


This article was downloaded by: [Massachusetts Institute of Technology], [Paul Muentener]
On: 24 January 2012, At: 09:38
Publisher: Psychology Press
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered
office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Language Learning and Development

Publication details, including instructions for authors and
subscription information:

<http://www.tandfonline.com/loi/hlld20>

What Doesn't Go Without Saying: Communication, Induction, and Exploration

Paul Muentener^a & Laura Schulz^a

^a Department of Brain & Cognitive Sciences, Massachusetts Institute of
Technology

Available online: 12 Jan 2012

To cite this article: Paul Muentener & Laura Schulz (2012): What Doesn't Go Without Saying:
Communication, Induction, and Exploration, *Language Learning and Development*, 8:1, 61-85

To link to this article: <http://dx.doi.org/10.1080/15475441.2011.616455>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any
substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing,
systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation
that the contents will be complete or accurate or up to date. The accuracy of any
instructions, formulae, and drug doses should be independently verified with primary
sources. The publisher shall not be liable for any loss, actions, claims, proceedings,
demand, or costs or damages whatsoever or howsoever caused arising directly or
indirectly in connection with or arising out of the use of this material.

What Doesn't Go Without Saying: Communication, Induction, and Exploration

Paul Muentener and Laura Schulz

Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology

Although prior research on the development of causal reasoning has focused on inferential abilities within the individual child, causal learning often occurs in a social and communicative context. In this paper, we review recent research from our laboratory and look at how linguistic communication may influence children's causal reasoning. First, we present a study suggesting that toddlers only treat spontaneously occurring predictive relationships as if they might support intervention if the events are described with causal language. Second, we show that presenting causal hypotheses as contrastive beliefs, rather than neutrally, improves kindergarteners' ability to provide evidence for their causal inferences. Third, we provide a rational analysis suggesting that stronger inductive inferences are licensed when evidence is presented pedagogically than nonpedagogically; preschoolers are sensitive to this with the consequence that, for better and worse, instruction constrains exploration. In each case study, we discuss the implication that language has a unique role in changing children's interpretation of evidence for causal relationships.

Here is something we learned the other day: porcupines float. Here is how we did *not* learn it: by seeing a floating porcupine. We learned it by reading a book. For all the controversy over the relationship between language and thought, one thing is uncontroversial: language (“porcupines float,” “this is a blicket,” “the block makes it go”) communicates information, including information you might not discover by any other means. The transmission of information is a trivially true property of language. However, in the discussion to follow, we will take on the “task of psychology to remove the veil of self-evidence” from things (Asch, 1952) and take a closer look at what it means that language is informative.

As researchers with primary interests outside of language per se, we will bring data to bear on the relationship between language and cognition not originally generated for this purpose. We became interested in language in part because one of the challenges of designing studies for young children is figuring out exactly what to say to them. A tempting strategy is to try to say as little as possible; in studies not focused on language itself, it is easy to think of language as a distraction. Nonetheless, to run experiments, we usually need to tell children something about what they have to do. In the course of thinking about what to say to children, we became interested in

how different kinds of communicative acts might serve as unique sources of information about causal relationships.

Here we will investigate three different kinds of communicative acts: describing, disagreeing, and demonstrating. We will look at how *describing* events can add evidence about causal relationships, supporting enriched (or even altogether new) conceptual representations. We will provide some preliminary data suggesting that *disagreeing* about causal relationships increases learners' attention to evidence, affecting their ability to report how they know what they know. Finally, we will suggest that what people say when *demonstrating* causal relationships changes learners' evaluation of evidence, licensing inferences not only about what is taught but also about what is not.

BACKGROUND

Our research has focused primarily on how children learn and develop coherent, structured causal theories of the world. A cartoon that we have sometimes used to talk about causal learning is displayed in Figure 1.

The top half of the cartoon suggests that evidence affects children's beliefs. These beliefs, together with core concepts rooted in innate, domain-specific knowledge (Carey, 2009; Spelke, Breinlinger, Macomber, & Jacobsen, 1992) affect children's interpretation of subsequent evidence. Thus, for instance, children, like the rest of us, require less evidence to accept belief-consistent causal relations than belief-violating ones. Studies from our own lab and many others suggest that children rationally integrate evidence and prior beliefs in this manner (Bonawitz, Fischer, & Schulz, in press; Kushnir & Gopnik, 2007; Schulz, Bonawitz, & Griffiths, 2007; Schulz & Gopnik, 2004; Sobel, Tenenbaum, & Gopnik, 2004).

The bottom half of the cartoon suggests that children's beliefs affect the actions they take. "Actions" here might mean many things, from performing an intervention that exploits existing knowledge to performing an action to check a prediction. However, we have been particularly interested in exploratory behavior, that is, the actions children take that can generate evidence previously unknown to the child. When the evidence children generate is compelling and inconsistent with their prior beliefs, children can learn from the evidence and change their minds. The arrow from action to evidence suggests a bootstrapping mechanism to support new learning.




FIGURE 1 Cartoon depiction of the relationship between induction and exploration: Evidence leads to beliefs, which affect the interpretation of subsequent evidence and support actions that can generate evidence for new learning. (Color figure available online.)

The claim that that children “learn by doing” is almost as uncontroversial as the idea that language “tells you things”; the centrality of exploratory play to learning is part of the public understanding of early childhood development and the cartoon itself could be taken as a depiction of Piagetian constructivism (Piaget, 1954). However, most studies of exploratory play have been descriptive rather than experimental, and most predate contemporary work on causal inference (Berlyne, 1954; Piaget, 1951, 1954). Thus, although it is obvious that children both play and learn, decades of research had left the link between the two relatively unspecified. In our lab, we have tried to bridge that gap. Several studies from our lab suggest that children do indeed engage in “rational play”: they explore more when there is something to be learned and they generate informative evidence in the course of their own exploration (e.g., Cook, Goodman, & Schulz, in press; Schulz & Bonawitz, 2007; Schulz, Hoopell, & Jenkins, 2008; Schulz, Standing, & Bonawitz, 2008).

Nonetheless, it is clear that this cartoon is incomplete. Evidence for causal relations is sometimes generated through an individual child’s observation and exploration, but it is also generated in social contexts with the help of knowledgeable others. Thus the classic Piagetian approach, in which a single child constructs knowledge through self-directed exploration, is often contrasted with that of Piaget’s contemporary, Lev Vygotsky, who consistently emphasized the importance of the sociocultural context (1934).¹

In contemporary developmental work, studies on the communicative context and causal reasoning have taken a variety of approaches. Some researchers have looked at interactions between event representations and the syntax of causal language (e.g., Ammon, 1980; Behrend, 1990; Bungler, 2008, Casasola & Cohen, 2000; Fausey & Boroditsky, 2009; Fisher, 1996, 2002; Gentner, 1978; Muentener & Lakusta, 2011; Naigles, 1990; Wolff, 2003). Other researchers have looked at the causal explanatory context in general, considering both how children’s own explanations, and requests for explanations (e.g., in the form of “why” questions), affect children’s causal reasoning (Chouinard, 2007; Hood & Bloom, 1979; Wellman & Liu, 2007). Finally, researchers have looked at how children’s causal representation of events is affected by testimony from informants (e.g., Harris, 2002; Harris & Koenig, 2006; Koenig, Clement, & Harris, 2004; Jaswal & Neely, 2006). We will briefly review these lines of work before turning to the implications of our own empirical data in considering how language may act as a distinctive source of information about causal relationships.

Many researchers have looked at the relationship between causal inferences and the grammar of causal constructions and the semantics of causal verbs (e.g., Ammon, 1980; Behrend, 1990; Bungler, 2008, Casasola & Cohen, 2000; Fisher, 1996, 2002; Gentner, 1978; Naigles, 1990, 1996; Wolff, 2003). Developmental work suggests that productive use of common causal syntactic structures emerges in early childhood. For instance, Bowerman (1982) reported in a diary study that both lexical (e.g., “He rolled the ball”) and periphrastic (e.g., “He made the ball roll”) causal syntactic structures are present in two-year-olds’ speech. Also, in a classic study, Naigles

¹The contrast should be qualified however. Although his writing on the topic is somewhat obscure, Piaget was well aware of the ways in which the individual child’s play and the communicative context might interact: “We have to attempt to determine the connection between the imitative image, ludic symbolism and representative intelligence, i.e., between cognitive representation and the representation of imitation and play. This very complex problem is still further complicated by the intervention of language, collective verbal signs coming to interfere with the symbols we have already analyzed, in order to make possible the construction of concepts” (Piaget, 1951).

(1990) found that 2-year-olds distinguish causal and noncausal descriptions, mapping transitive sentences, (“the duck is kradding the rabbit”) onto causal events (the duck pushing the rabbit into a squatting position), and intransitive sentences (“The duck and rabbit are kradding”) onto noncausal events (the duck and the rabbit each waving their arms).

