

The invisible hand: Toddlers connect probabilistic events with agentive causes

Yang Wu

Massachusetts Institute of Technology

Paul Muentener

Tufts University

Laura E. Schulz

Massachusetts Institute of Technology

Author Note

This research was funded by a National Science Foundation CAREER award to L.S. and the Center for Minds, Brains and Machines (CBMM), funded by NSF STC award CCF-1231216. Warm thanks to the Boston Children's Museum and participating parents and children. Thanks to Patience R. Stevens, Faith O'Brian, Emily Tsang, and Veronika Jedryka for help with data collection, to Rachel Magid and Lutong T Cheng for help with coding the data, and to Hyowon Gweon, Rebecca Saxe, and Julian Jara-Ettinger for helpful feedback on the paper.

Correspondence concerning this article should be addressed to Yang Wu, MIT Department of Brain and Cognitive Sciences, 77 Massachusetts Avenue, Cambridge, MA 02139. Phone: 617-324-4859 Email: yangwu@mit.edu,

Abstract

Children posit unobserved causes when events appear to occur spontaneously (e.g., Gelman & Gottfried, 1996). What about when events appear to occur probabilistically? Here toddlers (mean: 20.1 months) saw arbitrary causal relationships (Cause A generated Effect A; Cause B generated Effect B) in a fixed, alternating order. The relationships were then changed in one of two ways. In the Deterministic condition, the event order changed (Event B preceded Event A); in the Probabilistic condition, the causal relationships changed (Cause A generated Effect B; Cause B generated Effect A). As intended, toddlers looked equally long at both changes (Experiment 1). We then introduced a previously unseen candidate cause. Toddlers looked longer at the appearance of a hand (Experiment 2) and novel agent (Experiment 3) in the Deterministic than the Probabilistic conditions, but looked equally long at novel non-agents (Experiment 4), suggesting that by two, toddlers connect probabilistic events with unobserved agents.

Keywords: causal learning, determinism, agency, toddlers, unobserved causes

The 19th century mathematician, Pierre Simon-LaPlace, speculated that if there were an intellect capable of analyzing all the forces operating in nature “to it nothing would be uncertain; the future, like the past, would be as the present before its eyes.” (1814/1951). Twentieth century physics has made this view untenable; we now know that our universe is comprised of irreducible uncertainties. Nonetheless, in everyday life, we may act as if the world were deterministic and believe that we could fully account for all our observations if we only had more information. Even Albert Einstein was loath to accept the implications of quantum mechanics, famously asserting: “God does not play dice with the universe”.

However misguided about the fundamental nature of the universe, a commitment to causal determinism may be advantageous for everyday inference and discovery. To the degree that people assume the world is deterministic, there are well-specified conditions under which they should posit the existence of hidden variables: if events appear to occur spontaneously, learners can assume that either an unobserved generative cause is present or that an inhibitory cause is absent; if events appear to occur stochastically, learners can infer that either an unobserved inhibitory cause is present or that a generative cause is absent.

Developmental evidence suggests that inferences of this nature are accessible even to children. Preschoolers resist both spontaneous and stochastic causation (at least for physical events) well before they receive formal science instruction. By the age of five, preschoolers posit hidden causes to account both for apparently uncaused events (Bullock, Gelman & Baillargeon, 1982; Chandler & Lalonde, 1994; Gelman, Coley, & Gottfried, 1994) and for caused events that occur some, but not all, of the time (Schulz & Sommerville, 2006; see also Piaget & Inhelder, 1975). Moreover, four and five-year-olds trade-off inferences about the kinds of unobserved

variables that might be present: if children know a generative cause is missing they are less likely to infer the existence of inhibitory cause, and vice versa (Schulz & Sommerville, 2006).

However, relatively little is known about the origins of deterministic beliefs earlier in development. The vast majority of studies looking at violations of causal determinism before the preschool years have focused only on the specific case of unexplained motion events (see Gelman & Gottfried, 1996; Gottfried & Gelman, 2005; Leslie, 1984; Luo & Baillargeon, 2005; Luo, Kaufman, & Baillargeon, 2009; Markson & Spelke, 2006; Muentener, Bonawitz, Horowitz, & Schulz, 2012; Premack, 1990; Saxe, Tenenbaum, & Carey, 2005; Saxe, Tzelnic, & Carey, 2007; Spelke, Phillips, & Woodward, 1995). If for instance, an inanimate object flies over a wall, infants look longer if a hand is revealed at the terminus of the movement than at the origin, suggesting that infants posit hidden causes when objects appear to move spontaneously (Saxe et al., 2005; 2007). Recent work has extended these findings to change of state events: infants infer causal relations not just when objects appear to move spontaneously but also when they appear to spontaneously break apart or play music (Muentener & Carey, 2010). Moreover, 2-year-old children both look predictively and selectively explore plausible candidate causes given evidence for otherwise spontaneous events (Muentener & Schulz, 2014).

However, although such findings suggest that very young children expect hidden causes when events appear to occur spontaneously, they do not say anything about whether children expect unobserved causes when events appear to occur probabilistically. Children might resist spontaneous causation while being untroubled by the possibility that observed causes can (for no reason at all) generate outcomes that occur only some of the time. Some evidence suggests that toddlers imitate deterministically effective actions more faithfully than probabilistically effective ones (Schulz, Hoopell, & Jenkins, 2008). However, although this suggests that probabilistic

evidence may encourage exploration, it does not bear directly on the question of whether toddlers connect probabilistic events to unobserved causes.

This is not to suggest that children are largely insensitive to probabilistic evidence. Even infants detect and generalize patterns in arbitrary stimuli based on the transitional probability between events (Fiser & Aslin, 2002; Johnson, et al., 2009; Kirkham, Slemmer, & Johnson, 2007; Marcus, Vijayan, Bandi Rao, & Vishton, 1999; Saffran, Aslin & Newport, 1996). Infants also distinguish probable and improbable relationships between samples and populations (e.g., Denison, Reed, & Xu, 2013; Denison & Xu, 2010a; Xu & Denison, 2009; Xu & Garcia, 2008; Teglas, Girotto, Gonzales, & Bonatti, 2007), and integrate their understanding of the statistical probability of event outcomes with their intuitive theories of the physical world (Denison & Xu, 2010b; Teglas, Vul, Girotto, Gonzales, Tenenbaum, & Bonatti, 2011). However, as noted, the only study suggesting that children infer unobserved causes given probabilistic evidence has focused on children already old enough to know about batteries, magnetism, and other unobservable physical causes (42 – 65 months; mean: 47 months; Schulz & Sommerville, 2006). Here, motivated both by theoretical considerations and related prior empirical work (Schulz, et al., 2008) we look at children two years younger: toddlers (18-24 months; mean: 20 months).

