, Carter (K, male) studies the board carefully to see if there is a pattern to what happens. In his first session, he guesses the hole will open at spot #10. Instead it opens at #13. He moves his finger around the board looking for soft spots. By narrowing his choice to the soft spots, he is narrowing the possibilities to nine posts instead of 26. He finds one at #16 and chooses that spot explaining that "it skips two and then it goes there." He turns once and it does not open there. He turns again and it does go there. He gets it wrong the next two times. The next time he gets it correct and then he gets it wrong. He gets it incorrect on the next try but is undeterred. "Oh, I know where I am going to pick now!" He is incorrect the next two times, but maintains "that you can know what will happen." Carter continues to focus on trying to find a pattern.

The stance that there is a pattern and it can be determined can be viewed as an adaptive one. The students are persistent in their attempts to figure out what is happening and do not start to consider an alternative to this possibility until the feedback continues to be discrepant with their predictions over multiple sessions.

2. Some students who held a deterministic stance at the outset began to make a shift towards the possibility that they could not predict the outcome in every case, "just most of the time." They predicted what a best guess would be even if it "would not always be right." When they were

unable to determine strategies for how to predict, some students redefined the cause (from a mechanistic to an anthropomorphizing one, for instance.) Eventually, they questioned whether it was possible to predict. Whether they thought the particular problem was too hard to figure out or were beginning to recognize the problem as having probabilistic characteristics was not clear. The strongest example of this is that of Rajon (Gr. 2, M).

The narrative below outlines this transition for Rajon. Notice how in the first session, Rajon discovers that the hole opens up and that it moves about the board. He quickly engages in looking for a pattern and initially assumes that it is like another game that he has played. When that fails to help him explain what happens, he looks for other evidence and zeroes in on the pattern of what moves. However, eventually over the course of five sessions, he struggles to explain what causes the hole to be where it is and shifts from a deterministic explanation. After five sessions, he questions whether there is a clear cause and effect pattern and if indeed, there is a knowable pattern. Instead of focusing on whether or not he (subjectively) knows, his language suggests that you just can't know. While these are clear shifts in how he reasons about the game, they may also indicate a "letting go" of an attempt to figure out the pattern of the game.

Rajon first engages by focusing on moving his rabbit. Early in the game, he is surprised when the hole opens up in the game board and he begins to look for patterns to explain it.

Within minutes of being introduced to Funny Bunny, he understands that the clicking of the carrot is related to a hole appearing on the board and that when a hole opens, a rabbit could fall through. This understanding is demonstrated by his response when the carrot is clicked for the first time, revealing the first hole (the game began with no holes visible):

R: Uh oh.
I: Uh oh. Why are you saying, "Uh oh?"
R: Because there's a hole!
I: Yeah. So what?
R: You're going to fall!
I: Who's going to fall?
R: The rabbit!
I: Oh! The hole is going to make the rabbit fall?
R: Yes!

He draws upon knowledge that he holds from another game. Believing that it is analogous to this game, a Piranha game, he attempts to map the features of the game to the game that he is familiar with. The next time a click card is turned, he uses his knowledge of another game, the "Piranha Game," to predict that the hole will move "to the next hole over" from the hole that was already open. He gives the following reason:

R: Because I have a um game that you have to [pause] that you have to catch a fish and it goes around and around and around and it's caught- reminds me of a game in that you- if you get some- and then- like this- like the piranha's game, because it keeps on going forward, forward, forward.

He indicates with his fingers that the holes will open one at a time, in a clockwise manner, which is similar to a pattern he observed in the piranha game. However, he also seemed to suggest that the pattern was not always consistent in the piranha game:

I: Uh huh. So in the piranha game does it go this one and then the next one and the next one, like that? Or does it go- Is that what you're saying when it goes?
R: Yes, but sometimes if you want to play- sometimes it goes different- like that- but sometimes it goes that way too.

