

Is it me or the world? 16-month-olds distinguish competing hypotheses about the cause of failed interventions

Hyowon Gweon (hyora@mit.edu) and Laura E. Schulz (lschulz@mit.edu)

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 021 USA

Abstract

When an agent fails to make an object function properly, there are two possibilities: the agent did something wrong or something is wrong with the object. As in all problems of confounding, these hypotheses can be disambiguated by varying one factor and holding the other constant: in this case, either by holding the object constant and varying the agent (e.g., by asking for help from others) or by holding the agent constant and varying the object (e.g., by trying another object). Here we show that 16-month-old infants engage in distinct patterns of behavior depending on the relative probability of the competing hypotheses: they ask for help more often when they (rather than the object) are the probable cause of failure; they reach for a new object more often when the object (rather than themselves) is the probable cause of failure.

Keywords: infants, confounding, exploratory behavior, ambiguous evidence, asking for help.

Imagine that you are trying to get into a new office but the key doesn't work. You jiggle it for a minute and then assess your options. What you choose to do next depends on what hypothesis you think is most probable: if you think you are having trouble positioning the key, you might ask a friend for help; if you think you picked up the wrong key, you will probably try a different one.

As intentional agents, we frequently plan and carry out goal-directed actions. Most of the time, these actions are successful. However, when we experience failure, we can experience not only the frustration of our intentions but also a problem of confounded evidence. Did we do something wrong or was something wrong in the world?

This problem of "me or the world" is perhaps the most common example of confounded causal variables that we encounter in everyday life. Of course, the variables "me" and "the world" can sometimes be more precisely defined. In the key example for instance, if you are the problem, you might have put the key in upside down, turned it in the wrong direction, or lack fine motor coordination. Alternatively, if you believe the problem lies in the world, the door might be jammed, or the lock might have been changed. However, these distinctions are subordinate to the primary problem of discovering whether you or the world is the culprit. When things go wrong, how do we identify the locus of failure?

As in any causally confounded situation, changing one variable at a time can disambiguate the evidence. Assuming that changing either variable is possible and equally costly,

a rational agent who wants to generate the effect should change the variable that seems most likely to be the source of the failure. If I think I'm the problem, I should hold the object constant and vary the agent (e.g., ask my friend to help); if I think the object poses the problem, I should hold myself constant and vary the object (e.g., try a new key).

There are many reasons to believe that recognizing, let alone solving, problems of confounding between the self and the world might require substantial expertise. Indeed, previous research suggests that both children and adults have difficulty recognizing when information is ambiguous (Penner & Klahr, 1996) and designing experiments that could generate informative evidence (Chen & Klahr, 1999; Koslowski, 1996; Kuhn, 1989). Such studies of formal scientific reasoning however, typically involve many hypotheses, including those that conflict with the learners' prior beliefs. In contexts where there are only two competing hypotheses and both are familiar and plausible, children seem to be sensitive to confounding at a much younger age (Gweon & Schulz, 2008; Kushnir & Gopnik, 2005; Schulz & Bonawitz, 2007; Sodian, Zaitchik, & Carey, 1991). Thus in simple cases, even very young children might be sensitive to competing hypotheses.

While a sensitivity to confounded variables might enable young children to recognize the "me vs. the world" problem, deconfounding the self and the world requires children to understand *how* to intervene on each variable. For example, imagine a simple confounded situation where a child tries a novel toy and fails to make it work; the child needs to understand that other people ('the agent' variable), and other toys of the same kind ('the world' variable), can both serve as useful sources of information. Here we briefly review some previous studies that suggest that even very young children might be capable of such an understanding.

A large body of literature on social referencing in infancy suggests that infants readily treat their caregivers as sources of information about the emotional valence of events (Sorce, Emde, Campos, & Klinnert, 1985; Walden & Ogan, 1988) and the referent of adults' attention (Baldwin, 1993; Carpenter, Nagell, Tomasello, Butterworth, & Moore, 1998). Moreover, infants use the information to regulate their own behavior. In particular, O'Neill (1996) showed that two-year-olds will request help from a knowledgeable (but not ignorant) parent in retrieving a hidden object, suggesting that toddlers not only look to parents for the information they might provide but also actively solicit such information.

