Psychological Review
2013, Vol. 120, No. 4, 779–797

0033-295X/13/$12.00 DOI: 10.1037/a0034191

Knowledge Matters: How Children Evaluate the Reliability of Testimony as a Process of Rational Inference

David M. Sobel
Brown University

Tamar Kushnir
Cornell University

Children's causal learning has been characterized as a rational process, in which children appropriately evaluate evidence from their observations and actions in light of their existing conceptual knowledge. We propose a similar framework for children's selective social learning, concentrating on information learned from others' testimony. We examine how children use their existing conceptual knowledge of the physical and social world to determine the reliability of testimony. We describe existing studies that offer both direct and indirect support for selective trust as rational inference and discuss how this framework may resolve some of the conflicting evidence surrounding cases of indiscriminate trust. Importantly, this framework emphasizes that children are active in selecting evidence (both social and experiential), rather than being passive recipients of knowledge, and motivates further studies that more systematically examine the process of learning from social information.

Keywords: cognitive development, social learning, selective trust, testimony, causal inference

Over the last 10 years, there has been a great deal of research on young children's causal learning from their own observations and actions (see e.g., Gopnik & Schulz, 2007, for a review). During the same time, there has been interest in another topic: young children's ability to learn selectively from other people's verbal testimony (see e.g., Harris, 2012, for a review). These research programs address a common question: What makes children such extraordinarily good learners? In answering this question, however, the two programs have taken different philosophical approaches (see Clément, Koenig, & Harris, 2004) and also have focused their empirical investigations on the acquisition of different sorts of knowledge.

Investigations of causal learning centrally examine how children acquire conceptual knowledge of physical, biological and psychological phenomena and emphasize how children actively explore and evaluate observed evidence for causal relations (e.g., Carey, 1995; Gopnik et al., 2004; Gopnik & Wellman, 1994; Schulz & Gopnik, 2004; Wellman & Liu, 2007). In keeping with Piagetian traditions, causal learning is viewed as a child-driven process, and the focal question has been the following: How do children learn about the causal structure of the world from patterns of evidence?

This article was published Online First September 9, 2013.

David M. Sobel, Cognitive, Linguistic, and Psychological Sciences Department, Brown University; Tamar Kushnir, Department of Human Development, Cornell University.

We were supported by National Science Foundation Grants 1223777 to David M. Sobel and 1023179 to Tamar Kushnir during the writing of this article. We would like to thank David Buchanan, Kathleen Corriveau, Fiery Cushman, Chris Erb, Mark Fedyk, Deanna Macris, Thomas Rucker, Julie Sedivy, Sarah Suarez, and Chris Vredenbugh for helpful discussion.

Correspondence concerning this article should be addressed to David M. Sobel, Cognitive, Linguistic, and Psychological Sciences Department, Box 1821, Brown University, Providence, RI 02912. E-mail: Dave_Sobel@brown.edu

The predominant approach to this question has been to show that children's causal learning is best described as rational inference (see e.g., Gopnik & Wellman, 2012; Xu & Kushnir, 2013). The term "rational" has a broad set of meanings when applied to human behavior, but when it comes to children's causal learning, it has been precisely defined as the process by which children interpret new evidence—in the form of probabilistic relations—in light of their existing knowledge—in the form of beliefs or "theories" about the world. Many studies have shown that children integrate domain-specific conceptual knowledge into their causal inferences in new situations (e.g., Denison & Xu, 2010; Schulz, Bonawitz, & Griffiths, 2007; Sobel & Munro, 2009).

Moreover, the *strength* of new evidence is critical to learning—weak evidence against existing knowledge does not overturn it, but strong evidence can lead children to revise their existing beliefs and ultimately to a deeper understanding of conceptual domains. Kushnir and Gopnik (2007), for example, examined the relation between preschoolers' assumptions about contact causation ("no action-at-a-distance") and probabilistic evidence. Children were initially assessed to have a "contact bias" as in previous work (e.g., Bullock, Gelman, & Baillargeon, 1982). One group of children saw strong evidence for *distance* causation: Action-at-a-distance was 100% efficacious and action by contact was never efficacious. Another group saw weaker evidence for distance causation: Both actions were efficacious, but action-at-a-distance was probabilistically more efficacious than contact (66% vs. 33%). In all cases, preschoolers correctly identified the more effective object. When later asked to activate the device with a new object, the group that saw strong evidence was overwhelmingly likely to change their initial contact bias (acting at a distance with the new object) while only a small minority in the group who saw weak evidence made the same change.

This approach also makes a commitment to describing children's learning as an active process and suggests that children gather evidence for new causal relations through exploratory play. For example, Schulz and Bonawitz (2007) demonstrated that preschoolers would explore a toy more frequently and for longer periods of time when its efficacy was uncertain (i.e., they did not know which of two activators would produce a novel effect) than when the child knew the exact efficacy of the toy. This active learning process has many benefits over passive observation. Preschoolers who actively discover the causal efficacy of novel objects learn causal relations more accurately (Schulz, Gopnik, & Glymour, 2007; Sobel & Sommerville, 2010). They also make stronger causal inferences from their own actions when evidence is ambiguous or uncertain (Kushnir & Gopnik, 2005; Kushnir, Wellman, & Gelman, 2009).

These are just a few of many examples showing rational, evidence-based causal inference and active causal learning in young children (see also, e.g., Bonawitz et al., 2011; Cook, Goodman, & Schulz, 2011; Griffiths, Sobel, Tenenbaum, & Gopnik, 2011; Gweon & Schulz, 2011; Sobel & Sommerville, 2009; Sobel, Tenenbaum, & Gopnik, 2004). Until recently, however, the role of social transmission in this process—in particular how explicit verbal testimony interacts with children's own actions and observations—has been unaddressed.

In parallel to these recent advances in causal learning, the "trust in testimony" (Harris & Koenig, 2006) literature has shown how children learn directly from others' explicit verbal statements. This line of research begins with the observation that children must learn about many abstract ideas that they cannot possibly gather evidence about through simple observation or solitary play (Harris & Koenig, 2006; Harris, Pasquini, Duke, Asscher, & Pons, 2006; see also Keil, 2010). Thus, this approach rejects the Piagetian assumption that young children are "stubborn autodidacts" (Faucher et al., 2002, p. 341) and instead emphasizes the role of social transmission in the learning process. Not surprisingly, this line of research has also emphasized the acquisition of social knowledge, such as language or cultural conventions. It highlights how children evaluate others' deliberate communicative intent, their trustworthiness, and their reliability as sources of information. In this line of work, the focal question has been the following: What degree of skepticism or credulity do children have about verbal testimony?

The predominant way this question has been examined recently is through studies of children's word learning. Because there is a conventional but arbitrary assignment of lexical items to concepts and objects, children must necessarily learn new words from others' testimony (e.g., Bloom, 2000; de Saussure, 1966; see Koenig & Harris, 2008, for a review). Contrary to classic and contemporary philosophical investigations (Coady, 1992; Reid, 1764/2000), a significant number of studies now show that children are not wholly credulous of testimony but rather are selective. Children as young as 2 are capable of selecting sources by tracking speakers' history of past accuracy (e.g., Clément et al., 2004; Koenig, Clément, & Harris, 2004; Koenig & Harris, 2005; Koenig & Woodward, 2010). Moreover, by the age of 4, children can infer others' accuracy based on probabilistic evidence (Pasquini, Corriveau, Koenig, & Harris, 2007). In addition to past accuracy, young children also register multiple cues that point to a speaker's potential reliability, such as their epistemic knowledge (e.g., Baldwin, 1991, 1993; Koenig & Echols, 2003; Sabbagh & Baldwin, 2001). This research argues that children's rapid learning—of social, cultural, and linguistic information, but also of causal and conceptual knowledge of the world—emanates from their ability to learn from others selectively (see e.g., Bergstrom, Moehlmann, & Boyer, 2006; Harris & Koenig, 2006; Mascaro & Sperber, 2009).

Throughout their daily lives, children learn in all of these ways—from their own observations and interactions with the world and also from the testimony of other people. It is unlikely that such everyday learning is governed by two separate processes—one for causal learning and the other for social learning. Children learn from playing alone but also from playing with others (e.g., Gauvain & Rogoff, 1989). They observe others but also ask questions and seek explanations (e.g., Callanan & Oakes, 1992; Chouinard, 2007; Frazier, Gelman, & Wellman, 2009; Mills, Legare, Bills, & Mejias, 2010; Mills, Legare, Grant, & Landrum, 2011). Even direct instruction is not an assurance that information is directly transmitted. Sometimes children learn ideas, concepts, and skills readily from direct instruction (e.g., Klahr & Nigam, 2004; Markson & Bloom, 1997); other times it takes longer, particularly when testimony runs contrary to direct observations that children themselves make (e.g., diSessa, 1993; Vosniadou & Brewer, 1992). Understanding children's learning, then, requires a unified framework that encompasses all of these situations. Importantly such an account should make systematic predictions regarding when and from whom children learn— through observations, actions, or testimony—and when and from whom they do not. Moreover, such an account should predict what happens when multiple sources of information interact or conflict.

In this article we present an argument for *rational learning from testimony*, resembling the argument that has been previously made for rational causal learning from observations and actions. We begin by noting a common theme in empirical research in both traditions: Children's causal and social learning involves being discriminant (i.e., selective) about new evidence—whether it be evidence for causal relations or evidence for reliability of testimony. We suggest that this is not a coincidental similarity, but rather that it indicates the presence of a shared inferential mechanism—a rational one. Our argument rests on the following related claims:

1. Knowledge Matters. Our central hypothesis is that children's inferences about and from others' testimony recruit their existing conceptual knowledge. This parallels the way in which children use their existing knowledge to infer the strength or quality of evidence for new causal relations (in physical systems, for example, e.g., Kushnir & Gopnik, 2007). In the first few sections of the article, we use examples from the literature on selective social learning to demonstrate this claim. The first set of examples shows how conceptual knowledge about individuals and groups leads children to different judgments of reliability at baseline, prior to observing any information about accuracy. The next examples show how children's existing conceptual knowledge about *kinds* and then about *psychological states* affects how they make use of accurate or inaccurate testimony to infer reliability. We further describe a few examples that suggest developmental differences in selective trust are due to developmental changes in conceptual knowledge.

2. Trust by Inference, Not Association. Some accounts of how children learn from testimony suggest that they associate accuracy or inaccuracy (or more generally, positive or negative valence) with informants and generalize based on those associations (see e.g., Brosseau-Liard & Birch, 2010). We argue against this idea. Making judgments of others' reliability from accuracy is an inferential process, not an associative one. Following from our previous claim, associative mechanisms do not account for the role that conceptual knowledge plays in children's inferences from accuracy to reliability. Throughout our examples, we show that associative learning also does not explain the level of selectivity and domain-specificity of children's generalizations, in particular their selective inferences that others' knowledge extends to conceptually clustered domains.

3. Epistemic Trust Is Rational, but Not All Trust Is Rational. Our discussion is focused on epistemic influences on trust, but there are other influences on trust that are non-epistemic. Emotional cues, for example, have been proposed to be separate but interacting with epistemic ones. To understand fully how and who we trust, these interactions must be considered (e.g., Corriveau, Meints, & Harris, 2009; Fiske, 2010; Koenig & Stephens, in press). We do not suggest that emotional trust is necessarily rational. Below we note empirical examples from the testimony literature that seem to rely on this interaction. These examples suggest that the influences of epistemic and emotional factors are complementary rather than contradictory.

4. Verifiability and Deception. For the most part, we restrict our claims to learning environments in which children are free to observe and explore the world in the presence of helpful and knowledgeable informants. In such an environment, information from each source (observations, actions, testimony) can be verified by information from the other sources to some extent. Cases in which information is never verifiable through direct observation (such as beliefs about supernatural entities, death, or angels) are currently beyond the scope of this approach and have been given more thorough theoretical analysis elsewhere (e.g., Harris, 2012; Woolley & Ghossainy, in press).

Related to this is the issue of what happens when informants are overtly deceptive. Here we address what happens when deception can be discovered (when information can be verified by some other source, or by the child's prior knowledge or direct perception). Critically, our approach to such situations stands in contrast to an alternative theoretical view that children are initially indiscriminate in their trust of testimony because of a default bias to trust verbal information (e.g., Jaswal, Croft, Setia, & Cole, 2010). As we will show below, our account often makes similar predictions as this framework, but for very different reasons. This contrast leaves open many interesting questions for future research.