He realizes that this game has different features but continues to determine what patterns account for when and where the hole will open up. After he makes this prediction, the carrot is accidentally clicked *twice* not once. Because it is clicked twice the hole that opens is a few holes over clockwise from the hole that was previously open. It was not quite what he expected because the next hole over did not open, but the hole did move clockwise as he predicted it would. When another click card is turned, Rajon expresses his understanding that the height of a spot is related to it opening up into a hole. He is still figuring out exactly how this works, initially thinking that holes go up and down but then figuring out that certain spots are consistently lower than others:

R: These only turn [pointing to the "low spots"] because- these ones don't [indicating the high spots], the high ones don't, only the low ones.
At the game's conclusion, Rajon continues to believe that holes open predictably, but only the low spots. He thinks the next spot to open will be the next low spot clockwise. His advice to someone playing the game was:
R: Try to stay off of the low one that is next to the other low one.

He sticks with this explanation in the second session, however, he isn't quite sure how to respond when his explanation does not deterministically result in an outcome or allow him to make consistent predictions

In session two, the first time the carrot is clicked, the hole opens up in a spot that he did not predict. He attempts to explain why it ended up there and not in a spot closer to the starting point.
I: So did you expect it to go to that one?
R: A little, because you had to click it until all the holes were blocked up.

Upon seeing that the hole did not open up where he expected it to, he tries to come up with a logical explanation why it did not. He tries to take the new information and fit it into his current thinking.

He then gets a click card. When asked what he thinks is going to happen, he points to the next low spot moving clockwise from the hole that is currently open. He claims that will be the spot that will open. His prediction is congruent with what he stated at the beginning of the session.

He then clicks the carrot and the hole does not open in the space he predicted it to open. The interviewer asked him what happened and why did he think it happened that way. He was quiet for a few seconds and looked at the game from different angles intently. He seemed surprised by the outcome and appeared to be trying to figure out what happened. He generates two possible explanations:

R: This one [pointing at the hole he thought was going to open] was a little too high so it went to this one [pointing to the hole that actually did open] or the carrot is trying to not let them get to it.

His first explanation suggests that he thinks that some low spots are higher than other low spots and that those "higher low spots" behave like regular high spots and therefore, do not open like the "lower low spots." This would imply that some low spots don't ever open. His second explanation clearly anthropomorphizes the carrot.

At minute 10:38, Rajon gets another click card and the interviewer asks for another prediction. The hole is at spot #25, just two away from the top of the carrot. He predicts the hole will open at #4, which is just to the bottom right of #26. His logic argues that if the carrot doesn't want the player to get to the top, it will open the next nearest spot, which is #26. But, as R points out, #26 is not a "low spot" so it won't open there. The choice of spot #4 is the next low spot over from #25, which matches his original pattern. But then he changes his guess to #7 explaining:
R: Every time, something gets lower and higher.
When he clicks, the hole opens on the opposite side of the board from #7. This time his only explanation anthropomorphizes the carrot:
R: It's still trying to protect itself.

It seems the more he is unsuccessful at guessing where the next hole will be, the more he leans toward explaining that the carrot is trying to prevent rabbits from getting to the top.

He is asked where he thinks it might open next. He says that any spot that makes a sound (any low spot) might open up. This is the first time he articulates the idea that *any* low spot could open up. He seems to have abandoned his idea that the next low spot clockwise will open up.

Increasingly in this session, he tries to predict where the hole will open up and is wrong much of the time. By the end of this session, we see him beginning to waiver about the game's predictability.

He says simply:

R: You never know where the hole is going to go.

This thinking is obviously very different from where he started at the beginning of this session.

In the third session, he begins in the same place where he left off in the second one and the terms "luck" and "chance" increasingly dominate his dialogue about the game:

I: Okay, so, before you click, where do you think the next hole is going to be? Could you show me where you think the next hole is going to be?
R: It's gonna be on one of these low ones. [He sweeps his fingers across the holes on the top part of the game.] You can't predict so good.

