CHAPTER FIVE

Learning about Causes from People and about People as Causes: Probabilistic Models and Social Causal Reasoning

Daphna Buchsbaum, Elizabeth Seiver, Sophie Bridgers, and Alison Gopnik[1]

Institute of Human Development, Department of Psychology, University of California at Berkeley, Tolman Hall, Berkeley, CA 94720, USA

[1]Corresponding author: E-mail: gopnik@berkeley.edu

Contents

Abstract

A major challenge children face is uncovering the causal structure of the world around them. Previous research on children's causal inference has demonstrated their ability to learn about causal relationships in the physical environment using probabilistic evidence. However, children must also learn about causal relationships in the social environment, including discovering the causes of other people's behavior, and understanding the causal relationships between others' goal-directed actions and the outcomes of those actions. In this chapter, we argue that social reasoning and causal reasoning are deeply linked, both in the real world and in children's minds. Children use both types of information together and in fact reason about both physical and social causation in fundamentally similar ways. We suggest that children jointly construct and update causal theories about their social and physical environment and that this process is best captured by probabilistic models of cognition. We first present studies

ISSN 0065-2407,
http://dx.doi.org/10.1016/B978-0-12-397919-3.00005-8
© 2012 Elsevier Inc.
All rights reserved.

showing that adults are able to jointly infer causal structure and human action structure from videos of unsegmented human motion. Next, we describe how children use social information to make inferences about physical causes. We show that the pedagogical nature of a demonstrator influences children's choices of which actions to imitate from within a causal sequence and that this social information interacts with statistical causal evidence. We then discuss how children combine evidence from an informant's testimony and expressed confidence with evidence from their own causal observations to infer the efficacy of different potential causes. We also discuss how children use these same causal observations to make inferences about the knowledge state of the social informant. Finally, we suggest that psychological causation and attribution are part of the same causal system as physical causation. We present evidence that just as children use covariation between physical causes and their effects to learn physical causal relationships, they also use covaration between people's actions and the environment to make inferences about the causes of human behavior.

1. INTRODUCTION

In the past 10 years, the probabilistic models approach to cognitive development, also known as rational constructivism, has begun to be applied to many aspects of children's development, particularly their causal inference and learning. In the first wave of this research, however, the focus was squarely on physical knowledge, such as the relation between blickets and blicket detectors (or the workings of other physical machines). In these types of studies, for example, an experimenter may place a series of blocks on top of a machine. Some blocks are "blickets" and make the machine produce an effect (e.g., lighting up and playing music), while other blocks do not. Children are then asked to make causal inferences from the evidence they see, such as which block was a blicket or which new block should make the machine go. In this vein, work from our lab and others has demonstrated that children possess sophisticated causal reasoning abilities, including making rational inferences from probabilistic input (e.g., Gopnik et al., 2004; Kushnir & Gopnik, 2005, 2007; Schulz, Bonawitz, & Griffiths, 2007; Schulz, Gopnik, & Glymour, 2007; Sobel & Kirkham, 2006; Sobel, Tenenbaum, & Gopnik, 2004).

These initial studies were generally limited to investigating how children learn by observing causal relationships in their physical environment and did not take the child's social environment into account. From an early age, children are exquisitely sensitive social beings and their causal learning takes place in a rich social context. A natural question is therefore how social interaction informs and influences children's causal learning and how causal reasoning influences children's social inferences.

Data about "purely physical" causes does not exist in a vacuum – blickets are not putting *themselves* on the machine, after all. There is a social and psychological component to the causal learning that results from our interactions with other people. Even in the relatively simple context of a blicket detector experiment, the child not only must consider the physical evidence of the machine's activation but also must make inferences about the experimenter's actions and mental states. Did she put the blicket on the machine in the right way? She says she knows what makes the machine go, but does she? Is she just trying to make the machine go or does she also want to teach me how it works? Children can use the physical blicket evidence to make social inferences (the block did not work, so she must not know what she is doing) or use the experimenter's testimony and actions to make inferences about the blickets (since she says she knows what she is doing, she must be teaching me about which blickets I should use, so I will pick the same one).

In general, social and physical causation will be inextricably linked in most real-life causal learning, especially since the goal-directed actions of others lead to many of the causal outcomes children observe. In fact, even infants and toddlers seem to expect that the causally relevant events they observe in the world will have been produced by the actions of social agents (Bonawitz et al., 2010; Meltzoff, Waismeyer, & Gopnik, in press; Saxe, Tenenbaum, & Carey, 2005; Saxe, Tzelnic, & Carey, 2007).

We argue that children jointly construct theories about both the physical and the social world, which in turn generate higher-order theories that shape children's interpretation of future events. This natural learning process parallels the scientific method, and thus, we can characterize children's learning with the metaphor of children as intuitive scientists.

This metaphor might suggest that children just learn on their own, but neither children nor scientists are solitary learners. Both scientists and children learn extensively from the actions, reports, and tuition of others.

Teachers serve a particularly important function in this regard, both formally in the classroom and informally in the world. Recent work on "natural pedagogy" (Csibra & Gergely, 2006, 2009; Gergely, Egyed, & Király, 2007) and children's understanding of testimony (e.g., Corriveau, Meints, & Harris, 2009; Jaswal, Croft, Setia, & Cole, 2010; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig, & Harris, 2007) has demonstrated that infants and young children are sensitively tuned to others and can learn from them in complex and subtle ways. The pedagogical intent of a social demonstrator can influence everything from children's exploration of a novel toy (Bonawitz et al., 2011) to their generalizations about objects' functional properties (Butler & Markman, in press). The expertise (e.g., Koenig & Jaswal, 2011; Kushnir, Vredenburgh, & Schneider, under review; Sobel & Corriveau, 2010) and past accuracy (e.g., Birch, Vauthier, & Bloom, 2008; Corriveau et al., 2009) of a social informant affects what children learn from this informant in the future.

At the same time that children learn from others, they also learn *about* others. In the past 10 years, "theory of mind" research has found not only more and more sophisticated psychological understanding at younger ages but also a strikingly consistent rational pattern of advances in that understanding as children get older (Wellman & Liu, 2004). More recently, there has been a renewed interest in children's social cognition and their understanding of social concepts such as in-groups and out-groups (Dunham, Baron, & Banaji, 2008; Kinzler, Dupoux, & Spelke, 2007; Rhodes & Gelman, 2008) and personality traits (Liu, Gelman, & Wellman, 2007). We suggest that the outcomes of other people's actions are not only informative about the causal systems they act on but also socially informative about the actors themselves. Furthermore, we argue that children's inferences about psychological causes of behavior such as traits are fundamentally causal inferences, relying on the same probabilistic learning mechanisms as their inferences about physical systems such as blicket detectors.

Other recent results further support the notion that we can apply probabilistic models to both the social context of causal understanding and the causal context of the social world. Schulz and Gopnik (2004) found that children inferred psychological causal relationships from covariation in much the same way that they inferred physical and biological relationships. Kushnir, Xu, and Wellman (2010) and Ma and Xu (2011) found that infants as young as 14 months old showed some capacity to infer an underlying desire from a person's pattern of nonrandom sampling behavior. Additionally, Kushnir, Wellman, and Gelman (2008) and Sobel, Sommerville, Travers, Blumenthal, and Stoddard (2009) found that children's causal inferences are sensitive to the social environment. On the computational side, Shafto and colleagues (Bonawitz et al., 2011; Shafto & Goodman, 2008; Shafto, Goodman, Gerstle, & Ladusaw, 2010) have modeled how pedagogical information may be used differently than nonpedagogical information in solving inductive problems.

How children learn from social sources of causal information becomes an especially interesting question when we move beyond artificial laboratory tasks such as blicket detectors. Much of the real-world causal evidence children receive involves complex statistical patterns of both actions and outcomes. Consider the case of learning which actions are necessary to open a door. Children might notice that people almost always grasp and then turn a doorknob before the door opens, but sometimes they pull a handle instead. They frequently insert a key into a lock and then turn it before trying the doorknob, but not always. Sometimes the sequence of actions must be repeated a couple of times (for instance, in the case of a jammed lock); other times, the sequence fails and is not followed by the door opening at all. Often, other actions precede the door opening as well – putting down groceries, fumbling around in a purse, ringing a doorbell, sliding a bolt – which of these are causally necessary and which are incidental? Does the order they were performed in matter? Finally, in addition to these observations, children might receive direct testimony about the door. For instance, someone who lived in the house might say that jiggling the key almost always works or someone unfamiliar with the door might guess that this is the case. How might children combine these statements with other sources of causal evidence?

In just this simple example of opening a door, we can see that there are not only many potential types of causal information available but also many different sources of statistical variation and ambiguity. There is variation in the physical data – actions (and other causes) may not always bring about their effects or may only lead to the desired outcome in certain combinations. There is variation in the action sequence – repeated demonstrations of bringing about the same outcome may include different actions. There is variation in people's behavior – some individuals might succeed at opening the door while others fail or might be successful with one door while failing to open another. There is even variation in direct testimony – people may express differing levels of certainty and causal knowledge, and the testimony of multiple people may even conflict. Finally, children must also take into account their own prior knowledge and expectations about not only the causal system in question but also the intentions, knowledgeability, and helpfulness of their social informant, all of which could vary widely across situations.