Children's imitation of object-directed actions is also often interpreted as an indication that children perceive others as agents like themselves (the 'like me' hypothesis; Meltzoff & Brooks, 2001) and use adult actions for information about how to interact with an object (Gergely, Bekkering, & Kiraly, 2002; Gopnik & Meltzoff, 1994). Notably, children are more likely to imitate an adult's goal-directed action if they themselves have previously failed to generate a target outcome than if they have succeeded (Williamson, Meltzoff, & Markman, 2008) which suggests that young children can use adult actions as evidence about the cause of their own failures, and modify their own actions accordingly.

Such studies speak to children's understanding of other agents as potential sources of information about objects in the world. What about children's understanding that one object can be informative about other members of the object kind? Previous research has shown that preschoolers generalize non-obvious properties (like squeaking or magnetism) from one member of a kind to others (Gopnik & Sobel, 2000; Nazzi & Gopnik, 2000). Moreover, children maintain this expectation even when one exemplar fails to function as expected (Schulz, Standing, & Bonawitz, 2008). Indeed, 9-month-old infants can generalize a property of an object to other identical-looking object after a single exposure (Baldwin, Markman, & Melartin, 1993), and by 15 months infants can even integrate information about how the exemplars are sampled in their inferences about object properties (Gweon, Tenenbaum, & Schulz, 2010). These studies establish that children expect object properties to generalize across similar-looking objects, maintain that expectation even when they themselves fail to elicit the expected property, and, having experienced failure, can both solicit help from caregivers and act on other similar objects. This study however, is the first to investigate the possibility that infants might be sensitive to competing hypotheses for why their actions fail, and might rationally trade-off actions directed towards agents and actions directed towards objects. Here we ask whether 16-month-old infants implicitly recognize the ambiguity in a failed attempt to activate a toy. We predict that children should be more likely to ask for help when they themselves are the probable source of failure and more likely to test another toy when the probable cause of failure lies in the toy itself.

Experiment

In the current study, we introduce infants to three identical-looking toys, differing only in color (Green, Yellow, and Red). The experimenter shows the child that she can push a button on the Green toy and the toy will make music. In the Agent condition, the experimenter then hands the child the Green toy; in the Object condition, the experimenter hands the child the Yellow toy. All the children are allowed to try to activate their toy. However, because the Green toy is actually activated by a hidden switch and the Yellow toy is inert, the toys never activate for the infants.

The condition manipulation is designed to affect the relative probability of the two hypotheses for why the toy fails to activate. In the Agent condition, the hypothesis that the toy is broken is relatively improbable given that the toy had just worked moments before; the hypothesis that the child herself is doing something wrong should seem more probable. By contrast, in the Object condition, where the child's toy has never activated, the hypothesis that the toy doesn't work should seem more probable than the hypothesis that the child herself is doing something wrong (given that the button is conspicuous and easy to press). Thus, we created a situation in which infants might differentially weigh the two hypotheses about the cause of the failure. In both cases however, infants had identical sources of information that they could use to resolve the ambiguity. All children were seated next to their parents. By turning and asking their parent for help with the toy they had, they could test the 'agent' variable. All children could also reach for the Red toy, which sat on the end of a piece of felt cloth. By pointing to the Red toy or pulling the piece of felt cloth they could try to retrieve a toy of the same kind and test 'the object' variable. (We placed the Red toy at a distance to ensure that all infants would initially attend to the toy they were given.) Because previous research indicates that infants reliably understand the intentional structure of action in a cloth-pulling sequence by the age of 12 months (e.g., Sommerville & Woodward, 2005), we recruited infants slightly older than this age and verified in a warm-up period that they can pull the cloth to retrieve a toy.

We hypothesized that the infants' behavior would be sensitive to the relative probability of the competing hypotheses. Therefore, we predicted that infants in the Agent condition should be more likely to appeal to their parents; infants in the Object condition should be more likely to reach for the other toy.