Rajon also continues to think that chance/luck play a role in the game. He gets a click card and predicts that the hole will open up under one of the rabbits. He turns that carrot and it does not open up in the spot he predicted:

R: Lucky! Lucky! [He laughs.]
I: Why do you think he was lucky?
R: 'Cause this one didn't fall.

By the fifth session, we see a dramatic increase in Rajon's use of the term risk. He gets a one-hop card. He moves his back bunny one space. He does not move his bunny on 20 to lower space 21.]

R: I'll risk moving this bunny [refers to back bunny] but I'm not risking that bunny [referring to bunny on 20].

His belief that his risk-aversion strategy is a good one is affirmed. Had Rajon figured out the pattern of how the holes open, he would have been able to place his bunny on lower spots without them falling and thus, advance his bunny to win the game. His strict adherence to his "no risk" strategy suggests that he believes there is not a deterministic pattern to how the holes open up. His moves suggest that he thinks that any lower spot can open up at any time and that where is completely unpredictable. This thinking matches the thinking he had at the end of sessions two and three.

Increasingly, he refers to luck as the reason for the particular outcomes in the game. Perhaps because so much time has passed without him getting a two-hop card, Rajon finally risks his lead bunny by placing it on lower space 21. His bunny doesn't fall on the next click and he claims it was luck that he didn't fall (19:42):

Rajon gets a one-hop card and moves his bunny from space 20 to lower space 21.
A click card comes up.
R: Oh, no! Never mind.
Rajon attempts to move his bunny on 21 back. The game is clicked and the interviewer's bunny falls.
R: Lucky! I'm lucky!
When asked what he thinks about this, he replies:
R: You just can't know.

We see a subtle shift here in the kind of language that Rajon uses over the course of the tasks. He has backed off looking for a pattern at this point (though he resumes searching for patterns in other tasks.)

The emergent coding of the games transcripts also revealed that kindergarteners commonly reasoned in what might be called “single frame analysis” where they described what happened in a given move, but did not necessarily try to reconcile this pattern with previous moves. It appeared as though the information in a given frame of the game was most salient for reasoning from at the moment. The second graders adopted a similar strategy that might be called “backwards reasoning.” This involved revising one’s ideas after a prediction failed to incorporate what was now known with the previous step. It led to a kind of reasoning with “20/20 hindsight” where the reasoner behaved as though they should have known information predictively that they could only know descriptively.

2. Two students did not approach the tasks from a highly deterministic stance from the outset. The variation in the data offers interesting insights. In each grade, there was one child who approached the tasks from a more varied stance. They seemed more open to the possibility of probabilistic causal structures at the outset.

For instance, Maia, a kindergartener, in the very first session, said that she didn’t know what would happen (subjective uncertainty), but maybe that you couldn’t know what was going to happen in terms of whether it was possible to know for certain (objective uncertainty) or not.

In her first experience with Funny Bunny, when the first hole opens, Maia refers to the order of the cards in explaining what made the rabbit fall.
M: (19:51) um... last time we played the game in that’s the way the cards got in order and the next card
She attends to the patterns in the game but holds a stance of uncertainty about what will happen.
M: (22:30) *Gets a Carrot*
I: Maia before you turn that, what do you think it’s going to make happen?
M: Either one of the other holes are going to open, or one of the rabbits is going to fall down
I: How do you know which?
M: (22:51) I don’t know which one. I cannot.
I: Cannot know… uh huh… do you know why you cannot?
M: No
I: (28:24) Do you have a prediction about what hole will open?
M: I don’t know which one

Throughout the game sessions, she focused on luck and possibility, but in balance with predictions and mechanisms. As she gets to know the game better, she considers the possibility that it contains patterns that can help her predict outcomes. However, she also maintains the possibility that there are not clear patterns to predict from.

M: I was very very lucky. (Gets a one-hop card.) One. I am ahead of you.
Interviewer gets a carrot card
M: Hope it’s not me, uhh. It’s not me, it’s not me!
At 5:59, the Interviewer gets a carrot card again.
M: That is far from us! Spin the carrot.
I: Spin the carrot. Alright, you ready?
M: No, no, no, no, not on mine, not on mine. Ooh.