Inferring Reliability From Accuracy Requires Conceptual Knowledge

In the selective trust literature, a basic paradigm has been to have children gather information from two informants, measure the accuracy of that information, and then make novel inferences about information those individuals subsequently generate as a function of the relative level of accuracy. When asked to extend a novel label to a novel object based on two informants' contrary labels, preschoolers typically rely on the informant who had previously shown a history of labeling familiar objects in a conventional manner (i.e., the "accurate" informant) than an informant who generated different labels from the child (i.e., the "inaccurate" informant, see e.g., Koenig et al., 2004). Across many studies, children prefer accurate informants when learning new words (e.g., Clément et al., 2004; Corriveau, Meints, & Harris, 2009; Koenig & Harris, 2005; Scofield & Berhend, 2008), when learning artifact functions (Birch, Vauthier, & Bloom, 2008), and even when learning the rules of games (Rakoczy, Warneken, & Tomasello, 2007).

The central hypothesis of our rational account is that children's existing conceptual knowledge determines the extent to which they infer an individual's future reliability from his or her past accuracy. On the surface, this argument might seem trivial, as, by necessity, children must rely on their knowledge to evaluate whether an individual's statements are "accurate" in the sense that they match the true state of the world. Indeed, in the Koenig et al.'s (2004) procedure (and many others), children were asked "explicit judgment" questions, in which they had to identify which of the two informants provided them with right or wrong information. This intuitively posits that children both understand the question and are able to infer the individual's general accuracy (which they do, as evidenced by their above chance performance) from some information about their accuracy in a few specific instances.

Inferring reliability, however, requires another step beyond appreciating accuracy. It involves inferring the extent to which future information coming from that individual will also be accurate. A stopped clock is accurate twice a day, but if the sample from which you base future judgments of time came only at those time points, you might be misguided in the future. Judgments of reliability must integrate accuracy with other pieces of conceptual knowledge. If you know how time is measured, then you would want to collect at least three data points (or two that are not exactly 12 hr apart) to determine the clock's reliability. Moreover, appreciating the accuracy of a clock is no guarantee that it will remain accurate. A functional clock's battery could die, rendering a previously thought "accurate" source now unreliable. By the same token, judgments of the reliability of an individual require that accuracy information is integrated with other information—characteristics of the individual or of the world that are relevant to determining reliability from accuracy. For this, children must use their conceptual knowledge.

To outline the next part of our argument, in the following sections we argue for three kinds of conceptual knowledge. The first kind we will call *baseline* knowledge—cases in which existing knowledge about the informant influences how children might interpret the data she or he generates. These cases do not involve integration between accuracy and conceptual knowledge but are simply driven by conceptual knowledge itself. Next, we will focus on the role that children's existing knowledge in two domains—about *kinds* and agents' *psychological states*—play in their interpretation of others' reliability. These cases involve integrating the child's existing knowledge with the data they observe to make novel inferences and critically examine how broadly the child should generalize an informant's information.

Baseline Judgments: How Conceptual Knowledge Influences Judgments of Reliability Independent of Accuracy

Our initial evidence for the role of children's knowledge in judgments of reliability comes from baseline cases—studies in which children evaluate the reliability of testimony when information about accuracy is ambiguous or unavailable. Our first example concerns what children know about the age of the person that is providing testimony. Following findings that suggest preschool children recognize that adults know more than they do (e.g., Taylor, Cartwright, & Bowden, 1991), Jaswal and Neely (2006) showed that preschoolers, all else being equal, tend to rely on testimony from previously accurate adults over previously accurate children. Similar findings hold true when children are given misinformation—when the informant is a child, preschoolers are less likely to be swayed by the misinformation than when the informant is an adult (Lampinen & Smith, 1995).

Another example involves what children know about social groups, specifically affiliation with other members of the child's linguistic community. Infants use accent to register differences among linguistic communities and prefer native to non-native speakers. For instance, 6-month-olds look longer at (i.e., show preference for) and 10-month-olds are more likely to take toys from speakers of their native language than speakers of another language (Kinzler, Dupoux, & Spelke, 2007). By the time children reach preschool, they appear to use accent as an indicator of membership in a linguistic community, which subsequently guides inferences about reliability of testimony. For example, Kinzler, Corriveau, and Harris (2011) showed that 4–5-year-olds relied on informants who spoke with a native accent for novel functions of objects over informants who spoke with a non-native accent, even though both individuals generated 100% accurate lexical information. This suggests that children used the familiarity of a speaker's accent to make inferences about the trustworthiness of their generalized knowledge.

More generally, these data suggest that in the absence of information about informants' accuracy, children rely on existing conceptual knowledge about individuals and social groups as a basis for trust. We also see evidence that developmental changes in this conceptual knowledge lead to changes in trust. By the time children are 6, they more broadly understand that children know more than adults about a range of child-specific activities, books, and television characters (Fitneva, 2010). As expected, these changes in children's knowledge about how domains of expertise cluster by age predict their subsequent evaluations of children and adults as informants (see Fitneva, 2010, Experiment 2). Similarly, VanderBorght and Jaswal (2009) showed that reliability judgments are elaborated by preschoolers' developing understanding of kind-relevant expertise. For example, children believe that other children are more reliable sources of knowledge about toys, but adults are more reliable sources of knowledge about nutrition. These differences reflect children's developing understanding that toys and nutrition fall into different domains of knowledge, which are acquired and attended to by different informants at different rates.

How Children's Knowledge of Kinds Influences Judgments of Reliability

We have suggested that children's knowledge about individuals influences how they assess the testimony of those individuals prior to having any information about accuracy. Here we begin our discussion of how children infer reliability from information about speaker accuracy, focusing first on the role of children's existing conceptual knowledge about *kinds*. In particular, the examples in this section show that reliability judgments depend on what children know about the relation between labels and nonobvious properties of objects and what children know about people with kind-relevant (i.e., domain-specific) expertise.

Labels and Nonobvious Properties

As part of their conceptual understanding, preschoolers understand that an object's label signifies its category membership, and thus having common labels implies the existence of a host of nonobvious commonalities among category members (e.g., Gelman, 2003; Gelman & Markman, 1986, 1987; Hall & Waxman, 1993). In particular, for object categories, knowing an objects label licenses inferences about its function (Bloom, 1996; Gelman & Bloom, 2000; Gopnik & Melzoff, 1997; Kemler Nelson, 1995; Kemler Nelson, Russell, Duke, & Jones, 2000; Nelson, 1973). This understanding that labels and functions "go together" has implications for children's judgments of reliability; one of the more ubiquitous findings in the testimony literature is that children take informants' past accuracy with words as a signal that the informant is a reliable source of both new labels and new object functions (Birch et al., 2008; Koenig & Harris, 2005).

Preschoolers' also expect labels to signal the presence of nonobvious commonalities among category members. For objects, common labels often indicate common causal properties of objects (e.g., Gelman, 2003; Gopnik & Sobel, 2000; Keil, 1989; Nazzi & Gopnik, 2000). This piece of conceptual knowledge also influences children's selective judgments of reliable testimony; Sobel and Corriveau (2010) found that 4-year-olds made inferences about others' expertise with objects based on the causal knowledge they demonstrated. Children were introduced to two confederates, each of whom had different privileged knowledge about a certain kind of unobservable causal property (i.e., one knew when objects would activate a machine in a certain way; the other knew when objects would activate the machine in another way). Children were then shown new objects that possessed one of these different causal properties and were asked to endorse one of the confederates' novel labels for the objects. Children tracked the informants' differential knowledge and relied on the appropriate knowledgeable informant's label. These data suggest that not only can children track different kinds of accuracy across informants but also that they use accuracy only when it is relevant to the inference they are asked to make.

Kinds and Experts

As an extension of their understanding of kinds, preschool children realize that, even among knowledgeable informants, it is possible to be ignorant about things in one domain but an expert in another. This has been demonstrated many times by Keil and colleagues (e.g., Danovitch & Keil, 2004; Erickson, Keil, & Lockhart, 2010; Keil, 2010; Keil, Stein, Webb, Billings, & Rozenblit, 2008; Lutz & Keil, 2002). Their studies have shown that, in addition to developing conceptual knowledge about kinds, children can apply this knowledge to understanding that different people have different expertise, and that expertise generalizes in kind-relevant ways. For instance, someone who possesses mechanical knowledge of cars might be more likely to possess similar mechanical knowledge about boats as opposed to knowledge about the digestive process (on which one would suppose they would be at baseline). Within our framework, we would expect to see studies that demonstrate children rationally and selectively infer reliability of testimony from accuracy only for areas of expertise, and not beyond.

Sobel and Corriveau (2010) showed that preschoolers could distinguish between domain-relevant and domain-irrelevant expertise even when two informants were equally accurate. They introduced children to a set of objects that had nonobvious internal properties—they had "red stuff" or "green stuff" inside. Children observed that one informant was knowledgeable when an object was made of red stuff but didn't know about objects made of green stuff, while the other informant showed the opposite understanding. Four-year-olds monitored these different histories of accuracy and relied on each informant regarding different types of knowledge—when shown a novel object made of red stuff, children relied on the "red stuff expert" for its label, but when shown a novel object made of green stuff, children relied on the "green stuff expert." In contrast, when the two informants could predict whether each object had an arbitrary, nonobvious, external property (a red or green sticker on the bottom) children did not selectively trust the informants with knowledge of the objects' labels. Children used domain-relevant accuracy to infer label knowledge only when it was about internal properties of objects, thus combining their understanding of object kinds with their understanding of kind-relevant expertise (e.g., Gelman & Wellman, 1991; Sobel, Yoachim, Gopnik, Meltzoff, & Blumenthal, 2007).

Koenig and Jaswal (2011) found similar effects of relevant expertise regarding specialized knowledge of animal kinds. Preschoolers were introduced to an informant who had more knowledge about dogs than another informant (who had knowledge equivalent to the child's own knowledge). Both informants were equally accurate, but one provided more specific information than the other (i.e., the expert labeled the dog by its subordinate category label, the nonexpert labeled the dog by color). Children relied on the dog expert when asked to label novel dogs but recognized that the dog expert's knowledge was domain-specific—the dog expert was no more knowledgeable about novel object labels than the non-expert. This point deserves emphasis, because it suggests that children are not treating all labels the same way. They recognize that expertise in animal labels (particularly, possessing knowledge of which the children themselves are unaware) does not license general knowledge beyond animal labels (and possibly is limited to just other dogs).

As a final example, Kushnir, Vredenburgh, and Schneider (2013) showed that children could use differential accuracy of informants to infer different domains of expertise. They introduced preschool children to an informant who demonstrated through his actions an ability to fix broken toys, while at the same time showed through verbal testimony that the informant was ignorant about the names of the tools he used to fix the toys. They then asked how far children would generalize the fixer's reliability. In one experiment, this evidence was contrasted with an informant who failed to fix toys but provided accurate labels of the tools. Children selectively relied on the accurate fixer for fixing new toys but relied on the accurate labeler for learning new words. In another experiment, after silent demonstrations of one informant successfully fixing contrasted with another who failed to fix toys, children endorsed the fixer's over the non-fixer's verbal testimony in the causal-mechanical domain—his explanations for a set of new mechanical failures—but did not endorse his testimony about labels for novel objects.

What these examples all show is that children's trust of others' testimony depends on both the accuracy of the information they generate and on the domain in which an informant has demonstrated accuracy. This sort of selectivity would not be possible without children having some prior knowledge of the boundaries of domains of expertise. Thus, children integrate domain-specific knowledge and accuracy in much in the same way that they integrate what they know about speaker characteristics (e.g., a speaker's accent) with accuracy information. More generally, these studies provide us with some insight as to how children make judgments about others' testimony when informants show equivalent levels of prior accuracy. In some cases, children's preference for one informant matches that person's expertise. These cases rely on children observing one informant demonstrate specialized knowledge germane to the problem they face. In other cases, children show no preference between informants. In these cases, the problem at hand usually requires no specialized knowledge or experience.

The current experimental evidence provides only course measures of generalization. Thus, there are open questions about how robustly children generalize expertise. Does generalization depends on how much knowledge children have about a particular domain? Does it change as they age and as their knowledge changes? More research on these questions is necessary. Guided by our rational framework, we would suggest starting with what we know about the knowledge children have across different domains (or even directly measuring that knowledge, as we will show in later examples, to capture individual differences). Children's knowledge of kinds, actions, biological properties, and mental states might each provide a basis from which to make inferences about how far to generalize expertise. An example for which there is a sufficiently large body of empirical evidence is children's understanding of mental states. We turn to this discussion in the next section.

How Children's Knowledge of Psychological States Influences Judgments of Reliability

Just as children's conceptual knowledge about kinds influences their judgments of reliable testimony, children's conceptual knowledge about people does as well. We have already shown several examples of this in our discussion of expertise—children appreciate that different people know different things. In fact, this is one example of children's developing understanding of individual behaviors in terms of underlying psychological states—their developing "theory of mind." In this section, we discuss how several other aspects of children's developing theory of mind—namely, their understanding of the situational and epistemic influences on others' behavior—influences judgments of reliability from others' accuracy.