Moreover, both adults and children describe causal events differently depending on the degree of animacy and intentionality present in the events. Adults tend to map lexical causal structures (e.g., “broke the window”) onto intentional actions and periphrastic structures (e.g., “made the window break”) onto unintentional actions and actions initiated by inanimate objects (Wolff, 2003). A similar pattern occurs in children’s descriptions of causal events. Children’s tendency to prefer causal to noncausal verbs depends on whether actions are intentional or accidental (e.g., they will say “she popped the balloon” for an intentional action but “she hit the balloon” for an accidental or object-initiated action; Muentener & Lakusta, 2011). This bias has also been shown to influence children’s conceptual representations of events. Children who hear unintentional events described with causal language (“she popped the balloon”) judge the actors as more responsible for the outcome than they do when the events are described with noncausal language (“the balloon popped”; Fausey & Boroditsky, 2009). Collectively, these studies suggest that children are sensitive to causal syntax from early in development and that the linguistic framing of events affects how children construe the evidence they observe. We were influenced by this work in thinking about how using causal language to describe events might shape children’s ability to represent relationships between prediction and intervention (Study 1 to follow).

A different line of research has looked at how children’s causal reasoning is affected, not by the specifics of causal language but by the practice of offering and requesting causal explanations more generally. Corpus work suggests that children both ask for and offer causal explanations from early in development (Callanan & Oakes, 1992; Dunn & Brown, 1993; Hickling & Wellman, 2001; Hood & Bloom, 1979; Inagaki & Hatano, 1993, 2002; Kelemen, Callanan, Casler, & Perez-Granados, 2005; Lagutta & Wellman, 2001; Schult & Wellman, 1997). Consistent with both adult and developmental work suggesting that the act of explaining can support new learning (Chi, Bassock, Lewis, Reinmanm, & Glaser, 1989; Chi, de Leeuw, Chiu, & LaVancher, 1994; Lombrozo, 2009; Crowley & Siegler, 1999; Pine & Siegler, 2003; Siegler, 1995; Trabasso & Suh, 1993; Williams & Lombrozo, 2010), research suggests that children benefit from explaining events and can sometimes offer informative explanations well before they can make accurate predictions (Amsterlaw & Wellman, 2006; Legare, Wellman, & Gelman, 2009; Wellman & Liu, 2007). By the age of three, children actively seek explanations for evidence that is anomalous with respect to their prior beliefs (Legare, *in press*; Legare, Gelman, & Wellman, 2010) and monitor the explanations they receive, resisting uninformative explanations (Baum, Danovitch, & Keil, 2006; Frazier, Gelman, & Wellman, 2009). In our discussion of Study 2, we consider how children’s understanding of the pragmatics of explanation may affect their tendency to generate sufficient evidence to justify their causal beliefs to others.

Finally, a growing body of work suggests that preschoolers are sensitive to the epistemic status of informants and preferentially learn from knowledgeable, reliable others (e.g., Harris, 2002; Harris & Koenig, 2006; Jaswal & Neely, 2006; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2005; Robinson & Whitcombe, 2003). If one person knows the label for a familiar object (e.g., a ball) and another person does not, children selectively accept a novel label for a novel object (and imitate a novel function for a novel object) only when the novel information is provided by the knowledgeable informant (Koenig & Harris, 2005). Furthermore, children

reject testimony from someone who has less access to information than they do. Children accept testimony that a red rather than a blue bug is hidden in a tunnel when the informant can see the toys and they can only feel the toys; however, even 3-year-olds are not so trusting of testimony that they defer to the informant when they can see the toys and the informant cannot (Robinson and Whitcombe, 2003; see also Koenig, Clement, & Harris, 2004). Such studies suggest that children's causal beliefs are sensitive to the quality of explanations and the epistemic status of informants. These findings motivate our own work on how children's causal judgments are affected by the testimony of others and the manner in which evidence is presented (Studies 2 and 3 to follow).

In summary, here we consider the relationship between the social-communicative context, construed broadly, and children's causal reasoning. We are interested in the ways in which information communicated through language changes the interpretation of evidence and affects the kinds of causal inferences children make. In particular we consider the ways that the mere fact of describing, disagreeing about, or demonstrating evidence can convey information that goes beyond the content itself.

STUDY 1: DESCRIBING CAUSAL RELATIONSHIPS

Taking our cartoon (Figure 1) as a point of reference, Study 1 looks at how different ways of talking about evidence affect children's beliefs and subsequent actions. In particular, we will look at how describing causal relationships supports children's ability to bind prediction and action.

A striking feature of adult causal cognition is that human adults make predictions about predictive relations. Specifically, adults believe that if event A reliably predicts event B, then acting to bring about A might bring about B. This ability is not trivial. Nonhuman animals generalize cues learned through classical conditioning (e.g., that a light predicts food) to behaviors learned through operant conditioning (they will push a lever for food more often when a light is present than when it is not; Estes, 1948). They also make different predictions about relationships learned under observation and intervention. If for instance, a rat learns that a light predicts both a tone and food, the rat expects food when it hears the tone, but the rat does not predict food if the rat triggers the tone itself (Blaisdell, Sawa, Leising, & Waldmann, 2006; Leising, Wong, Waldmann, & Blaisdell, 2008). However, there is no evidence that nonhuman animals generalize directly from observation to intervention: having learned that A predicts B, nonhuman animals do not seem to spontaneously infer that they might intervene on A to try to generate B (Tomasello & Call, 1997).

In recent work, we wondered whether the ability to form integrated representations of prediction and action emerged relatively late developmentally as well as phylogenetically. Given decades of research on the sophistication of children's causal reasoning (e.g., Bullock, Gelman, & Baillargeon, 1982; Gopnik & Sobel, 2000; Gopnik et al., 2004; Kushnir & Gopnik, 2005, 2007; Schulz & Bonawitz, 2007; Schulz, Goodman, Tenenbaum, & Jenkins, 2008; Schulz & Sommerville, 2006; Shultz, 1982; Sobel, 2004; Sobel & Kirkham, 2006; Williamson, Meltzoff, & Markman, 2008), it might seem surprising even to wonder whether children might fail to infer that predictive relations could support effective intervention. However, two features common to experimental studies of children's causal reasoning might have masked children's limitations. First, in almost every experimental setting testing children's causal reasoning, the causal event is initiated by an intentional agent: an adult, a puppet, or the child herself rolls a ball down a tube

(Bullock et al., 1982), puts a block on a toy (Gopnik & Sobel, 2000), or flips a switch (Schulz & Sommerville, 2006). In principle, children might be able to recognize predictive relationships as potentially causal when they are initiated by goal-directed actions (e.g., a person or a puppet initiates an event) but unable to represent causal relationships otherwise. That is, children might fail to recognize that predictive relations are potentially causal when they occur spontaneously (e.g., a ball just happens to roll down a tube and an effect occurs). Second, and critically for the current discussion, most experiments on children's causal reasoning have been accompanied by causal language (e.g., "The block makes the toy go!"). Children might succeed in representing predictive relations as causal (and thus supporting intervention) when they are told that the events are causal, but not otherwise.

To look at whether, in the absence of dispositional agents or linguistic cues, young children generalize from prediction to action, we showed children a block that slid (apparently spontaneously) across a stage towards a base (Bonawitz et al., 2010) (Figure 2). When the block contacted the base, a toy airplane, connected to the base by an orange wire, immediately lit up and began to spin. We occluded the stage and then showed children the event sequence again three more times. We wanted to know whether children would 1) predict that the plane would activate when the block contacted the base; 2) be willing and able to move the block to the base; and, critically, 3) whether when children themselves moved the block into the base, they would look up to see if the plane had activated. Note that of course children might not, and indeed should not, infer that the block *definitely* causes the plane to activate. (The events might be spuriously associated or there might be a hidden common cause.) However, if children simply recognize that events that predict each other sometimes cause each other, we would expect them to perform the action and look to see whether the effect occurs.

After children saw the predictive relation four times, we inserted a catch trial, in which the plane did not activate, to see if the children learned the predictive relation between the block and the plane. Specifically, we coded whether the children immediately looked up toward the plane after the block contacted the base. We tested both preschoolers ($n = 14$; mean age: 47 months;




FIGURE 2 Stimuli used in Study 1 (Bonawitz et al., 2010). The red block slides across the stage towards the yellow block (base). The base is attached to the toy airplane (top left) by an orange wire. When the block contacts the yellow base, the plane lights up and spins. (Color figure available online.)

range: 37–60 months) and toddlers ($n = 14$; mean age: 24 months; range: 18–30 months). There were no differences between age groups: almost all of the children (89% in each age group) learned the predictive relation. Because we were interested in whether children would generalize from prediction to action, we excluded from subsequent analyses those children who failed to perform the predictive look. After the catch trial, children saw a final successful outcome, in which the block contacted the base and the plane activated.

Immediately following the final presentation of the predictive relation, we gave the block to the children, pointed to the plane, and said, “Okay, now it’s your turn. Can you make it go?” Children were given 60 seconds to initiate the action. Here there was a striking age difference: 62% of preschoolers spontaneously slid the block into the base; not one of the 14 toddlers did (although all the children readily manipulated the block). Note, however, that this in itself does not establish that preschoolers succeed in representing the causal relationship or that toddlers fail to do so. The preschoolers might have slid the block to the base without any expectation that this might cause the plane to activate, and the toddlers might have failed to move the block to the base because they were more interested in the block than the plane, or because they were simply unable to perform the target action.