…but she also begins to detect a possible pattern.

M: I get it, what is happening.

I: What's happening?
M: When you’re in the beginning, then, the end numbers, are are, are have holes. When you're in the end, the beginning numbers have holes, and when you're in the middle, the middle numbers have holes.
I: Yeah?
M: And that's the easiest way to get dropped.
I: The easiest way to get dropped?
M: Yeah. Is when you're um, in the, um, middle, because then the, they go in the middle holes.
I: They go in the middle holes when you're in the middle of the board?
M: Yeah. Not, not right here, but somewhere in the middle of the numbers. Maybe on twelve, eleven, right?
I: mmhmm
M: Is in the middle. Twenty-six and one.
I: Are you in the middle or am I in the middle right now?
M: Um I'm getting close to the middle. Maybe, two more steps I guess.
I:: Okay. And then you think there will be a hole where?
M: In the middle.
I: Okay.
M: Well maybe not…

In a later session, she shifts back to viewing the game as probabilistic.
M: My carrot. Whoah (no holes open) But where did it go? Oh, just probably part of the pattern. Well I don’t know what is going to happen. We spin it so little and it comes no, it must not be a pattern.
Next, the interviewer and Maia play Last Bunny Standing” with 3 bunnies each (32:25)
K – (Loses one of her bunnies)
M – Why do you think that happened?
K – Because it’s not a pattern, not anymore.

Given her age and her openness to both possibilities from early in the sessions, Maia’s approach differs from most of the kindergartners and second graders.

Another student who is an interesting outlier is Layla. In generating her own examples of instances of probabilistic causation, she offers two:

“My shower, because sometimes when you turn the thing nothing comes out. And I’m like mom the shower stopped working! And she just tells me to get in the shower, so I go in the shower and all of a sudden water starts popping out and its cold and sometimes it will be piping hot. Oh and another thing, my baby brother does that. I’ll be looking for him and he’ll be standing in one place. And I’ll be like Tyrone[1] don’t move, that’s my brother’s name, and then I’ll go over there, and I’ll come back and he’s nowhere in sight. I’ll go all the way around the house. Then I go in the hall way and I open the door and then boo out of nowhere. He pops out like card machine and my shower sometimes doesn’t work.”

Given their ages and openness to both deterministic and probabilistic possibilities from early in the sessions, students like Maia and Layla are particularly interesting case studies because they introduce a different growth pattern and perhaps, points of leverage that can be used in service of instruction.

3. Most students responded similarly in a similarly deterministic manner to mechanical devices. For instance, when turning the handle on a gumball machine lets out candy all but every fourth time or that lets out five candies in a row five times and then let out 4, 6, or 7 in a sixth turn, most children looked for nonobvious mechanisms that allowed for the pattern, for instance, that “the turning device is jammed” or “there’s an opening inside that got stuck and let more candies down”

[1] All names have been changed to pseudonyms.

or “I can make it move slowly and then more come down.” This echoes earlier findings (e.g. Schulz & Sommerville, 2006). However, even if children gave a deterministic, mechanistic explanation for the outcome, some of them allowed for probabilistic outcomes in their predictions, recognizing that the mechanism behaved stochastically even if there was a deterministic reason for doing so. For instance, a second grader said, “Five is the best number to predict because it comes up the most times, but you will be wrong some of the time if you guess it.”

Kaylee, a second grader, appeared to make a breakthrough went she recounted in great detail the indignance of a gumball machine that did not deliver.

K: One time we went to a store. They had a gumball machine. My brother put one quarter in and he got NOTHING.

I: And he got nothing? So he didn’t get a gumball that time. No? So, um, well how often when you put a quarter in does it come out compared to when it doesn’t come out? Which do you think happens more?

K: Um. It comes out.