Toddlers are a particularly important age group to study with respect to their reasoning about latent causes of probabilistic events. By the age of two, children have rich intuitive theories of the physical and psychological world that support inferences about unobserved causal variables (Onishi & Baillargeon, 2005; Song, Onishi, Baillargeon, & Fisher, 2008; Southgate, Senju, & Csibra, 2007; Wellman & Gelman, 1992; Wellman, 1990; Wellman & Wooley, 1990). In the second year of life, toddlers are also able to learn novel, seemingly arbitrary causal relationships from patterns of covariation evidence (e.g., Gweon & Schulz, 2011; Sobel &

Kirkham, 2006; Walker & Gopnik, 2014). However, as noted, there is as yet no evidence that 18-24-month-olds selectively represent the presence of unobserved variables when observed causes behave probabilistically. Indeed, most studies of causal reasoning in this age group (and even in older children) have focused on children's inferences from deterministic data (Bonawitz, et al., 2010; Meltzoff, Waismeyer, & Gopnik, 2012; Sobel & Kirkham, 2006; Vredenburg, Kushnir, & Casasola, 2014; Walker & Gopnik, 2014; Yu & Kushnir, 2014). However, outside the laboratory, toddlers are very likely to observe instances of apparently probabilistic causation, given both the complexity of real world physical events and the potential for noise and error in children's interventions. If before the age of two, children already connect novel probabilistic events to hidden causes, this might support their ability to reason about the latent structure of novel events even before they know much about specific causal mechanisms.

Here we test children's inferences about probabilistic causation by introducing them to novel, arbitrary causal relationships. Toddlers are shown familiarization trials in which first Cause A generates Effect A and then Cause B generates Effect B. Children are shown this sequence three times. We then introduce a Switch Trial in which we change the events in one of two ways, designed to match for salience (which we verify in Experiment 1). In the Deterministic conditions, we switch the order of the causal events so that the B event happens before the A event; the relationship between each cause and its effect remains unchanged. In the Probabilistic conditions, we switch the causal relationship so that Cause A generates Effect B and Cause B generates Effect A; the order of the causal events remains unchanged (Cause A always occurs first and Cause B always occurs second). See Figure 1.

We predicted that toddlers would find both of these changes surprising. Previous studies have found that even very young infants selectively attend to changes in the temporal order of

audiovisual events (e.g., Lewkowicz, 2004a; 2004b) and toddlers are particularly likely to encode the temporal order of causally relevant events (e.g., Bauer & Mandler, 1989; 1992). This is valuable because changes in causal order affect the construal of an event, even if each event is causally independent of the other. (Consider the relatively innocuous event of applying the breaks on your car and then the accelerator versus the potentially more dramatic event of pushing the accelerator and then the breaks). Indeed, changing the order of causal events may be diagnostic of whether two events are independent or not. (If a toddler pushes the “on” button on the computer and then the letter “D” she may think the two causal events are independent until she tries pressing the letter “D” before pushing the “on” button.) Even when an agent engages in a sequence of arbitrary, causally unnecessary actions, children appear to construe the order of events as normatively important and they imitate the sequence faithfully (e.g., Lyons, Young, & Keil, 2007). Thus both because changing the order of causal events changes the event construal and because the order of causal events often has functional or normative relevance, we believed toddlers would attend to the change in causal order.

For different reasons, we predicted that the change in the causal relationship would attract toddlers’ attention. The critical point about deterministic causal relationships is that they obtain 100% of the time; as soon as that fails to be the case (e.g., when Cause A fails to generate Effect A, even on a single trial) the relationship between Cause A and Effect A is (apparently) no longer deterministic. Cause A now appears to generate Effect A probabilistically: only 75% of the time. Of course, we do not usually view such events as irreducibly probabilistic; insofar as we are intuitive determinists, we assume either that there is a hidden cause explaining the change in the events or an unobserved source of noise. However, to “rescue” determinism in the face of apparently probabilistic evidence, the observer must posit an unobserved latent cause.

Consistent with a large body of research on infant looking time suggesting that infants attend to events that are novel or violate their expectations (see Aslin, 2007 for review) we expected infants to attend to both Switch Trial events, and the events were designed to be well matched for perceptual similarity and overall interest. Nonetheless, they represent two different kinds of changes to the events and differ in any number of respects that might affect toddlers' looking time. Because any subsequent looking time differences would be difficult to interpret if children looked differently at these two events, Experiment 1 tests whether toddlers, as intended, look equally long at these two kinds of changes.

However, our primary question of interest is whether children selectively represent unobserved candidate causes given evidence for probabilistic causation. In the subsequent studies, we introduce a previously unseen, additional candidate cause after the Switch Trial. We hypothesized that toddlers would look longer at the appearance of the previously unobserved cause in the Deterministic condition than in the Probabilistic condition. In the Deterministic condition, the events in the Switch Trial can be explained as a change in the experimenter's goal-directed action (i.e., she changed her mind and acted on B first rather than A). This is a novel action with implications for how to construe the event as a whole but there is no reason to posit an additional cause. By contrast, in the Probabilistic condition, the appearance of the additional unobserved cause can account for the apparent violation of determinism; whether Cause A generates Effect A or B could depend on the state of the unobserved cause. Thus we predicted that although toddlers would be equally attentive to a change in causal order and a change in causal relationships, they would look longer at the unobserved cause in the Deterministic condition than the Probabilistic condition.

Finally, consistent with abundant research suggesting that young children preferentially represent intentional agents (rather than objects) as causes of events (Bonawitz, et al., 2010; Muentener, Bonawitz, Horowitz, & Schulz, 2012; Muentener & Carey, 2010; Newman, Keil, Kuhlmeier, & Wynn, 2010; Saxe et al., 2005; 2007), we predicted that toddlers would be more likely to accept a hidden agent as a candidate cause than a hidden object. To test this we manipulated the ontological status of the candidate cause across conditions: in Experiment 2, we introduced a human hand; in Experiment 3, we introduced a novel agent puppet (with a face but no hands), and in Experiment 4, we introduced a novel non-agent (identical to the novel agent except that the face was scrambled). We predicted that toddlers would look longer at the appearance of an agent (Experiments 2 and 3) in the Deterministic condition than the Probabilistic condition, but would show no looking time differences given the appearance of an object (Experiment 4).