Situational Constraints

By the time children are 3, they appreciate that other people's perspective—and by extension their ability to know things—can be limited by situational factors. For instance, children know that individuals who cannot see the insides of a box do not know its contents (e.g., Brooks & Meltzoff, 2002; Flavell, 1988; Flavell, Everett, Croft, & Flavell, 1981; Hogrefe, Wimmer, & Perner, 1986; Meltzoff & Brooks, 2008; Pillow, 1989; Pratt & Bryant, 1990; Senju, Southgate, Snape, Leonard, & Csibra, 2011).

Recent studies have demonstrated that children's inferences about reliability from accuracy information are influenced by their knowledge of such situational constraints. For example, Nurmsoo and Robinson (2009) showed that 4- and 5-year-olds (but not 3-year-olds) would forgive past inaccuracy if the informant had been blindfolded. Similarly, Einav and Robinson (2011) found that children take into account informational access when evaluating accurate testimony. They showed children two accurate informants, one of whom was getting information from a third person who whispered words into their ear. Later, in the absence of the whisperer, 4-year-olds (but not 3-year-olds) only trusted new words coming from the independently accurate informant, showing that they could understand when informational access was available or blocked. The inferences children make in these cases involve situational factors that potentially influence whether an informant has access to particular information, which would affect whether they can generate accurate or relevant information. Children are not merely tracking accuracy as a surface feature of each informant but are making genuine inferences about the causes of each informants' behavior.

Combining Situational Constraints With Epistemic States

Between the ages of 3 and 5, children undergo a further shift in their social cognition and begin to be explicitly aware that knowledge and knowing, and by extension, ignorance and not knowing, are subjective epistemic states of individuals (see, e.g., Kuhn, Chenney, & Wainstock, 2000; Perner, 1991; Wellman & Liu, 2004). Sabbagh and Baldwin (2001) directly examined the effect of such epistemic state understanding on word learning by asking whether children learn words better from knowledgeable speakers than from ignorant ones. They showed that both 3- and 4-year-olds prefer to learn words from a knowledgeable speaker as opposed to an ignorant one. They also showed that developmental differences between ages 3- and 4-years in children's trust of knowledgeable over ignorant speakers directly corresponded to developmental changes in their understanding of cues that signal ignorance, such as hesitancy.

Furthermore, preschool children can combine this understanding of epistemic states with their understanding of situational constraints (such as blindfolds) in their selective trust. Kushnir, Wellman, and Gelman (2008) varied whether a source was knowledgeable or ignorant about a novel toy (an epistemic constraint) and also whether the source was permitted to use that knowledge in performing an action (a situational constraint). Specifically, 3- and 4-year-olds saw two puppets, and these children were told that one of them knew which blocks activated a machine and the other did not. Children then observed the puppets in three conditions: the puppets intentionally picked blocks to make the machine go, the child picked blocks and gave them to each puppet, or the puppets picked blocks while blindfolded. In all conditions, the puppets placed the blocks on the machine together, causing it to activate (play music). Then children were asked which block was more likely to be efficacious. The results showed that only when the puppets picked intentionally did preschoolers state that the knowledgeable puppet's block was more likely to be efficacious. They responded at chance in the other conditions. These data suggest that by the end of the preschool years, children integrate both situational and epistemic constraints when tracking others' accuracy.

Differentiating Among Epistemic States

There are many reasons speakers could be inaccurate. They could hold a genuine false belief. They could be guessing and generate the wrong response. They could also know they do not know and state their ignorance. A critical ability for learning from others requires children to differentiate between possible causes for being inaccurate when generating information. Preschoolers tend not to understand guessing as independent from knowledge; if one does not possess knowledge, but must respond and does so correctly, 4-year-olds (and most 5-year-olds as well) state that they knew the information (as well as were guessing, see C. N. Johnson & Wellman, 1980). More sophisticated inferences, like understanding the difference between one's own deductive inference and guessing does not emerge until well after the preschool years (Pillow, 2002).

But even if preschoolers do not have explicit understanding of guesswork, they might recognize that there are important differences between ignorance and inaccuracy (e.g., Friedman & Petrashek, 2009). If a speaker says, "I don't know" only sometimes and only when appropriate, this demonstrates an ability to monitor one's own knowledge accurately. Inaccurate statements, on the other hand, could reflect a range of psychological states, such as an inability to monitor one's own knowledge, an intent to deceive the listener, or simply strange, inexplicable behavior. Intuitively, then, ignorant speakers should be more generally trusted over inaccurate ones. Four-year-olds seem to have this intuition—when they are introduced to only one informant who is either consistently ignorant or consistently inaccurate, they will endorse new labels from the previously ignorant informant only (Kushnir, 2009).

Evidence of whether (and when) children can reason about the underlying causes of outright inaccurate testimony are mixed. For example, 6-year-olds seem to take confidence as a sign of trustworthiness and don't understand that inaccuracies stated with confidence are actually worse than those stated with some hesitancy (Tenney, Small, Kondrad, Jaswal, & Spellman, 2011). At that same age, however, children know that people sometimes fabricate testimony to reflect their own self-interest (Mills & Keil, 2005). Interesting issues remain, then, about how children can interpret behaviors as signals of underlying mental states (such as lack of knowledge or intent to deceive), while at the same time applying this understanding to decide from whom they should or should not learn.

Interim Summary

To review our argument so far, we claim that the same principles of rational inference that inform children's causal learning also inform children's social learning from others' testimony. Our first (and most central) hypothesis is that children's existing conceptual knowledge allows them to make inferences about the reliability of testimony—both in advance of and interacting with evidence of accuracy. We reviewed studies that show how children's conceptual knowledge allows them to make inferences about others' reliability in the absence of any information about accuracy. We also reviewed studies showing that children's conceptual knowledge of individuals, kinds, and psychological states plays a role in translating an understanding of an individual's past accuracy into a judgment about future reliability. We have tried to show through our review that the interaction between conceptual knowledge and accuracy information can be quite subtle. For example, conceptual knowledge about psychological states allows children to excuse past inaccuracies when they are based on a lack of epistemic access.

Critical to our argument is the interaction of knowledge with evidence (usually, evidence about informants' accuracy). In particular, knowledge that would lead children to make one inference about reliability can be overturned by accuracy evidence that leads to the opposite inference. For instance, although 4–5-year-olds treat a speaker with a native accent as more reliable than a speaker with a non-native accent when both informants generate accurate information, children treat an accurate speaker with a non-native accent as a more reliable source of knowledge than an inaccurate speaker with a native accent (Corriveau, Kinzler, & Harris, 2013). Similarly, Jaswal and Neely (2006) show that preschoolers treat an inaccurate adult as less reliable than an accurate child. Both of these examples show how children can revise their baseline judgments (based on prior knowledge) when other information (new evidence of accurate testimony), warrants a different conclusion. Together, they further support the idea that trust in testimony can be viewed as a form of rational inference.

Trusting Others When Evidence Conflicts With Testimony

So far, we have shown many examples in which children make rational inferences about the reliability of testimony based on their own prior knowledge and the strength of new evidence about the informants they observe. There are, however, some data that suggest children are overly trusting when there is direct evidence to the contrary. How do we reconcile this with our claims that children are making rational inferences?

First, it is important to remember that children's rational inferences about the reliability of testimony is best suited to environments in which children can explore the world in the presence of helpful and knowledgeable informants. Thus, the rational framework is consistent with observations that children are generally trusting of adults and that general trust often leads to them to infer epistemic (and testimonial) competence (e.g., Corriveau, Meints, & Harris, 2009; Jaswal & Neely, 2006). Through their own exploration, children can observe conflicting evidence when it arises and use it to verify the accuracy of testimony. The prediction of the rational account is that the extent to which children will be influenced by testimony will depend critically on the opportunity to observe conflicting evidence and on the strength of that evidence.

Jaswal (2010) nicely demonstrates the shift from trust to mistrust when children are faced with different kinds of conflicting evidence. Children observed a version of the "tubes task" (based on Hood, 1995), in which they watched objects roll down clear chimneys, such that they did not land directly beneath the location where they were dropped. When given verbal testimony from an observer that the object landed in a different location (specifically, in the location suggested by the "gravity bias"—directly below where it was dropped), children often instructed the experimenter to look in that location when the cups (i.e., where the objects landed) were opaque (such that children could not see the objects at rest). They ignored this misinformation when the cups were transparent (i.e., they could see the objects at rest). These data suggest that when children had direct perception of the location of the object, they simply evaluated the informant as inaccurate. When they lacked that information, they appeared swayed by that informant's testimony.

However, relevant experience can override children's reliance on testimony. In a subsequent experiment, when children were given relevant experience with the opaque cups, their reliance on the misleading information decreased over the course of the experiment. Jaswal's (2010) conclusion was that "although children in the relevant experience condition were willing to 'try out' E2's misleading testimony on the early conflict trials, they came to weight their own self-generated belief more heavily as the session went on." (p. 265). This suggests that children could use experiences they just had to make judgments about the reliability of conflicting testimony, even in a brief lab session (see also Ma and Ganea, 2010, who demonstrate a similar finding for 3-year-olds). This simple example, however, might have powerful implications for everyday contexts such as home and classroom environments; opportunities to explore on their own may lead to similar shifts in children's tendency to believe others when testimony is deceptive, inappropriate, or uninformed.

What about cases where multiple informants generate information that is inconsistent with direct perception? Following a procedure described by Asch (1956), Corriveau and Harris (2010) found that preschoolers would sometimes defer to a majority of individuals who provided them with testimony that was distinct from their direct perception (i.e., a group would say that one line was the longest of three lines in a display when a longer line was present).

There are several aspects of these results that seem quite consistent with the rational view we advocate. First, the strength of the conflicting evidence matters; using a slightly different paradigm than the Asch procedure, DiYanni, Nini, Nagrini, Kurkel, and Corriveau (2013) found that children do not faithfully imitate a questionable action performed by one informant alone, but they do when the action is performed by three different informants individually. Second, experience—and thus prior knowledge—matters. Suggestive evidence comes from several studies showing cultural differences in children's responses to conformity, for example, between Asian American and Caucasian children (Corriveau & Harris, 2010; Corriveau, Kim, Song, & Harris, in press, although the direction of the difference depends on the study, see Chan & Tardiff, 2013). Lastly, the immediate context matters. Here evidence comes from studies showing children are more conformist when asked questions in public with others watching than when asked in private with no one watching (Corriveau et al., in press; Haun & Tomasello, 2011). Together these studies suggest that conformity is not necessarily irrational and may even be the best course of action when it is low cost (though there are still open questions about what experiences and situations contribute to this tendency).

Cases Where Broadly Generalizing Reliability Is Warranted

The advantage of having (and using) conceptual knowledge for inferring reliability is that it allows for strong inductive inferences even with very little data (note, this is also an advantage in causal learning, see Schulz, Kushnir, & Gopnik, 2007). On the other hand, this can be a limitation of rational inference; a process that requires knowing *something* cannot operate without any knowledge at all. This can be seen in studies that examine children's inferences about reliability from information unrelated to accuracy. For example, children judge that a strong person will be a more accurate source of knowledge about novel object labels than a weak person (Fusaro, Corriveau, & Harris, 2011) and that a social person will be a more accurate labeler of novel objects than someone who is less social (Brosseau-Liard & Birch, 2010). These cases are relevant because children presumably do not have conceptual knowledge relating strength or prosocial behavior to kinds or labels. Children show *halo effects* in these circumstances—they respond based on valence, treating the positive (e.g., stronger) informant as more knowledgeable overall. These cases present dissimilar findings from those presented in the section on expertise above, in which children only generalize from others' accuracy judiciously. Why do children generalize selectively in some cases but broadly in others?

In our view, halo effects demonstrate what happens when the knowledge children have is not applicable to the inference at hand. Consider generalizing from strength to labels. There is research showing that preschoolers have little knowledge of strength as related to intrinsic properties of an object or person (e.g., Wenham, 2004). Without a way to bind knowledge to the inference at hand, children are faced with a choice: either make no inferences about whom to trust or resort to other (perhaps more associative or emotion-based) mechanisms, such as valence. Such an explanation is consistent with Fusaro et al.'s (2011) other finding: When shown that individuals are accurate and inaccurate labelers, children know not to generalize that accuracy information to a judgment about the strength of the informants. Children know (through their understanding of kinds) that only certain properties are related to kinds, and strength of the informant is not one of them.

The degree to which children generalize unreliability or incompetence is subject to the same analysis. Koenig and Jaswal (2011) describe cases of *pitchfork effects*, in which children treat some kinds of inaccuracy as a lack of specific knowledge and other kinds of inaccuracy as fatally indicating incompetence. If one is inaccurate about a piece of specialized knowledge (e.g., if one does not know that a Cocker Spaniel is distinct from a Water Spaniel, and just labels both as "dog"), such "inaccuracy" should be unrelated to the informant's capacity to generate other kinds of knowledge. This is because the "inaccurate" informant here is not incorrect but simply lacks specialized knowledge. In contrast, when shown the two Spaniels, if an informant *is* incorrect—if he calls one a "car" and the other a "shirt"—such inaccuracy should generalize more broadly and such an informant's knowledge should no longer be considered the same as that of a baseline informant. Koenig and Jaswal (2011) demonstrated that preschoolers make exactly these kinds of inferences.