S: So you think more often than not, you’ll get a gumball if you put your quarter in the machine? Okay, what do you think happens when you turn the crank? So you put the quarter in right there and you turn it. You turn that little lever and then the gumball comes out down here. What do you think is happening?

K: Um, maybe…I had a gumball machine at my house and I turned the knob, and there was like a little hole and it filled in and it came through there.

S: Okay so you turned the knob and it goes through a little hole, the gumball does, and you see it come out the bottom.

K: Uh-huh.

S: Okay. So again, we’re going to think back to the Sleeping Bear game. Is there any way that the bear waking up in Don’t Wake the Sleeping Bear is like the candies coming out in this gumball machine? So, again, let’s think about the gumball machine. You put the coin in. You spin the knob. And what happens?

S: The gumball comes out. And, you said, sometimes you don’t get one but most of the time you do.

K: Yes.

She was not able extend this experience to her reasoning beyond the gumball example. However, in this case, the interviewer’s question focused her on the mechanism and may have distracted her from the focus on the less-than-reliable pattern.

4. Most of the students also brought a strongly deterministic stance to the biology examples. This included growing of seeds, which the second graders had done as part of a classroom activity during the month prior to the interviews. The experience of planting seeds in their classrooms and actually witnessing that not all of the seeds grew was not enough to shift their stance. When asked to write down how many seeds they would plant to have bean plants in a few weeks, most of the second graders (three out of four) said that they would plant the exact number of seed that they wanted plants for. Rajon said that he would plant seven seeds for seven plants and five seeds for five plants. In the class activity, Rajon said that he planted eight seeds for eight plants. He then recounted planting twelve once and only getting seven plants. When pushed to tell more, he said that something could have happened, “his little sister destroyed them, or a dog or lizard could eat them.” He also mentions birds breaking them in half.

Q: um, say birds, cause they, they have beaks and they can break the seeds in half.

E: oh yeah, and what do you would happen then?

Q: they won't grow.

E: the plant won't grow...

Similarly with the incubator, Rajon drew exactly eight chickens for the eight eggs. He did add that the “incubator could be too hot and you would have breakfast.”

Carlos discussed planting seeds in class. He said that if he planted four, that four would grow and that he had planted four in class and four grew. He also said that if he planted three, all three would grow because they had enough space. In the third example, he said that it was possible that if he planted five, may not have enough space, may be that only four would grow. With the incubator example, he said that some eggs might not hatch, but then clarified that eventually they would, the rate would be different.

I: Uh-huh. Okay. Tell me about that, for example. So if it lays six eggs and sometimes only three hatch, there's a lot that don't hatch, right? Can you know if an egg is going to hatch?
C: No.
I: No?
C: It could be a day earlier or a day later.
I: Uh-huh. If I were to give it 30 days, do you think they would all hatch?
C: nods

Two of the kindergarteners, Jordan and Carter had never planted before and they each predicted a one to one correspondence between seed planting and the number of plants that they would get (16-16; 4-4 and 3-3; 21-21, and 5-5, respectively). However, Carter purposely chose not to plant a cracked bean in the package because he said, "it wouldn't grow." Maia said that she had "planted many times before in pots in her yard and in her garden. Maia planted five and predicted three would grow, two and that two would grow; six and that two or three would grow. She could not verbalize by some would and some wouldn't, but aid that if there were "more seeds, it would take longer."

5. The social examples appeared to invite probabilistic responses from the two students who already offered probabilistic examples (Maia and Layla), but not necessarily from the others. For example, Carlos appeared to be searching for a clear pattern:

C: (referring to the video) She called her.
I: Uh-huh.
C: And called her and called her.
I: How many times did she have to call her?
C: The last one four. The third one five. The first one, one time. The second one three times.
I: Today she's doing her homework. Totally different day. And she called her mom. How many times do you think she'll have to call her mom today?
C: Probably six or seven.
I: Six or seven? Why six or seven?
C: Because I think it's going in a pattern, like one, two, three, four, five. And then it went back down.
I: Is there a pattern when you call your mom and she comes?
C: No.
C: No? Tell me about that.
I: Every time I call my mom, and she says, "What?"
I: She says, "What?"
C: Yeah. And then I come to her and tell her I need help on my homework.
I: She's always only one time?
C: Nods