Experiment 1: Manipulation check

In this manipulation check experiment, we familiarize children with the same novel causal relationship across conditions, and then introduce one of two changes. In the Deterministic condition, we change the order of the causal events; in the Probabilistic condition we change the causal relationship between the events. As discussed, we designed the stimuli so that the changes were perceptually well matched and both kinds of changes might be equally interesting to the children. The order change should be interesting to the degree that changing the order of causal relationships can have functional or normative significance; the causal change should be interesting to the degree that children notice the change in the causal relationship. If children's looking times are equivalent at this point, we can then investigate toddlers' looking time when the hidden causes are revealed with less concern that any subsequent differences in

looking times are due to different initial looking at the Switch Trials. Because our test participants are toddlers (18-24 months old) rather than young infants, and we wanted to hold their attention, we opted for familiarization and a single test trial, rather than habituation and multiple test trials (following the example of other toddler violation of expectation studies in the literature, e.g., Onishi & Baillargeon, 2005; Surian, Caldi, & Sperber, 2007). Given previous research suggesting that toddlers can learn distinct patterns of causal relations given just a few trials of evidence (e.g., Gweon & Schulz, 2011; Sobel & Kirkham, 2006) there was a priori reason to believe that toddlers should be able to encode the kinds of causal relations we showed them here.

Methods

Participants

Thirty-six toddlers were recruited at a Children's Museum (range: 18.0 to 24.0 months, mean age = 21.1 months, 21 male). They were randomly assigned to either the Deterministic condition (mean age = 20.9 months, $SD = 2.1$) or Probabilistic condition (mean age = 21.2 months, $SD = 2.0$, $n = 18$ toddlers per condition). Although most of the children were white and middle class, a range of ethnicities and socioeconomic backgrounds reflecting the diversity of the local population were represented. An additional 7 toddlers were recruited but not included in the final sample due to: experimenter error ($n = 2$), fussiness ($n = 4$), or parent interference ($n = 1$).

Materials

Toddlers were introduced to a purple box (37.6 cm × 29.2 cm × 20.3 cm) with two handles (21.6 cm in length). The left handle was red with black stripes. The right handle was green with white spots. (See Fig. 1.) The box was placed in front of a black foam board screen

(117.9 cm × 97.8 cm). The experimenter could hide behind the screen and observe the child through pinholes in the screen. Two openings in the screen on either side of the box allowed the experimenter to reach her hands through to manipulate the handles. The box had an opening in the back and the top so that the experimenter could conceal her hand in the box and lift objects out of the box. When a handle was pressed the experimenter lifted either a lollipop (9.4 cm in diameter) or a toy cake (7.6 cm in height, 7.6 cm in width) out of the box. An MP3 player was also used: the red handle was always accompanied by the sound of an ascending scale on a xylophone; the green handle was always accompanied by the sound of a descending scale on a xylophone. The experimenter also used earphones and a metronome to track the timing of the experiment.

Procedure

Toddlers were tested in a private room located in the children's museum. The child was placed in a high chair approximately 100 cm in front of the box. The child's parent sat to the right of the high chair, out of the child's direct line of sight.

Before the experiment, the experimenter shook a set of keys at three regions of the display—the red handle, the top of the box and the green handle—to attract the child's attention to the display. These were used to calibrate the offline coder to the locations that counted as looking at the stimuli. Then the experimenter went behind the screen. The experimenter knocked on the center of the screen and said, "Hi, [child's name]! Watch this box!" She began the **Familiarization Trials** by putting her hand out of the left hole and waving at the child. She waved at the child in order to catch the child's attention. When the child was looking at the hand, she then pressed the red handle (Cause A) and, with her other hand concealed in the box, lifted the lollipop out of the box (Effect A) and at the same time triggered the ascending scale. She

held the lollipop up for 2 seconds and then released the red handle and simultaneously returned the lollipop to the box. She brought her hand back behind the screen. She then put her hand out of the right hole and waved at the child. When the child looked at the hand, she pressed the green handle (Cause B) and, with her other hand concealed in the box, lifted the cake out of the box (Effect B) and at the same time triggered the descending scale. She held the cake up for 2 seconds and then released the green handle and simultaneously returned the cake to the box. Critically, the hand that pushed the toy out of the box was never visible to the child; pilot work established that to an adult observer, it looked like the handle caused the toys to emerge from the box. The experimenter repeated the familiarization trials a total of three times so that children saw three instances of each causal relationship. The experimenter wore earphones attached to an electronic metronome and synchronized her actions with the metronome throughout.

On the **Switch Trial**, the experimenter said, “[child’s name], watch!” In the Deterministic condition, she switched the order of events, repeating the events in the Familiarization Trials except that she pressed the green handle first (Cause B) and the cake popped up (Effect B); she then pressed the red handle (Cause A) and the lollipop popped up (Effect A). In the Probabilistic condition, the experimenter switched the relation between events, repeating the events in the Familiarization Trials except that when the experimenter pressed the red handle (Cause A), the cake popped up (Effect B); when she pressed the green handle (Cause B), the lollipop popped up (Effect A). In both conditions, the experimenter froze and held the position with the handle down and the lollipop up. The experimenter ended the experiment when the child looked away from the stage for at least 2 consecutive seconds.

Results and discussion

All experiments were coded blind to condition, offline from videotape. One coder coded toddlers' cumulative looking time to the events from the start of the experiment to the end of the Switch Trial when the scene froze. Toddlers were equally attentive to the entire sequence of events in the two conditions (Deterministic condition: $M = 45.7s$, $SD = 7.0$; Probabilistic condition: $M = 48.1s$, $SD = 6.8$; $t(34) = -1.015$, $p = .317$, $d = -.35$, 95% CI [-7.06, 2.36]).