Methods

Participants Thirty infants (mean: 16 months, 10 days; range: 14 – 20 months; 47% girls) were recruited from a local children's museum; infants were randomly assigned to an Agent condition or an Object condition ($n = 15/\text{condition}$). Six infants were replaced due to parental interference or experimental error. Two additional infants were replaced because they did not pull the cloth to retrieve a toy during the warm-up procedure. (See Procedure.) Finally, two infants (one in each condition) were excluded from analyses because they never showed any of the target behaviors. (See Results.)

Materials One commercially available toy (a plastic fish) was used during the warm-up period. Three similar-looking novel toys were built by attaching a wooden stick (10 cm in length) to a round plastic container (4 inches in diameter). The toys resembled small hand drums with handles. A square-shaped button (2 x 2 x 1 cm) was attached to the top of the container. This button was inert. Each object was covered with green, red, or yellow electrical tape and felt




Figure 1: Schematic drawing of the experimental setup and procedure. R (Red), G(Green), Y (Yellow) refer to the color of the toys.

cloth. The Green toy had a small battery-powered circuit that was operated by a hidden switch at the bottom of the container: when the toy was laid flat on a hard surface and the fake button was pressed down, the real switch depressed and the toy played a musical tune (creating the appearance that pushing the fake button activated the toy). Children sat in a highchair. The tray on the high chair was covered with white felt, creating a surface that was too soft to activate the real switch at the bottom of the Green toy. The Green toy never worked on this tray when the fake button was pressed. The Red and Yellow toys did not have a musical mechanism inside, but contained play-dough so that all three toys were matched in approximate weight.

Procedure All children were tested individually in a quiet room inside the museum. The children sat in the highchair and the parents sat next to them on a chair. (See Figure 1 for experimental setup and stimuli.) Parents were instructed not to interact with the toys and only to smile and nod if the child addressed them. They were given a brochure about the study and asked to read it during the experimental procedure. Once the child was positioned in the highchair, the experimenter put a piece of orange felt cloth (approx. 20 x 75 cm) on the table and placed the warm-up toy on one end of the cloth. She pulled the cloth towards herself and retrieved the toy. Then she encouraged the infant to pull the cloth. Infants who did not pull the cloth and retrieve the toy after two demonstrations were excluded from analysis and replaced.

The experimenter removed the warm-up toy and introduced the child to a basket containing the Green, Red, and Yellow toys. She took the Green toy out, put it on the table, and pressed the button on top of the toy to play the music. She demonstrated this three times. Then she showed the child the basket containing the other two toys. She took out the Red toy and placed it on one end of the felt cloth. The toy was approximately 70 centimeters away from the

child and was not within direct reach of the child's hands. She placed the other end of the felt cloth on the child's tray within easy reach of the child. Then, the experimenter handed the child either the Green toy (Agent condition) or the Yellow toy (Object condition) and said, "Here you go, you can go ahead and play!" She took the basket with the remaining toy (the Yellow toy in the Agent condition: the Green toy in the Object condition) out of the child's line of sight. The child's behavior was videotaped for 90 seconds (24 children) or until the child fussed-out (6 children); all but one of the infants who stopped playing before 90 seconds played for at least 60 seconds. The remaining infant was in the Object condition and played for 35 seconds. There was no difference between conditions in children's mean length of free play (Agent Condition: mean 89 seconds; Object Condition: mean 84 seconds, $p = ns$).

Results

For our preliminary analyses, we looked at whether all the children imitated the experimenter's action on the toy and whether they were equally persistent in the Agent condition (where they were given the same toy on which the action had been modeled) and the Object condition (where they had to make an inductive generalization from the Green toy to the Yellow toy). Given previous research suggesting that even 9-month-olds readily make such generalizations (Baldwin et al., 1993), we did not expect any difference in their button-pushing behavior. Indeed, all but one infant immediately (within two seconds) pressed the inert button on the toy in front of them. There was no difference in the frequency of children's button-pushing attempts in the two conditions (Agent Condition: mean 3.0 times; Object Condition: mean 3.2 times, $p = ns$).