On the other hand, while all this ambiguity can make the causal inference problem children face more challenging, there are times when the presence of statistical variation can actually be quite illuminating and aid inference. Actions that do not consistently precede outcomes are less likely to be causally necessary. Actions that reliably appear together and, in fact, predict each other, are more likely to be coherent units, corresponding to intentional, goal-directed action. Variations in the certainty and accuracy of a social informant can facilitate our judgments of the trustworthiness of the information they provide. Variation in the success and failure of individuals might help us infer situational or psychological causes for their behavior such as personality traits.

Formal modeling can be extremely helpful in disentangling these complex inferences. First, formal models give us a way of precisely characterizing hypotheses about what the child thinks and knows. Rational constructivism, and probabilistic computational models in particular, is a natural way to approach understanding how social information, along with other evidence, contributes to children's causal reasoning, because they allow us to systematically represent both beliefs and evidence. Intuitively, this can be seen as a formal version of the approach developmental psychologists have used historically. The method is to hypothesize that children have particular beliefs or conceptions of the world and to assume that children's answers and actions follow rationally from those beliefs. For example, if children initially have a non-representational theory of mind, we would expect them to rationally infer in a false-belief task that a person will immediately search for an object in the location where it actually resides rather than where she last saw it. The classic developmental methodology, then, is to work backward and infer children's current theories from their answers and actions, by assuming that they are operating under the theory that is most consistent with their behavior.

Describing the child's current conception of the world as a particular rational model gives us a more exact way of both characterizing the child's beliefs and working out the predictions that should rationally follow from those beliefs. It also lets us make predictions about how children should rationally update those beliefs with new evidence. By specifying a model, we make explicit our hypotheses about the prior biases and information children bring to a problem and how these biases should be combined with new information in order to update beliefs or even potentially change models or theories. Conversely, we can compare different possible models of the children's beliefs and see which models are most congruent with children's behavior. This approach allows us to give a more precise justification for attributing particular theories to the children, theories that may or may not be like adult theories.

Second, probabilistic models give us a way of more precisely combining and weighting how different factors interact in the child's mind to bring about a particular response. It is common in developmental psychology to see children make different judgments in different contexts. This inconsistency has sometimes been taken to mean that all children's cognition is variable and context dependent and that there is no coherent conceptual structure to be found (e.g., Greeno, 1998; Lave & Wenger, 1991; Thelen & Smith, 1994). At other times, it has led to unresolved debates, for example, about whether early imitation is rational or not. As we will see, probabilistic models allow one to precisely show how multiple sources of evidence, reflecting different contexts, can be rationally combined and integrated to lead to a particular response.

In this chapter, we report two lines of research that apply the ideas of probabilistic modeling to social cognition and explore the complex and interdependent relationship between social and causal learning. In the first set of studies, we examine how the social context, in the form of both demonstrations and testimony, influences children's causal learning. We also examine how causal learning can influence the understanding and segmentation of action and how observed statistical structure in human action can affect causal inferences. In the second set of studies, we examine how children might use covariation in human behavior to infer and attribute mental traits to others, in the same way that they use covariation in cause and effect data to infer physical causal structure. Both lines of research extend probabilistic models from reasoning about purely physical causes to include children's social cognitive development, while also characterizing the distinctive aspects of psychological and physical causal reasoning.

2. THE SOCIAL CONTEXT OF CAUSAL REASONING

2.1. Jointly Inferring Causal Structure and Action Structure

As we discussed in the introduction, many if not most of the causal outcomes children witness are the result of intentional human action. Children must be able to distinguish the unique actions they see other people performing and recognize their effects in order to understand the reasons behind others' behavior and in order to potentially bring about those effects themselves. But before we can interpret actions, we first must parse a continuous stream of motion into meaningful behavior (Byrne, 2003). What cues do we use to do this? How might infants and young children begin to break into the behavior stream in order to identify intentional, goal-directed actions?

Could the causal relationships between actions and their outcomes in the world help children understand action structure itself? How might children identify reaching, grasping, and turning and then group them into the action "opening the door"?

One way that infants might be able to segment actions is by using statistical regularities in human motion. There is now a lot of evidence that both infants and adults use statistical patterns in spoken language to help solve the related problem of segmenting words from continuous speech (e.g., Aslin, Saffran, & Newport, 1998; Pelucchi, Hay, & Saffran, 2009; Saffran, Aslin, & Newport, 1996; Saffran, Newport, & Aslin, 1996). In these experiments, infants (and adults) listen to an artificial language constructed of made-up words, usually created from English syllables (e.g., dutaba, patubi, pidabu). The words are assembled into a continuous speech stream (e.g., dutabapatubipidabu …), with other potential segmentation cues such as intonation and pauses removed. In these experiments, as in many words in real languages, syllables within a word have higher transitional probabilities than syllables between words – you are more likely to hear ta followed by ba (as in dutaba) than to hear bi followed by pi (as in patubi pidabu). Both infants and adults are able to use these transitional probabilities in order to distinguish words in these artificial languages (dutaba, patubi, pidabu), from part-words – combinations of syllables that cross a word boundary (e.g., tabapa, tubipi), and from nonwords, combinations of syllables that do not appear in the artificial language at all (e.g., dupapi, babibu). Infants have also been shown to succeed at statistical language segmentation even when more naturalistic language stimuli are used (Pelucchi et al., 2009).

More recently, a similar sensitivity to statistical regularities has been shown to play a role in action segmentation in both adults (Baldwin, Andersson, Saffran, & Meyer, 2008) and infants (Roseberry, Richie, Hirsh-Pasek, Golinkoff, & Shipley, 2011). Intriguingly, there is also evidence that children can successfully map words learned through this type of segmentation to meanings (Estes, Evans, Alibali, & Saffran, 2007) and, conversely, can use words they already know to help find segment boundaries and discover new words (Bortfeld, Morgan, Golinkoff, & Rathbun, 2005). Similarly, a recent study shows that, in the visual domain, children use statistical patterns to infer the boundaries between objects and then use that information to make further predictions about how objects will behave (Wu, Gopnik, Richardson, & Kirkham, 2011). So children do not just detect the statistics and then segment the streams accordingly. They actually treat those statistical units as if they were meaningful.

In the same way that words have meanings, intentional actions usually lead to causal outcomes. This suggests that just as identifying words assists in mapping them to meanings, segmenting human action may bootstrap learning about causation and vice versa. Recent work has demonstrated that adults can segment videos of common everyday behaviors into coherent actions (Baldwin et al., 2008; Hard, Tversky, & Lang, 2006; Meyer, Decamp, Hard, Baldwin, & Roy, 2010; Newtson, Engquist, & Bois, 1977; Zacks & Tversky, 2001; Zacks, Speer, Swallow, & Maley, 2010) and that both children and adults can infer causal relationships from conditional probabilities (Cheng, 1997; Gopnik et al., 2004; Griffiths, Sobel, Tenenbaum, & Gopnik, in press; Griffiths & Tenenbaum, 2009). However, researchers have not yet explored whether action parsing and causal structure can be learned jointly.

In our work (Buchsbaum, Griffiths, Gopnik, & Baldwin, 2009, 2012), we adapted a Bayesian word segmentation model (Goldwater, Griffiths, & Johnson, 2009), with actions composed of individual small motion elements (SMEs) taking the place of words composed of phonemes or syllables, and extended this model to incorporate causal information. The key intuition behind this model is that action segmentation and causal structure are jointly learned, taking advantage of statistical evidence in both domains. In the model, sequences of motion that correspond to known actions are considered more likely to be causes, and sequences of motion that appear to be causal (they predict outcomes in the world) are considered more likely to be actions. The inferred action boundaries help determine the inferred causal structure and vice versa. This corresponds to our hypothesis that people believe intentional actions and causal effects go hand in hand. If statistical action structure is a cue to causal relationships then, like our model, people should think statistically grouped actions are more likely to be potential causes than other equivalent sequences. Additionally, if people believe that causal sequences of motion are also likely to be actions, then adults should find causal sequences to be more meaningful and coherent than other sequences with equivalent statistical regularities. Finally, if action segmentation and causal relationships are truly jointly learned, then we should see cue combination and cue conflict effects emerge, as in other cases of joint perceptual inference (Ernst & Banks, 2002).

We tested all these predictions in a set of experiments using "artificial action grammars" as in Baldwin et al. (2008). Just as a sentence is composed of words, which are in turn composed of phonemes or syllables, here an action sequence is composed of actions, which are themselves composed of SMEs. Similar to Baldwin et al., we used video clips of object-directed motions to create three-motion actions, which we then combined to create continuous videos of a person manipulating an object.

Just as people can recognize words from an artificial language, and distinguish them from nonwords and part-words, we also know that they can recognize artificial actions grouped only by statistical relationships and can distinguish these sequences from nonactions (motions that never appeared together) and part-actions (motion sequences that cross an action boundary) (Baldwin et al., 2008). We wanted to see whether people think these statistically defined actions are meaningful sequences that can help them understand and interpret others' behavior and whether they believe that these actions are likely to be causal.

In the first experiment, after watching a video, adult participants rated actions, part-actions, and nonactions on how coherent the sequences seemed to be. They were given the example of removing a pen cap and then writing with the pen as motions "going together" and of removing a pen cap and then tying your shoes as motions "not going together." Participants also rated sequences on how likely they thought those motions were to be causal. In this case, we gave participants a cover story. They were told that some of the sequences of motion they were observing would make the manipulated object play music, but there was no sound in the video, so they would just have to guess how likely each sequence was to cause music.