Thus if children failed to initiate the action spontaneously in 60 seconds, we prompted them: we pushed the block towards the base but stopped just short of contacting it. We then returned the block and gave children another 60 seconds to perform the action. After prompting, all the remaining preschoolers and 87% of the toddlers successfully performed the target action. Because we wanted to ensure that all children were willing and able to perform the target action, children were excluded from subsequent analyses if they failed to perform the target action after prompting.

The critical question was whether, having learned the predictive relations and having performed the target action, children predicted the outcome of their own intervention. Did children look predictively towards the plane after performing the action themselves? Again, the age effect was striking: 87% of the preschoolers succeeded and 100% of the toddlers failed (Figure 3). Although the toddlers learned the predictive relation between the block’s movement and the plane’s activation as well as the preschoolers, in a full minute of free play, no toddler moved the block to the base spontaneously, and when prompted to move the block to the base, no toddler looked up to see if their action had activated the plane.

Toddlers’ failure to move from prediction to action, although in dramatic contrast with the behavior of preschoolers, is consistent with research on earlier conceptual representations of causality. Research investigating causal representations in infancy has shown that infants represent only a limited range of events causally. For instance, although 8.5-month-old infants establish causal representations of Michottean launching events (where one ball strikes another and appears to set it in motion; Cohen & Oakes, 1993; Leslie & Keeble, 1987), they fail to do so if the target event is not a motion event but a change of state event (e.g., a ball strikes a box and the box breaks apart or plays music; Muentener & Carey, 2010). Interestingly, infants succeed if the causal agent is also a dispositional agent (e.g., a hand strikes the box; Muentener & Carey, 2010).

If, absent intentional action, infants and toddlers initially construe only a limited set of physical sequences as causal (e.g., events initiated by dispositional agents and motion events involving direct contact), how do they eventually develop enriched causal representations such that children just a few years older spontaneously realize that novel predictive relations might support




FIGURE 3 Results of Study 1 (Bonawitz et al., 2010). Percentage of children in each condition who looked towards the airplane following their own spontaneous or prompted intervention.

intervention? We now turn to the hypothesis that causal language facilitates children's causal understanding.

There are two possible roles that causal language might play in changing children's conceptual representations. One possibility is that, in using the same words to describe causal events that do and do not involve dispositional agents (e.g., "You made the branch shake"; "The wind made the branch shake") causal language draws attention to the commonalities between events that result from agent actions and observed correlations. Another possibility is that causal language simply testifies that an observed sequence *is* a direct causal relationship (i.e., rather than a spurious association or the consequence of a hidden common cause). The fact that almost no toddlers tried to see whether the block would activate the toy seems to suggest that they failed to construe the relationship even as *potentially* causal. However, we cannot rule out the possibility that toddlers were simply less certain about the causal relationship than the preschoolers. On either account, introducing causal language might improve children's ability to move from prediction to intervention.

To test this, we replicated the procedure described above but added a verbal description during the initial predictive learning phase of the experiment. In one condition (the *Identical Causal Language* condition), children ($n = 16$; mean age: 24 months; range: 19–29 months) were told "The block can make it go!" as they watched the block move towards the base and the toy activate. This is, of course, the same language ("make it go") then repeated during the test phase when the experimenter gives the child the block and says, "Can you make it go?" In a second condition (the *Different Causal Language* condition), we used different but semantically equivalent language during the predictive learning and test phases. During the predictive learning phase, children ($n = 16$; mean age: 24 months; range: 18–30 months) were told "The block

can make it go,” and during the test phase they were asked, “Can you turn it on?” Finally, in a *Noncausal Language* control condition, we simply attracted children’s ($n = 16$; mean age: 24 months; range: 18–30 months) attention to the events during the predictive learning phase (“Let’s watch my show. Here it goes!”).

There were no differences between the causal and noncausal language conditions in children’s ability to learn the initial predictive relation, suggesting that children attended equally to the events across conditions. As in the initial study, only children who performed the predictive look and the target action were included for analysis. The critical question was whether children would spontaneously look from the block to the airplane following the target intervention. Replicating the original finding, toddlers in the *Noncausal Language* condition failed to perform the action spontaneously and, with the exception of a single child (6%), failed to look predictively when prompted to perform the target action. By contrast, 50% of the children in the *Identical Causal Language* condition and 62% of the children in the *Different Causal Language* condition succeeded at the task. (The *Causal Language* conditions did not differ significantly from one another.)

How then does language help children bind prediction and action? We believe that the answer depends on what the initial state of children’s knowledge is like. Children might have the wrong concept of causation altogether. Alternatively, they might be right about the concept of causation but fail to recognize many events in the world as instances of that concept.

The idea that children have the wrong concept of causation suggests a genuine discontinuity in the causal representations of younger and older children. Specifically, infants and toddlers might recognize many features associated with causal events (predictive relations, spatiotemporal contiguity, and the ability to support intervention) but, unlike older children and adults, fail to recognize that these features predict one another; that is, they might not have an integrated, adult-like concept of “cause.” For instance, one might construe the evidence that 8-month-old infants treat hands engaged in goal-directed actions but not hands acting accidentally as potential participants in contact causality (Muentener, 2009) as evidence that infants’ concept of causation is too narrow; it represents only event outcomes that immediately follow goal-directed actions.

If this is the case, then causal language might help children integrate representations that are originally distinct. Arguments for this kind of linguistic bootstrapping have been suggested in the domain of number (for integrating the small exact and large approximate number systems into the adult number concept; e.g., Carey, 2009) and space (for integrating distance, angle, and directional relationships into Euclidian geometry; Spelke, Lee, & Izard, 2010; see also Spelke, 2003). Language might play a comparable role in forging an integrated concept of causation out of evolutionarily ancient systems for learning predictive relations among events (as in classical conditioning), the outcomes of agent actions (as in operant conditioning), and spatiotemporal relations associated with contact causality (as in launching events). Even if the initial representations are not modular, the use of common words could help children integrate representations that they might otherwise deploy in different contexts for different tasks. This kind of role for language has been suggested for instance, in supporting children’s categorization by causal rather than perceptual properties of objects (Gopnik & Nazzi, 2003; Nazzi & Gopnik, 2000, 2001; Schulz, Standing, & Bonawitz, 2008). In this way, causal language might help children develop a theoretical construct of “causation” that integrates predictive relations, information about interventions, and spatiotemporal cues appropriate to the domain.

An alternative possibility however, is that infants and toddlers implicitly have an integrated, adult-like representation of causation. Statisticians and philosophers (e.g., Hitchcock, 1997; Lewis, 1973; Pearl, 2000; Salmon, 1998; Spirtes, Glymour, & Scheines, 1993; Suppes, 1970; Woodward, 2003) argue about exactly what a representation of causation is; however, one useful supposition is that if there is some context in which you can intervene on one variable to change another, the first variable is a cause of the second (Woodward, 2000). Under this construal, toddlers may also believe that a predictive relation might also be a causal relation but only for a limited set of events. Thus, the fact that 8-month-old infants treat hands engaged in goal-directed actions but not hands acting accidentally as potential participants in contact causality (Muentener, 2009) is evidence not that they have the wrong concept of causation but only that they fail to recognize hands acting accidentally as a potential candidate cause. In particular, infants may apply the concept of causation at first only to the relationship between their own actions and immediate outcomes and the goal-directed actions of others (Gergely, Bekkering, & Király, 2002; Meltzoff, 2007; Rovee-Collier, 1987; Watson & Ramey, 1987). This early restricted nature of causal reasoning may be the product of the types of evidence infants receive over the first couple months of life or might be related to the kinds of events infants' core knowledge prepares them to detect (Carey, 2009; Muentener & Carey, 2010). In either case, children might only gradually come to recognize events without dispositional agents or visible transformations across the path of transmission as causal. If so, causal language would not transform children's concept of causation per se but would allow children to recognize that superficially very different sequences of events were all nonetheless instances of causal relations.

We do not think that it is an easy matter to distinguish these accounts; there are many tasks in which infants and toddlers could behave identically because they lack an adult-like concept of causation or because they recognize a restricted class of events as causal. However, we can suggest two follow-up experiments. First, we might want to know whether toddlers simply have difficulty moving from prediction to action or if they fail to *represent* predictive relations as causal in the absence of dispositional agents or causal language. One way to test this is to see whether toddlers fail not only to act on predictive relations but also fail to expect contact causality for these relations — an expectation that has been shown to be a default assumption for causal events in infancy (Muentener & Carey, 2010) and early childhood (Kushnir & Gopnik, 2007). Suppose, for instance, a block goes towards the base but disappears behind an occluder before children see whether it makes contact; the plane then activates. If children represent the predictive relation as causal, then toddlers should expect contact if the plane activates but not if it doesn't; if they fail to represent the predictive event sequence as causal, they should make no differential prediction. Some preliminary data from our lab (Muentener, Bonawitz, Horowitz, & Schulz, 2010; see also Bonawitz, Ferranti, Horowitz, & Schulz, 2009) suggest that the results described here for action replicate in toddlers' predictions about contact causality: toddlers expect the appropriate contact relations when nonagentive events are described in causal language but not otherwise.