Layla (Grade 2) first maps the social example by its surface features, then she maps to the deeper probabilistic aspects:

I: Okay.
L: She's probably tired of having to help her with her homework.
I: Okay. So how do you think the girl calling her mom- how do you think that is like UNO attack?
L: Because it's like surprising. It's like surprising like in UNO Attack when you're playing, you press the button and then in the middle of the game maybe you press it twice and then cards just start popping out and more than two cards come out sometimes, which is all you need but you still have to take them all and like her, she's doing a math problem and then all of a sudden she's calling her mom out of nowhere.

I: So the mom eventually showed up, was it the same number of times in each video that the girl called?
L: No.
I: No?
L: And that's like UNO Attack.
I: Oh?.
L: Not the same amount of cards come out every time.
I: Okay. So you think it's like UNO Attack because you can't predict the number of cards that are going to come out. So that's how they're the same. How are they different, do you think?
L: How are they different?
I: Yeah.
L: Because if the girl is doing her homework, she's calling a person. If you're playing UNO Attack, you're just pressing, you don't really have to talk for you to get cards.
I: How is that the same?
L: Because when I'm calling my friend, sometimes she comes over, sometimes she doesn't. And then UNO Attack, sometimes cards come out, sometimes they don't. : Sometimes when I'm doing stuff, I might call my baby brother and then sometimes he comes, then other times, if he gets mad at me or something, he'll- he won't say anything. He just won't come and then I'll have to go get him myself.
I: How can you predict if he's going to come or not?
L: You can't really predict because our house- sometimes you can predict if you have a house on the floorboards, sometimes you can hear the floorboards moving when they're walking around.
I: Okay.
L: But in our house you can't so you can't predict.

Kindergarten (Maia): could respond to but not generate examples:
M: So when you raise your hand in class, do you get called on all the time?
I: No!
M: No?
I: No.
M: What do you mean?
I: If you don't get called on then the other person has to get called.
M: The other person. Can you- if you raise your hand can you know for sure if you're ever going to get called on?
I: No. You can't.
M: You can't?
I: There are many people who raise their hand so Ms. D can pick on someone else.
M: Someone else.
I: Sometimes she'll come. Sometimes she'll not come right.
M: Sometimes she'll come and sometimes she'll not come.

About the video:
M: So can you think about any times in your life when this has happened?
I: Happens couple times but I don't remember what was going on.
M: So when you call for your mom, does she always come?
I: No. Not always.
M: Can you know when she's gonna come?
I: No.

For these two students, their responses appear to take a much less deterministic stance than their classmates.

In looking across the tasks, the path for most of the students seemed to be to push hard on deterministic explanations and to seek out patterns. In some cases, they eventually adopted anthropomorphic explanations when they struggled to explain the pattern in other ways, for instance, arguing that the carrot didn't want them to win or that the game knew where they were. After playing the Funny Bunny Game long enough, two of the students, a kindergartner, Carter and a second grader, Rajon appeared to let go of their deterministic stance to some extent. However, it seemed to be accompanied, at least for Rajon, with a type of giving up on trying to detect a pattern, at least on the Funny Bunny game. As researchers, we acknowledged that there is a level of admirable perseverance, a sense that the world is ultimately entirely predictable, that may be adaptive and enable resilience for learners in taking such a deterministic stance. Thus we noted that Rajon resumed looking for a pattern with "Don't Wake the Sleeping Bear" with some sense of pleasure.

I: You're only not too worried. Now why did you choose that?

R: 'Cause I'm really sure it won't pop up on me this time.

I: You're really sure? Can I ask why you're really sure this time?

R: 'Cause it only needs 6 and 6 more to wake up the bear.

I: Okay. So you're saying that if we press it six times it'll definitely go up but in the next two times it won't. Is that your idea?