Adults rated the sequences corresponding to actions as both more coherent and more likely to be causal than the nonactions and part-actions. In fact, after the experiment, some of the participants commented on how much more "sense" some of the action sequences made, often coming up with post hoc intentional explanations for the actor's behaviors ("she shook it to see if anything was inside, then emptied it, then looked inside to check"). This is striking because the "nonactions" of one video were in fact the "actions" of another, meaning that people found the very same sequences of motion to be more meaningful based purely on how frequently they appeared together and how well the component motions predicted each other.

These results show that people's sensitivity to statistical patterns in action is not just an artifact of the impoverished stimuli but plays a real role in their understanding of the structure of observed human behavior. The fact that people found the statistically grouped actions to be more coherent suggests that they do not experience the sequences they segment out as arbitrary but assume that they are meaningful groupings that play some (possibly intentional) role. This is further supported by the fact that, even without being presented with overt causal structure, people believe that the statistically grouped actions are more likely to be causally effective, suggesting that inference of action structure and causal structure really are linked.

People seem to use statistical action structure to infer causal relationships but can they use causal relationships to identify meaningful actions? We hypothesized that when statistical cues to action segmentation are unavailable, adults will be able to use causal structure to identify coherent units of action. In a second experiment, we had adults watch specially constructed videos where all possible combinations of three motions appeared equally often together, so that joint and transitional probabilities could not be used to identify groupings. However, one particular sequence of three motions was chosen to be causal and was always followed by the manipulated object playing music, and this time the sound in the video was on. Adults easily identified the correct set of causal motions from within the longer sequence, one of the first demonstrations of causal variable discovery from a continuous stream of events. Additionally, even though there were no statistically grouped actions in this experiment, participants perceived the causal sequence as being more meaningful (going together better) than the other sequences, suggesting that they had nonetheless segmented it out as a coherent action based on its causal efficacy.

In a third experiment, we looked at the inferences people make when both types of cues – statistical action structure and causal relationships – are present. Can people combine information from both these sources of evidence, even when they conflict? This type of cue integration is often used as evidence of true joint inference, for instance, when visual and haptic information about the same stimuli are combined in inferences about an object's size (Ernst & Banks, 2002). As in our first experiment, we showed adults videos of statistically grouped actions, but now we selected a part-action (a set of motions crossing an action boundary) as the causal sequence that leads to music. Adults appeared to take both the causal relationships and the statistical structure into account, correctly identifying the part-action as the most likely cause, but continuing to rate actions as more likely to also be causal when compared to other part-actions and nonactions. Similarly, they judged the causal part-action to be very cohesive, even though it violated the statistical regularities of the action sequence, suggesting that its causal properties led to it being considered a coherent unit of human action.

Together, these three studies demonstrate that adults, at least, can combine statistical regularities and causal structure to divide observed human behavior into meaningful actions. They can also use this inferred segmentation to help them identify likely causal actions. Additionally, the parallels between people's word segmentation and action segmentation abilities support the possibility of a more general statistical learning mechanism. These results also provide a demonstration that causal and social information can be jointly used to infer goal structures. In the following section, we will look at whether young children make similar types of inferences. Can children identify causal subsequences of action from within a longer action sequence when deciding which actions to imitate?

2.2. Causal Imitation from Social Demonstrations

Imitation is a characteristic and pervasive behavior of human children and so it seems like a natural mechanism for identifying and learning causal actions. How do children choose what to imitate from all the actions they see performed around them? When they see a sequence of behaviors preceding an interesting outcome, can they choose the relevant actions? Do they imitate different portions of sequences when given different evidence about their effectiveness?

Recent studies of children's imitation have produced varying answers to the question of whether children are in fact capable of inferring causal action sequences from observed demonstrations. Children can use information about an actor's prior intentions to help them identify causally effective actions (Carpenter, Call, & Tomasello, 2002). Similarly, when children observe unsuccessful demonstrations, they reproduce the actor's intended goals rather than the unsuccessful actions themselves (Hamlin, Hallinan, & Woodward, 2008; Meltzoff, 1995). In some cases, they vary the precision and faithfulness of their imitation with apparent causal relevance (Brugger, Lariviere, Mumme, & Bushnell, 2007; Harnick, 1978; Williamson & Markman, 2006) and selectively imitate actions based on how causally effective they appear to be (Schulz, Hooppell, & Jenkins, 2008; Want & Harris, 2001; Williamson, Meltzoff, & Markman, 2008). At other times, however, children will "overimitate," reproducing apparently unnecessary parts of a causal sequence (Horner & Whiten, 2005; Lyons, Young, & Keil, 2007; Lyons, Damrosch, Lin, Macris, & Keil, 2011; McGuigan & Whiten, 2009; McGuigan, Whiten, Flynn, & Horner, 2007) or copying an actor's precise means (Meltzoff, 1988) even when this makes them less efficient at accomplishing their goal.

There are even cases where children do both in the same study. In the "rational imitation" studies by Gergely, Bekkering, and Király (2002), children saw an experimenter whose hands were either free or confined activate a machine using their forehead. Children both produced exact imitations of the actor (touching their head to the machine to make it go) and produced more obviously causally efficient actions (touching the machine with a hand), though the proportion of such actions differed in the different intentional contexts. In fact, finding a distribution of imitative responses is the norm across all these studies. Even in the most intriguing demonstrations of overimitation, it is not the case that all children blindly mimic the demonstrator's actions, and similarly, even in experiments where children show an overall appreciation for causal efficacy, some children still imitate unnecessary or ineffective actions.

We are interested in reconciling these results by suggesting that perhaps all these imitative choices are the result of rational imitation using a combination of social, physical, and statistical evidence as well as prior knowledge. In particular, evidence for which actions are causally necessary includes more than just the immediately observed demonstration. It also includes children's previous experiences with causal systems and objects, their prior observations of bringing about the same effect, and social information including the adult's knowledge state, intentions, and pedagogical stance (we know that observing a helpful teacher versus a neutral [Bonawitz et al., 2011; Brugger et al., 2007], ineffective [Schulz et al., 2008; Want & Harris, 2001; Williamson et al., 2008], or naïve [Bonawitz et al., 2011; Butler & Markman, in press] demonstrator changes children's inferences). If different imitative choices are the result of different evidence, then we should be able to manipulate these choices and get children to imitate different portions of the same action sequences by changing the combination of social and physical evidence they receive.

Moreover, in many real-world situations, the causal structure of a demonstrated sequence of actions is not fully observable, and which actions are necessary and which are superfluous may be unclear. Therefore, there is often no single "right answer" to the question of what to imitate. After all, a longer "overimitation" sequence might actually be necessary to bring about an effect, though that might initially seem unlikely. One way in which children may overcome this difficulty is by using statistical evidence provided by repeated observations of bringing about the effect. By watching someone unlock and open a door or turn on a light bulb on multiple occasions, children can detect which actions consistently predict the desired outcome and which do not.

To test this prediction, we ran an experiment that manipulated the statistical evidence children received from a series of demonstrated action sequences (Buchsbaum, Griffiths, Gopnik, & Shafto, 2011). We used a Bayesian model to help us construct demonstration sequences that normatively predict selective imitation in some cases and "overimitation" in others. If children make rational inferences from variations in the action sequences they observe, then their choice of whether to imitate only part of an action sequence versus the complete sequence should similarly vary with the evidence.

In this study, children watched a naïve informant (who claimed to have no knowledge of how the toy worked) demonstrate five sequences of three actions each on a toy (e.g., the experimenter squishing the toy, then shaking it, and then rolling it would be one sequence). Some of these sequences but not others led to the toy playing music. In the "ABC" data condition, the same three actions (e.g., knock, shake, pull) always made the toy play music, while in the "BC" data condition, the first action of the successful sequences varied while the final two actions preceding the music stayed the same (e.g., knock, shake, pull, or squish, shake, pull, or roll, shake, pull would all be followed by the toy playing music). Children either could exactly reproduce one of the three-action sequences that had caused the toy to activate or could just produce the final two actions in isolation.

Intuitively, it is more likely that all three-actions are necessary in the "ABC" condition, while perhaps only the final two-actions are necessary in the "BC" condition. However, both three-action and two-action sequences reflect potentially correct hypotheses about what caused the toy to activate in either condition. It could be that the last two-actions by themselves cause the toy to activate in the "ABC" condition and the first is superfluous or it could be that three-actions are necessary in the "BC" condition, but the first action can vary. It is just the probability of these hypotheses that changes between the two conditions. Our Bayesian model predicts just those differences in probability.

If children automatically encode the adult's successful actions as causally necessary, then they should exclusively imitate three actions in both conditions. However, if children are also using more complex statistical information, then we expect that children in the "ABC" condition should reproduce three actions more often than children in the "BC" condition and that children in the "BC" condition might imitate the two-action subsequence by itself. This is, in fact, what we found – children imitated all three actions almost exclusively in the "ABC" condition, while children in the "BC" condition imitated much more variably, with a number of them imitating the two-action subsequence, even though they had never seen it performed on its own and even though three actions would have also activated the toy. Like adults in our first set of experiments, preschool children used statistical patterns to identify causal subsequences within longer sequences of action.

The particular model parameters that best fit children's performance also tell us something about children's expectations going into this task. The model suggested that children employ a causal Occam's razor, assuming that simpler hypotheses, which require fewer unique causal sequences to explain the data, are more likely than more complex hypotheses. The model also suggested that children were biased to imitate the adult's complete action sequence (though this bias could be overcome), perhaps indicating a preexisting belief that adults usually do not perform extraneous actions.