We also used two different measures to look at whether parents differentially cued the infants to ask for help in the two conditions. Two coders, blind to hypotheses and conditions, coded from videotape; the monitor was partially covered with a cardboard occluder so that only the parent was visible. One coder was asked to make Yes/No judgments about whether the parent ever encouraged the child to ask for help. A second coder rated parents' attempts to initiate communication on a scale from 1 (no attempts to communicate) to 7 (repeated attempts to communicate). There was no difference between conditions on either measure (% of Yes judgments = 7% in both conditions; mean rating: 1.71 (Agent) vs. 1.79 (Object), $p = ns$).

The primary measure of interest was whether children's first response to failure was directed towards their parents or to the other toy. To determine this, we coded three target behaviors: Ask, Point, and Pull. The criteria for coding a behavior as Ask was that the child turned to the parent¹ and

¹ The infants could also have asked the experimenter for help. However, the experimenter stood behind the high chair during the free play period and acted busy (i.e., by writing something on a clipboard). Therefore, although there were a few cases where the infants looked as if they wanted the experimenter's attention, we did not include these attempts as one of our target behaviors.

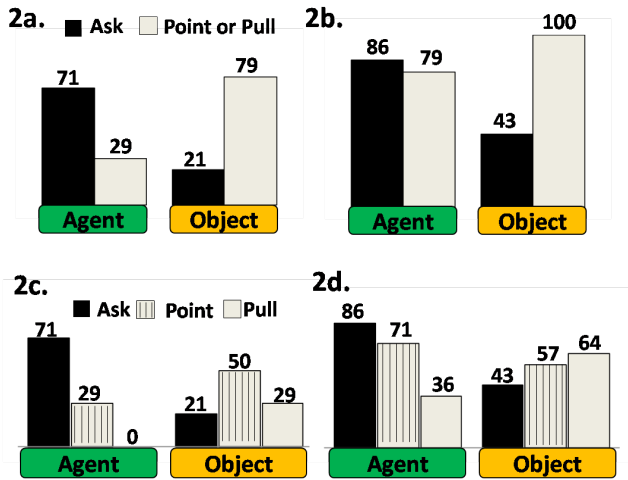


Figure 2: Experimental results. 2a. % infants Asking *first* or Pointing/Pulling *first* in each condition. 2b. % infants *ever* Asking or Pointing/Pulling in each condition (infants can perform more than one behavior, so the combined percentage within a condition can exceed 100). 2c and 2d. Same as 2a and 2b, respectively, but showing Point and Pull separately.

tried to give her the inoperative toy or grabbed the parent's hand and tried to bring it towards the inoperative toy (the Green toy in the Agent condition; the Yellow toy in the Object condition). We coded Point as a finger point to the Red toy or a direct reach for the Red toy. We coded Pull as pulling on the cloth and successfully retrieving the Red Toy. We predicted that children's first response to failure would be to Ask for help in the Agent condition but to Point or Pull in the Object condition. Additionally we coded whether infants ever showed the three target behaviors during free play, the latency to their first target behavior, and what they did with the Red toy if they retrieved it. The data for three target actions were originally coded by the experimenter, but also coded by an observer blind to hypotheses and conditions. The inter-coder reliability was high (Cohen's kappa = 92.6); using the data from the blind-coder did not change the results. All but two infants (one in each condition) performed at least one of the three target behaviors during the course of their free play; these two children were eliminated from subsequent analyses.

As predicted, children were significantly more likely to Ask First than Point or Pull first in the Agent condition than the Object condition ($\chi^2(1, N = 28) = 7.04, p < 0.01$). In the Agent condition, 10 infants (71%) Asked first and 4 infants (29%) Pointed or Pulled first; in the Object condition, 3 infants Asked first (21%); 11 infants (79%) Pointed or Pulled first. Within conditions, children were marginally more likely to Ask first than Point/Pull in the Agent condition ($p < 0.10$ by binomial test) and more likely to Point/Pull first than Ask in the Object condition ($p < 0.05$ by binomial test). See Figure 2a.