A second prediction is that if children genuinely fail to have an adult-like concept of causality, then their failures should be robust across tasks and domains. If instead, children simply fail to recognize some events as instances of causal relationships, then they should show a more mixed pattern of success and failure. Suppose, for instance, that a block moves towards a base but instead of activating a plane, the block contacting the base causes a puppet to giggle; we can then ask children "Can you make the puppet giggle?" If toddlers again fail to perform the

action spontaneously and fail to look towards the puppet when prompted to perform the action, this would provide additional support for the hypothesis that children's causal representations are discontinuous from adults. However, if toddlers succeed, this would suggest that they can, at least in some domains, recognize that predictive relations support intervention. This would suggest that their difficulties are not in knowing what causation is, but in recognizing particular events *as* causal. We are currently investigating these possibilities in our lab.

STUDY 2: DISAGREEING ABOUT CAUSAL RELATIONS

In Study 1, we focused on linguistic communication about evidence, suggesting that describing events with causal language affects children's ability to integrate prediction and action. Returning to our cartoon (Figure 1), Study 2 looks at the arrow between belief and action. We will discuss how different ways of talking about beliefs can affect children's interpretation of evidence, subsequent actions, and the new evidence they generate.

One of the striking features of cognitive development is that although children learn rapidly, even from sparse noisy data, they often fail to remember the evidence for their beliefs. Very young children may not even recognize that they learned anything at all. Three-year-olds who look into a candy box only to discover that it contains pencils report seconds later that they always thought that there were pencils in the box (Gopnik & Astington, 1988). Similarly, children who are taught a new fact (e.g., that cats use their whiskers to decide whether they can fit into narrow openings) deny that they just learned the information and report that they have known it for a long time (Taylor, Esbensen, & Bennett, 1994).

Failures of source memory are not restricted to early childhood (see Johnson, Hashtroudi, & Lindsay, 1993, for review); at all ages, we are better at remembering new information than its source (e.g., Cycowicz, Friedman, Snodgrass, & Duff, 2001; Ferguson, Hashtroudi, & Johnson, 1992; McIntyre & Craik, 1987; Schacter, Kaszniak, Kihlstrom, & Valdiserri, 1991; Simons, Dodson, Bell, & Schacter, 2004; Spencer & Raz, 1995; Wegesin, Jacobs, Zubin, Ventura, & Stern, 2000). However, outside of the rare embarrassing moment (telling a joke to the friend who taught it to you; Taylor et al., 1994), failures to remember the sources of your beliefs may not seem particularly costly.² As long you are reasonably certain that it is raining outside, it may not matter whether you learned the information by reading it, hearing it, or observing it directly. We can comfortably adopt a stance of naïve realism, assuming that the true state of the world is accessible to anyone with access to sufficient perceptual input; the particular means by which we acquire information is arguably irrelevant.

There are a few domains of expertise however (e.g., law and science) where how you came by the evidence for your beliefs *does* matter. You need not only to know things but also to be able to

²This is particularly true given that adults seem to have better source memory when the integrity, rather than merely the identity, of the source is at stake (Rahhal, May, & Hasher, 2002). Interestingly, research suggests that this is true of preschoolers as well. Although preschoolers are generally poor at source memory, they encode the difference between knowledgeable and ignorant informants, and remember this information up to a week later in deciding which informant to trust (Birch, Vauthier, & Bloom, 2004; Corriveau & Harris, 2009; Scofield & Behrend, 2008). Importantly, such studies of selective trust involve conflicting claims about the world. Thus consistent with the current proposal, preschoolers seem most likely to check the source of evidence in the face of contrastive beliefs. We are grateful to an anonymous reviewer for bringing this to our attention.

articulate how you know what you know. Unsurprisingly, given the failures of source memory in general, children and lay adults tend to be poor at this kind of metacognition. Even when people successfully use evidence to make accurate judgments, they typically fail to reproduce evidence sufficient to support their conclusions (Chen & Klahr, 1999; Dunbar & Klahr, 1989; Koslowski, 1996; Kuhn, 1989; Kuhn, Amsel, & O’Laughlin, 1988; Inhelder & Piaget, 1958).

We were struck by the fact that the adversarial system — having individuals present alternative views of a case and allowing a jury of peers to decide — has been a societal solution to improving attention to sources of evidence in both law and science. We wondered whether there was something about embedding competing hypotheses in a context of contrastive beliefs that might affect learners’ relationship to evidence. Notably, this context is universally present in real science but almost uniformly absent in experimental tests of scientific reasoning. Children are routinely asked to distinguish competing hypotheses from each other, but the hypotheses are presented neutrally as alternative possibilities, not as individuals’ differing beliefs.

Why should a context of contrastive beliefs affect learners’ attention to evidence? Consider the difference between A) hearing your friend Bob say that it is sunny and B) hearing your friend Bob say that it is sunny and your friend Emily say that it is raining. In context A, there is no threat to naïve realism; the information that it is sunny is presumably accessible to anyone with the requisite input, and the specific source of evidence is irrelevant (even if in principle, many alternative hypotheses abound: it could be sunny, snowing, etc.). Thus in context A, we can retain the information, “it is sunny”, and drop the source of information, “Bob says” (let alone any further inquiry into how Bob got his facts). Context B is not so simple. If Bob and Emily disagree about the world as it really is then you could assume that one of them is crazy (i.e., you could abandon the assumption that they are both rational agents), you could assume that one of them is engaged in deliberate deception (i.e., you could abandon the assumption that they are both moral agents), or you could assume that there is no objective fact of the matter (i.e., you could abandon naïve realism). Alternatively, you could preserve rationality, morality, and realism by considering the possibility that Bob and Emily have different sources of evidence: Bob heard a weather report from his two-year-old daughter; Emily just came in soaked from a rainstorm.

If contrastive beliefs support attention to sources of evidence, then phrasing competing hypotheses as contrastive beliefs might improve learners’ ability to provide evidence for their inferences in scientific reasoning tasks. We have recently begun investigating this possibility (Cook, Bonawitz, & Schulz, 2010). We showed five and six-year-olds (mean: 68 months) a V-shaped ramp and two different balls, one white and one black (see Figure 4). The ramp was modeled after an apparatus used in other studies of scientific reasoning (e.g., Klahr & Nigam, 2004); the height of the downside ramp could be adjusted to two settings by manipulating a wooden handle below the ramp. On the upside ramp there was a high red notch and a lower blue notch, where a ball might land after rolling down the ramp. A plastic tube affixed to the top of the downside ramp ensured that balls were released uniformly from the same location. The balls differed in appearance but generated identical outcomes on the ramp. The experimenter told the child, “We can change the kind of ball we use: we can use this (black) kind of ball or this (white) kind of ball. And we can change the height of the ramp like this: we can make it low or high.” (Order counterbalanced throughout.) The experimenter then drew the child’s attention to the red and blue areas on the upside of the ramp and told the child that sometimes a ball lands in the higher red area, and sometimes it lands in the lower blue area.



FIGURE 4 Ramp set-up used in Study 2. The height of the downside ramp (left side of Figure) can be raised and lowered. Either a white or black ball can be dropped through the tube to roll down the ramp. The ball can land in either the blue notch or red notch on the upside ramp (right side of figure). (Color figure available online.)

Thirty-two children were randomly assigned to a Neutral condition or a Contrastive Beliefs condition ($n = 16/\text{condition}$). In the Neutral condition, the competing hypotheses were presented simply as alternative possibilities. Children were told:

I wonder what makes a difference for where the ball lands. There's two ways it could happen. One is called the Height Way: the height of the ramp is what makes a difference for where the ball lands. The other way is called the Ball Way: the kind of ball you use is what makes a difference for where it lands. We want to know which way's right and which way's wrong.

In a Contrastive Beliefs condition, children were told:

I wonder what makes a difference for where the ball lands. I have two friends. My friend Bob thinks it is the height of the ramp is what makes a difference for where the ball lands. My friend Emily thinks it is the kind of ball you use is what makes a difference for where it lands. We want to know which way's right and which way's wrong.

In both conditions, the experimenter then told the child, "Go ahead and play and see if you can figure it out. When I come back I want you to show me if the height of the ramp matters and if the kind of ball matters for where the ball lands." The experimenter left the child's line of sight and the child was allowed to play freely for 90 seconds. When the experimenter returned, she asked, "Can you show me that the height of the ramp makes a difference for where the ball lands, and that the kind of ball does not make a difference for where the ball lands?" The child was allowed to demonstrate anything she liked. When the child indicated that she was finished, the experimenter asked, "Is there anything else you want to show me?" The demonstration phase stopped when the child indicated that she was done.

There are four relevant things that children might do (see Figure 5): they might drop the white ball from the high height, the white ball from the low height, the black ball from the high height, or the black ball from the low height. Children could choose to do any one, any two, any three, or all four of these. Critically, however, a correct response to the test question ("Can you show me that the height of the ball makes a difference for where the ball lands, and that the kind of ball does not?") requires performing at least three of the four relevant demonstrations.