Children might make this "rational actor" assumption because they are using information about the adult's knowledgeability (e.g., Jaswal, 2006; Kushnir et al., 2008), reliability (e.g., Koenig, Clément, & Harris, 2004; Zmyj, Buttelmann, Carpenter, & Daum, 2010), and intentional stance (Bonawitz et al., 2011; Butler & Markman, in press). For instance, children might notice that the experimenter always performs three actions and infer that the experimenter, while not knowing the exact causal sequence, knows that it must be three actions long. We explored this possibility in a next experiment, where we manipulated the intentional state of the demonstrator rather than the statistics of the demonstration.

In our original study, the experimenter acted clueless, as if she did not know anything about how the machine worked. In the next study, the experimenter became a knowledgeable teacher. She told the children that she was showing them how the machine worked – and then showed them exactly the same sequences of actions as in the original "BC" condition. Now, children were much more likely to "overimitate;" almost all of them reproduced a complete sequence of three actions. So children made different causal inferences depending on the social context. When it was their turn to bring about the effect, children chose to reproduce more of the demonstrated actions when the demonstrator was a knowledgeable teacher than when she was naïve about the workings of the toy. Intuitively, children, like our model, understood that a helpful teacher would only be demonstrating all these extra actions if they were in fact necessary to make the toy work (see, e.g., Shafto & Goodman, 2008, for more details on Bayesian models of inference from pedagogically versus non-pedagogically selected data).

These studies suggest that causal learning is informed by both social knowledge and statistical information. Children are sensitive to probabilities, knowledge state, and pedagogical intent when deciding which actions to imitate. These studies also suggest a rational account of "overimitation." In particular, imitating three actions in these studies can be thought of as a kind of overimitation, reproducing parts of a causal sequence that are not actually demonstrably necessary for the effect. These results suggest that this behavior varies depending on the statistics of the data and the probability of various hypotheses concerning them. "Overimitation" also varies depending on the social demonstrator. By explicitly representing the contributions of these different sources of evidence and using them to assign probabilities to causal hypotheses, a Bayesian model can predict these behaviors quite precisely.

Many of the studies of imitation we discussed earlier in this section did not provide the child with either clearly pedagogical or nonpedagogical demonstrators. These demonstrators may have used cues such as directed gaze and pointing (Csibra & Gergely, 2009; Gergely et al., 2007; Senju, Csibra, & Johnson, 2008), leading children to assume that they were in a teaching situation. In general, these studies also showed children only one way to bring about the desired effect and used causal systems where children's prior expectations were unclear. These differences may help explain why children's imitative choices seem so varied across studies. This work also suggests that despite appearances, such behavior is a rational response to different combinations of social, statistical, and physical information. In situations where causal structure is ambiguous, children not only take advantage of social demonstrations, they use relevant information about the demonstrators themselves to make causal inferences.

2.3. Causal Inference from Social Testimony

The previous experiments show that social observation influences children's causal reasoning. Children used the demonstrator's intentional state to help infer which actions were causally necessary to produce an effect. Presumably, children in these experiments were learning not only about the causal system but also about the causal demonstrator. What assumptions might children have made about the value of the demonstrator as a social informant and how might these assumptions guide children's future interactions with that person? In this section, we address these questions by investigating the influence of a different type of social information: verbal testimony. What can we learn from other people's causal statements about the world and what might the world tell us about the reliability of those statements and thus, other people?

Much of what we know about the world we learn from what other people tell us to be true. Our parents, teachers, and peers are continually providing us with information about the causal structure and mechanisms of our environment (e.g., Callanan & Oakes, 1992). However, the role of verbal testimony in causal learning is not obvious a priori. Cause and effect relationships can often be inferred from direct observation without explicit instruction. Again, a child could learn that turning a key in the lock makes a door open because someone said this is so, but a child could just as successfully learn this causal link simply by observing someone turning a key and seeing the door open. So how might children use informant testimony in the context of causal inference?

When what we hear corroborates what we observe, then testimony should facilitate children's causal understanding. However, as we discussed in the introduction, the real world is stochastic and unpredictable, and informants might be ignorant, mistaken, or even deceptive. What would happen if the testimony children receive conflicts with what they see? Would children choose one source to rely on or integrate information from each to inform their causal judgments? To use testimony effectively, we must know when it is prudent to trust others and when they are likely giving inaccurate information. What would a conflict between testimony and observation tell children about the credibility of their informant?

Young children have a strong bias to trust the testimony of others (Jaswal et al., 2010). However, children are not entirely credulous. Just as they can use patterns of evidence to make sophisticated judgments about the relative strengths of different causes (e.g., Kushnir & Gopnik, 2005), children can use patterns of past accuracy to make sophisticated judgments about the relative credibility of different informants. Preschoolers are more likely to trust future testimony from informants who have demonstrated that they tend to be knowledgeable and accurate over that of informants who have demonstrated ignorance and inaccuracy (e.g., Corriveau et al., 2009; Koenig & Harris, 2005; Pasquini et al., 2007). This phenomenon is referred to as selective trust (Koenig & Harris, 2005). Past accuracy is not the only cue children rely on, however. As we saw in our studies of children's imitation, children also take the expressed confidence of the informant into account and are more likely to trust the testimony of informants who speak with confidence than informants who indicate that they are unsure (e.g., Jaswal & Malone, 2007; Tenney, Small, Kondrad, Jaswal, & Spellman, 2011).

Informants are not only reporters about the world but also reporters of their own knowledge states. Informants may be unreliable because they hold mistaken beliefs about the world or because they hold mistaken beliefs about the extent of their own knowledge. Thus, yet another cue children might use to evaluate testimony is an informant's level of self-knowledge: how well their confidence predicts their accuracy (what Tenney et al., 2011, refer to as calibration). Research on eyewitness testimony has suggested that though children are sensitive to an informant's confidence and past accuracy, they are not sensitive to an informant's level of self-knowledge, whereas adults are attuned to all three cues (Tenney et al., 2011).

Some recent research has explored how children combine social information with their observations when making causal inferences. This research, including the studies on imitation in our lab, has found that just as children consider testimony from certain informants more informative than others based on past reliability, children find the interventions of certain causal demonstrators more informative than others based on the social information the demonstrators offer. For example, children favor the causal interventions of a demonstrator who claims to be knowledgeable about the causal system over those of a demonstrator who claims to be naïve (Kushnir et al., 2008). Children also learn more from a disambiguating intervention when the demonstrator supplies an explanation relevant to the causal problem at hand than when the demonstrator supplies an irrelevant rationale (Sobel & Sommerville, 2009). Additionally, children are better able to infer causal strength from probabilistic data when the demonstrator acts surprised by the anomalous outcomes (Sobel et al., 2009). Together this research shows that children's causal inferences are not solely determined by the statistical evidence children observe but are also mediated by the social information communicated by the demonstrator.

What happens, though, when the social information explicitly contrasts with children's observations? How might children handle a conflict between what observed statistical data show and what an informant says? Furthermore, what inferences do children make about the reliability of an informant based on the information they provide and the causal evidence children see? In the following experiment (Bridgers, Buchsbaum, Seiver, Gopnik, & Griffiths, 2011, 2012), we explored children's causal and social inferences when they were presented with a disagreement between an informant's statements and their own causal observations.

The experiment involved four between-subject conditions: the knowledgeable conflict condition, the naïve conflict condition, the knowledgeable baseline condition, and the naïve baseline condition. In the conflict conditions, we investigated how 3-, 4-, and 5-year-olds resolve a conflict between the information provided by either a knowledgeable or a naïve informant and by probabilistic causal demonstrations. In the baseline conditions, we explored preschoolers' baseline trust in a knowledgeable and a naïve informant's testimony in the absence of conflicting data. We describe the conflict conditions first.

The knowledgeable and naïve conflict conditions had two within-subject phases: the causal phase and the generalization phase. First, in the causal phase, children were introduced to two blocks (the "causal" pair) and a machine that lit up and played music when certain blocks were placed on top. An informant explained to the children that one block was better at activating the machine than the other. In the knowledgeable conflict condition, the informant claimed to really know which block was better, while in the naïve conflict condition, the informant said she was just guessing. The informant then left the room, and a second, neutral experimenter demonstrated the blocks on the machine, providing probabilistic evidence that in both conditions, challenged the informant's statement: The block endorsed by the informant was actually less causally efficacious, statistically speaking, than the unendorsed block. The endorsed block only activated the machine two out of six times, while the unendorsed block activated it two out of three times (past research has shown that children can correctly infer causal strength from this pattern of activation; see Kushnir & Gopnik, 2007). Children were then asked to choose which block they thought was better at activating the machine. Children were thus confronted with an ambiguous situation in which they had to decide whether or not the informant or perhaps their own observations were unreliable. In the generalization phase, the informant returned with two novel blocks (the "generalization" pair) and, in both conditions, claimed that she *knew* which block was better at activating the machine. Last, children were asked by the neutral experimenter to choose one of these new blocks to make the machine go.

One might expect that when provided with contradictory verbal and visual information, children would always trust what they directly see over what they hear. However, children might instead rationally combine their prior beliefs about the reliability of these types of sources with the evidence to make a joint causal and social inference. In doing so, children would also update their beliefs about the validity of both the informant's testimony and the observed causal data, which would affect their later inferences from these sources. Therefore, as in our previous experiments looking at children's causal imitation, we expected children's inferences to vary with both the social and the causal evidence.