A second coder coded the toddlers' looking times from the end of the Switch Trial when the scene froze until the child looked away from the stage for 2 consecutive seconds; an additional coder blind to conditions coded this measure on one third of the test clips. Inter-coder reliability was high, $r^2 > .9$. This coding corroborated the experimenter's decision about the endpoint of the experiment in all cases. There was no difference in toddlers' looking times to the last scene of the Switch Trial between conditions (Deterministic condition: $M = 9.9s$, $SD = 6.4$; Probabilistic condition: $M = 9.0s$, $SD = 3.5$; Welch's unequal variances t -test: $t(26) = .541$, $p = .593$, $d = .21$, 95% CI [-2.59, 4.45]).

These results suggest (1) that toddlers in the two conditions were equally engaged in the experiment through the end of the Switch Trial and (2) that they looked equally long at the final outcome of the Switch Trial. Experiment 1 thus serves as a manipulation check, suggesting that, as intended, the stimuli were matched for perceptual similarity and interest well enough to equate looking times across the two conditions. This allows us to ask whether toddlers nonetheless have different responses to hidden agents in the two conditions. If toddlers look longer at the appearance of the agent in the Deterministic condition than the Probabilistic condition, this manipulation check mitigates against the possibility that the looking time difference is due to the difference in the Switch Trials across conditions; instead it would suggest that, as hypothesized,

toddlers are more likely to accept the appearance of an agent given the probabilistic event than the deterministic event. We investigate this hypothesis in Experiment 2.

Experiment 2

In Experiment 1, we established that toddlers looked equally long given a change in the order of the causal events (the Switch Trial in the Deterministic condition) and the causal relationship between events (the Switch Trial in the Probabilistic condition). In Experiment 2, we look at whether toddlers are more likely to accept the existence of hidden agents given probabilistic than deterministic events. Experiment 2 is identical to Experiment 1 except that rather than freezing the scene at the end of the Switch Trial until the toddlers look away for two seconds, we pause on the last scene of the Switch Trial for two seconds and then proceed directly to the Test trial.¹ At the Test Trial, a previously concealed hand emerges from the box, holding the lollipop. We predict that (as in Experiment 1) there should be no difference in children's attention to the events through the end of the Switch Trial, but that toddlers should look longer at the hand in the Deterministic condition (when there was no reason to assume a hidden agent) than the Probabilistic condition (given that an unobserved agent might be inferred).

Methods

Participants

Thirty-two toddlers were recruited at a Children's Museum (range: 18.0 to 23.5 months, mean age = 19.8 months, 15 male). Toddlers were randomly assigned to either the Deterministic condition (mean age = 20.0 months, $SD = 1.8$) or Probabilistic condition (mean age = 19.7

¹ An alternative would have been to replicate Experiment 1 in its entirety (including terminating the Switch Trial after the 2 second look away) and only then introduce the Test Trial. However, because we were working with active, ambulatory toddlers we were concerned that we would fail to recover the children's attention after losing it. Thus we opted to proceed directly to the Test Trial.

months, $SD = 1.3$, $n = 16$ toddlers per condition). Although most of the children were white and middle class, a range of ethnicities and socioeconomic backgrounds reflecting the diversity of the local population were represented. An additional 7 toddlers were recruited but not included in the final sample due to: experimenter error ($n = 4$), fussiness ($n = 1$), or parent interference ($n = 2$).

Materials and procedure

We used the same materials and procedure as those in Experiment 1 with one exception. Instead of freezing at the last scene of the Switch Trial and terminating the experiment after the child looked away for 2 consecutive seconds, we paused at the last scene of the Switch Trial for 2 seconds and then continued immediately onto the Test Trial. On the Test Trial, the experimenter put her hand out of the left hole and waved to the child. She then said “Aha!”, and at the same time pressed the red handle, and lifted her hand holding the lollipop all the way out of the box so that both her hand and the lollipop were now visible to the child. The experimenter remained stationary in this position until she judged that the child looked away from the stage for at least 2 consecutive seconds.

Results and discussion

A coder blind to conditions coded toddlers’ cumulative looking time to the ongoing events from the start of the experiment to the end of the Switch Trial, after the scene froze for 2 seconds. Toddlers were equally attentive to the entire sequence of events in the two conditions up to the end of the Switch Trial (Deterministic condition: $M = 52.6s$, $SD = 3.5$; Probabilistic condition: $M = 50.9s$, $SD = 4.2$; $t(30) = 1.225$, $p = .230$, $d = .45$, 95% CI [-1.12, 4.48]). We will refer to this measure as the toddlers’ Baseline looking in each condition.

A second coder blind to conditions coded off-line from videotape the children’s looking times to the test trial, from the beginning of the “Aha!” sound until the child looked away for 2

consecutive seconds. An additional coder blind to conditions coded one third of the test clips. Inter-coder reliability was high, $r^2 > .9$. The off-line coding from videotape corroborated the experimenter's decision to end the experiment in all but three cases (one in the Deterministic condition, two in the Probabilistic condition); these three children were dropped from the analysis and replaced due to premature termination of the test trial.

Next, we examined the question of interest: the effect of the condition manipulation on toddlers' looking time to the test trial. (See Fig. 2.) As predicted, toddlers looked longer at the test trial in the Deterministic condition ($M = 13.7s$, $SD = 7.3$) than the Probabilistic condition ($M = 8.2s$, $SD = 4.8$; $t(30) = 2.507$, $p = .018$, $d = 0.92$, 95% CI [1.02, 9.97]).

As noted, the timing of the experimenter's actions was synchronized with a metronome throughout. However, to ensure that the experimenter did not differentially cue the children (i.e., when she said "Aha!") between the Deterministic and Probabilistic conditions, a coder blind to conditions rated the enthusiasm of the experimenter's vocalization from videotape on a Likert scale from 1 (not enthusiastic at all) to 5 (extremely enthusiastic); there were no differences between conditions (Deterministic condition: $M = 2.6$, $SD = 0.6$; Probabilistic condition: $M = 2.4$, $SD = 0.6$; $t(30) = .850$, $p = .402$, $d = .31$, 95% CI [-.26, .64]).

Both the failure to observe any condition differences in the manipulation check in Experiment 1 and the Baseline looking here suggest that the results of the test condition are unlikely to be due to toddlers' differential attention to looking to the Familiarization and Switch Trials across the two conditions. The results can be predicted from the assumption that toddlers had no reason to expect an unobserved agent in the Deterministic condition, whereas in the Probabilistic condition, the apparently probabilistic relationship between each observed cause and the effect could be explained away by the appearance of the previously hidden cause.