This discussion brings up a broader question we have raised previously, which is not directly addressed by any of these investigations above: When children observe an informant generating accurate information, what can they infer about the extent of his or her knowledge? We know of at least two phenomena in the existing literature that suggest children make broad but appropriate generalizations from accuracy information. One example is linguistic proficiency. Membership in a linguistic community generally indicates proficiency with language. Children know this, so we expect children's generalizations from evidence or proficient language use to extend to all types of linguistic information.

Recent work from adult language processing (Grodner & Sedivy, 2011) suggests that evidence of speaker unreliability (simply being told that utterances come from an individual with aphasia) can lead to the suspension of real-time inferences about reference. Corriveau, Pickard, and Harris (2011) showed that 4-year-olds may also make this particular inference; they used a speaker's history of accurately generating labels or morphological structure (particularly familiar plural morphology) when making inferences about novel labels or novel morphology (specifically generating irregular past tense forms). Similarly, Sobel and Macris (2013) found that when older 4-year-olds were asked to make inferences about what novel objects were labeled, they were less likely to rely on a speaker whose lexical information was accurate but used subject-verb agreement inappropriately than a speaker who generated equivalently accurate lexical information but used proper subject-verb agreement in her utterances. Critically, even these inferences appear to be based on the child's existing knowledge. Children do not simply ignore past regularities present in language—in fact they rely more on this statistical information than speaker accuracy (Jaswal, McKercher, & VanderBorght, 2008), presumably because of the degree of difference in the strength of their beliefs about language use. Generally, these data suggest that by children might register language as a domain in which one should make broad generalizations based on the competence of an individual informant.

Another example is one that we have already discussed: Children are sensitive to differences in the knowledge that particular experts possess (e.g., Danovitch & Keil, 2004; Keil et al., 2008). That is, children recognize that different people know different things, but what a person knows coheres around domains. This is why children extend expertise in labeling to object functions (e.g., Koenig & Harris, 2005), why children recognize that an informant who knows the labels for tools does not necessarily know how to fix things with them (Kushnir et al., 2013), but also why individuals with distinct expertise about novel causal properties do not have the same knowledge (Sobel & Corriveau, 2010).

We conclude this section with another open question: Are children sensitive to the degree of difficulty necessary to acquire expertise, and can they use that difficulty when making inferences? We know of no research investigating this question, but it generally relates to what kinds of trait judgments children might make from observing informants' behaviors. It might be the case that two informants can both be experts in different domains, but because one domain might require more depth of knowledge, children generalize more expertise or knowledge more broadly from the expert in that domain. This and many other questions about generalization are left for future research.

A Comparison to Other Theories of Selective Trust

How does this account differ from other accounts of learning from selective trust? Here, we attempt to consider the accounts that have been proposed and how our rational framework compares with each of them.

Associative Learning

One possible way an individual's past history of accuracy could be generalized to inferences about them as sources of future information is through a very simple process of calculating the associative strength of that individual's accuracy. Early accounts of "trust in testimony" seem to take this position (e.g., Clément et al., 2004), and this idea appears to be the basis of some findings in generalizability of accuracy (e.g., Birch et al., 2008) and more generally, halo effects (Brosseau-Liard & Birch, 2010). On this account, children generalize from one form of accuracy to another piece of knowledge because one informant is more associated with being correct than the other in general.

Associative learning is not able to explain children's inferences about reliability, which rest on more than simply accuracy information, as we have already shown. Evidence against this account comes from many studies reviewed above, in particular those which show how children generalize from expertise selectively (e.g., Koenig & Jaswal, 2011; Sobel & Corriveau, 2010) and how children discount accurate information that is caused by situational rather than epistemic states (e.g., Einav & Robinson, 2011; Nurmsoo & Robinson, 2009).

Social Learning Theory

We do not know of a particular theory of selective trust that emanates specifically from Social Learning Theory (e.g., Bandura, 1977), but such an account could clearly be integrated with the literature on selective trust. Generally speaking, Social Learning Theory argues that people learn from observing others, and what is learned is governed by the individual's existing knowledge. These ideas are consistent with our rational account. Social Learning Theory, however, posits a strong reliance on associative learning, particularly for generalization. Thus, many of our arguments about the scope (and limits) of generalization which argue against associative learning would apply to Social Learning Theory in its current form. In order for social learning theory to account for children's selective trust, it would require some updated learning mechanism, with the most likely candidate being a rational—rather than an associative—one.

Default Bias to Trust

Gilbert (1991) suggested that adults tend to believe that new information they encounter is true. Jaswal et al. (2010) suggested that in some cases, young children also possess such a belief system early in development. In particular, children repeatedly trusted an informant who verbally informed them of the location of the object when that information was consistently incorrect but learned quickly to ignore that information if the informant used an arrow to indicate the location. They conclude from this that "a specific bias to trust testimony develops out of a generally trusting disposition that would be adaptive in infancy (e.g., Baier, 1986)" (Jaswal et al., 2010, p. 1545) and that "the specific, highly robust bias to trust testimony could emerge from the accumulation of evidence that what people say is normally true (Hume, 1748/2004)" (Jaswal et al., 2010, p. 1546).

A default bias to trust testimony would be adaptive in infancy, where children are highly reliant on adults. It would also be necessary for certain cognitive endeavors (such as language learning, see Coady, 1992; these are both points made by Jaswal et al., 2010). Jaswal et al. (2010) also argue that the bias to trust others' testimony changes with experience; initially it is strong, but perhaps later in development it weakens because children learn that informants can be inaccurate and that they themselves have experience verifying testimony through their own play and exploration of the world.

It is hard to reconcile this claim, however, with much of the evidence we have reviewed (and that is reviewed elsewhere—e.g., Harris, 2012; Mills, 2013) that even very young children are selective in trusting others. In general, across studies, there is a conflict that needs to be resolved. In the study above, 3-year-olds were overly trusting of misleading informants, but in many other studies they are not (e.g., Corriveau, Meints, & Harris, 2009; Koenig & Harris, 2005). In even younger children, we see a similar conflict: Krogh-Jespersen and Echols (2012) found 24-month-olds engaged in indiscriminate trust for a novel label regardless of accuracy, but Koenig and Woodward (2010), showed that 24-month-olds were selective and preferred accurate informants.

We see an opportunity here to resolve this conflict by examining how both the knowledge that children have and the evidence that children are presented varies between these experimental contexts. In the studies with 24-month-olds, knowledge of words as conventions might be critical. Toddlers appreciate that labels deemed accurate by both the child and informant not only need to agree, but need to agree with the conventions of the linguistic community—that is other speakers must use those labels in the same manner (Diesendruck & Shemer, 2006; Henderson & Woodward, 2012; Markson & Bloom, 1997; Moll & Tomasello, 2007). Thus, we speculate that an important difference between Krogh-Jespersen and Echols (2012) and Koenig and Woodward (2010) may be the context with which errors are presented. When only one informant is present and generates inaccurate lexical information, such errors might be treated less severely when there are more speakers present (and particularly if another speaker generates the accurate label). In Krogh-Jespersen and Echols's study, children only ever interacted with a single informant. In most of Koenig and Woodward's studies, children interact with an accurate informant and an inaccurate informant, and the inaccurate informant's errors are generated in front of all speakers. Thus, the errors made by the inaccurate informant in Koenig and Woodward's study might be considered more severe than the errors made by the informant in Krogh-Jespersen and Echols' study. Interestingly, in the one experiment reported by Koenig and Woodward in which children only interact with one informant, they find similar results to Krogh-Jespersen and Echols.

Moreover, errors involving conventional knowledge (like the meaning of words) might be treated more severely than other forms of inaccurate information (such as knowledge about specific events, like the location of a hidden object). An error in conventional knowledge is an error for any individual who follows that convention. An error in event knowledge is an error only for individuals with knowledge of that event. This might help explain various differences among Jaswal et al.'s (2010) own findings that 3-year-olds were less likely to rely on consistently misleading arrows than on similar verbal information. Because 3-year-olds hold weaker knowledge about the conventionality of arrows than the conventionality of verbal information, they might overcome this default bias more easily. Jaswal et al. (2010) also found that presenting just a disembodied voice as opposed to a person produced a weaker effect (i.e., 3-year-olds were more likely to recognize that the informant was deceptive). Even for adults, live presentation of information is more persuasive than hearing the same information in a "disembodied" form (e.g., Milgram, 1965).

In addition, we suggest a subtly different reason for *why* children interpret different kinds of errors (and more generally, different kinds of intentionally deceptive evidence) differently. Children might not be trusting by default because generally everyone tells the truth; rather, they fail to appreciate interactions in which they are being lied to. While young children engage in lying at very early ages (e.g., Newton, Reddy, & Bull, 2000; Piaget, 1932), it is only by age 4–5 that children judge one is lying based on the factuality of the statement or action (Strichartz & Burton, 1990). These authors found that younger children's (i.e., 3-year-olds) judgments of whether someone was lying were not based on the truth-value of the statement and were unsystematic. Moreover, while 4-year-olds recognized that lying had a negative valence (Taylor, Lussier, & Maring, 2003), whether younger children understand the negative connotations of lying has not been investigated (although there is evidence suggesting that 3-year-olds understand that there are negative connotations for deontic violations, see Cummins, 1996; Harris & Nunez, 1996). We return to our claim that *knowledge matters*; if children do not know they are being intentionally deceived by another person, they do not treat that individual as a liar. This is also potentially why 3-year-olds can learn that an arrow can be deceptive—they do not have to make a judgment about the arrow's intentions.

Finally, combining these ideas with evidence reviewed above that 3-year-olds *can* learn to rely on their direct perception (e.g., Jaswal, 2010; Ma & Ganea, 2010) potentially explains a difference between the results of Jaswal et al. (2010) and some recent findings. Heyman, Sritanyaratana, and Vanderbilt (2013) presented 3-year-olds with a similar task to Jaswal et al. (2010) in that children were exposed to a misleading informant in a search game. A critical difference between these procedures was that children were first shown evidence of the informant intentionally misinforming the experimenter, and before each trial, the experimenter reminded children that the informant was "tricky and mean." Unlike Jaswal et al. (2010), who consistently found that 3-year-olds mostly relied on the informant's testimony, Heyman et al. found that 3-year-olds responded at chance levels (i.e., they rejected the informant's testimony approximately half the time) and 4- and 5-year-olds were better than chance at rejecting the misleading testimony. This difference indicates that even a small amount of initial evidence showing deceptive intentions gives 3-year-olds some ability to reject misleading testimony.

Reasons and Rationales

We see our rational approach as most consistent with two recent theoretical accounts of selective trust. First, Harris and Corriveau (2011) recently argued that children's selective trust is based on an individual's actions and his or her social status in comparison to the child. We agree. At issue is how children evaluate an individual's actions and social status. We would argue that children assess the relevance of social status to trust in the same way as they do when assessing the relevance of an individual's knowledge and that this is done by integrating what they already know with whatever new data they observe.

Second, Koenig (2012) has demonstrated that children not only use others' accuracy as a basis for learning but use others' rationales for holding particular beliefs. These findings are nicely consistent with our rational account—in particular the incorporation of theory of mind into reliability judgments. If children know that reasonable beliefs are more likely to be true, they should treat individuals with good reasons for their beliefs as reliable sources of knowledge.

Pedagogy

Another theoretical account that is highly consistent with our framework is not specifically about selective trust but about how individuals learn from others more generally. Some have argued that in a learning interaction between a child and adult, children are sensitive to the pedagogical information inherent in what they communicate (Bonawitz et al., 2011; Gweon, Tenenbaum, & Schulz, 2010; see also Shafto & Goodman, 2008, for similar work with adults). The focus of this research has mostly been on the communication between teachers and children in learning contexts, where children clearly make inferences about the teacher's mental states with regards to who will be taught and how things should be taught (e.g., Strauss, Ziv, & Stein, 2002; Ziv & Frye, 2004). More generally, a theory of "natural pedagogy" might be inherent in all of human communication, and the inferences children make from what others say and do not say might be available very early in infancy (see e.g., Csibra & Gergely, 2009). The idea behind this hypothesis is that when children learn from others, they treat others as teachers, who have nonrandomly selected (i.e., chosen) what they demonstrate in order to be the most beneficial for the learner.