We coded the children's actions both during the free play period and in response to the test question. There were no differences between conditions in the number of actions or kinds of actions taken during free play and during free play. This suggests that children were equally motivated by the task in both conditions. There were also of course no differences in children's




FIGURE 5 Possible target responses in Study 2. Children could place either the black or white ball at either the high or low heights. At least three of the four demonstrations are necessary to show that the height affects the outcome and that the kind of ball does not. (Color figure available online.)

understanding of how the ramp worked (given that the test prompt itself provided the answer – and post-tests confirmed that children indeed understood this). However, at test, only 4 of the 16 children (25%) in the Neutral condition successfully confirmed the height hypothesis and disconfirmed the ball hypothesis.³ By contrast, 13 of the 16 (81%) children in the Contrastive Beliefs condition successfully answered the test question. Children were significantly more likely to succeed in the Contrastive Beliefs condition than the Neutral condition ($\chi^2(1, N = 32) = 10.17, p < .005$). Moreover, the effect was specific to the introduction of contrastive beliefs, rather than the inclusion of a social context more generally: in a follow-up condition, Bob and Emily disagreed over their preferences rather than their beliefs (“Bob likes to change the height of the ramp”; “Emily likes to change the kind of ball”), and children performed no better than they did in the Neutral condition (Fisher’s exact = 1.00, $p = ns$).

Note that children could successfully confirm the height hypothesis and disconfirm the ball hypothesis simply by demonstrating all the permutations available to them with two balls and two heights. Thus disagreement might exert a relatively crude effect on children’s reasoning about evidence: children might simply realize that, when two people disagree, it is a good idea to provide as much evidence for any claim as possible. Alternatively in the context of disagreement, children might simply be more motivated to show the experimenter everything they themselves know. Critically, however, if children’s success reflects a genuine understanding of the relationship between the evidence and the hypotheses, then children ought to be able to select all *and only* the evidence relevant to any particular claim.

To see whether contrastive beliefs affected children’s ability to generate relevant evidence and not merely their ability or motivation to generate evidence generally, we subjected children to a more rigorous test: we replicated the original study but separated the test question into two

³The children who failed to produce sufficient evidence showed no distinct pattern of behavior. Of the 12 children who failed the task in the Neutral condition, 2/12 produced only a single demonstration. The remaining children produced two demonstrations: 4/12 varied only the ball, 1/12 varied only the height, and 5/12 changed both variables between demonstrations.

parts. After the free play period, children were asked (order counterbalanced), “Can you show me that the height of the ramp makes a difference for where the ball lands?” and separately, “Can you show me that the kind of ball you use doesn’t make a difference for where the ball lands?” For each question, we counted as a correct response only the “target pair” of demonstrations that minimally confirmed/disconfirmed the relevant hypothesis (i.e., showing one ball at both heights to confirm the height hypothesis; showing both balls at one height to disconfirm the ball hypothesis). If children performed superfluous (3 or 4) demonstrations, demonstrated evidence relevant to the other hypothesis, generated confounded evidence (varying both the ball and height across a pair of demonstrations), or gave only a single demonstration, they were counted as failing.

This task was much more difficult for the kindergarteners; nonetheless the effect of condition was striking. In the Neutral condition, children performed at floor: no child answered both questions correctly, 8% answered one of the two target questions correctly, and 92% failed to answer either question correctly. In the Contrastive Beliefs condition, 25% of children answered at ceiling, an additional 50% answered one of the questions correctly, and only 25% of children failed both questions. Children were more likely to succeed in the Contrastive Beliefs condition than the Neutral condition ($\chi^2(1, N = 24) = 10.97, p < .001$).

These findings suggest that framing hypotheses as contrastive beliefs facilitates kindergarteners’ ability to generate evidence to distinguish competing hypotheses. In the absence of a belief context, children were unable to coordinate evidence to distinguish hypotheses; given contrastive beliefs, kindergarten children succeeded at a type of scientific reasoning task thought to require years of formal schooling.

However, this study also raises many questions. Our task was relatively simple: children were given only two competing hypotheses and were told the correct answer (the height makes a difference; the ball does not); they were asked only to demonstrate the evidence that would support this. We do not know whether a belief context would facilitate children’s success on more complicated problems. We also do not know how competing beliefs would affect children’s performance if the children themselves had differential prior beliefs about the relevant variables. (Pilot work established that the kindergarteners initially accepted both the ball and the height as plausible variables.) Future work might investigate the interaction between contrastive belief contexts and the children’s own naïve theories.

For the purpose of this paper, however, the critical questions center on the relationship among contrastive beliefs, evidential reasoning, and language. Organisms seem to be able to understand that beliefs can be mistaken without language; both chimpanzees and preverbal infants pass false belief tasks in at least some contexts (Baillargeon, Scott, & He, 2010; Hare, Call, Agnetta, & Tomasello, 2000; Southgate, Senju, & Csibra, 2007). However, language (“Bob thinks this but Emily thinks that”) makes it possible to convey disagreement between agents without requiring any knowledge about their informational access (e.g., what each agent can see or hear). Disagreement expressed through language thus poses a direct threat to naïve realism. The observer is confronted with two different representations of reality without any obvious grounds for this difference; the learner must herself infer that the disagreement might be explained by access to different sources of evidence. Indeed, research suggests that conversation and, in particular, the practice of engaging multiple perspectives, improves children’s ability to represent competing mental states (see Harris, de Rosnay, & Pons, 2005, for review). Framing competing hypotheses as contrastive beliefs might thus support increased attention to the particular facts that license competing claims.

In addition to the role of pragmatic enrichment, there is some reason to speculate that the syntax, as well as the semantics, of disagreement might promote children's attention to evidence. Many empirical studies suggest that mastery of linguistic complement structures predicts children's performance on false belief tasks (de Villiers, Burns, & Zurer Pearson, 2003; de Villiers & Pyers, 2001, 2002; de Villiers, 2005; Hale & Tager-Flusberg, 2003; Lohmann & Tomasello, 2003; Pyers, 2003; Schick, de Villiers, de Villiers, & Hoffmeister, 2007). A hypothesis expressed with a mental state complement (e.g., "Bob thinks it's the height that matters") requires the hearer to evaluate the act of representation ("Does this accurately reflect Bob's beliefs?") independently of the truth-value of the represented content ("Does the height matter?"). That is, a listener who understands complement structures must be able to recognize that a true sentence can contain false claims. All the children in the current study were old enough to understand complement structures. Thus children might have performed better when competing hypotheses were presented as contrastive beliefs (rather than contrastive possibilities) because the complement structure of mental state verbs itself set up the possibility that the embedded claim may be false. This may have helped draw children's attention to the evidence for the claim.

At present, these ideas are preliminary; future research is needed to understand precisely how framing hypotheses as beliefs might affect learners' relationship to evidence. For now, we have joked that if a Martian were to read our journal papers, she might be struck not so much by humans' scientific insight as by our habit of following every claim with proper nouns: parenthetical references to the individuals who advanced it. It would be interesting if this apparently arbitrary social convention actually helped make us better scientists.⁴

STUDY 3: DEMONSTRATING CAUSAL RELATIONSHIPS

Many causal relationships are demonstrated by action: we can teach children how something works by pushing a button or flipping a switch. In Study 3, we look at the arrow in our cartoon between action and evidence (Figure 1), and consider how what we say in the course of functional actions might affect children's causal learning.

We alluded earlier to the gap between what we might loosely characterize as the Piagetian view of learning (the individual child exploring her environment) and the Vygotskeyan view (the child learning in a social context, from helpful, informative others). Strikingly, however, relatively little research has looked at how these two forms of learning interact. In recent work, we have begun to look specifically at how instruction affects exploration (Bonawitz et al., in press).

Imagine a toy with four functions (see Figure 6). Suppose you want to teach the child how the toy works. What actions should you take? In considering this question, we have been motivated by recent computational research (Shafto & Goodman, 2008) advancing a rational analysis of pedagogy. The analysis assumes that a learner will update her belief in a hypothesis given new data and that a teacher will choose data likely to increase the learner's belief in the correct

⁴Mercier & Sperber (in press) have recently turned this claim on its head. Rather than assuming that arguments from evidence support accurate reasoning, they have suggested that explicit reasoning, in the form of arguments from evidence, serves as a costly signal — advertising to conversational partners the contestants' epistemic reliability. We think this is an interesting possibility. Note however, that if explicit reasoning did not actually improve epistemic reliability it is not clear that it would function effectively as a reliable signal.

hypothesis (Shafto & Goodman, 2008). In this case, for instance, the analysis predicts that a knowledgeable, helpful teacher should demonstrate all four functions because a learner is more likely to infer the existence of four functions from evidence for four functions than from evidence for three. Similarly, given evidence for four functions, a learner should infer the existence of four functions (rather than five or more) because a rational teacher would have been unlikely to demonstrate four functions if five or more were present.

Thus pedagogically generated evidence not only provides evidence for demonstrated properties, it also suggests the nonexistence of additional properties; that is, in pedagogical contexts, absence of evidence for additional functions is strong evidence for their absence. Moreover, this inference is specific to pedagogical contexts. If an explicitly naïve adult demonstrates the functions or if the learner herself discovers the functions, there is no reason for the learner to infer that other functions are not present as well. Thus evidence generated pedagogically and nonpedagogically provides equivalent support for the existence of demonstrated functions but pedagogically generated evidence provides stronger evidence for the absence of additional functions.