Our results suggest just such an interaction. In the causal phase of the naïve conflict condition, children overwhelmingly trusted the data and chose the unendorsed block as better at activating the machine, while children in the knowledgeable conflict condition were torn between the two blocks. Thus, when there was strong evidence supporting the causal efficacy of each block (the knowledgeable informant's testimony for the endorsed block and the causal observations for the unendorsed block), children were at chance between inferring the endorsed or the unendorsed block as the better cause. When the testimony was weaker (because the informant was naïve), children favored the block that the causal data suggested was better. Though the causal evidence was constant across conditions, children put more confidence in the knowledgeable informant's claim than in the naïve informant's guess and so were willing to believe the informant over their own observations when she expressed certainty but not when she expressed uncertainty. As predicted and again, as demonstrated in our imitation experiments, children's causal inferences about the same pattern of causal data differed depending on the social context. These results suggest that children combine informant testimony and causal data to infer causal relations even when these cues conflict.

In the conflict conditions, the informants expressed different levels of knowledge about the causal blocks initially but both were wrong in their endorsements. Thus, in addition to different levels of claimed knowledgeability, the informants also had different levels of self-knowledge. Even though both informants made incorrect predictions about the causal blocks' relative causal strengths, the naïve informant actually demonstrated more self-knowledge because she was aware that she did not know about the blocks. The knowledgeable informant, on the other hand, was oblivious to the fact that she was mistaken in her beliefs. Therefore, when both informants later say that they "know" about the generalization blocks, it is more judicious to trust the previously naïve informant because she is more likely to actually know about the causal system when she says that she does. Would children be sensitive to this difference?

At first glance, the answer appears to be no. In the generalization phase, there was no difference in performance across the two conditions; children were equally likely to extend trust to the informant who was correctly uncertain in her prior testimony as to the informant who was incorrectly certain. In both the knowledge and the naïve conflict conditions, children were willing to trust the informant and intervened with the generalization block she endorsed more often than the unendorsed block. Children's failure to selectively trust the naïve informant in the generalization phase implies that, as earlier research has suggested, children are more sensitive to an informant's expressed knowledge level about the general world than to her level of self-knowledge (e.g., Tenney et al., 2011).

However, there may be alternative explanations for children's performance in the generalization phase that do not assume that children entirely lack a concept of self-knowledge. In the causal phase, the informant endorses one block, while the causal data "endorse" the other. However, in the generalization phase, though the informant endorses one of the new blocks, there is no evidence about the second block. Since there is no evidence directly contradicting the informant's claim, there is a relatively low cost in choosing to intervene with the endorsed generalization block over the unendorsed one. Furthermore, the conflicting data observed in the causal phase are probabilistic, suggesting that perhaps the causal data, rather than the informant, are the unreliable information source. Maybe the informant was correct about the relative efficacies of the causal blocks and the particular pattern of causal data the children observed were merely a fluke (for instance, the result of faulty wiring or battery failure) and unrepresentative of the actual causal system. Additionally, children's strong tendency to trust testimony may have further convinced them to trust the informant about the generalization blocks regardless of the conflict observed and the informant's prior knowledge state. In summary, given children's bias to believe testimony, their beliefs about the reliability of stochastic data, and the low cost of intervening with the endorsed generalization block, children could be sensitive to self-knowledge and still rationally trust both informants more or less equally.

To help us better understand children's performance in the generalization phase of the conflict conditions, we need to consider the level of trust that children place in a knowledgeable and a naïve informant's causal testimony when no conflicting causal data are present. In the knowledgeable and naïve baseline conditions, children were told by either a knowledgeable or a naïve informant, respectively, that one block was better than another at making a machine go, without seeing any causal demonstrations of either one. These two conditions are identical in structure to the generalization phase of the two conflict conditions, but since there is no preceding causal phase, children only have the informant's current testimony to guide their intervention choice. We predicted that in general across both baseline conditions, children would trust the informant but would be more likely to do so when the informant expressed certainty than when she expressed uncertainty. And that is basically what we found: A majority of children chose to intervene with the endorsed block across both conditions though slightly fewer did so when the informant was naïve.

We can compare children's performance at baseline to their performance in the generalization phase of the conflict conditions. In both situations, children are presented with informant testimony endorsing one of two blocks but are not given a chance to observe the blocks on the machine. However, in the baseline conditions, children have no prior experience with the informants, while in the generalization phase of the conflict conditions, children have witnessed a disagreement between the informants' earlier statements and the causal data. If children's trust in the informants is influenced by this conflict, we would expect children to be less trusting of the informants in the generalization phase of the conflict conditions than in the baseline conditions. When we make this comparison, we find that overall, more children intervened with the endorsed block in the baseline conditions than in the generalization phase of the conflict conditions, suggesting that children were in fact more willing to trust the informants before observing a conflict between their testimony and the causal data than afterward. Moreover, there was a greater decline in children's trust in the knowledgeable informant than in the naïve one. This may indeed suggest that children are potentially sensitive to self-knowledge though further research is necessary to test this claim.

This experiment confirms that children's causal judgments are informed by both social knowledge (in this case, testimony) and statistical data. Additionally, these experiments provide us with further insight into how children combine information from these sources. They demonstrate that children do not entirely discount one source and privilege another when the information from each conflicts. Rather children are evaluating, weighting, and integrating information from both social and physical cues to guide their inferences about both the causal system and the informant.

This situation lends itself particularly well to Bayesian modeling since children are being asked to combine information from two probabilistic sources and the disagreement between the two only adds to the complexity and ambiguity of the information each source provides. We are developing just such a model (Buchsbaum et al., in press) to better understand how children might be relating testimony and direct observation in their social and causal inferences (for a related model, see Eaves & Shafto, this volume).

3. USING CAUSAL INFERENCE TO LEARN ABOUT PEOPLE

The studies described in the preceding sections demonstrate how the social domain can inform our inferences about physical causation and how the causal outcomes of people's actions can be used to make inferences about both the causal structure of the physical world and the intentions, knowledge, and reliability of the social demonstrator or informant. But as adults, we not only make causal inferences about physical systems, we also make extensive inferences about the causes of people's behavior, a process termed "attribution." How do children reason about these psychological causes? Does their reasoning about the causes of human behavior proceed along the same lines as their reasoning about physical causation?

Even in the earliest years of life, babies are already making attributions about other people and figuring out the causes of their behavior. For example, even infants expect there to be different sources of movement for physical objects and people (Saxe et al., 2005; Schult & Wellman, 1997; Woodward, Phillips, & Spelke, 1993). In other studies, infants are capable of even more sophisticated reasoning about others' behavior; for example, they expect that agents who help others reach their goal will be treated differently than agents who hinder others' progress and they also treat helpers and hinderers differently themselves (Hamlin, Wynn, & Bloom, 2007; Kuhlmeier, Wynn, & Bloom, 2003).

As children grow older, they show increasingly sophisticated understanding of such social constructs as in-groups and out-groups. Some more sophisticated aspects of social cognition, however, do not seem to emerge until much later in children's development, in some cases not until the school-age period (Rholes & Ruble, 1984; Ruble & Dweck, 1995). This includes the propensity to causally explain human behavior in terms of personality traits.

In particular, people explain the causes of human actions in different ways. First, they may attribute a person's actions to internal, individual, and enduring characteristics (i.e., traits). An internal attribution places the cause of behavior in the mind of the acting agent. To revisit the case of interpreting the opening of a door, we might see someone open a door and think she did it because she is the kind of curious person who enjoys opening doors. Second, people may attribute actions to external situations, circumstances, or other objects in the environment. An external attribution for opening the door might be that it was a hot day outside or that the cat needs to be let in. These different styles of attribution have far-reaching consequences; social psychologists have found that a preference for one type of causal explanation and attribution affects other kinds of social cognition and behavior, such as motivation, achievement, blame, mental health, and general emotional well-being (e.g., Levy & Dweck, 1998), even in children (Levy & Dweck, 1999; Patrick, Skinner, & Connell, 1993). Especially in Western cultures, many adults have a bias to attribute the actions of others to individual enduring traits of the person rather than to external situations (Jones & Harris, 1967; Ross, 1977). Some researchers have suggested that this is because these adults have developed an intuitive theory that explains action in terms of such traits (Molden, Plaks, & Dweck, 2006; Morris & Peng, 1994; Rosati et al., 2001). That theory might then bias the observer's interpretation of behavioral evidence toward favoring internal causes.

Where do these attributions come from? It is unclear when and why children begin to explain action in terms of internal, individual, and enduring traits. Even very young children explain action in terms of internal mental states (Flavell, Flavell, Green, & Moses, 1990). However, trait explanations include two additional factors beyond mental states themselves – traits are specific to particular individual people, and they are constant over time and across situations. Many researchers have demonstrated that children do not spontaneously explain actions in terms of traits or endorse trait explanations for a single instance of behavior until middle childhood (Alvarez, Ruble, & Bolger, 2001; Peevers & Secord, 1973; Rholes & Ruble, 1984; Shimizu, 2000). However, other studies show that when preschoolers are given trait labels or behavioral frequency information, they can use that information to make inferences about future behavior and that they can infer the right trait label from frequent behaviors (Boseovski & Lee, 2006; Ferguson, Olthof, Luiten, & Rule, 1984; Heyman & Gelman, 1999; Liu et al., 2007; Matsunaga, 2002). On the other hand, these preschoolers still did not spontaneously construct trait explanations; rather they simply matched the frequency of behaviors to trait labels that were provided for them. This suggests that the failure to attribute traits more broadly is not simply a problem with word comprehension or conceptual development.