Note however, that human hands are very familiar to toddlers and have the requisite affordances for performing myriad different tasks. The toddlers might have accepted the human hand as a candidate cause only because of its familiarity and particular affordances. A more interesting possibility is that, consistent with previous work (e.g., Saxe, et al., 2007) toddlers represent the hand as an extension of a dispositional agent, and that children treat dispositional agents as plausible candidate causes of a wide range of events. We test this claim in Experiment 3 by replicating the design of Experiment 2 varying only the candidate cause that emerges during the test trial. In Experiment 3, we introduce a novel agent without hands.

Experiment 3

In contrast to studies of children's inferences about unobserved causes of apparently spontaneous events, the evidence children observed in Experiment 2 does not immediately suggest a unique target action that generated it. That is, a tossed beanbag implies an agent who tossed it (Saxe, et al., 2005) and stacked blocks, an agent who stacked them (Newman et al., 2010). In Experiments 1-2, however, the hand might have done any of a number of different things to generate the probabilistic evidence: it might have switched the cake and the lollipop on the ends of the handles, it might have lifted the cake and lollipop on the Switch Trial (coincident with the handle presses and creating the illusion that the handle caused the events), or it might have lifted the cake and lollipop all along.² However, human hands have the affordances requisite to performing any one of these tasks (e.g., the ability to grasp and lift objects) and hands are of course very familiar to toddlers. Nonetheless, given the various possibilities for

² Note that the inference is not that the hand moved the object per se. (No additional cause was needed merely to explain the movement of the objects. The observed intervention on the lever was a sufficient plausible cause of the object popping out of the box, as is clear from children's failure to impute additional causes in the Deterministic condition.) The inference is that in the Probabilistic condition the changing relationship between the levers and the objects could be explained away by considering the role of the hand.

how the evidence was generated, it seems unlikely that the toddlers envisioned a unique action underlying the probabilistic data.

Instead, toddlers might have accepted the human hand as a candidate cause, not because of its particular affordances but because of its abstract ontological status as an extension of a dispositional agent. A large body of work suggests that infants and toddlers attribute unique causal powers to dispositional agents including the ability to engage in self-generated movement (Luo & Baillargeon, 2005; Premack, 1990), the ability to resist being moved (Wang, Kaufman, & Baillargeon, 2003), the ability to resist gravity (Leslie, 1984), the ability to cause objects to change state (Muentener & Carey, 2010), the ability to create order (Newman, et al., 2010), and the ability to draw non-random samples (Ma & Xu, 2013). Rather than inferring a particular action underlying the events, toddlers may simply assume that dispositional agents can freely engage in goal-directed actions. In this sense, the appearance of a dispositional agent transforms the apparently stochastic physical events in the Probabilistic condition to a plausibly deterministic action governed by changes in goal-directed action (much like the Deterministic condition, except that of course the appearance of the hidden agent in the Deterministic condition is less likely given its redundancy with the observed cause.)

In Experiment 3 we look at whether toddlers only accept familiar, plausible candidate causes of probabilistic evidence or whether they have a more abstract expectation that any entity with the ontological status of a dispositional agent can intervene on the causal structure of the events. Following the approach used in previous studies (Muentener & Carey, 2010; Newman et al., 2010; Saxe et al., 2005; 2007), we test whether children have an abstract commitment to intentional agents as candidate causes by replacing the familiar, biomechanically plausible human hand with a novel, apparently self-moving, agent puppet with a face but no hands. Note

that although these studies involved infants, abundant work suggests that, well through preschool, children readily accept puppets as dispositional agents (e.g., Garvin & Woodward, 2015; Jara-Ettinger, Tenenbaum, & Schulz, 2015; Kushnir, Xu, & Wellman, 2010; Wimmer & Perner, 1983). However, if toddlers do not accept unfamiliar agents, or agents without relevant affordances as candidate causes, then the children should look equally long at the appearance of the puppet in both conditions. By contrast, if children accept any dispositional agent as a candidate cause, they should look longer at the appearance of the agent in the Deterministic condition than in the Probabilistic condition.

Methods

Participants

Thirty-two toddlers were recruited at a Children's Museum (range: 18.1 to 24.0 months, mean age = 21.1 months, 19 male). Toddlers were randomly assigned to either the Deterministic condition (mean age = 20.7, $SD = 1.9$) or Probabilistic condition (mean age = 21.4 months, $SD = 1.9$, $n = 16$ per condition). Although most of the children were white and middle class, a range of ethnicities and socioeconomic backgrounds reflecting the diversity of the local population were represented. An additional 14 toddlers were recruited but not included in the final sample due to: experimenter error ($n = 5$), fussiness ($n = 5$), or parent interference ($n = 4$).

Materials

We created a novel agent puppet (9.1 cm. in diameter) out of a ball, pink yarn, two paper eyes and a mouth. (See Fig. 1, *Test Trial* part). Earphones and a metronome were also used to track the timing of the puppet's motion. Additionally, all of the materials used in Experiment 2 were used here.

Procedure

The procedure was similar to Experiment 2 except as follows. Following the procedure in Saxe et al., 2005, 2007, children were first familiarized with the novel agent. Before the experiment began, the experimenter went behind the screen and said to the child, “We have a friend here today!” Then a puppet popped up from the top of the screen, swayed to the left and to the right three times, and then went behind the screen. This was repeated a total of three times. Throughout the familiarization phase, the experimenter controlled the puppet; however, to the child it appeared that the puppet was moving on its own.

On the test trial, when the red handle was pressed, the puppet, instead of a human hand, emerged from the box with the lollipop, and at the same time the experimenter said “Aha!”. (The experimenter held the stick of the lollipop and the puppet together so it looked like the puppet was “holding” the lollipop.) The experimenter swayed the puppet and lollipop to the left and to the right three times (just like how we familiarized the child with the puppet) and then held the puppet and lollipop upright and still. The experimenter was cued by a metronome that she listened to over earphones to control the pace of the puppet’s motion to ensure that it was constant across all the participants. The experimenter remained stationary until she judged that the child looked away from the stage for at least 2 consecutive seconds.