R: Yes. And if you land on 8, I think it would be 2 more and then- or 4 more and then you'll wake up the bear.

I: Okay. It sounds like there's some kind of pattern you're talking about. Can you tell me what-

R: The higher numbers, the closer you get to wake up the bear. The lower numbers, the farther you don't get to wake up the bear.

[R presses the alarm clock twice and the bear does not wake up.]

R: Woohoo!

How to encourage reflective knowledge while protecting this stance might well be the right instructional question. As we explore this terrain further, this might involve drawing upon the power of narrative, metacognition, and analogy to introduce new ideas about the reliability of scientific findings while protecting this fundamental, in many ways potentially adaptive, belief in determinism.

While Rajon made subtle shifts during the interview sessions with each task, the rate of the shifts was slow and uneven. This was generally true for all of the students with a return to generally deterministic explanations with each new task. The breadth of the change was also narrow for most of the students, for instance Kaylee realizing that bubble gum machines do not always work but then not extending this to other tasks. This particular experience seems to have stuck with her, perhaps for the injustices that it entailed and with further scaffolding may be the future source of insights for her. However, it has not yet helped to broaden her thinking about other stochastic tasks. Across the students, there was important variability as witnessed by the responses of Maia and Layla. Their responses suggest a different reflective sense of the reliability between causes and effects. While they are outliers in comparison to their classmates, they may offer points of leverage in instruction with the language that they introduce and the examples that they understand.

We are continuing to analyze the rich data set that we have collected and to mine it for further patterns. In addition, we are studying the other end of the age range and are looking at the reasoning of fourth and sixth graders. This data hints that by sixth grade, determinism is a strong factor in how students reason about scientific explanation—that it is an explicit criterion for whether a causal relationship holds and that this may impact how the students interpret the results of science experiments and data that they learn about.

One might ask, "Towards what instructional end is children's developing reflective understanding of probabilistic covariation because cause and effects important?" We are interested in developing children's reflective awareness that a lack of reliability between causes and effects does not necessarily signal a lack of causal relationship and that noticing of causes and effects may be influenced by the patterns between causes and effects and what we believe about those patterns. In analyzing the component understandings of probabilistic causation and how they are elicited by particular contexts, it is our goal to scaffold children's understanding and to leverage these understandings in service of their science learning and the understanding that causal relationships are often not entirely reliable. Thus lack of a perfect correlation is not necessarily grounds for rejecting a causal relationship and that we may not always notice relevant patterns because their patterns of co-variation may be complex.

While few phenomena are truly probabilistic, the information that we receive often has probabilistic characteristics. The existing research argues that children can handle probabilistic data but that they hold a deterministic stance and engage in a search for non-obvious causes that is ultimately adaptive (Schulz & Sommerville, 2006). We agree, but believe that a reflective stance is important to the broader problem space of detecting causality despite instances with the appearance of probabilistic characteristics in our world. The findings confirm children's deterministic stance, however, also suggest the ability amongst some children to reason reflectively about probabilistic causation. Of particular interest moving forward is whether it is possible to broaden students' reflective capacities to some extent. Our ability to detect stochastic event structures matters because in a complex world, it is easy to lose sight of the impact of particular choices, individual and societal, when there is no observable effect following an action. Such patterns characterize many of the most recalcitrant and imminent problems of our time, such as climate change, ecosystems decline, and global disease transmission and may well be important to developing thinkers for the future.

ACKNOWLEDGEMENTS

The authors wish to thank Saida L. Solis, Molly Levitt, Erika Spangler, Adi Flesher, Samantha Marengell for their assistance in collecting and analyzing this data. We also wish to thank Suzannah Bukenya, Maureen Danahy, Lauren Mueller, and Mark Breneman for enabling us to work with their classes.

This work is supported by National Science Foundation, Grant No. NSF#0845632 to Tina Grotzer. All opinions, findings, conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

Borton, R.W. (1979, March). *The perception of causality in infants.* Paper presented at the meeting of the Society of Research in Child Development (SRCD), San Francisco.