More significantly, we do not know the learning mechanisms that underlie the course of attribution in childhood and beyond. Kelley was one of the first psychologists to suggest that person and situation covariation evidence might play an important role in internal versus external attributions (Kelley, 1967; Plaks, Grant, & Dweck, 2005). Empirical studies confirm that adults use statistical information tracking multiple people in multiple situations to make behavioral attributions (Cheng & Novick, 1990; Hewstone & Jaspars, 1987; Morris & Larrick, 1995; Orvis, Cunningham, & Kelley, 1975; Sutton & McClure, 2001). However, adults already have intuitive theories of action they can apply to the covariation data to interpret and predict behavior. Could covariation play a role in the development of trait attribution itself?

As we discussed earlier in this chapter, Bayesian causal learning theories, in particular, suggest that children make new rational inferences by systematically combining prior knowledge and current covariation evidence to arrive at the right causal hypothesis. This suggests a potential mechanism for the development of attribution. Children may begin by observing person and situation covariation evidence that confirms a particular type of hypothesis, particularly the hypothesis that internal traits cause actions. Once that theory has been highly confirmed, it will be more difficult to overturn in the future, though it might still be overturned with sufficient evidence. Eventually, in adulthood, this may result in a consistent "trait bias" that is difficult and thus requires a larger amount of contrary evidence to overcome.

In a series of studies, we examined the developmental origins of Kelley's social schemas. We integrated research on the development of causal inference and trait attribution to see if the same domain-general machinery children used to learn about physical causation in our experiments on causal imitation and causal testimony might also underlie their reasoning about psychological causation.

3.1. Reasoning about Psychological Causes

First, Seiver, Gopnik, and Goodman (in press) conducted a study where 4- and 6-year-old children observed a scenario of two dolls playing on two activities (chosen from a bicycle, trampoline, and diving board). Children were either in the doll condition (where the two doll characters acted consistently on the two activities and differently from each other) or in the toy condition (where both dolls played on one toy activity and did not play on the other). The children in each condition received different covariation information about the person and situation while still observing the same overall frequency of playing and not playing. At the end, we asked the children to explain the dolls' actions (e.g., "Why did Josie play on the bicycle?") and predict their behavior in a future situation.

In the doll condition, one doll always plays and the other doll never plays. This evidence suggests that something about the individual rather than the situation is responsible for their behavior. In the toy condition, the two characters never play on one toy and always play on the other, suggesting instead that the situation or the toy itself is responsible for their actions. So how would children explain the dolls' behavior in these two different conditions? Four-year-olds more closely tracked the behavioral data than 6-year-olds and offered explanations that matched the data. For example, in the doll condition, when the overall pattern of behaviors indicated that something about the person was responsible for the dolls' behavior, both 4- and 6-year-olds gave internal explanations for their behavior – explanations about the person, including physical characteristics such as age or height or mental states such as desires and beliefs. However, in the toy condition, when the data indicated that the situations were driving the dolls' actions (i.e., they both played on one activity and did not play on the other), 4-year-olds appropriately gave more external explanations – explanations involving the environment or the specific toy activity – but 6-year-olds persisted in giving internal explanations. This difference in attribution style between the two age groups in the toy condition suggests that the 4-year-olds were more sensitive to the covariation data than the 6-year-olds. Further evidence included a control condition where children were asked to explain why a single doll did or did not play on a single activity. In this case, the data are ambiguous about the possible cause of the behavior. In the control condition, 6-year-olds gave internal explanations significantly more often than chance and 4-year-olds were, correctly, at chance.

The prediction question provided additional evidence for 6-year-olds' preference for internal causes. In the doll condition, children were asked to predict whether each doll would play or not play on a new toy. Both 4- and 6-year-olds generalized from the previous pattern of data and said that the doll who had played before would play on the new toy and the doll who did not play before would continue to refrain from playing. In the toy condition, children were asked to predict whether a new doll would play or not play on the same two toys. Four-year-olds again accurately assessed the data and said she would play on the toy the other dolls played on but would not play on the one that they backed away from. Six-year-olds, on the other hand, did not predict consistent behavior in this case.

This provides further evidence for age differences in children's behavioral causal inference. Four-year-olds predicted the pattern of playing and not playing would be consistent in the future, irrespective of whether it favored internal or external attributions. Six-year-olds thought that only the behavioral data that supported internal attributions would generalize to future behavior.

This pattern of results suggests that 6-year-olds have developed a specific prior attributional theory that the internal qualities of a person, rather than the situation, drive their behavior. Six-year-olds seem to have developed expectations about the source of people's actions, so when they are asked to explain the cause of a person's behavior, 6-year-olds use both the actual evidence at hand and their prior beliefs to arrive at a conclusion. The 4-year-olds, in contrast, seem to use a more general "bottom-up" data-based strategy and only use the most immediately available data to draw conclusions.

How domain specific or how general is this higher-order bias? Would it only be applied to the case of psychological causation or would children reason similarly about internal versus external causes of physical outcomes? Studies with adults suggest that there is a relationship between cultural attributional biases and seemingly unrelated views about physical causes. Some studies have shown in adults that culturally based attributional biases affect scientific reasoning (e.g., Morris & Peng, 1994). Even though culturally based differences may be rooted in social cognition, they also cause differences in reasoning about simple Newtonian physics. Westerners, for example, who have a stronger social trait bias, are also more likely to attribute causal power to individual physical objects rather than to relationships or forces.

3.2. From People to Magnets

To explore potential attributional biases in understanding physical causation, we replicated the previous study with children but changed the outcome of interest to a physical rather than psychological one – "stickiness" instead of willingness to play. Without changing the task in any other way, we altered the cover story to implicate physical instead of psychological causation. Thus, rather than saying that the doll character was playing on the scooter, we would say that the doll was sticking to the scooter. The relevant explanatory question then became "Why did the Josie doll stick to the scooter?" We again divided children's responses into two categories. In "internal" responses, children talked about the properties of the doll. In "external" responses, they talked about properties of the toy.

When we made this small modification, changing the language used to describe the dolls to be physical rather than agentive, 6-year-olds lost their overall preference for internal explanations. Moreover, for predicting future sticking or not sticking, 6-year-olds now reliably extended the data pattern in both conditions. That is, they were willing to use the most recently available data to make causal inferences instead of relying on their prior beliefs.

Four-year-olds, however, still gave more accurate explanations than 6-year-olds – they continued to follow the data pattern more closely and gave more internal explanations in the doll condition and external explanations in the toy condition. Closer examination of the results suggests that once again the 6-year-olds had shifted from largely relying on the data to relying instead on a prior bias. Unlike in the original psychological case, however, the 6-year-olds gave explanations in terms of a rather different everyday causal theory – namely, magnetism. They often appealed to the scientific properties of magnetism, such as the relationship between magnets and metal, in their explanations. They also were more likely to give interactive causal explanations that implicated both the doll and the toy as causes for the outcome (e.g., "she has metal shoes and the skateboard is a magnet"). Children never produced these interactive explanations in the social case, and 4-year-olds rarely produced them in the physical case. These explanations suggest that the 6-year-old children relied on a deeper and more scientifically based causal framework about stickiness and magnetism, in particular, rather than relying on the data. Four-year-olds tended to give more vague answers such as "she has sticky stuff on her feet" and were less sophisticated in terms of appealing to physical mechanisms and magnets. However, they again tracked the data more accurately, perhaps due to this less sophisticated understanding of magnetism's interactive properties, and therefore weaker prior beliefs.

3.3. Cross-Cultural Studies

We are also conducting versions of these studies in Beijing, China, to compare children's beliefs about psychological causes across cultures. Research with adults shows that broadly speaking, the cultural preference for trait explanations in the United States is not present in Chinese culture (e.g., Morris & Peng, 1994). If there is truly a cultural difference in prior expectations about the causes of people's actions, then 6-year-olds in a non-trait-biased culture should not perform the same way as the American

6-year-olds; that is, 6-year-olds in China should not show a bias favoring internal explanations. Although the data collection is still ongoing, preliminary results suggest that 6-year-olds in Beijing have a similar expectation in the control condition, where the data does not support an internal or an external explanation, to American (and Chinese) 4-year-olds. They are at chance for preferring internal or external explanations. These findings suggest that by the age of 6, children's prior beliefs about others' behavior are influenced by culture and these attributional styles shape their interpretation of new behavioral information.

4. CONCLUSIONS

Taken together, the studies in this chapter show how the tools of probabilistic modeling and Bayesian learning can be applied to the social as well as the physical domain and how the physical and social domains can jointly inform each other. When children learn about causes from other people, whether through demonstration or testimony, they appear to integrate their prior hypotheses about pedagogy, cues to informant reliability, and the statistical evidence they observe from people's actions. Children are sensitive to the pedagogical intent of a demonstrator and can use this information to aid their decisions about which of the demonstrator's actions to imitate in order to bring about an effect. Similarly, children can use an informant's statements to help them evaluate the data produced by a causal system and likewise can use these data to evaluate the credibility of the testimony produced by an informant. They also can use covariation information to decide whether a situational or psychological cause is a better explanation for a person's behavior and take into consideration whether the event involves people or just physical objects. Together, these studies demonstrate that causal reasoning and social reasoning are linked, both in the real world and in children's minds. When children reason about physical causal systems, they are incorporating social information, and when children reason about seemingly purely social issues like personality trait attribution, they are applying the same causal reasoning that they use for physical systems.