The hypothesis that children infer that teachers sample from a different hypothesis space of actions (or words or ideas) is rational. But where do these assumptions come from? Some (e.g., Csibra & Gergely, 2006, 2009) have suggested that these assumptions are "natural" in the sense that children access them directly from ostensible cues in communication (e.g., joint attention, child directed speech). On these accounts, sampling assumptions are the basis of trust. This may be true initially, in particular when infants are faced with social information whose truth value they cannot evaluate. In fact, this is another way of explaining children's default bias, which is compatible with their inability to detect deception—they not only assume people are telling the truth but make a stronger assumption that adults are intentionally teaching (see discussions of "strong" vs. "weak" sampling—e.g., Xu & Tenenbaum, 2007). However, we believe that, rather quickly, conceptual knowledge itself becomes the basis from which children derive sampling assumptions, rather than doing so by default.

Recent evidence supports our view. For instance, in order to make the inference that another's false belief states potentially creates a non-useful "sample" of information, children need to understand false belief. Sobel, Sommerville, Travers, Blumenthal, and Stoddard (2009) showed that only 3–4-year-olds who passed standard false belief measures used an informant's false beliefs to make causal inferences about objects' efficacy. Such evidence suggests that such inferences are not available to very young children but rather emerge as children's conceptual knowledge (in this case, about belief) emerges. Here is a case where future research is needed to attempt to link children's existing conceptual knowledge with their understanding of whether a sample generated by a teacher is useful for learning (see Rhodes, Gelman, & Brickman, 2010, as an interesting starting point).

Proposals for Future Investigations

While we have made some suggestions for future investigations throughout this article, in this penultimate section, we turn our attention to other open questions that are inspired by our discussion—places where little to no empirical results exist that potentially could shed light on the framework we have described. This is not meant to be an exhaustive list but hopefully will lead to interesting future research.

Younger Children?

In our review, we have mostly focused on preschool-age children, but our rational account suggests that conceptual knowledge plays a role in learning from others, regardless of age. So, are younger children capable of tracking others' accuracy and integrating that information with their own existing knowledge?

We would posit that the answer is "yes." Young children have statistical learning capacities at very early ages (e.g., Haith, 1993; Kirkham, Slemmer, & Johnson, 2002), which influence their interpretation of causal information (e.g., Sobel & Kirkham, 2006, 2007), social knowledge (e.g., S. C. Johnson, Slaughter, & Carey, 1998; Kaye & Fogel, 1980; Trevarthen, 1979), and linguistic data (e.g., Goldstein, Schwade, & Bornstein, 2009; Saffran, Aslin, & Newport, 1996). Moreover, statistical learning capacities in infancy are related to infants' ability to generalize (Denison, Reed, & Xu, 2013; Xu & Garcia, 2008). These learning capacities integrate with infants' existing physical (e.g., Denison & Xu, 2010) and social (e.g., Gweon & Schulz, 2011; Wu, Gopnik, Richardson, & Kirkham, 2011) knowledge. It is certainly possible that the origins of children's capacity to track others' accuracy involve these statistical learning mechanisms.

But we might expect development between the first and second birthday, as integrating statistical information with conceptual knowledge requires access to memory systems and connections among memory and reasoning systems. Young infants might be unable to do this, but by the time the child is 12–14 months old this connection is evident in their imitation, which at this point is heavily influenced by their developing social cognition (e.g., Carpenter, Akhtar, & Tomasello, 1998; Gergely, Bekkering, & Király, 2002; Schwier, van Maanen, Carpenter, & Tomasello, 2006). For example, when shown an action that produced a novel, nonobvious effect when one object, but not another was used, 15-month-olds, but not 12-month-olds could generalize the efficacy based on the differences between the configuration of the parts of the objects (Elsner & Pauen, 2007). Critically, 12-month-olds categorized these stimuli differently when their efficacy was observed (Träuble & Pauen, 2007), suggesting that infants' generalization abilities from data are developing during the second year. What is necessary is a better description of the cognitive and social processes children are using to engage in these inferences (although see Yang, Bushnell, Buchanan, & Sobel, in press, for one such example).

During the second year of life, infants understand that individuals can be reliable or unreliable sources of information in specific domains. For instance, 14-month-olds appreciate that an informant is a reliable source of information about attentional information if they look appropriately, given the attentional information inherent in the environment (Chow, Poulin-Dubois, & Lewis, 2008; Poulin-Dubois & Chow, 2009). Fourteen-month-olds also distinguish between an actor who acts in a consistent or conventional manner (puts his shoes on his feet, acts enthusiastically in responses to finding an object) from an actor who acts inconsistently or in an unconventional manner (puts his shoes on his hands or acts enthusiastically in response to an empty box) when imitating novel acts (Poulin-Dubois, Brooker, & Polonia, 2011; Zmyj, Buttelmann, Carpenter, & Daum, 2010). Selective statistical sampling (i.e., *non-random* sampling) of objects leads 18–24 month-old infants to infer that the individual has a subjective preference for those objects (Kushnir, Xu, & Wellman, 2010; Ma & Xu, 2011). Thus, by age 2, children can use the regularity in other people's choices to generalize their preferences to categories of objects (Fawcett & Markson, 2010). Tracking accuracy of desire, preference, or goal-directed intentional action might be the function of this privileged, intuitive system that also tracks statistical regularity in the physical domain. This is potentially why even though 2-year-olds can distinguish completely accurate from completely inaccurate speakers and show signs of selective learning based on past accuracy, the linguistic mappings children make at this age are extremely fragile (Koenig & Woodward, 2010).

How Much Data?

In most demonstrations of learning from others' information, children are usually shown several examples of an informant's accuracy. What is not clear is how much data children need in order to make inferences about those individuals' future information. Preschoolers do generalize inferences based on a single data point (Fitneva & Dunfield, 2010). Similarly, very young children respond to the truth value of another's individual statements (Koenig & Echols, 2003; see also Pea, 1982).

Critically, our rational account makes predictions regardless of the amount of data children observe. That is, unlike associative mechanisms, which require large amounts of data to make appropriate inferences (see e.g., Danks, 2003), a rational approach simply considers whatever data are available when making inferences. Such a mechanism, thus, can account for inferences with very little data but also inferences based on large quantities of information. For example, 5-year-olds trust familiar adults more than unfamiliar ones, even though these sources might generate equivalently accurate information in small samples (Corriveau & Harris, 2009). Similarly, 5-year-olds trust their mother differently from a stranger, mediated by the attachment relationship they have with her (Corriveau, Harris, et al., 2009). Both of these inferences are consistent with our rational account. We have much more evidence for familiar adults' accuracy than inaccuracy, and the ways in which children relate to their parents presumably are based on large amounts of evidence.

Note that this last point is slightly different from the one made by Jaswal et al. (2010) regarding a default bias to trust others. Jaswal et al. (2010) argue that children start with a relatively generic bias to trust others' information, regardless of the nature of the information or who the informant is. What they believe is learned is that others' verbal information is almost always accurate, so children should develop a bias to trust verbal information. This argument says nothing directly about the familiarity of the speaker. The argument we are making here is that children have much more data from familiar than unfamiliar informants (and presumably from their parents or caregivers more than anyone else) and given the extent of those data, children might treat any single accurate or inaccurate event from a familiar informant differently than such an event from an unfamiliar one.

Critically, when presented with very little data, one must consider what mechanisms children have for generalization from small sample sizes. This question has been considered in the causal reasoning literature. When presented with deterministic evidence, very young children make inferences from small samples, but when presented with probabilistic evidence, it is not until age 4–5 that children reliably generalize based on differences in probability (e.g., Kushnir & Gopnik, 2005; Sobel et al., 2009). When presented with accuracy evidence, children might be using a similar mechanism to track an individual's accuracy. Such a hypothesis might explain why Pasquini et al. (2007) demonstrated that 3-year-olds (but not older children) struggle with informants who differ in their degree of accuracy.

Pasquini et al. (2007) did find one case in which 3-year-olds appeared to use probabilistic data—if one informant is always accurate (100% vs. 25%). Where 3-year-olds struggle is when both informants erred, just at different rates (75% vs. 0% or 25% accurate). They explain this asymmetry by suggesting that young children believe that once an informant generates inaccurate information, she or he is unreliable. A concern with this conclusion is that in everyday life, even the most trustworthy and well-intentioned adults that young children encounter are not 100% accurate. Adults mislabel objects, misremember information or goals, and act in ways that are clearly misinformed. We (like all parents) sincerely hope our own 3-year-olds do not end up treating us as unreliable sources of knowledge forevermore when we make such errors!

Our rational hypothesis posits an alternative explanation to this asymmetry. There is evidence that children (Vaish, Grossman, & Woodward, 2008) and adults (Fiske, 1980; Kahneman & Tversky, 1984) sometimes make broad generalizations of unreliability from even a little bit of inaccuracy—that is, they might show a "negativity bias": Negative evidence might be weighed more heavily than positive evidence. When presented with probabilistic data, if children (i.e., 4-year-olds and older, Sobel et al., 2009) can reason about such data, they do so. When they cannot (i.e., when they lack knowledge to assess probabilistic evidence), they rely on other factors, like a negativity bias, to make inferences. When one informant is 100% accurate, that informant has never presented negative evidence, so she or he is more reliable than the one who has. When both err, but at different rates, children might succumb to this bias for both informants and have no basis for choosing between the informants as they believe both are unreliable.

Finally, investigating how much data children need to encounter before they change their beliefs about the reliability of an informant or about a piece of knowledge might shed insight into the degree of difficulty children have learning various scientific principles. For instance, Shtulman and Schulz (2008) found that understanding within-species variation was key for adults to appreciate a selection-based understanding of evolution. The problem is that understanding within-species variation is contrary to essentialist beliefs about natural kinds, which children possess at very young ages (e.g., Gelman, 2003), and robustly persist throughout development (e.g., Keil, 1989; Newman & Keil, 2008). Potentially, one reason why it is so difficult to learn a concept like natural selection is that the existing knowledge one has about evolution contradicts what others tell us about how evolution works.

A Processing Account of Gauging Reliability

We have argued that similar learning processes are recruited for both causal and social learning. To support this idea, studies that include more focus on processing (like memory or attention, e.g., Sabbagh & Shafman, 2009; Sabbagh & Henderson, 2013b) are needed. In the causal reasoning literature, many researchers who advocate a rational hypothesis emphatically state that they work on the "computational" level of description (cf. Marr, 1982)—that is, they describe what the system is doing. The rational account we have provided here is potentially at the same descriptive level of analysis. A focus on processing—the "algorithmic" level, in Marr's (1982) terms—would bring us one step closer to understanding how learning actually occurs.

One way of addressing this is to use methods that tap children's online processing of others' reliability in real time. For instance, Sobel, Sedivy, Buchanan, and Hennessey (2012) presented preschoolers with a live speaker who accurately or inaccurately generated labels for familiar objects. Children were then shown two objects and informed that a third object was hidden in a secret location, visible to the speaker but not to the child. On some trials, both of the objects in common ground were familiar to the child; on others, one was familiar and the other was novel. When a speaker with a history of accurately labeling familiar objects asked for an object with a novel label, children often gave them the novel object when it was visible to them and looked in the secret location when only familiar objects were visible to them. When an inaccurate speaker made this request, children were more likely to give familiar objects in both cases. Critically, children's patterns of eye gaze to the array paralleled their inferences: Children looked at the secret location first more often when a reliable speaker asked for an object with a novel label when only familiar objects were visible than when a novel object was visible or when an unreliable speaker made the same request. This example suggests that children's selective trust is not just reflected in their inferences, but the cognitive mechanism responsible for those inferences potentially affects their online language processing. While this investigation only makes minor inroads toward answering the processing question, eye-tracking might be a good method to investigate how children process information from others over time (see e.g., Sabbagh & Henderson, 2013a, for a promising new line of work using this method).

A second way of considering this question is to examine computational models of selective trust. There are currently several computational models of selective trust, some of which have rational components that describe children's inferences on a computational level (e.g., Butterfield, Jenkins, Sobel, & Schwertfeger, 2009; Eaves & Shafto, 2012; Shafto, Eaves, Navarro, & Perfors, 2012; Yang et al., in press). These models nicely explain existing findings, including many that we have described here, although we do believe that no computational model explains all of the empirical data. For instance, the model articulated by Butterfield et al. (2009) explains why children generalize the novel labels of a previously seen accurate speaker over a previously seen inaccurate speaker (e.g., Koenig & Harris, 2005) but cannot explain why children rely on certain kinds of expertise and not others, even though accuracy is held constant (e.g., Koenig & Jaswal, 2011; Sobel & Corriveau, 2010). This limitation was pointed out by Sobel, Buchanan, Butterfield, and Jenkins (2010), who suggested ways in which the computational framework could be expanded (some of these suggestions are instantiated in the modeling work of Shafto and colleagues). Critically, the primary usage of such models is to motivate direct empirical tests of different assumptions about the processing of social information. Our approach here is not to provide the formal details of these models but rather to describe the importance of a rational framework to psychological theories of selective trust.