We also looked at how many infants in each condition exhibited each of the target behaviors at least once during




Figure 3: Data excluding Point behavior. 3a. % infants who Asked or Pulled *first* in each condition. 3b: % infants who *ever* Asked or Pulled during free play in each condition.

the course of their free play. Children were more likely to Ask for help over the course of their free play in the Agent condition than the Object condition ($\chi^2(1, N = 28) = 5.6, p < 0.05$, see Figure 2b.). Twelve infants (86%) in the Object condition Asked at some point; only 6 infants (43%) did so in the Object condition. Similarly, there was a trend for children to be more likely to Point to or Pull the red object in the Object condition than the Agent condition: 11 infants (79%) Pointed/Pulled at some point in the Agent condition whereas all 14 infants (100%) did so in the Object condition ($\chi^2(1, N = 28) = 3.36, p = .07$). See Figures 2c and 2d for the first target action and any instance of the target actions broken down by each of the three target behaviors.

There was no difference between conditions in the mean latency to the first target action (Agent condition, mean: 19.3 s; Object condition, mean: 25.4 s, $p = ns$). This suggests that the children in the two conditions were approximately matched in their motivation to act. There was also no difference in latency between the agent-directed and object-directed actions (Ask: 20.5 s; Point/Pull: 26.6 s, $p = ns$). This suggests that the agent-directed and object-directed actions were equivalently easy for the children to perform.

Although there was no overall latency difference, the Pulling action occurred (non-significantly) later than the Ask or Point actions (because most children in both conditions pointed before they pulled). Prima facie, Point is a less complex action than either Ask or Pull. Point required only a finger movement whereas Ask required the child to try to hand the object to the parent or to try to place the parent's hand on the object and Pull required a means/end sequence. We believe the collapsed Point/Pull measure is the correct measure of children's interest in the distal object as there is little doubt from the videotapes that infants coded as Pointing were unambiguously asking for the Red toy. However, to match for the overall complexity of the action sequence, we looked at whether infants were more likely to Ask or Pull first if the Point measure is excluded. Under this analysis, and excluding infants whose only target behavior was pointing (one child in the Agent condition; three children in the Object condition) infants were more likely to Ask than Pull in the Agent condition compared to the Object condition ($\chi^2(1, N = 24) = 11.7, p < 0.001$; see Figure 3a). Within conditions, infants in the Agent condition were more likely to Ask than Pull (12 Ask first, 1 Pull first; $p < 0.01$ by binomial test); infants in the Object condition were equally

likely to Ask and Pull (4 Ask first, 7 Pull first $p = ns$, by binomial test). Looking at any instance of Asking or Pulling over the course of free play, infants in the two conditions again tended to show different patterns of behavior ($\chi^2(1, N = 24) = 3.03, p = .08$; Agent condition: 12 Ask overall, 5 Pull overall; Object condition: 6 Ask overall, 9 Pull overall; See Figure 3b).

Finally, of those infants who pulled the cloth and successfully retrieved the Red toy (5 infants in the Agent condition; 9 agents in the Object condition), all but one immediately (within 2 seconds) pressed the button on the Red toy, suggesting that infants did indeed retrieve the toy in order to see whether they could make the toy go.

Discussion

These results suggest that when an object fails to function as expected, 16-month-old infants entertain competing hypotheses about the cause of the failure and act on the most probable hypothesis. Not only did almost every child (28 out of 30) actively try to elicit information from the available sources (another agent or another object), they selectively accessed different sources of information given different evidence about the likely cause for the failure. When the hypothesis that the agent caused the failure was more probable than the hypothesis that the toy was broken (because infants were given a toy that worked for the experimenter), the majority of infants asked their parents for help. In this condition, varying the 'agent' variable is the most effective strategy: if you're doing something wrong, doing the same thing with a different object will not solve the problem. By contrast, when the hypothesis that their toy was broken was more probable than the hypothesis that they were doing something wrong (because infants were given a similar but non-identical toy), the majority of infants reached for a new object. In this condition, trying another exemplar is the most effective strategy: if a toy is broken, asking someone else to act on the broken toy will not solve the problem. These results suggest that infants rationally trade-off help-seeking and object-exploration behaviors depending on the relative probability of the two hypotheses.