Results and discussion

All results were coded offline from videotape. As in Experiment 2, a coder blind to conditions coded the toddlers’ looking times up to the end of the Switch Trial. Toddlers’ Baseline looking to the events did not differ between conditions (Deterministic condition: $M = 51.7s$, $SD = 6.5$; Probabilistic condition: $M = 54.7s$, $SD = 7.6$; $t(30) = -1.215$, $p = .234$, $d = -.44$, 95% CI [-8.16, 2.07]).

A second coder blind to conditions coded the children's looking times to the test trial, from the beginning of the "Aha!" sound to the start of the 2-second look away. An additional coder blind to conditions coded one third of the clips. Inter-coder reliability was high, $r^2 > .9$. The off-line blind coding from videotape corroborated the experimenter's judgment about the endpoint of the experiment in all but one case (in the Probabilistic condition); this child was dropped from the analysis and replaced due to premature termination of the test trial.

Next we examined our question of interest: the effect of the condition manipulation on toddlers' looking time to the test trial. (See Fig. 2.) The results replicated the pattern in Experiment 2: the toddlers looked longer at the test trial in the Deterministic condition ($M = 16.9s$, $SD = 7.0$) than the Probabilistic condition ($M = 9.9s$, $SD = 3.8$; Welch's unequal variances t -test: $t(23) = 3.490$, $p = .002$, $d = 1.46$, 95% CI [2.90, 11.09]). Consistent with Experiment 2, this result suggests that toddlers were more likely to accept the presence of a previously unobserved agent in the Probabilistic condition than the Deterministic condition.

An additional coder blind to conditions rated the enthusiasm of the experimenter's vocalization (i.e. "Aha") from videotape on a Likert scale from 1 (not enthusiastic at all) to 5 (extremely enthusiastic); there were no differences between conditions (Deterministic condition: $M = 1.6$, $SD = 0.6$; Probabilistic condition: $M = 1.6$, $SD = 0.6$; $t(30) = .000$, $p = 1.000$, $d = .00$, 95% CI [-.4844, .4844]).

As in Experiment 2, the manipulation check from Experiment 1 and the results of the baseline coding here suggest that the difference in looking at the test trial is unlikely to be due to artifacts of the design. Rather, the results of Experiment 3 suggest that children's representations of the kinds of unobserved causes that might explain probabilistic causation are abstract: toddlers accept handless puppets as well as familiar hands as plausible candidate causes. Would toddlers

accept non-agents as well? Previous work suggests that infants and toddlers are more constrained than adults in the kinds of candidate causal relations they accept. Infants and toddlers do not seem to believe, for instance, that objects can cause other objects to move (Saxe et al., 2005), change states (Bonawtitz et al., 2010; Muentener & Carey, 2010), create order (Newman et al., 2010), or engage in selective sampling (Ma & Xu, 2013). Thus we might expect that toddlers would similarly be reluctant to accept that a non-agent could account for probabilistic data. In Experiment 4, we replaced the novel agent puppet with a perceptually similar non-agent (object) to see whether toddlers accept this as a candidate cause.

Experiment 4

In Experiment 4, we show toddlers probabilistic or deterministic evidence and then present them with an inanimate object as a candidate cause. If, given evidence for probabilistic causation, toddlers broadly infer the existence of *any* additional cause, then toddlers should look longer at the object in the Deterministic condition than in the Probabilistic condition. However, if toddlers only attribute causal agency to dispositional agents, then the appearance of the object should be equally likely in the Deterministic and Probabilistic conditions. Additionally, Experiment 4 allows us to test a possible methodological concern. Because the toddlers were familiarized to the events, rather than habituated, we cannot definitively rule out the possibility that although the toddlers looked equally at the events through the Switch Trials across conditions, they may nonetheless have failed to fully encode the change in the causal relationship. If so longer looking in the Deterministic condition than the Probabilistic condition might reflect a familiarity preference (for the simpler causal order change) rather than a violation of expectation (given the unexpected appearance of a hidden agent). If toddlers fail to encode

the causal Switch Trial and merely show a preference for the events in the Deterministic condition, then they should – contra our hypothesis – continue to do so in Experiment 4.

Methods

Participants

In this Experiment, we predicted no difference between conditions, thus we conducted a power analysis to ensure that we had sufficient power to detect a difference if one existed. We assumed a strong effect size consistent with the previous results (i.e. aggregating all the data from Experiments 2 and 3 generates an effect size of $d = 1.17$), resulting in a sample of $n = 18$ /condition, sufficient to detect a difference 93% of the time. Thirty-six toddlers were recruited at a Children's Museum (range: 18.1 to 23.8 months, mean age = 21.3 months, 18 male). They were randomly assigned to either the Deterministic condition (mean age = 21.1 months, $SD = 1.6$) or Probabilistic condition (mean age = 21.6 months, $SD = 1.4$, $n = 18$ toddlers per condition). Although most of the children were white and middle class, a range of ethnicities and socioeconomic backgrounds reflecting the diversity of the local population were represented. An additional 6 toddlers were recruited but not included in the final sample due to: experimenter error ($n = 2$), fussiness ($n = 2$), or parent interference ($n = 2$).

Materials

We used the same materials as in Experiment 3, except that the materials used to create the face on the novel puppet were scrambled such that the puppet no longer looked like an animate, intentional agent. (See Fig. 1, bottom panel.)

Procedure

The procedure was similar to that in Experiment 3 with two exceptions. Following the procedure in Saxe et al., 2005, 2007, the experimenter familiarized the child with the novel

object before the experiment began. She took out the object and said to the child, “We have a new toy here today!” She put it in front of the child and pointed at it, saying, “Look!” She kept the pointing position for 3-4 seconds. She then moved it to another place in front of the child and pointed at it, saying, “Look!” Again she kept this position for 3-4 seconds. She repeated the moving and pointing procedure for a total of three times.

On the test trial, while the experimenter said “Aha!”, the red handle was pressed, and the object, attached to the lollipop, popped up from the box. The experimenter froze the scene until the experimenter judged that the child looked away from the stage for at least 2 consecutive seconds.