Finally, we would like to see more of an emphasis placed on examining individual differences in children's existing knowledge as a predictor of performance in measures of selective trust. What information children trust will change as their knowledge develops. Our approach predicts that when new knowledge is acquired (e.g., via the output of social or causal learning processes), it can be used to evaluate new data (social, causal or otherwise) or revise existing beliefs about others' reliability.

As an example discussed earlier, Einav and Robinson (2011) showed that 3-year-olds treated an accurate informant who was fed information by a third party as equally accurate to an informant who generated that information independently, while older children relied more on the independent informant. Such development might not be surprising. Four-year-olds have a better understanding of information and perceptual access (e.g., Wellman, Cross, & Watson, 2001; Wellman & Liu, 2004), and this specific understanding presumably underlies children's ability to distinguish an informed informant from an ignorant one who is being fed the right answers by a more knowledgeable individuals. We would predict that as children's understanding of perceptual access develops, that comprehension, and not age, would predict how children treat these informants.

Similarly, there are now a number of findings demonstrating individual differences in children's reliance on others' testimony based on their existing knowledge about the inference at hand (e.g., Chan & Tardiff, 2013; E. E. Chen, Corriveau, & Harris, 2013; Lane, Harris, Gelman, & Wellman, in press; Robinson, Einav, & Fox, 2013; Van Reet, Green, & Sobel, 2013). For instance, Lane et al. (in press) showed that preschoolers' performance on standard appearance-reality tests predicted how they would interpret testimony about the reality status of objects that did not look like what they were (i.e., other stimuli that could be used to measure Appearance-Reality). Children who succeeded on Appearance-Reality measures were willing to trust testimony that did not conform with the appearance of the objects, while children who did not succeed were less credulous.

We believe that these kinds of studies offer perhaps the strongest evidence for the rational approach we articulated being used in everyday learning. After all, the general trust in testimony paradigm presents an artificial setting—rarely outside of a lab do children observe two informants provide different information. However, children often observe informants generate information that is consistent or inconsistent with their own beliefs or knowledge. Demonstrating that that knowledge affects how children learn from those informants is critical to describing the way in which children are making these online judgments.

Concluding Remarks

We want to conclude by emphasizing that the parallels—and the direct links—between the child as an active causal learner and the child as an active seeker of social information (Baldwin & Moses, 1996). For instance, in many theories of causal learning, explanation is thought to be critical for learning causal structure (e.g., Gopnik, 1998; Koslowski, 1996). This may be equivalent to the influence of explanation-seeking as part of social learning. Toddlers in a "naming spurt" often generate a lexical item soliciting an object's name (Nelson, 1973). Children begin to generate "why" questions around the time they themselves offer causal explanations (Hickling & Wellman, 2001; Hood & Bloom, 1979; Wellman & Liu, 2007). More generally, children's questions appear based on their existing knowledge (Chouinard, 2007), and children use their existing conceptual knowledge to learn from different experts by asking them different kinds of questions (Mills et al., 2010, 2011). Further, children explore through play what they seek to explain (Legare, 2012) or ask for more information when they are not getting good explanations from adults (Frazier et al., 2009). Our framework lends itself to further integrating active social learning with active causal learning through more systematic observations of the trade-offs children make between exploration on their own and relying on the knowledge of others (Gweon & Schulz, 2011; Vredenburgh & Kushnir, 2013).

These ideas are consistent with recent anthropological investigations, which have considered how human societies evolved. Given that few individuals have all the knowledge needed to survive in a single environment (let alone in a set of diverse environments), Boyd, Richerson, and Henrich (2011) argued that any one individual in a community must rely on others as sources of knowledge. They use historical evidence to show the presence of division of cognitive labor across time and culture. Their conclusion is that human societies evolved not only because of species-unique causal reasoning capacities, but more so due to our human ability to learn from others. Our proximate explanation of children as rational learners is logically consistent with these ultimate claims (Fedyk, in press). Together, they provide a coherent explanation of both the mechanisms and ends of human social learning.

Thus, rather than viewing a trade-off between causal and social learning, each with its own advantages and disadvantages (e.g., Bonawitz et al., 2011; Z. Chen & Klahr, 1999; Sobel & Sommerville, 2010), we conclude by suggesting that both are part of the same active learning process. *All* of children's learning involves attempting to ascertain the relevance and informativeness of evidence in a rational way such that they fill gaps in their existing knowledge to learn more about the world.

References

Asch, S. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychology Monographs: General and Applied, 70*(9, Serial no. 416).

Baier, A. C. (1986). Trust and antitrust. *Ethics, 96,* 231–260. doi:10.1086/292745

Baldwin, D. A. (1991). Infants' contribution to the achievement of joint reference. *Child Development, 62,* 875–890. doi:10.2307/1131140

Baldwin, D. A. (1993). Early referential understanding: Infants' ability to understand referential acts for what they are. *Developmental Psychology, 29,* 832–843. doi:10.1037/0012-1649.29.5.832

Baldwin, D. A., & Moses, L. J. (1996). The ontogeny of social information gathering. *Child Development, 67,* 1915–1939. doi:10.2307/1131601

Bandura, A. (1977). *Social learning theory*. Oxford, England: Prentice-Hall.

Bergstrom, B., Moehlmann, B., & Boyer, P. (2006). Extending the testimony problem: Evaluating the truth, scope, and source of cultural information. *Child Development, 77,* 531–538. doi:10.1111/j.1467-8624.2006.00888.x

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,* 1018–1034. doi:10.1016/j.cognition.2007.12.008

Bloom, P. (1996). Intention, history, and artifact concepts. *Cognition, 60,* 1–29. doi:10.1016/0010-0277(95)00699-0

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge, MA. MIT Press.

Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition, 120,* 322–330. doi:10.1016/j.cognition.2010.10.001

Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Science of the United States of America, 108,* 10918–10925. doi:10.1073/pnas.1100290108

Brooks, R., & Meltzoff, A. N. (2002). The importance of eyes: How infants interpret adult looking behavior. *Developmental Psychology, 38,* 958–966. doi:10.1037/0012-1649.38.6.958

Brosseau-Liard, P. E., & Birch, S. J. (2010). "I bet you know more and are nicer too!" What children infer from others' accuracy. *Developmental Science, 13,* 772–778. doi:10.1111/j.1467-7687.2009.00932.x

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York, NY: Academic Press.

Butterfield, J., Jenkins, O. C., Sobel, D. M., & Schwertfeger, J. (2009). Modeling aspects of theory of mind with Markov Random Fields. *International Journal of Social Robotics, 1,* 41–51. doi:10.1007/s12369-008-0003-1

Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development, 7,* 213–233. doi:10.1016/0885-2014(92)90012-G

Carey, S. (1995). On the origin of causal understanding. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 268–3082). New York, NY: Oxford University Press.

Carpenter, M., Akhtar, N., & Tomasello, M. (1998). Fourteen- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior & Development, 21,* 315–330. doi:10.1016/S0163-6383(98)90009-1

Chan, C. C. Y., & Tardiff, T. (2013). Knowing better: The role of prior knowledge and culture in trust in testimony. *Developmental Psychology, 49,* 591–601. doi:10.1037/a0031336

Chen, E. E., Corriveau, K. H., & Harris, P. L. (2013). Children lose trust in a consensus composed of outgroup members—but do not retain that trust. *Child Development, 84,* 269–282. doi:10.1111/j.1467-8624.2012.01850.x

Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development, 70,* 1098–1120. doi:10.1111/1467-8624.00081

Chouinard, M. M. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development, 72,* 1–108.

Chow, V., Poulin-Dubois, D., & Lewis, J. (2008). To see or not to see: Infants prefer to follow the gaze of a reliable looker. *Developmental Science, 11,* 761–770. doi:10.1111/j.1467-7687.2008.00726.x

Clément, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language, 19,* 360–379. doi:10.1111/j.0268-1064.2004.00263.x

Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford, England: Clarendon Press.

Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition, 120,* 341–349. doi:10.1016/j.cognition.2011.03.003

Corriveau, K., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and past accuracy. *Developmental Science, 12,* 426–437. doi:10.1111/j.1467-7687.2008.00792.x

Corriveau, K. H., & Harris, P. L. (2010). Young children's trust in what other people say. In K. Rotenberg (Ed.), *Interpersonal trust during childhood and adolescence* (pp. 87–109). Cambridge, United Kingdom: Cambridge University Press. doi:10.1017/CBO9780511750946.005

Corriveau, K. H., Harris, P. L., Meins, E., Ferneyhough, C., Arnott, B., Elliott, L., . . . de Rosnay, M. (2009). Young children's trust in their mother's claims: Longitudinal links with attachment security in infancy. *Child Development, 80,* 750–761. doi:10.1111/j.1467-8624.2009.01295.x

Corriveau, K. H., Kim, E., Song, G., & Harris, P. L. (in press). Young children's deference to a consensus varies by culture and judgment setting. *Journal of Cognition and Culture*.

Corriveau, K. H., Kinzler, K. D., & Harris, P. L. (2013). Accuracy trumps accent in children's endorsement of object labels. *Developmental Psychology, 49,* 470–479. doi:10.1037/a0030604

Corriveau, K. H., Meints, K., & Harris, P. L. (2009). Early tracking of informant accuracy and inaccuracy by young children. *British Journal of Developmental Psychology, 27,* 331–342. doi:10.1348/026151008X310229

Corriveau, K. H., Pickard, K., & Harris, P. L. (2011). Preschoolers trust particular informants when learning new names and new morphological forms. *British Journal of Developmental Psychology, 29,* 46–63. doi:10.1348/2044-835X.002009

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 xxi* (pp. 249–274). Oxford, England: Oxford University Press.

Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences, 13,* 148–153. doi:10.1016/j.tics.2009.01.005

Cummins, D. D. (1996). Evidence of deontic reasoning in 3- and 4-year-old children. *Memory & Cognition, 24,* 823–829. doi:10.3758/BF03201105

Danks, D. (2003). Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology, 47,* 109–121. doi:10.1016/S0022-2496(02)00016-0

Danovitch, J. H., & Keil, F. C. (2004). Should you ask a fisherman or a biologist? Developmental shifts in ways of clustering knowledge. *Child Development, 75,* 918–931. doi:10.1111/j.1467-8624.2004.00714.x

Denison, S., Reed, C., & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: Evidence from 4.5- and 6-month-olds. *Developmental Psychology, 49,* 243–249.

Denison, S., & Xu, F. (2010). Integrating physical constraints in statistical inference by 11-month-old infants. *Cognitive Science, 34,* 885–908. doi:10.1111/j.1551-6709.2010.01111.x

de Saussure, F. (1966). *Course in general linguistics*. New York, NY: McGraw-Hill.

Diesendruck, G., & Shemer, G. (2006). Young children's expectation of competence in word learning. *Journal of Child Language, 33,* 321–328.

diSessa, A. (1993). Towards an epistemology of physics. *Cognition and Instruction, 10,* 105–225. doi:10.1080/07370008.1985.9649008

DiYanni, C., Nini, N., Nagrini, N., Kurkel, J., & Corriveau, K. H. (2013). *The role of consensus and culture in children's imitation of questionable actions*. Manuscript submitted for publication.

Eaves, B., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in Child Development and Behavior, 43,* 295–319. doi:10.1016/B978-0-12-397919-3.00011-3

Einav, S., & Robinson, E. J. (2011). When being right is not enough: Four-year-olds distinguish knowledgeable from merely accurate informants. *Psychological Science, 22,* 1250–1253. doi:10.1177/0956797611416998

Elsner, B., & Pauen, S. (2007). Social learning of artifact function in 12- and 15-month-olds. *European Journal of Developmental Psychology, 4,* 80–99. doi:10.1080/17405620601051220

Erickson, J. E., Keil, F. C., & Lockhart, K. L. (2010). Sensing the coherence of biology in contrast to psychology: Young children's use of causal relations to distinguish two foundational domains. *Child Development, 81,* 390–409. doi:10.1111/j.1467-8624.2009.01402.x

Faucher, L., Mallon, R., Nazer, D., Nichols, S., Ruby, A., Stich, S., & Weisberg, J. (2002). The baby in the lab-coat: Why child development is not an adequate model for understanding the development of science. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis of science* (pp. 335–362). Cambridge, United Kingdom: Cambridge University Press. doi:10.1017/CBO9780511613517.019

Fawcett, C. A., & Markson, L. (2010). Children reason about shared preferences. *Developmental Psychology, 46,* 299–309. doi:10.1037/a0018539

Fedyk, M. (in press). How (not) to bring biology and psychology together. *Philosophical Studies*.

Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology, 38,* 889–906. doi:10.1037/0022-3514.38.6.889

Fiske, S. T. (2010). *Social beings: Core motives in social psychology* (2nd ed.). New York, NY: Wiley.