The rational constructivist approach can help us understand how children resolve, and even benefit from, multiple sources of ambiguous and probabilistic data, and social data, in particular, in order to solve challenging causal learning problems. And because these data are often probabilistic, Bayesian models help us describe the complex, uncertain, joint inferences about the nature of both other people and the world that underlie our ability to learn from others. At the same time, the work on attribution shows how similarly complex integrations between prior knowledge and current statistics can lead children to understand the actions of others in a new way. In fact, we can construe the information we get from people, either in the form of testimony or observable actions, as causal information. These studies suggest that children use covariation evidence to construct abstract causal schemas that they then employ to explain the behavior of both the people and the objects around them.

The studies on imitation and pedagogy, in particular, suggest that we would be wise to fully consider the social environment when looking at children's physical causal reasoning. The degree of confidence that the social demonstrator has, and the level of authority they convey to the child, might not just socially influence the child to feel pressured to respond in a certain way but also might actually change their inferences about the physical causal events they are observing. In fact, incorporating this social evidence into causal reasoning is a rational response, especially in the face of uncertainty. Therefore, to get a complete picture of how children understand the causal landscape of both the physical and the social worlds we need to understand how they use the entire rich set of data they encounter in the real world. Studies directly manipulating social information, such as how pedagogically the demonstrator is behaving and how much certainty she expresses, integrate the human element into experiments that model causal understanding.

Future research should computationally address how children develop priors about the causes and results of people's behavior and of the social information they provide. What leads children to believe that a person is an expert, and what process guides their assumptions based on that attribution? What are the components of children's pedagogical understanding, and what prior beliefs do children have about the likely causes and effects of pedagogical behavior? How do children integrate data about people's beliefs (via testimony) and actions when making attributions about people's behavior? How do children conceptualize people causing changes in other people's beliefs or actions? What are children's prior beliefs about person-to-person causes, and how would they parse these events? Furthermore, how would they integrate physical causes into those judgments?

The studies in this chapter begin to show how we can move beyond basic laboratory problems, like determining the causal structure of blicket detectors, to more complex inferences that more closely mirror the real world. The probabilistic models approach can be applied to real and ecologically significant kinds of conceptual change. It sheds new light on classic topics in cognitive development such as the nature of imitation and trait attribution. Instead of looking at how children evaluate individual or isolated events, we can more appropriately study how children *learn* in and from the complex social–physical environment that makes up the world around them.

ACKNOWLEDGMENTS

This work is supported by National Science Foundation Graduate Research Fellowships to D.B. and E.S. and National Science Foundation Grant BCS-1023875 to A.G. and by the McDonnell Foundation Causal Learning Initiative and Grant.

REFERENCES

Alvarez, J. M., Ruble, D. N., & Bolger, N. (2001). Trait understanding or evaluative reasoning? An analysis of children's behavioral predictions. *Child Development, 72*(5), 1409–1425.

Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science, 9*(4), 321–324.

Baldwin, D., Andersson, A., Saffran, J., & Meyer, M. (2008). Segmenting dynamic human action via statistical structure. *Cognition, 106*(3), 1382–1407.

Birch, S. A. J., Vauthier, S. A., & Bloom, P. (2008). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition, 107*(3), 1018–1034.

Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., et al. (2010). Just do it? Investigating the gap between prediction and action in toddlers' causal inferences. *Cognition, 115*(1), 104–117.

Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. E. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition, 120*(3), 322–330.

Bortfeld, H., Morgan, J. L., Golinkoff, R. M., & Rathbun, K. (2005). Mommy and me: Familiar names help launch babies into speech-stream segmentation. *Psychological Science, 16*(4), 298–304.

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioral prediction. *Developmental Psychology, 42*(3), 500–513.

Bridgers, S., Buchsbaum, D., Seiver, E., Gopnik, A., & Griffiths, T. L. (2011). *Which block is better at making the machine go?: How children balance their trust in an informant vs. the data.* Poster presented at Biennial Meeting of the Cognitive Development Society. October.

Bridgers, S., Buchsbaum, D., Seiver, E., Gopnik, A., & Griffiths, T. L. (2012). *Which block is better at making the machine go?: How children balance their trust in an informant vs. the data.* Manuscript in preparation.

Brugger, A., Lariviere, L. A., Mumme, D. L., & Bushnell, E. W. (2007). Doing the right thing: Infants' selection of actions to imitate from observed event sequences. *Child Development, 78*(3), 806–824.

Buchsbaum, D., Bridgers, S., Whalen, A., Seiver, E., Griffiths, T. L., & Gopnik, A. (in press). Do I know that you know what you know? Modeling testimony in causal inference. In *Proceedings of the 34th Annual Cognitive Science Society*.

Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition, 120*(3), 331–340.

Buchsbaum, D., Griffiths, T. L., Gopnik, A., & Baldwin, D. (2009). Learning from actions and their consequences: Inferring causal variables from a continuous sequence of human action. In N. A. Taatgen, & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2493–2498). Austin, TX: Cognitive Science Society.

Buchsbaum, D., Griffiths, T. L., Gopnik, A., & Baldwin, D. (2012). *Inferring action structure and causal relationships in continuous sequences of human action*. Manuscript in preparation.

Butler, L. P., & Markman, E. M. (in press). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Development.*

Byrne, R. W. (2003). Imitation as behaviour parsing. *Philosophical Transactions of the Royal Society London B, 358*(1431), 529–536.

Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development, 7*(2), 213–233.

Carpenter, M., Call, J., & Tomasello, M. (2002). Understanding "prior intentions" enables two-year-olds to imitatively learn a complex task. *Child Development, 73*(5), 1431–1441.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*(2), 367–405.

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*(4), 545–567.

Corriveau, K. H., Meints, K., & Harris, P. L. (2009). Early tracking of informant accuracy and inaccuracy. *British Journal of Developmental Psychology, 27*(2), 331–342.

Csibra, G., & Gergely, G. (2006). Social learning and social cognition: The case for pedagogy. In Y. Munakata, & M. H. Johnson (Eds.), *Processes of change in brain and cognitive development. Attention and performance*, Vol. 21 (pp. 249–274). Oxford: Oxford University Press.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences, 13*(4), 148–153.

Dunham, Y., Baron, A. S., & Banaji, M. R. (2008). The development of implicit intergroup cognition. *Trends in Cognitive Sciences, 12*(7), 248–253.

Eaves, B., & Shafto, P. (this volume). Unifying pedagogical reasoning and epistemic trust. In F. Xu & T. Kushnir (Eds.), *Advances in child development and behavior.*

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*, 429–433.

Estes, K. G., Evans, J. L., Alibali, M. W., & Saffran, J. R. (2007). Can infants map meaning to newly segmented words? *Psychological Science, 18*(3), 254–260.

Ferguson, T. J., Olthof, T., Luiten, A., & Rule, B. G. (1984). Children's use of observed behavioral frequency versus behavioral covariation in ascribing dispositions to others. *Child Development, 55*(6), 2094–2105.

Flavell, J. H., Flavell, E. R., Green, F. L., & Moses, L. J. (1990). Young children's understanding of fact beliefs versus value beliefs. *Child Development, 61*(4), 915–928.

Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature, 415*, 755.

Gergely, G., Egyed, K., & Király, I. (2007). On pedagogy. *Developmental Science, 10*(1), 139–146.

Goldwater, S., Griffiths, T. L., & Johnson, M. (2009). A Bayesian framework for word segmentation: Exploring the effects of context. *Cognition, 112*(1), 21–54.

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.

Greeno, J. G. (1998). The situativity of knowing, learning, and research. *American Psychologist, 53*(1), 5–26.

Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (in press). Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science*.
Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review, 116*(4), 661–716.
Hamlin, J. K., Hallinan, E. V., & Woodward, A. L. (2008). Do as I do: 7-month-old infants selectively reproduce others' goals. *Developmental Science, 11*(4), 487–494.
Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature, 450*, 557–559.
Hard, B. M., Tversky, B., & Lang, D. S. (2006). Making sense of abstract events: Building event schemas. *Memory and Cognition, 34*(6), 1221–1235.
Harnick, F. S. (1978). The relationship between ability level and task difficulty in producing imitation in infants. *Child Development, 49*(1), 209–212.
Hewstone, M., & Jaspars, J. (1987). Covariation and causal attribution: A Logical model of the intuitive analysis of variance. *Journal of Personality and Social Psychology, 53*(4), 663–672.
Heyman, G. D., & Gelman, S. A. (1999). The use of trait labels in making psychological inferences. *Child Development, 70*(3), 604–619.
Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal Cognition, 8*(3), 164–181.
Jaswal, V. K. (2006). Preschoolers favor the creator's label when reasoning about an artifact's function. *Cognition, 99*(3), B83–B92.
Jaswal, V. K., Croft, A. C., Setia, A. R., & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science, 21*(10), 1541–1547.
Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development, 8*(3), 263–283.
Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology, 3*(1), 1–24.
Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation*, Vol. 15 (pp. 192–238). Lincoln: University of Nebraska Press.
Kinzler, K. D., Dupoux, E., & Spelke, E. S. (2007). The native language of social cognition. *Proceedings of the National Academy of Sciences of the United States of America, 104*(30), 12577–12580.
Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science, 15*(10), 694–698.
Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development, 76*(6), 1261–1277.
Koenig, M. A., & Jaswal, V. K. (2011). Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Development, 82*(5), 1634–1647.
Kuhlmeier, V., Wynn, K., & Bloom, P. (2003). Attribution of dispositional states by 12-month-olds. *Psychological Science, 14*(5), 402–408.
Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science, 16*(9), 678–683.
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.
Kushnir, T. et al. "Who can help me fix this toy?:" The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology*.
Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition, 107*(3), 1084–1092.