Results and discussion

As in the previous experiments, a coder blind to conditions coded the toddlers’ looking times from videotape up to the end of the Switch Trial. Toddlers’ Baseline looking to the events up to the end of the Switch Trial did not differ between conditions (Deterministic condition: $M = 53.6s$, $SD = 4.1$; Probabilistic condition: $M = 54.5s$, $SD = 3.6$; $t(34) = -.678$, $p = .502$, $d = -.23$, 95% CI [-3.47, 1.73]). Despite the difference in the initial familiarization (to a puppet agent versus an object), toddlers’ Baseline looking also did not differ between Experiments 3 and Experiment 4 (Experiment 3: Deterministic condition: $M = 51.7s$, $SD = 6.5$; Probabilistic condition: $M = 54.7s$, $SD = 7.6s$; no main effect of experiments: $F(1, 64) = .367$, $p = .547$, $\eta_p^2 = .01$; no main effect of conditions: $F(1, 64) = 2.067$, $p = .155$, $\eta_p^2 = .03$; no interaction: $F(1, 64) = .640$, $p = .427$, $\eta_p^2 = .01$).

A second coder blind to conditions coded the children’s looking times to the test trial, from the beginning of the “Aha!” sound to the start of the 2-second looking away off-line from videotape. An additional coder blind to conditions coded 100% the video clips. Inter-coder

reliability was high, $r^2 > .9$. The offline blind coding from videotape corroborated the experimenter's online judgment about the end point of the experiment in all cases.

Next we examined the effect of the condition manipulation on toddlers' looking time to the test trial. (See Fig. 2.) In contrast to the results of Experiments 2 and 3, in Experiment 4 toddlers looked equally long to the non-agent in the Deterministic ($M = 14.3s$, $SD = 8.4$) and Probabilistic conditions ($M = 15.3s$, $SD = 8.4$; $t(34) = -.374$, $p = .711$, $d = -.13$, 95% CI [-6.72, 4.63]). We also conducted an ANOVA on toddler's looking time to the test events in Experiments 3 and 4 with Agency (agent vs. object) and Condition (Deterministic vs. Probabilistic) as the between-subjects factors. There was no main effect of agency ($F(1, 64) = .683$, $p = .412$, $\eta_p^2 = .01$) or condition ($F(1, 64) = 2.867$, $p = .095$, $\eta_p^2 = .04$). However, there was an interaction between agency and condition ($F(1, 64) = 5.233$, $p = .025$, $\eta_p^2 = .08$). Follow-up analyses revealed that when toddlers viewed the Deterministic events, they looked equally long at the non-agent ($M = 14.3$ s, $SD = 8.4$) and the novel agent ($M = 16.9$ s, $SD = 7.0$; $t(32) = -.962$, $p = .343$, $d = -.34$, 95% CI [-2.87, 8.00]). When they viewed the Probabilistic events, however, toddlers looked longer at the non-agent puppet ($M = 15.3$ s, $SD = 8.4$) than at the novel agent ($M = 9.9$ s, $SD = 3.8$; Welch's unequal variances t -test: $t(24) = 2.490$, $p = .020$, $d = 1.02$, 95% CI [.94, 10.00]). These results suggest that the appearance of an object in Experiment 4 was no more likely in the Probabilistic condition than the Deterministic condition, and toddlers specifically accept the presence of agents, but not objects, given probabilistic events.

As in Experiments 2 and 3, an additional coder blind to conditions rated the enthusiasm of the experimenter's vocalization (i.e. "Aha") from videotape on a Likert scale from 1 (not enthusiastic at all) to 5 (extremely enthusiastic); there were no differences between conditions

(Deterministic condition: $M = 3.4$, $SD = 0.9$; Probabilistic condition: $M = 3.3$, $SD = 0.7$; $t(34) = .415$, $p = .680$, $d = .14$, 95% CI [-.43, .65]).

Experiment 4 further corroborates the manipulation check in Experiment 1 and the baseline measures in Experiments 2 and 3 in ruling out a class of alternative explanations for the test results in Experiments 2 and 3. If the toddlers' looking times were driven by ancillary features of the design, or by a failure to fully encode the events and an attendant familiarity preference for the (potentially simpler) causal order change in the Deterministic condition, we should have found similar patterns of looking in Experiment 4 as in Experiment 2 and 3. The fact that toddlers' Baseline looking in Experiments 3 and 4 were equivalent, but, in contrast to Experiment 3, toddlers did not look longer at the Deterministic condition than the Probabilistic condition in Experiment 4, suggests that these kinds of explanations cannot account for the results. Rather the results across all four experiments suggest that toddlers are more likely to accept the appearance of previously hidden agents given probabilistic evidence than deterministic evidence, and that they selectively treat dispositional agents as plausible candidate causes.

General Discussion

We presented toddlers with novel causal relationships followed by either a change in the order of the causal events (Deterministic condition) or by a change in the causal relationship between events (Probabilistic condition). Toddlers' looking times to the two types of changes were indistinguishable (Experiment 1). However, when a previously hidden human hand (Experiment 2) or a novel handless agent puppet (Experiment 3) was revealed, toddlers looked longer at the agent in the Deterministic condition than the Probabilistic condition. Toddlers' inferences were specific to dispositional agents: they did not distinguish the conditions when a

hidden object was revealed as the candidate cause (Experiment 4). These results are consistent with the possibility that toddlers expect physical causes to behave deterministically; when probabilistic events occur, they treat hidden agents as plausible candidate causes.

The finding that toddlers selectively accept latent causes given otherwise unexplained probabilistic causation is consistent with previous work suggesting that children are resistant to the idea that events can happen unpredictably; many studies suggest that children have difficulty understanding and accepting apparently random events (e.g. Piaget & Inhelder, 1975; Kuzmak & Gelman, 1986). However, there are at least two interpretations consistent with toddlers' looking behavior. One possibility is that before the age of two, children expect physical causes to behave deterministically. Thus when causal relationships change for no apparent reason, children infer the presence of unobserved agents. This interpretation would suggest that the inferences of 18-24-month-old may be continuous with the inferences of children two years older, who actively infer that necessary generative causes are missing, or that inhibitory causes are present, when they observe probabilistic events (Schulz & Sommerville, 2006).

However, it is not clear from this data that toddlers actively posited the existence of a hidden cause. Instead, children might have understood the events retroactively: when the agent appeared, the appearance may have made more sense given the probabilistic event than the deterministic event. Indeed, a learner trying to maximize information gain should accept additional candidate causes up until the point that a set of candidate causes deterministically accounts for the evidence. At a computational level (Marr, 1982) this is arguably equivalent to a belief in determinism: a learner who maximizes information gain by endorsing latent variables until 100% of the data is explained may be functionally indistinguishable from a learner committed to determinism.