Fitneva, S. A. (2010). Children's representation of child and adult knowledge. *Journal of Cognition and Development, 11,* 458–484. doi:10.1080/15248371003700023

Fitneva, S. A., & Dunfield, K. A. (2010). Selective information seeking after a single encounter. *Developmental Psychology, 46,* 1380–1384. doi:10.1037/a0019818

Flavell, J. H. (1988). The development of children's knowledge about the mind: From cognitive connections to mental representations. In J. W. Astington, P. L. Harris, & D. R. Olson (Eds.), *Developing theories of mind* (pp. 244–267). New York, NY: Cambridge University Press.

Flavell, J. H., Everett, B. A., Croft, K., & Flavell, E. R. (1981). Young children's knowledge about visual perception: Further evidence for the Level 1–Level 2 distinction. *Developmental Psychology, 17,* 99–103. doi:10.1037/0012-1649.17.1.99

Frazier, B. N., Gelman, S. A., & Wellman, H. M. (2009). Preschoolers' search for explanatory information within adult-child conversation. *Child Development, 80,* 1592–1611. doi:10.1111/j.1467-8624.2009.01356.x

Friedman, O., & Petrashek, A. R. (2009). Children do not follow the rule "ignorance means getting it wrong". *Journal of Experimental Child Psychology, 102,* 114–121. doi:10.1016/j.jecp.2008.07.009

Fusaro, M., Corriveau, K. H., & Harris, P. L. (2011). The good, the strong, and the accurate. Preschoolers' evaluations of accurate and strong informants. *Journal of Experimental Child Psychology, 110,* 561–574. doi:10.1016/j.jecp.2011.06.008

Gauvain, M., & Rogoff, B. (1989). Collaborative problem solving and children's planning skills. *Developmental Psychology, 25,* 139–151. doi:10.1037/0012-1649.25.1.139

Gelman, S. A. (2003). *The essential child*. New York, NY: Oxford University Press.

Gelman, S. A., & Bloom, P. (2000). Young children are sensitive to how an object was created when deciding what to name it. *Cognition, 76,* 91–103. doi:10.1016/S0010-0277(00)00071-8

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition, 23,* 183–209. doi:10.1016/0010-0277(86)90034-X

Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development, 58,* 1532–1541.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understanding of the non-obvious. *Cognition, 38,* 213–244. doi:10.1016/0010-0277(91)90007-Q

Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature, 415,* 755–756. doi:10.1038/415755a

Gilbert, D. T. (1991). How mental systems behave. *American Psychologist, 46,* 107–119. doi:10.1037/0003-066X.46.2.107

Goldstein, M. H., Schwade, J. A., & Bornstein, M. H. (2009). The value of vocalizing: Five-month-old infants associate their own noncry vocalizations with responses from adults. *Child Development, 80,* 636–644. doi:10.1111/j.1467-8624.2009.01287.x

Gopnik, A. (1998). Explanation as orgasm. *Minds and Machines, 8,* 101–118. doi:10.1023/A:1008290415597

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,* 3–32. doi:10.1037/0033-295X.111.1.3

Gopnik, A., & Melzoff, A. N. (1997). *Words, thoughts, and theories*. Cambridge, MA: MIT Press.

Gopnik, A., & Schulz, L. (2007). *Causal learning: Psychology, philosophy, and computation*. New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780195176803.001.0001

Gopnik, A., & Sobel, D. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Development, 71,* 5, 1205–1222. doi:10.1111/1467-8624.00224

Gopnik, A., & Wellman, H. M. (1994). The theory theory. In L. Hirschfield & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 257–293). New York, NY: Cambridge University Press. doi:10.1017/CBO9780511752902.011

Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological Bulletin, 138,* 1085–1108. doi:10.1037/a0028044

Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive Science, 35,* 1407–1455. doi:10.1111/j.1551-6709.2011.01203.x

Grodner, D., & Sedivy, J. (2011). The effects of speaker-specific information on pragmatic inferences. In N. Pearlmutter & E. Gibson (Eds.), *The processing and acquisition of reference* (pp. 239–272). Cambridge, MA: MIT Press.

Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science, 332,* 1524. doi:10.1126/science.1204493

Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences of the United States of America, 107,* 9066–9071. doi:10.1073/pnas.1003095107

Haith, M. M. (1993). Future-oriented processes in infancy: The case of visual expectations. In C. Granrud (Ed.), *Visual perception and cognition in infancy* (pp. 235–264). Hillsdale, NJ: Erlbaum.

Hall, D. G., & Waxman, S. R. (1993). Assumptions about word meaning: Individualization and basic-level kinds. *Child Development, 64,* 1550–1570. doi:10.2307/1131552

Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge, MA: Harvard University Press. doi:10.4159/harvard.9780674065192

Harris, P. L., & Corriveau, K. H. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society: B. Biological Sciences, 366,* 1179–1187. doi:10.1098/rstb.2010.0321

Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child Development, 77,* 505–524. doi:10.1111/j.1467-8624.2006.00886.x

Harris, P. L., & Nunez, M. (1996). Understanding of permission rules by preschool children. *Child Development, 67,* 1572–1591. doi:10.2307/1131719

Harris, P. L., Pasquini, E. S., Duke, S., Asscher, J. J., & Pons, F. (2006). Germs and angels: The role of testimony in young children's ontology. *Developmental Science, 9,* 76–96. doi:10.1111/j.1467-7687.2005.00465.x

Haun, D. B., & Tomasello, M. (2011). Conformity to peer pressure in preschool children. *Child Development, 82,* 1759–1767. doi:10.1111/j.1467-8624.2011.01666.x

Henderson, A. M. E., & Woodward, A. (2012). Nine-month old infants generalize object labels, but not object preferences across individuals. *Developmental Science, 15,* 641–652. doi:10.1111/j.1467-7687.2012.01157.x

Heyman, G. D., Sritanyaratana, L., & Vanderbilt, K. E. (2013). Young children's trust in overtly misleading advice. *Cognitive Science, 37,* 646–667. doi:10.1111/cogs.12020

Hickling, A., & Wellman, H. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology, 37,* 668–683. doi:10.1037/0012-1649.37.5.668

Hogrefe, G. J., Wimmer, H., & Perner, J. (1986). Ignorance versus false belief: A developmental lag in attribution of epistemic states. *Child Development, 57,* 567–582. doi:10.2307/1130337

Hood, B. M. (1995). Gravity rules for 2–4-year-olds? *Cognitive Development, 10,* 577–598. doi:10.1016/0885-2014(95)90027-6

Hood, L., & Bloom, L. (1979). What, when, and how about why: A longitudinal study of early expressions of causality. *Monographs of the Society for Research in Child Development, 44* (6, Serial No. 181).

Hume, D. (1748/2004). *An enquiry concerning human understanding*. New York, NY: Dover.

Jaswal, V. K. (2010). Believing what you're told: Young children's trust in unexpected testimony about the physical world. *Cognitive Psychology, 61,* 248–272. doi:10.1016/j.cogpsych.2010.06.002

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,* 1541–1547. doi:10.1177/0956797610383438

Jaswal, V. K., McKercher, D. A., & VanderBorght, M. (2008). Limitations on reliability: Regularity rules in the English plural and past tense. *Child Development, 79,* 750–760. doi:10.1111/j.1467-8624.2008.01155.x

Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science, 17,* 757–758. doi:10.1111/j.1467-9280.2006.01778.x

Johnson, C. N., & Wellman, H. M. (1980). Children's developing understanding of mental verbs: Remember, know, and guess. *Child Development, 51,* 1095–1102. doi:10.2307/1129549

Johnson, S. C., Slaughter, V., & Carey, S. (1998). Whose gaze will infants follow? The elicitation of gaze-following in 12-month-olds. *Developmental Science, 1,* 233–238. doi:10.1111/1467-7687.00036

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist, 39,* 341–350. doi:10.1037/0003-066X.39.4.341

Kaye, K., & Fogel, A. (1980). The temporal structure of face-to-face communication between mothers and infants. *Developmental Psychology, 16,* 454–464. doi:10.1037/0012-1649.16.5.454

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Keil, F. C. (2010). The feasibility of folk science. *Cognitive Science, 34,* 826–862. doi:10.1111/j.1551-6709.2010.01108.x

Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science, 32,* 259–300. doi:10.1080/03640210701863339

Kemler Nelson, D. G. (1995). Principle-based inferences in young children's categorization: Revisiting the impact of function on the naming of artifacts. *Cognitive Development, 10,* 347–380. doi:10.1016/0885-2014(95)90002-0

Kemler Nelson, D. G., Russell, R., Duke, N., & Jones, K. (2000). Two-year-olds will name artifacts by their functions. *Child Development, 71,* 1271–1288. doi:10.1111/1467-8624.00228

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native accented speakers. *Developmental Science, 14,* 106–111. doi:10.1111/j.1467-7687.2010.00965.x

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,* 12577–12580. doi:10.1073/pnas.0705345104

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence of a domain general learning mechanism. *Cognition, 83,* B35–B42. doi:10.1016/S0010-0277(02)00004-5

Klahr, D., & Nigam, M. (2004). The equivalence of learning paths in early science instruction: Effects of direct instruction and discovery learning. *Psychological Science, 15,* 661–667. doi:10.1111/j.0956-7976.2004.00737.x

Koenig, M. A. (2012). Beyond semantic accuracy: Preschoolers evaluate a speaker's reasons. *Child Development, 83,* 1051–1063. doi:10.1111/j.1467-8624.2012.01742.x

Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science, 15,* 694–698. doi:10.1111/j.0956-7976.2004.00742.x

Koenig, M. A., & Echols, C. H. (2003). Infants' understanding of false labeling events: The referential role of words and the people who use them. *Cognition, 87,* 179–206. doi:10.1016/S0010-0277(03)00002-7

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development, 76,* 1261–1277. doi:10.1111/j.1467-8624.2005.00849.x

Koenig, M. A., & Harris, P. L. (2008). The basis of children's trust: Reliable testimony or reliable sources? *Episteme, 4,* 264–284.

Koenig, M. A., & Jaswal, V. K. (2011). Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Development, 82,* 1634–1647. doi:10.1111/j.1467-8624.2011.01618.x

Koenig, M., & Stephens, E. (in press). Characterizing children's responsiveness to cues of speaker trustworthiness: Two proposals. In L. Robinson & S. Einav (Eds.), *Current issues in developmental psychology: Children's trust in testimony*. Oxford, England: Psychology Press.

Koenig, M. A., & Woodward, A. L. (2010). 24-month-olds' sensitivity to the prior accuracy of the source: Possible mechanisms. *Developmental Psychology, 46,* 815–826. doi:10.1037/a0019664

Koslowski, B. M. (1996). *Theory and evidence: The development of scientific reasoning*. Cambridge, MA: MIT Press.

Krogh-Jespersen, S., & Echols, C. E. (2012). The influence of speaker reliability on first versus second label learning. *Child Development, 83,* 581–590.

Kuhn, D., Chenney, R., & Wainstock, M. (2000). The development of epistemic understanding. *Cognitive Development, 15,* 309–328. doi:10.1016/S0885-2014(00)00030-7

Kushnir, T. (2009, October) *Preschooler's trust of others in causal learning*. Paper presented at the 2009 Biennial Meeting of the Cognitive Development Society, San Antonio, TX.

Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science, 16,* 678–683. doi:10.1111/j.1467-9280.2005.01595.x

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,* 186–196. doi:10.1037/0012-1649.43.1.186

Kushnir, T., Vredenburgh, C., & Schneider, L. A. (2013). "Who can help me fix this toy?" The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology, 49,* 446–453. doi:10.1037/a0031649

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition, 107,* 1084–1092. doi:10.1016/j.cognition.2007.10.004

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in children's causal inferences. *Developmental Psychology, 45,* 597–603. doi:10.1037/a0014727

Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science, 21,* 1134–1140. doi:10.1177/0956797610376652

Lampinen, J. M., & Smith, V. L. (1995). The incredible (and sometimes incredulous) child witness: Child eyewitnesses' sensitivity to source credibility cues. *Journal of Applied Psychology, 80,* 621–627. doi:10.1037/0021-9010.80.5.621

Lane, J. D., Harris, P. L., Gelman, S. A., & Wellman, H. M. (in press). More than meets the eye: Young children's trust in claims that defy their perceptions. *Developmental Psychology*.

Legare, C. H. (2012). Exploring explanation: Explaining inconsistent information guides hypothesis-testing behavior in young children. *Child Development, 83,* 173–185. doi:10.1111/j.1467-8624.2011.01691.x

Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development, 73,* 1073–1084. doi:10.1111/1467-8624.00458

Ma, L., & Ganea, P. A. (2010). Dealing with conflicting information: Young children's reliance on what they see versus what they are told. *Developmental Science, 13,* 151–160. doi:10.1111/j.1467-7687.2009.00878.x

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

Markson, L., & Bloom, P. (1997). Evidence against a dedicated system for word learning in children. *Nature, 385,* 813–815. doi:10.1038/385813a0

Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.

Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition, 112,* 367–380. doi:10.1016/j.cognition.2009.05.012

Meltzoff, A. N., & Brooks, R. (2008). Self-experience as a mechanism for learning about others: A training study in social cognition. *Developmental Psychology, 44,* 1257–1265. doi:10.1037/a0012888

Milgram, S. (1965). Some conditions of obedience and disobedience to authority. *Human Relations, 18,* 57–76. doi:10.1177/001872676501800105

Mills, C. M. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology, 49,* 404–418. doi:10.1037/a0029500

Mills, C. M., & Keil, F. C. (2005). The development of cynicism. *Psychological Science, 16,* 385–390. doi:10.1111/j.0956-7976.2005.01545.x

Mills, C. M., Legare, C. H., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development, 11,* 533–560. doi:10.1080/15248372.2010.516419

Mills, C. M., Legare, C. H., Grant, M., & Landrum, A. R. (2011). Determining whom to question, what to ask, and how much information to ask for: The development of inquiry in young children. *Journal of Experimental Child Psychology, 110,* 539–560. doi:10.1016/j.jecp.2011.06.003

Moll, H., & Tomasello, M. (2007). How 14- and 18-month-olds know what others have experienced. *Developmental Psychology, 43,* 309–317. doi:10.1037/0012-1649.43.2.309

Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science, 3,* 389–396. doi:10.1111/1467-7687.00133

Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development, 38*(1–2, Serial No. 149).

Newman, G. E., & Keil, F. C. (2008). Where is the essence? Developmental shifts in children's beliefs about internal features. *Child Development, 79,* 1344–1356. doi:10.1111/j.1467-8624.2008.01192.x

Newton, P., Reddy, V., & Bull, R. (2000). Children's everyday deception and performance on false-belief tasks. *British Journal of Developmental Psychology, 18,* 297–317. doi:10.1348/026151000165706

Nurmsoo, E., & Robinson, E. J. (2009). Identifying unreliable informants: Do children excuse past inaccuracy? *Developmental Science, 12,* 41–47. doi:10.1111/j.1467-7687.2008.00750.x

Pasquini, E. S., Corriveau, K., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology, 43,* 1216–1226. doi:10.1037/0012-1649.43.5.1216

Pea, R. D. (1982). Origins of verbal logic: Spontaneous denials by two- and three-year-olds. *Journal of Child Language, 9,* 597–626. doi:10.1017/S0305000900004931

Perner, J. (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press.

Piaget, J. (1932). *The moral judgment of the child*. London, England: Routledge & Kegan-Paul.

Pillow, B. H. (1989). Early understanding of perception as a source of knowledge. *Journal of Experimental Child Psychology, 47,* 116–129. doi:10.1016/0022-0965(89)90066-0

Pillow, B. H. (2002). Children's and adults' evaluation of the certainty of deductive inferences, inductive inferences, and guesses. *Child Development, 73,* 779–792. doi:10.1111/1467-8624.00438

Poulin-Dubois, D., Brooker, I., & Polonia, A. (2011). Infants prefer to imitate a reliable person. *Infant Behavior & Development, 34,* 303–309. doi:10.1016/j.infbeh.2011.01.006

Poulin-Dubois, D., & Chow, V. (2009). The effect of a looker's past reliability on infants' reasoning about beliefs. *Developmental Psychology, 45,* 1576–1582. doi:10.1037/a0016715

Pratt, C., & Bryant, P. (1990). Young children understand that looking leads to knowing (so long as they are looking into a single barrel). *Child Development, 61,* 973–982. doi:10.2307/1130869

Rakoczy, H., Warneken, F., & Tomasello, M. (2007). "This way!" "No! That way!"—3-year olds know that two people can have mutually incompatible desires. *Cognitive Development, 22,* 47–68. doi:10.1016/j.cogdev.2006.08.002

Reid, T. (2000). *An inquiry into the human mind on the principles of common sense*. Edinburgh, Scotland: Edinburgh University Press. (Original work published 1764)

Rhodes, M., Gelman, S. A., & Brickman, D. (2010). Children's attention to sample composition in learning, teaching, and discovery. *Developmental Science, 13,* 421–429. doi:10.1111/j.1467-7687.2009.00896.x

Robinson, E. J., Einav, S., & Fox, A. (2013). Reading to learn: Prereaders' and early readers' trust in text as a source of knowledge. *Developmental Psychology, 49,* 505–513. doi:10.1037/a0029494

Sabbagh, M. A., & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development, 72,* 1054–1070. doi:10.1111/1467-8624.00334

Sabbagh, M. A., & Henderson, A. (2013a). *Eye-tracking evidence regarding how preschoolers manifest selective learning in different contexts*. Paper presented at the 2013 Biennial Meeting of the Society for Research in Child Development, Seattle, WA.

Sabbagh, M. A., & Henderson, A. (2013b). Preschoolers are selective word learners. In M. Banaji & S. A. Gelman (Eds.), *Navigating the social world* (pp. 177–180). New York, NY: Oxford University Press.

Sabbagh, M. A., & Shafman, D. (2009). How children block learning from ignorant speakers. *Cognition, 112,* 415–422. doi:10.1016/j.cognition.2009.06.005

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926–1928. doi:10.1126/science.274.5294.1926

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers play more when evidence is confounded. *Developmental Psychology, 43,* 1045–1050. doi:10.1037/0012-1649.43.4.1045

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

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40,* 162–176. doi:10.1037/0012-1649.40.2.162

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science, 10,* 322–332. doi:10.1111/j.1467-7687.2007.00587.x

Schulz, L. E., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Interventions and causal inference. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy and computation* (pp. 67–86). New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780195176803.003.0006

Schwier, C., van Maanen, C., Carpenter, M., & Tomasello, M. (2006). Rational imitation in 12-month-old infants. *Infancy, 10,* 303–311. doi:10.1207/s15327078in1003_6

Scofield, J., & Behrend, D. A. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development, 23,* 278–290. doi:10.1016/j.cogdev.2008.01.003

Senju, A., Southgate, V., Snape, C., Leonard, M., & Csibra, G. (2011). Do 18-month-olds really attribute mental states to others? A critical test. *Psychological Science, 22,* 878–880.

Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science, 15,* 436–447. doi:10.1111/j.1467-7687.2012.01135.x

Shafto, P., & Goodman, N. (2008). *Teaching games: Statistical sampling assumptions for learning in pedagogical situations*. Proceedings of the 30th annual conference of the Cognitive Science Society, Washington, DC.

Shtulman, A., & Schulz, L. E. (2008). The relationship between essentialist beliefs and evolutionary reasoning. *Cognitive Science, 32,* 1049–1062. doi:10.1080/03640210801897864

Sobel, D. M., Buchanan, D. W., Butterfield, J., & Jenkins, O. C. (2010). Interactions between causal models, theories, and social cognitive development. *Neural Networks, 8-9,* 1060–1071. doi:10.1016/j.neunet.2010.06.004

Sobel, D. M., & Corriveau, K. H. (2010). Children monitor individuals' expertise for word learning. *Child Development, 81,* 669–679. doi:10.1111/j.1467-8624.2009.01422.x

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42,* 1103–1115. doi:10.1037/0012-1649.42.6.1103

Sobel, D. M., & Kirkham, N. Z. (2007). Bayes nets and blickets: Infants' developing representations of causal knowledge. *Developmental Science, 10,* 298–306. doi:10.1111/j.1467-7687.2007.00589.x

Sobel, D. M., & Macris, D. M. (2013). Children's understanding of speaker reliability between lexical and semantic knowledge. *Developmental Psychology, 49,* 523–532. doi:10.1037/a0029658

Sobel, D. M., & Munro, S. A. (2009). Domain generality and specificity in children's causal inferences about ambiguous data. *Developmental Psychology, 45,* 511–524.

Sobel, D. M., Sedivy, J., Buchanan, D. W., & Hennessey, R. (2012). The role of speaker reliability in children's inferences about the meaning of novel words. *Journal of Child Language, 39,* 90–104. doi:10.1017/S0305000911000018

Sobel, D. M., & Sommerville, J. A. (2009). Rationales in children's causal learning from others' actions. *Cognitive Development, 24,* 70–79. doi:10.1016/j.cogdev.2008.08.003

Sobel, D. M., & Sommerville, J. A. (2010). The importance of discovery in children's causal learning from interventions. *Frontiers in Developmental Psychology, 1,* 176–183.

Sobel, D. M., Sommerville, J. A., Travers, L. V., Blumenthal, E. J., & Stoddard, E. (2009). Preschoolers' use of others' beliefs to make causal inferences from probabilistic data. *Journal of Cognition and Development, 10,* 262–284. doi:10.1080/15248370903389416

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science, 28,* 303–333.

Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N., & Blumenthal, E. J. (2007). The blicket within: Preschoolers' inferences about insides and causes. *Journal of Cognition and Development, 8,* 159–182. doi:10.1080/15248370701202356

Strauss, S., Ziv, M., & Stein, A. (2002). Teaching as a natural cognition and its relations to preschoolers' developing theory of mind. *Cognitive Development, 17,* 1473–1487. doi:10.1016/S0885-2014(02)00128-4

Strichartz, A. F., & Burton, R. V. (1990). Lies and truth: A study of the development of the concept. *Child Development, 61,* 211–220. doi:10.2307/1131060

Taylor, M., Cartwright, B. S., & Bowden, T. (1991). Perspective taking and theory of mind: Do children predict interpretive diversity as a function of differences in observers' knowledge? *Child Development, 62,* 1334–1351. doi:10.2307/1130810

Taylor, M., Lussier, G. L., & Maring, B. L. (2003). The distinction between lying and pretending. *Journal of Cognition and Development, 4,* 299–323. doi:10.1207/S15327647JCD0403_04

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,* 1065–1077. doi:10.1037/a0023273

Träuble, B., & Pauen, S. (2007). The role of functional information for infant categorization. *Cognition, 105,* 362–379. doi:10.1016/j.cognition.2006.10.003

Trevarthen, C. (1979). Communication and cooperation in early infancy: A description of primary intersubjectivity. In M. Bullowa (Ed.), *Before speech: The beginning of interpersonal communication* (pp. 321–348). Cambridge, MA: Cambridge University Press.

Vaish, A., Grossmann, T., & Woodward, A. (2008). Not all emotions are created equal: The negativity bias in social-emotional development. *Psychological Bulletin, 134,* 383–403. doi:10.1037/0033-2909.134.3.383

VanderBorght, M., & Jaswal, V. K. (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development, 18,* 61–71. doi:10.1002/icd.591

Van Reet, J., Green, K., & Sobel, D. M. (2013). *Preschoolers' existing theory of mind knowledge influences whom they trust about others' theory of mind*. Manuscript submitted for publication.

Vosniadou, S., & Brewer, W. F. (1992). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology, 24,* 535–585. doi:10.1016/0010-0285(92)90018-W

Vredenburgh, C., & Kushnir, T. (2013). Help-seeking as a cause of young children's collaboration. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th annual conference of the Cognitive Science Society* (pp. 3705–3710). Austin TX: Cognitive Science Society.

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory of mind development: The truth about false belief. *Child Development, 72,* 655–684. doi:10.1111/1467-8624.00304

Wellman, H. M., & Liu, D. (2004). Scaling of theory-of-mind tasks. *Child Development, 75,* 523–541. doi:10.1111/j.1467-8624.2004.00691.x

Wellman, H. M., & Liu, D. (2007). Causal reasoning as informed by the early development of explanations. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 261–279). New York, NY: Oxford University Press. doi:10.1093/acprof:oso/9780195176803.003.0017

Wenham, M. W. (2004). *Understanding primary science*. London, England: Paul Chapman.

Woolley, J. D., & Ghossainy, M. (in press). Revising the fantasy-reality distinction: Children have naïve skeptics. *Child Development*. Advance online publication.

Wu, R., Gopnik, A., Richardson, D., & Kirkham, N. Z. (2011). Infants learn about objects from statistics and people. *Developmental Psychology, 47,* 1220–1229. doi:10.1037/a0024023

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences of the United States of America, 105,* 5012–5015. doi:10.1073/pnas.0704450105

Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science, 22,* 28–32. doi:10.1177/0963721412469396

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review, 114,* 245–272. doi:10.1037/0033-295X.114.2.245

Yang, D. J., Bushnell, E., Buchanan, D. W., & Sobel, D. M. (in press). Infants' use of contextual cues in the generalization of effective actions from imitation. *Journal of Experimental Child Psychology*.

Ziv, M., & Frye, D. (2004). Children's understanding of teaching: The role of knowledge and belief. *Cognitive Development, 19,* 457–477. doi:10.1016/j.cogdev.2004.09.002

Zmyj, N., Buttelmann, D., Carpenter, M. J., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. *Journal of Experimental Child Psychology, 106,* 208–220. doi:10.1016/j.jecp.2010.03.002

Received November 29, 2012
Revision received June 19, 2013
Accepted June 26, 2013 ■