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science, 21*(8), 1134–1140.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.

Levy, S. R., & Dweck, C. S. (1998). Trait-versus process-focused social judgment. *Social Cognition, 16*(1), 151–172.

Levy, S. R., & Dweck, C. S. (1999). The impact of children's static versus dynamic conceptions of people on stereotype formation. *Child Development, 70*(5), 1163–1180.

Liu, D., Gelman, S. A., & Wellman, H. M. (2007). Components of young children's trait understanding: Behavior-to-trait inferences and trait-to-behavior predictions. *Child Development, 78*(5), 1543–1558.

Lyons, D. E., Damrosch, D. H., Lin, J. K., Macris, D. M., & Keil, F. C. (2011). The scope and limits of overimitation in the transmission of artefact culture. *Philosophical Transactions of the Royal Society of London B, 366*(1567), 1158–1167.

Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences of the United States of America, 104*(50), 19751–19756.

Ma, L., & Xu, F. (2011). Young children's use of statistical sampling evidence to infer the subjectivity of preferences. *Cognition, 120*(3), 403–411.

Matsunaga, A. (2002). Preschool children's inferences about personality traits. *Japanese Journal of Developmental Psychology, 13*, 168–177.

McGuigan, N., & Whiten, A. (2009). Emulation and "overemulation" in the social learning of causally opaque versus causally transparent tool use by 23- and 30-month-olds. *Journal of Experimental Child Psychology, 104*(4), 367–381.

McGuigan, N., Whiten, A., Flynn, E., & Horner, V. (2007). Imitation of causally opaque versus causally transparent tool use by 3- and 5-year-old children. *Cognitive Development, 22*(3), 353–364.

Meltzoff, A. N. (1988). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology, 24*(4), 470–476.

Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology, 31*(5), 838–850.

Meltzoff, A. N., Waismeyer, A. S., & Gopnik, A. (in press). Learning about causes from people: Observational causal learning in 24-month old infants. *Developmental Psychology*.

Meyer, M., Decamp, P., Hard, B., Baldwin, D., & Roy, D. (2010). Assessing behavioral and computational approaches to naturalistic action segmentation. In S. Ohlsson, & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Cognitive Science Society* (pp. 2710–2715). Austin, TX: Cognitive Science Society.

Molden, D. C., Plaks, J. E., & Dweck, C. S. (2006). "Meaningful" social inferences: Effects of implicit theories on inferential processes. *Journal of Experimental Social Psychology, 42*(6), 738–752.

Morris, M. W., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review, 102*(2), 331–355.

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology, 67*(6), 949–971.

Newtson, D., Engquist, G., & Bois, J. (1977). The objective basis of behavior units. *Journal of Personality and Social Psychology, 35*(12), 847–862.

Orvis, B. R., Cunningham, J. D., & Kelley, H. H. (1975). A closer examination of causal inference: The roles of consensus, distinctiveness, and consistency information. *Journal of Personality and Social Psychology, 32*(4), 605–616.

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43*(5), 1216–1226.

Patrick, B. C., Skinner, E. A., & Connell, J. P. (1993). What motivates children's behavior and emotion? Joint effects of perceived control and autonomy in the academic domain. *Journal of Personality and Social Psychology, 65*(4), 781–791.

Peevers, B. H., & Secord, P. F. (1973). Developmental changes in attribution of descriptive concepts to persons. *Journal of Personality and Social Psychology, 27*(1), 120–128.

Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Statistical learning in a natural language by 8-month-old infants. *Child Development, 80*(3), 674–685.

Plaks, J. E., Grant, H., & Dweck, C. S. (2005). Violations of implicit theories and the sense of prediction and control: Implications for motivated person perception. *Journal of Personality and Social Psychology, 88*(2), 245–262.

Rhodes, M., & Gelman, S. A. (2008). Categories influence predictions about individual consistency. *Child Development, 79*(5), 1270–1287.

Rholes, W. S., & Ruble, D. N. (1984). Children's understanding of dispositional characteristics of others. *Child Development, 55*(2), 550–560.

Rosati, A. D., Knowles, E. D., Kalish, C. W., Gopnik, A., Ames, D. R., & Morris, M. W. (2001). The rocky road from acts to dispositions: Insights for attribution theory from developmental research on theories of mind. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition* (pp. 287–303). Cambridge, MA: The MIT Press.

Roseberry, S., Richie, R., Hirsh-Pasek, K., Golinkoff, R. M., & Shipley, T. F. (2011). Babies catch a break: 7- to 9-month-olds track statistical probabilities in continuous dynamic events. *Psychological Science, 22*(11), 1422–1424.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology*, Vol. 10 (pp. 173–220). New York: Academic Press.

Ruble, D. N., & Dweck, C. S. (1995). Self-perceptions, person conceptions, and their development. In N. Eisenberg (Ed.), *Social development: Review of personality and social psychology*, Vol. 15 (pp. 109–139). Thousand Oaks, CA: Sage Publications.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*(5294), 1926–1928.

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language, 35*(4), 606–621.

Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science, 16*(12), 995–1001.

Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental Psychology, 43*(1), 149–158.

Schult, C. A., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition, 62*(3), 291–324.

Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology, 43*(5), 1124–1139.

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

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science, 10*(3), 322–332.

Schulz, L. E., Hooppell, C., & Jenkins, A. C. (2008). Judicious imitation: Children differentially imitate deterministically and probabilistically effective actions. *Child Development, 79*(2), 395–410.

Seiver, E., Gopnik, A., & Goodman, N. D. (in press). Did she jump because she was the big sister or because the trampoline was safe? Causal inference and the development of social attribution. *Child Development*

Senju, A., Csibra, G., & Johnson, M. H. (2008). Understanding the referential nature of looking: Infants' preference for object-directed gaze. *Cognition, 108*(2), 303–319.

Shafto, P., & Goodman, N. D. (2008). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1632–1637). Austin, TX: Cognitive Science Society.

Shafto, P., Goodman, N. D., Gerstle, B., & Ladusaw, F. (2010). Prior expectations in pedagogical situations. In S. Ohlsson, & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Cognitive Science Society* (pp. 2182–2187). Austin, TX: Cognitive Science Society.

Shimizu, Y. (2000). Development of trait inference: Do young children understand the causal relation of trait, motive, and behavior? *Japanese Journal of Educational Psychology, 48*(3), 255–266.

Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development, 81*(2), 669–679.

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42*(6), 1103–1115.

Sobel, D. M., & Sommerville, J. A. (2009). Rationales in children's causal learning from others' actions. *Cognitive Development, 24*(1), 70–79.

Sobel, D. M., Sommerville, J. A., Travers, L. V., Blumenthal, E. J., & Stoddard, E. (2009). The role of probability and intentionality in preschoolers' causal generalizations. *Journal of Cognition and Development, 10*(4), 262–284.

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.

Sutton, R. M., & McClure, J. (2001). Covariational influences on goal-based explanation: An integrative model. *Journal of Personality and Social Psychology, 80*(2), 222–236.

Tenney, E. R., Small, J. E., Kondrad, R. L., Jaswal, V. K., & Spellman, B. A. (2011). Accuracy, confidence, and calibration: How young children and adults assess credibility. *Developmental Psychology, 47*(4), 1065–1077.

Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: The MIT Press.

Want, S. C., & Harris, P. L. (2001). Learning from other people's mistakes: Causal understanding in learning to use a tool. *Child Development, 72*(2), 431–443.

Wellman, H. M., & Liu, D. (2004). Scaling of Theory-of-Mind tasks. *Child Development, 75*(2), 523–541.

Williamson, R. A., & Markman, E. M. (2006). Precision of imitation as a function of preschoolers' understanding of the goal of the demonstration. *Developmental Psychology, 42*(4), 723–731.

Williamson, R. A., Meltzoff, A. N., & Markman, E. M. (2008). Prior experiences and perceived efficacy influence 3-year-olds' imitation. *Developmental Psychology, 44*(1), 275–285.

Woodward, A. L., Phillips, A. T., & Spelke, E. S. (1993). Infants' expectations about the motion of animate versus inanimate objects. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 1087–1091). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wu, R., Gopnik, A., Richardson, D. C., & Kirkham, N. Z. (2011). Infants learn about objects from statistics and people. *Developmental Psychology, 47*(5), 1220–1229.

Zacks, J. M., Speer, N. K., Swallow, K. M., & Maley, C. J. (2010). The brain's cutting-room floor: Segmentation of narrative cinema. *Frontiers in Human Neuroscience, 4*(168), 1–15.

Zacks, J. M., & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin, 127*(1), 3–21.

Zmyj, N., Buttelmann, D., Carpenter, M., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. *Journal of Experimental Child Psychology, 106*(4), 208–220.