This account suggests a trade-off between instruction and exploration. If a knowledgeable teacher shows a child a single function of a toy, the child should assume there is relatively little else to learn. She should thus engage in relatively little exploration and be unlikely to discover other functions of the toy. By contrast, in nonpedagogical contexts children should take all of the toy's affordances into account and explore broadly.

To test this hypothesis we randomly assigned preschoolers (mean: 4 years, 10 months) to a Pedagogical condition and each of three non-Pedagogical conditions: a condition where the pedagogical instruction was interrupted (Interrupted), a condition where the “teacher” was explicitly ignorant about the toy (Naïve), and a baseline condition. In all conditions, children were introduced to a novel toy (see Figure 6). In all but the baseline condition, children were shown a single function of the toy (that pulling the tube made the toy squeak) but the language that accompanied the demonstration varied (see below). The toy had three additional functions (a strip of tape that played music when pressed, a hidden button that made a light turn on, and a reversing mirror inside the black tubes) that were never demonstrated.




FIGURE 6 Novel toy used in Study 3 (Bonawitz et al., in press). (Color figure available online.)

In the Pedagogical condition, the experimenter said, "Look at my toy! This is my toy. I'm going to show you how my toy works. Watch this!" The experimenter pulled the tube so that it squeaked. She then repeated the demonstration. The Interrupted condition was exactly like the Pedagogical condition, except that the experimenter interrupted herself immediately after the second demonstration. ("I just realized I have to stop because I forgot to write down something over there. I have to go take care of it right now!") Although the Pedagogical and Interrupted conditions were otherwise identical, the interruption explains the absence of additional evidence and should disrupt the inference that the demonstrated property is the only property. In the Naive condition, the experimenter acted as if she had never before seen the toy. She said, "Look at this toy. I just found this toy." As she set the toy on the table, she squeaked the tube by chance and said, "Huh! Did you see that? Let me try to do that!" and performed the action. Finally, in the Baseline condition, children were introduced to the toy but not shown any of its functions. In all conditions the children were then left to explore freely for as long as they liked. The experiment was terminated when a child said she was finished or stopped interacting with the toy for 5 consecutive seconds.

To analyze the results, we coded children's total time playing, the number of unique actions children performed (e.g., two pulls on the tube were counted as one action; one pull on the tube and one activation of the music were counted as two actions), the proportion of children's play time spent only on the squeaking tube (excluding the Baseline condition where the squeaking tube was not demonstrated), and the total number of functions discovered in the course of free play.

We performed planned linear contrasts throughout, formalizing the prediction that the Pedagogical condition would differ from the three nonpedagogical conditions (and that the nonpedagogical conditions would not differ from each other). The four conditions were thus assigned weights of 3, -1, -1, and -1, respectively.

All the linear contrasts were significant. Children in the Pedagogical condition played with the toy for significantly less time ($M = 119.2$ s) than children in the Interrupted ($M = 179.6$ s), Naïve ($M = 132.7$ s), or Baseline ($M = 205.7$ s) conditions ($F(1,81) = 4.52$, $p < 0.05$). Children in the Pedagogical condition also performed fewer different kinds of actions on the toy ($M = 4.00$) than children in the Interrupted ($M = 5.30$), Naïve ($M = 5.90$) or Baseline ($M = 6.15$ s) conditions ($F(1,81) = 9.39$, $p < 0.01$). Children in the Pedagogical condition spent more of their play time with the squeaker ($M = 68\%$) than children in the other conditions where it was demonstrated: Interrupted ($M = 53\%$) or Naïve ($M = 38\%$) conditions ($F(1, 62) = 13.91$, $p < 0.001$). Finally, children in the Pedagogical condition discovered fewer of the other target functions ($M = 0.72$) than children in the Interrupted ($M = 1.3$), Naïve ($M = 1.2$), or Baseline ($M = 1.15$) conditions ($F(1,81) = 4.58$, $p < 0.05$).

These results suggest that pedagogical instruction constrains exploration and that well before children begin formal schooling, children are sensitive to pedagogical contexts. Four- and five-year-olds taught a single function of a toy assumed that the demonstrated function was the only function and thus failed to discover all the other things the toy could do. We predicted this constraint as the result of the rational inductive inference that teachers select samples of evidence in proportion to the probability that a learner will infer the target hypothesis from those data. Because the learner assumes that if more functions had been present, the teacher would have demonstrated them, evidence for a single function is treated as evidence against additional functions.

Of course, for the purpose of the experiment, the experimenter demonstrated only the squeaking function although the toy could also light up, make music, and show mirror images. A truly helpful and knowledgeable teacher would have shown the children all four functions. If she had, the inference that no additional functions were present would be not only rational but also correct: there would have been little else for the children to discover. Critically, however, even an expert teacher cannot know what currently unknown information the child might discover in the absence of instruction. In the infamous but apt dictum of former Secretary of Defense Donald Rumsfeld, there are “unknown unknowns — the ones we don’t know we don’t know.” Because unexpected properties might always be present, the inductive trade-off introduced by instruction is a general one. Pedagogical communication is effective precisely because it constrains the hypothesis space; this increases the probability that the learner will learn the instructed information but necessarily makes it less likely that the learner will discover anything else.

With respect to the connection between pedagogical instruction and language, we were struck by the subtlety of the distinction between pedagogical and nonpedagogical contexts. The difference between the Pedagogical and Interrupted conditions rested on whether the teacher was or was not acting without external constraints; the difference between the Pedagogical and Naïve conditions rested on whether the teacher knew or did not know how the toy worked. Other distinctions between pedagogical and nonpedagogical contexts may be subtler still. Agents can act intentionally without acting pedagogically: an adult might make a tube squeak simply because she likes the squeaking sound, not because she is trying to show a child how the toy works. If an adult is sampling evidence in proportion to her own preferences, not in proportion to the probability that it will induce any given hypothesis in the learner, demonstration of a single property should not lead the learner to infer the absence of other properties. Thus the difference between pedagogical and intentional conditions may be very subtle indeed: is she doing this for her own sake or for mine?

Given the potential ambiguity among these contexts, we believe that linguistic cues may be the most efficient, if not the only mechanism, by which a pedagogical context can be uniquely specified. As other researchers have noted (e.g., Csibra & Gergeley, 2009), cues as understated as simply saying the child’s name before a demonstration may impact whether the child thinks a pedagogical (rather than, for instance, intentional) context has been distinctively cued (see Bonawitz et al., in press, for discussion). Given that information communicated by pedagogical sampling licenses, for better and for worse, stronger inferences than identical evidence sampled by other processes, accurate cueing of these contexts may be critical and language may offer the best vehicle for this precision. From this perspective, it may be no coincidence that humans are both the only animals that speak and the only animals that teach (though see Thornton & McAuliffe, 2006).

CONCLUSION

At the start of this paper, we presented a cartoon depiction (Figure 1) of the relationship between induction and exploration, suggesting a bootstrapping mechanism by which the development of children’s theories about the world results from the interaction between their prior beliefs, the actions in which they engage, and the new evidence produced by these actions. The studies we have reviewed in the current paper suggest that linguistic communication can influence all three

arrows in our cartoon. Study 1 showed that the ways in which evidence is described (using causal language or not) influences how children interpret predictive relationships. Only children who heard causal language spontaneously intervened to make the event occur. Study 2 showed the impact of different descriptions of beliefs on children's reasoning. Children who heard competing hypotheses described as contrastive beliefs were more likely to generate evidence sufficient to support causal claims than children who heard the hypotheses described neutrally. Finally, Study 3 suggested that different descriptions of functional actions (e.g., as actions performed by a knowledgeable, freely acting teacher, as interrupted actions, or as random actions by a naïve adult) affect children's interpretation of subsequent data and the extent to which they engaged in additional exploration. Collectively, these studies suggest that the communicative context has a pervasive effect on children's exploration and learning whether what is described is the evidence itself, competing beliefs about the evidence, or the nature of the actions taken to demonstrate the evidence.

Thus, although the idea that linguistic communication transmits information is not a controversial claim, neither is it a trivial one. Certainly we rely on language to tell us things that we have not observed ourselves, ranging from the contents of our friend's dream to information about the buoyancy of porcupines. However, the current studies suggest that language serves as a unique source of evidence even about immediately perceptible events, providing information over and above what can be learned from observation of the events themselves and intervention on those events. Children's representations of otherwise equivalent events can be significantly influenced by the linguistic context in which they are presented. As research continues to bridge the gap between learning and cognition in the individual child and information provided by the social communicative context, we look forward to what future studies will tell us.

ACKNOWLEDGMENT

The authors thank Elizabeth Bonawitz for feedback on this article. This work was supported by the National Science Foundation, the John Templeton Foundation, and the James S. McDonnell Foundation.

REFERENCES

- Ammon, M. (1980). Development in the linguistic expression of causal relations: Comprehension of features of lexical and periphrastic causatives. Unpublished dissertation, University of California, Berkeley, CA.
- Amsterlaw, J., & Wellman, H. (2006). Theories of mind in transition: A microgenetic study of the development of false belief understanding. *Journal of Cognition & Development, 7*, 139–172.
- Asch, S. (1952). *Social psychology*. New York, NY: Prentice-Hall.
- Baillargeon, R., Scott, R., & He, Z. (2010). False-belief understanding in infants. *Trends in Cognitive Sciences, 14*, 110–118.
- Baum, L., Danovitch, J., & Keil, F. (2006). Children's sensitivity to circular explanations. *Journal of Experimental Psychology, 100*, 146–155.
- Behrend, D. A. (1990). The development of verb concepts: Children's use of verbs to label familiar and novel events. *Child Development, 61*, 681–696.
- Berlyne, D. E. (1954). A theory of human curiosity. *British Journal of Psychology, 45*, 180–191.