Are there alternative ways of accounting for the results? One possibility is that infants' differential behavior across conditions might reflect different affective responses to differentially frustrating situations rather than active requests for information. The manipulation was set up so that infants in the Agent condition would have a stronger expectation that their toy should work than infants in the Object condition. Arguably, infants in the Agent condition might have been more frustrated by their failure, and more likely to turn to their parents than infants in the Object condition. Conversely, infants in the Object condition arguably had a more "boring" toy than infants in the Agent condition (because they had never seen their toy activated). They thus may have been more motivated to discard it and reach for a new toy than infants in the Agent condition.

Further research is needed to definitively rule out these accounts but we believe that the current data renders both

explanations unlikely. First, differential frustration or boredom might be indicated by a difference in children's overall playtime between conditions but children played just as long in the Agent condition as the Object condition. Second, infants in the Agent condition who asked their parents for help did not show any signs of upset and did not look for comfort. They handed their parents the toy or tried to place their parents' hands on the toy but they did not cling to their parents or fuss out. Similarly, there was no indication that infants in the Object condition were more bored by the toy than infants in the Agent condition. Infants in the Object condition were just as likely to push the button on the toy as infants in the Agent condition, and they pushed the button just as persistently. Moreover, infants in the Object condition who retrieved the red toy immediately tried the button on the red toy. These behaviors suggest that infants in the Object condition expected that the toys would work and were strongly motivated to try to activate them. Thus the alternative accounts are inconsistent with how infants used the two different means: rather than reflecting frustration or boredom, infants' behavior is consistent with an attempt to generate an effective intervention. As noted however, conclusively distinguishing these possibilities requires further research. If for instance, an irrelevant distracter toy (rather than the Red toy) is placed at the end of the cloth, there should be no differences between the two conditions. We are currently running this control.

We note that the current study falls short of looking at whether children *learn* from the source of information they choose. That is, we cannot distinguish between the possibility that infants are taking the most rational steps to try to generate an outcome and the possibility that infants are (additionally) using the disambiguating evidence to determine the cause of the initial failure. In the current study, we deliberately asked the parents not to touch the toy, and the Red toy on the cloth was always inert. In future studies, we aim to look at whether learning occurs by studying infants' responses to different information that other people or other toys might provide. Imagine for instance, that if children retrieve the Red toy, it works for half the children and is inert for the other half. This evidence should give the children different information about the Yellow toy: if the Red toy works, the yellow toy is probably broken; if the Red toy does not work, it is now more probable that the child herself is the source of the failure. Thus if the Red toy is removed and the Yellow toy is returned to the children, they should be more likely to discard it if the Red toy worked, and more likely to ask for help if the Red toy failed. This would suggest that the infants' interventions not only serve the purpose of helping them make things happen but also help children disambiguate evidence to support causal learning. This research is also currently underway.

The study proposed above might also help clarify whether infants actually entertain the two competing hypotheses simultaneously, or whether only one of the two is considered in a given context. In the current study, some children who pulled the cloth and confirmed that the Red

toy did not work subsequently turned to their caregivers and asked for help. This suggests the possibility that children might indeed consider both hypotheses and act on the one favored by the evidence. Whether infants' choices of actions change dynamically given evidence that favors different hypotheses is an exciting topic for future research.

The current results however, already reveal impressive abilities in 16-month-old infants. There is abundant evidence that young children both ask adults for help (Dunham, Dunham, & O'Keefe, 2000; O'Neill, 1996) and explore objects in the world (Bonawitz, Shafto, Gweon, Spelke, & Schulz, submitted; Gweon & Schulz, 2008; Piaget, 1930; Gweon, Tenenbaum, & Schulz, 2010). This study goes beyond previous work in suggesting that infants' actively trade-off these two alternatives. Infants not only consider competing hypotheses about the failure of goal-directed actions, but also choose different means to resolve the ambiguity depending on which hypothesis is more probable. In the face of failure to achieve a goal, 16-month-old infants do not simply look to their parents nor do they simply move on to a new toy. Instead, they are able to infer the likely cause for their failure, and flexibly and rationally adjust their behavior. In solving the problem of assigning causal responsibility to themselves or the world, infants might lay the earliest foundations for scientific inquiry.