Brown, A. L. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions. *Journal of the Learning Sciences, 2* (2), 141-178.

Bullock, M. (1979). *Aspects of the young child's theory of causation*. Unpublished doctoral dissertation. University of Pennsylvania.

Bullock, M. (1985). Causal reasoning and developmental change over the preschool years. *Human Development*, *28,* 169-191.

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.

Carey, S. & Spelke, L. (1994). Domain-specific knowledge and conceptual change. In L.A. Hirschfeld & S.A. Gelman (Eds.). *Mapping the mind: Domain specificity in cognition and culture* (pp 169-200). New York: Cambridge University Press.

Collins, A. (1999).The changing infrastructure of education research. In E. Lagemann & L. Shulman (Eds.) *Issues in education research* (pp 289-298). San Francisco: Jossey-Bass

Einhorn, H. & Hogarth, R. (1986). Judging probable cause. *Psychological Bulletin, 99(1)*, 3-19.

Gopnik, A., Glymour, C., Sobel, D.M., Schulz, L.E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review.* 111(1) 3-32.

Gopnik, A. & Schulz, L. (2007). Causal learning: Psychology, philosophy, and computation. New York: Oxford University Press.

Grotzer, T.A. (2003). Learning to understand the forms of causality implicit in scientific explanations. *Studies in Science Education, 39,* 1-74.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, MA: Cambridge University Press.

Kalish, C.W. (1998). Young children's predictions of illness: Failure to recognize probabilistic causation. *Developmental Psychology, 34*(5), 1046-1058.

Keil, F.C. (1994). The birth and nurturance of concepts by domains. In L.A. Hirschfield & S.A. Gelman (Eds.), *Domain specificity in cognition and culture* (pp 234-254). New York: Cambridge University Press.

Keil, F.C. (1995). The growth of causal understandings of natural kinds. In D. Sperber, D. Premack, & A.J. Premack (Eds.), *Causal cognition: A multidisciplinary debate.* (pp 234-262). Oxford: Clarendon Press.

Kurtz, K.J., Miao, C.H., & Gentner, D. (2001). Learning by analogical bootstrapping. *The Journal of the Learning Sciences*, 10(4), 417-446.

Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions, *Developmental Psychology*, 43(1), 186-196.

Leslie, A.M. (1982). The perception of causality in infants. *Perception, 11*, 173-86.

Leslie, A.M. (1984). Spatiotemporal continuity and the perception of causality in infants. *Perception, 13*, 287-305.

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

Oakes, L.M. (1993, March). *The perception of causality by 7- and 10-month-old infants.* Meeting of the Society for Research in Child Development (SRCD), New Orleans, LA.

Opfer, J. E. & Siegler, R.S. (2004). Revisiting preschoolers' *living things* concept: A microgenetic analysis of conceptual change in basic biology. *Cognitive Psychology*, *49*, 301-332.

Schulz, L & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology*, 40(2), 162-176.

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

Shultz, T.R. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47*(1, 194), 1-51.

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

Siegler, R.S. (1976). The effects of simple necessity and sufficiency relationships on children's causal inferences. *Child Development, 47*, 1058-1063.

Siegler, R., & Liebert, R. (1974). Effects of contiguity, regularity, and age on children's causal inferences. *Developmental Psychology, 10*(4), 574-579.

Sobel, D.M.. & Buchanan, D.W. (2009). Bridging the gap: Causality-at-a-distance in children's categorization and inferences about internal properties. *Cognitive Development, 24*, 274-283.

Spelke, E.S., Phillips, A., & Woodward, A.L. (1995). Infants' knowledge of object motion and human action. In D. Sperber, D. Premack, & A.J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp 44-78). Clarendon Press: Oxford.

Van de Walle, G., & Spelke, E.S. (1993, March). *Integrating information over time: Infant perception of partly occluded objects.* Biennial meeting of the Society for Research in Child Development (SRCD), New Orleans.