- Birch, S., Vauthier, S., & Bloom, P. (2004). Three- and four-year-olds spontaneously use others' past performance to guide their learning. *Cognition*, *107*, 1018–1034.
- Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, *311*, 1020–1022.
- Bonawitz, E., Fischer, A., & Schulz, E. (in press). Teaching three-and-a-half year olds to reason about ambiguous evidence. *Journal of Cognition and Development*.
- Bonawitz, E., Ferranti, D., Horowitz, A., & Schulz, L.E. (2009). The block makes it go: Causal language helps toddlers integrate prediction, action, and expectations about contact relations. *30th Annual Proceedings of the Cognitive Science Society*.
- Bonawitz, E., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A., Woodward, J., & Schulz, L. (2010). Just do it? Investigating the gap between prediction and action in toddlers' causal inference. *Cognition*, *115*, 104–117.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N., Spelke, E., & Schulz, L. E. (in press). The double-edged sword of pedagogy: Teaching limits children's spontaneous exploration and discovery. *Cognition*.
- Bowerman, M. (1982). Evaluating competing linguistic models with language acquisition data: Implications of developmental errors with causative verbs. *Quaderni di Semantica*, *3*, 5–66.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. F. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York, NY: Academic Press.
- Bunger, A. (2008). How we learn to talk about events: Linguistic and conceptual constraints on verb learning. *Language Acquisition*, *15*, 69–71.
- Callanan, M., & Oakes, L. (1992). Preschoolers' questions and parents' explanation: Causal thinking in everyday activity. *Cognitive Development*, *7*, 213–233.
- Carey, S. (2009). *The origin of concepts*. New York, NY: Oxford University Press.
- Casasola, M., & Cohen, L. B. (2000). Infants' associations of linguistic labels with causal actions. *Developmental Psychology*, *36*(2), 155–168.
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, *70*, 1098–1120.
- Chi, M., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145–182.
- Chi, M., de Leeuw, N., Chiu, M., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, *18*, 439–477.
- Chouinard, M. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development*, *72*.
- Cohen, L., & Oakes, L. (1993). How infants perceive a simple causal event. *Developmental Psychology*, *29*, 421–433.
- Cook, C., Bonawitz, E., & Schulz, L. (2010). *Inference in a social context: What social and nonsocial reasoning have to teach each other*. Paper presented at the biennial meeting of the Cognitive Development Society, San Antonio, TX.
- Cook, C., Goodman, N., & Schulz, L. (in press). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*.
- Corriveau, K., & Harris, P. (2009). Choosing your informant: Weighing familiarity and recent accuracy. *Developmental Science*, *12*, 426–437.
- Crowley, K., & Siegler, R. (1999). Explanation and generalization in young children's strategy learning. *Child Development*, *70*, 304–316.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Science*, *13*, 148–153.
- Cycowicz, Y., Friedman, D., Snodgrass, J., & Duff, M. (2001). Recognition and source memory for pictures in children and adults. *Neuropsychologia*, *39*, 255–267.
- de Villiers, P. (2005). The role of language in theory of mind development: What deaf children tell us. In J. Astington & J. Baird (Eds.), *Why language matters for theory of mind*. New York: Oxford University Press.
- de Villiers, J., & Pyers, J. (2002). Complements to cognition: A longitudinal study of the relationship between complex syntax and false-belief-understanding. *Cognitive Development*, *17*, 1037–1060.
- de Villiers, P., Burns, F., & Zurer Pearson, B. (2003). The role of language in the theory of mind development of language-impaired children: Complementing theories. *BUCLD 27: Proceedings of the 27th Boston University Conference on Language Development*. Somerville, MA: Cascadilla Press.
- de Villiers, P., & Pyers, J. (2001). Complementation and false belief representation. In M. Almgren, A. Barrena, M.-J. Ezeizabarrena, I. Idiazabal, & B. MacWhinney (Eds.), *Research on Child Language Acquisition: Proceedings of the 8th Conference of the International Association for the Study of Child Language*. Somerville, MA: Cascadilla Press.

- Dunbar, K., & Klahr, D. (1989). Developmental differences in scientific discovery strategies. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon. Proceedings of the 21st Carnegie Mellon Symposium on Cognition* (pp. 109–144). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Dunn, J., & Brown, J. (1993). Early conversations about causality: Contents, pragmatics, and developmental change. *British Journal of Developmental Psychology, 11*, 107–123.
- Estes, W. (1948). Discriminative conditioning. II. Effects of a Pavlovian conditioned stimulus upon a subsequently established operant response. *Journal of Experimental Psychology, 38*, 173–177.
- Fausey, C., & Boroditsky, L. (2009). *How preschoolers integrate linguistic and nonlinguistic information when reasoning about causal agents*. Paper presented at the biennial meeting of the Society for Research in Child Development, Denver, CO.
- Ferguson, S., Hashtroudi, S., & Johnson, M. (1992). Age differences in using source-relevant cues. *Psychology and Aging, 7*, 443–452.
- Fisher, C. (1996). Structural limits on verb mapping: The role of analogy in children's interpretations of sentences. *Cognitive Psychology, 31*, 41–81.
- Fisher, C. (2002). Structural limits on verb mapping: The role of abstract structure in 2.5-year-olds' interpretations of novel verbs. *Developmental Science, 5*, 55–64.
- Frazier, B., Gelman, S., & Wellman, H. (2009). Preschoolers' search for explanatory information within adult-child conversation. *Child Development, 80*, 1592–1611.
- Gentner, D. (1978). On relational meaning: The acquisition of verb meaning. *Child Development, 49*, 988–998.
- Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature, 415*, 755.
- Gopnik, A., & Astington, J. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development, 59*, 26–37.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*, 1–31.
- Gopnik, A., & Nazzi, T. (2003). Words, kinds, and causal powers: A theory perspective on early naming and categorization. In D. Rakison & L. Oakes (Eds.), *Early category and concept development: Making sense of the blooming, buzzing confusion* (Vol. xxi, pp. 303–329). New York, NY: Oxford University Press.
- Gopnik, A., & Sobel, D. M. (2000). Detectingblickets: How young children use information about causal properties in categorization and induction. *Child Development, 71*, 1205–1222.
- Hale, C., & Tager-Flusberg, H. (2003). Social communication in children with autism: The relationship between theory of mind and discourse development. *Autism, 9*, 157–178.
- Hare, B., Call, J., Agnetta, B., & Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour, 59*, 771–786.
- Harris, P. L. (2002). What do children learn from testimony? In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis of science* (pp. 316–334). Cambridge, England: Cambridge University Press.
- Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child Development, 77*, 505–524.
- Harris, P. L., de Rosnay, M., & Pons, F. (2005). Language and children's understanding of mental states. *Current Directions in Psychological Science, 14*, 69–73.
- Hickling, A., & Wellman, H. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology, 37*, 668–683.
- Hitchcock, Christopher, "Probabilistic Causation", *The Stanford Encyclopedia of Philosophy* (Fall 1997 Edition), Edward N. Zalta (Ed.). Available at: <http://plato.stanford.edu/archives/fall1997/entries/causation-probabilistic/>
- Hood, L., & Bloom, L. (1979). What, when, and how about why: A longitudinal study of early expressions of causality. *Monographs for the Society for Research in Child Development, 44*.
- Inagaki, K., & Hatano, G. (1993). Young children's understanding of the mind-body distinction. *Child Development, 64*, 1534–1549.
- Inagaki, K., & Hatano, G. (2002). *Young children's thinking about the biological world*. New York, NY: Psychology Press.
- Inhelder, B., & Piaget, J. (1958). *The growth of logical thinking from childhood to adolescence*. New York, NY: Basic Books.
- Jaswal, V., & Neely, L. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science, 17*, 757–758.