Acknowledgments

Thanks to Stephanie Tong for help with data collection. This research was supported by an NSF Faculty Early Career Development Award, a John Templeton Foundation Award, and James S. McDonnell Foundation Collaborative Interdisciplinary Grant on Causal Reasoning to L.S.

References

- Baldwin, D. A. (1993). Infants' ability to consult the speaker for clues to word reference. *Journal of Child Language*, 20(2), 395-418.
- Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants' ability to draw inferences about nonobvious object properties: evidence from exploratory play. *Child Development*, 64(3), 711-728.
- Bonawitz, E. B., Shafto, P., Gweon, H., Spelke, E. S., & Schulz, L. E. (submitted). The double-edged sword of pedagogy: Teaching limits children's spontaneous exploration and discovery.
- Carpenter, M., Nagell, K., Tomasello, M., Butterworth, G., & Moore, C. (1998). Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63(4).
- Chen, Z., & Klahr, D. (1999). All other things being equal: Acquisition and transfer of the control of variables strategy. *Child Development*, 70(5), 1098-1120.
- Dunham, P., Dunham, F., & O'Keefe, C. (2000). Two-year-olds' sensitivity to a parent's knowledge state: Mind reading or contextual cues? *British Journal of Developmental Psychology*, 18(4), 519-532.

- Gergely, G., Bekkering, H., & Kiraly, I. (2002). Rational imitation in preverbal infants. *Nature*, 415(6873), 755.
- Gopnik, A., & Meltzoff, A. N. (Eds.). (1994). *Minds, bodies, and persons: Young children's understanding of the self and others as reflected in imitation and "theory of mind" research*. New York: Cambridge University Press.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (in press). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*.
- Gweon, H., & Schulz, L. E. (2008). Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play. *Proceedings of the 30th annual meeting of the Cognitive Science Society*, 570-574.
- Koslowski, B. (1996). *Theory and evidence: The development of scientific reasoning*. Cambridge, MA: The MIT Press.
- Kuhn, D. (1989). Children and adults as intuitive scientists. *Psychological Review*, 96(4), 674-689.
- Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, 16(9), 678-683.
- Meltzoff, A., & Brooks, R. (2001). "Like me" as a building block for understanding other minds: bodily acts, attention, and intention. In B. F. Malle, L. J. Moses & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition*. Cambridge MA: The MIT Press.
- O'Neill, D. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development*, 67(2), 659-677.
- Penner, D., & Klahr, D. (1996). The interaction of domain-specific knowledge and domain-general discovery strategies: A study with sinking objects. *Child Development*, 67(6), 2709-2727.
- Piaget, J. (1930). *The child's conception of physical causality*. New York: Harcourt, Brace.
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43(4), 1045-1050.
- Sodian, B., Zaitchik, D., & Carey, S. (1991). Young children's differentiation of hypothetical beliefs from evidence. *Child Development*, 62(4), 753-766.
- Sommerville, J., & Woodward, A. (2005). Pulling out the intentional structure of action: the relation between action processing and action production in infancy. *Cognition*, 95(1), 1-30.
- Sorce, J., Emde, R., Campos, J., & Klinnert, M. (1985). Maternal emotional signaling: Its effect on the visual cliff behavior of 1-year-olds. *Developmental Psychology*, 21(1), 195-200.
- Walden, T., & Ogan, T. (1988). The development of social referencing. *Child Development*, 1230-1240.
- Williamson, R., Meltzoff, A., & Markman, E. (2008). Prior experiences and perceived efficacy influence 3-year-olds' imitation. *Developmental Psychology*, 44(1), 275-285.