- Johnson, M., Hashtroudi, S., & Lindsay, D. (1993). Source monitoring. *Psychological Bulletin*, *114*, 3–28.
- Kelemen, D., Callanan, M., Casler, K., & Perez-Granados, D. (2005). Why things happen: Teleological explanation in parent-child conversations. *Developmental Psychology*, *41*, 251–264.
- 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.
- Koenig, M., Clement, F., & Harris, P. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, *10*, 694–698.
- Koenig, M., & Harris, P. (2005). The role of social cognition in early trust. *Trends in Cognitive Science*, *9*, 457–459.
- Koslowski, B. (1996). *Theory and evidence: The development of scientific reasoning*. Cambridge, MA: MIT Press.
- Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, *96*, 674–689.
- Kuhn, D., Amsel, E., & O'Laughlin, M. (1988). *The development of scientific thinking skills*. Orlando, FL: Academic Press.
- Kushnir, T., & Gopnik, A. (2005). Children infer causal strength from probabilities and interventions. *Psychological Science*, *16*, 678–683.
- Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, *44*, 186–196.
- Lagattuta, K., & Wellman, H. (2001). Thinking about the past: Early knowledge about links between prior experience, thinking, and emotion. *Child Development*, *72*, 82–102.
- Legare, C. (in press). Exploring explanation: Explaining inconsistent information guides hypothesis-testing behavior in young children. *Child Development*.
- Legare, C., Gelman, S., & Wellman, H. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child Development*, *81*, 929–944.
- Legare, C., Wellman, H., & Gelman, S. (2009). Evidence for an explanation advantage in naïve biological reasoning. *Cognitive Psychology*, *58*, 177–194.
- Leising, K. J., Wong, J., Waldmann, M. R., & Blaisdell, A. P. (2008). The special status of actions in causal reasoning in rats. *Journal of Experimental Psychology: General*, *137*, 514–527.
- Leslie, A., & Keeble, S. (1987). Do six-month-old infants perceive causality? *Cognition*, *25*, 265–288.
- Lewis, D. (1973). Causation. *Philosophical Papers*, *2*, 159–213.
- Lohmann, H., & Tomasello, M. (2003). The role of language in the development of false belief understanding: A training study. *Child Development*, *74*, 1130–1144.
- Lombrozo, T. (2009). Explanation and categorization: How “why?” informs “what?”. *Cognition*, *110*, 248–253.
- McIntyre, J., & Craik, F. (1987). Age differences in memory for item and source information. *Canadian Journal of Psychology*, *31*, 175–192.
- Meltzoff, A. (2007). Infants' causal learning: Intervention, observation, imitation. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 37–47). Oxford, England: Oxford University Press.
- Mercier, H. & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, *34*, 57–111.
- Muentener, P., Bonawitz, E., Horowitz, A., & Schulz, E. (2010). *From prediction to expectations of contact causality: Causal language and dispositional agency facilitate toddlers' causal representations*. Poster presented at the 2nd Annual Workshop on Cognition and Evolution, Rovereto, Italy.
- Muentener, P. (2009). The origin and development of causal reasoning. (Unpublished doctoral dissertation). Harvard University, Cambridge, MA.
- Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology*, *61*, 63–86.
- Muentener, P., & Lakusta, L. (2011). The intention-to-CAUSE bias: Evidence from children's causal language. *Cognition*, *119*, 341–355.
- Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, *17*, 357–374.
- Naigles, L. (1996). The use of multiple frames in verb learning via syntactic bootstrapping. *Cognition*, *58*(2), 221–251.
- Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science*, *3*, 389–396.
- Nazzi, T., & Gopnik, A. (2001). Linguistic and cognitive abilities in infancy: when does language become a tool for categorization? *Cognition*, *80*, B11–B20.
- Pearl, J. (2000). *Causality, models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.

- Piaget, J. (1951). *Play, dreams and imitation in childhood* (C. Gattegno & F. M. Hodgson, Trans.). New York, NY: W. W. Norton.
- Piaget, J. (1954). *The construction of reality in the child*. New York, NY: Basic Books.
- Pine, K., & Siegler, R. (2003). *The role of explanatory activity in increasing the generality of thinking*. Paper presented at the biennial meeting of the Society for Research in Child Development, Tampa, FL.
- Pyers, J. (2003). The expression of mental states in American Sign Language. In A. Baker, B. van den Bogaerde, & O. Crasborn (Eds.), *Crosslinguistic perspectives on sign language research*. Hamburg, Germany: Signum Press.
- Rahhal, T., May, C., & Hasher, L. (2002). Truth and character: Sources that older adults can remember. *Psychological Science, 13*, 101–105.
- Robinson, E. J., & Whitcombe, E. L. (2003). Children's suggestibility in relation to their understanding about sources of knowledge. *Child Development, 74*, 48–62.
- Rovee-Collier, C. (1987). Learning and memory in infancy. In J. D. Osofsky (Ed.), *Handbook of infant development* (2nd ed., pp. 98–148). New York, NY: Wiley.
- Rumsfeld, D. (2002). Donald Rumsfeld press conference as US Secretary of Defense. [WWW Document]. Available at: <http://www.defense.gov/transcripts/transcript.aspx?transcriptid=2636>
- Salmon, W. (1998). *Causality and explanation*. Oxford, England: Oxford University Press.
- Schacter, D., Kaszniak, A., Kihlstrom, J., & Valdiserri, M. (1991). The relation between source memory and aging. *Psychology & Aging, 6*, 559–568.
- Schick, B., de Villiers, P., de Villiers, J. G., & Hoffmeister, R. (2007). Language and false belief reasoning: A study of deaf children. *Child Development, 78*(2), 376–398.
- Schult, C., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition, 62*, 291–324.
- Schulz, L., & Bonawitz, E. B. (2007). Serious fun: Preschoolers play more when evidence is confounded. *Developmental Psychology, 43*, 1045–1050.
- Schulz, L., 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., Goodman, N., Tenenbaum, J., & Jenkins, A. (2008). Going beyond the evidence: Preschoolers' inferences about abstract laws and anomalous data. *Cognition, 109*(2), 211–223.
- Schulz, L., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology, 40*, 162–176.
- Schulz, L., Hoopel, K., & Jenkins, A. (2008). Judicious imitation: Young children imitate deterministic actions exactly, stochastic actions more variably. *Child Development, 79*, 395–410.
- Schulz, L., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers' causal inferences. *Child Development, 77*, 427–442.
- Schulz, L., Standing, H., & Bonawitz, E. (2008). Word, thought, and deed: The role of object labels in children's inductive inferences and exploratory play. *Developmental Psychology, 44*, 1266–1276.
- Scofield, J., & Behrend, D. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development, 23*, 278–290.
- Shafto, P., & Goodman, N. (2008). Teaching games: statistical sampling assumptions for learning in pedagogical situations. In *Proceedings of the Thirtieth Annual Conference of the Cognitive Science Society* (pp. 1632–1637). Austin, TX: Cognitive Science Society.
- Shultz, T. (1982). Rules of causal attribution. *Monographs of the Society for Research in Child Development, 47*.
- Siegler, R. (1995). How does change occur: A microgenetic study of number conservation. *Cognitive Development, 28*, 225–273.
- Simons, J., Dodson, C., Bell, D., & Schacter, D. (2004). Specific and partial source memory: Effects of aging. *Psychology and Aging, 19*, 689–694.
- Sobel, D. (2004). Exploring the coherence of young children's explanatory abilities: Evidence from generating counterfactuals. *British Journal of Developmental Psychology, 22*, 37–58.
- Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42*, 1103–1115.
- 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.
- Southgate, V., Senju, A., & Csibra, G. (2007). Action anticipation through attribution of false belief by two-year-olds. *Psychological Science, 18*, 587–592.

- Spelke, E. (2003). What makes us smart? Core knowledge and natural language. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the investigation of language and thought*. Cambridge, MA: MIT Press.
- Spelke, E., Breinlinger, K., Macomber, J., & Jacobsen, K. (1992). Origins of knowledge. *Psychological Review*, 99, 605–632.
- Spelke, E., Lee, S. A., & Izard, V. (2010). Beyond core knowledge: Natural geometry. *Cognitive Science*.
- Spencer, W., & Raz, N. (1995). Effects of aging on memory for content and context: A meta-analysis. *Psychology and Aging*, 10, 527–539.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. New York, NY: Springer.
- Suppes, P. (1970). *A probabilistic theory of causality*. Amsterdam: North Holland.
- Taylor, M., Esbensen, B., & Bennett, R. (1994). Children's understanding of knowledge acquisition: The tendency for children to report that they have always known what they have just learned. *Child Development*, 65, 1581–1604.
- Thorton, A., & McAuliffe, K. (2006). Teaching in wild meerkats. *Science*, 313, 227–229.
- Tomasello, M., & Call, J. (1997). *Primate cognition*. New York, NY: Oxford University Press.
- Trabasso, T., & Suh, S. (1993). Understanding text: Achieving explanatory coherence through online inferences and mental operations in working memory. *Discourse Processes*, 16, 3–34.
- Vygotsky, L. (1934/1962). *Thought and language*. Cambridge, MA: MIT Press.
- Watson, J. S., & Ramey, C. T. (1987). Reactions to response-contingent stimulation in early infancy. In J. Oates & S. Sheldon (Eds.), *Cognitive development in infancy*. Hove, England: Erlbaum.
- Wegesin, D., Jacobs, D., Zubin, N., Ventura, P., & Stern, Y. (2000). Source memory and encoding strategy in normal aging. *Journal of Clinical and Experimental Neuropsychology*, 22, 455–464.
- Wellman, H., & Liu, D. (2007). Causal reasoning as informed by the early development of explanations. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy, computation* (pp. 261–279). New York, NY: Oxford University Press.
- Williams, J., & Lombrozo, T. (2010). The role of explanation in discovery and generalization: Evidence from category learning. *Cognitive Science*, 34, 776–806.
- Williamson, R., Meltzoff, A., & Markman, E. (2008). Prior experiences and perceived efficacy influence 3-year-olds' imitation. *Developmental Psychology*, 44, 275–285.
- Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition*, 88, 1–48.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford, England: Oxford University Press.