However, it is noteworthy that both older children given instances of probabilistic causation (Schulz & Sommerville, 2006), and 2-year-old toddlers given instances of apparently spontaneous causation (Muentener & Schulz, 2014) do not merely accept the possibility of unobserved causes; given a well-constrained search space, they will incur costs (i.e. give up the chance to exploit a known cause) to explore a previously hidden potential cause. This provides at least suggestive evidence that children do not merely accept, but assume, the existence of unobserved causes when they observe violations of determinism. In the history of science, the idea of a genuinely non-deterministic universe is both a recent discovery and a counter-intuitive one, suggesting that a belief in causal determinism may be central to our intuitive theories of the physical world.

Of course, children may assume that artifacts behave deterministically without extending this assumption to the physical world more broadly. We do not know to what extent toddlers might endorse previously unobserved causes to account for probabilistic causal relations between naturally occurring events. Additionally, we do not know to what extent either adults or children extend a belief in causal determinism beyond the physical world, to psychological and social events. Note that the contrast between the causal change in the Probabilistic conditions and the order change in the Deterministic condition relied on the intuition that physical causal relationships do not change themselves but that people can freely change their goal-directed actions. The results of Experiment 1, and the baseline measures throughout, suggest that toddlers found the change in the order of the causal events as interesting as the change in the causal relationships. However, toddlers did not appear to treat the appearance of any additional agentive or non-agentive candidate cause as in any way expected given that the change in the causal order can be attributed entirely to a change in the experimenter's goal-directed actions.

One possibility is that children may simply accept that psychological causality is probabilistic. That is, they may infer that changes in agents' behavior can happen without causes. Alternatively, toddlers may assume that psychological events have causes but that such causes are intrinsically hidden, insofar as they are internal to the agent. The distinction between causes (for physical events) and reasons (for agent's behavior) has attracted considerable interest in the philosophical literature (e.g., Davidson, 1963) and it would be interesting to know to what extent even young children distinguish physical and psychological events on this basis.

Consistent with many other studies (Bonawitz, et al., 2010; Muentener & Carey, 2010; Newman et al., 2010; Max & Xu, 2013; Saxe et al., 2005; 2007) toddlers in our study appeared to be willing to accept dispositional agents, but not non-agents, as plausible candidate causes. Moreover, in our study toddlers inferred the presence of an unobserved agent even though the range of possible agent actions the agent might have performed was indeterminate: there were many things the agent might have done to alter the outcome and the child could not know what specific actions the agent might have performed. This is consistent with the possibility that toddlers treat dispositional agents as causal placeholders: given an otherwise unexplained event, children may infer that "an agent did it" even if they do not know precisely how or what the agent might have done.

In this sense, toddlers' inferences are consistent with a tendency to refer to agents as placeholder causes quite broadly. Gods, angels, and demons are invoked as causes of ill-understood phenomena cross-culturally (Barrett & Keil, 1996; Barrett, 2000; Boyer, 1996; Guthrie, 1993). Piaget noted a similar phenomenon in school-age children. When asked to explain events that were too complicated for them to explain with respect to physical causality, children often appealed to psychological causation, sometimes attributing agency to the

phenomenon itself (e.g., “The river flows because it wants to”; Piaget, 1951). In such cases, the appeal to agency betrays the absence of a genuine understanding of the phenomenon. However, the appeal to agentic causes also reflects a commitment to explanation itself: to the idea that there is some underlying cause of the events, even if the details are left to conjecture. The current results extend the previous literature in suggesting that, even if in all other respects events appear to have known observable candidate causes, the class of events that children treat as “unexplained” includes events that generate outcomes probabilistically. This suggests that toddlers’ causal reasoning goes well beyond the evidence they observe. Given sparse data for a novel probabilistic causal relation, toddlers selectively attend to hidden, latent variables that might help them better explain events in the world.

References

- Aslin, R. N. (2007). What's in a look?. *Developmental Science*, *10*(1), 48-53.
- Barrett, J. L. (2000). Exploring the natural foundations of religion. *Trends in Cognitive Sciences*, *4*(1), 29-34.
- Barrett, J. L., & Keil, F. C. (1996). Conceptualizing a nonnatural entity: Anthropomorphism in God concepts. *Cognitive Psychology*, *31*(3), 219-247.
- Bauer, P. J., & Mandler, J. M. (1989). One thing follows another: Effects of temporal structure on 1-to 2-year-olds' recall of events. *Developmental Psychology*, *25*(2), 197.
- Bauer, P. J., & Mandler, J. M. (1992). Putting the horse before the cart: The use of temporal order in recall of events by one-year-old children. *Developmental Psychology*, *28*(3), 441.
- Bonawitz, E. B., Ferranti, D., Gopnik, A., Meltzoff, A. Woodward, J., & Schulz, L.E. (2010). Just do it? Toddlers' ability to integrate prediction and action in causal inference. *Cognition*, *115*, 104-117.
- Boyer, P. (1996). What makes anthropomorphism natural: Intuitive ontology and cultural representations. *Journal of the Royal Anthropological Institute*, 83-97.
- Bullock, M., Gelman, R., & Baillargeon, R. (1982). The development of causal reasoning. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 209–254). New York, NY: Academic Press.
- Chandler, M. J., & Lalonde, C. E. (1994). Surprising, miraculous, and magical turns of events. *British Journal of Developmental Psychology*, *12*, 83–95.

- Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23), 685-700.
- Denison, S., Reed, C., & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: Evidence from 4.5-and 6-month-olds. *Developmental Psychology*, 49(2), 243.
- Denison, S., & Xu, F. (2010a). Twelve-to 14-month-old infants can predict single-event probability with large set sizes. *Developmental Science*, 13(5), 798-803.
- Denison, S., & Xu, F. (2010b). Integrating Physical Constraints in Statistical Inference by 11-Month-Old Infants. *Cognitive Science*, 34(5), 885-908.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99(24), 15822-15826.
- Garvin, L. E., & Woodward, A. L. (2015). Verbal framing of statistical evidence drives children's preference inferences. *Cognition*, 138, 35-48.
- Gelman, S. A., Coley, J. D., & Gottfried, G. M. (1994). Essentialist beliefs in children: The acquisition of concepts and theories. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture*. New York: Cambridge University Press.
- Gelman, S. A., & Gottfried, G.M. (1996). Children's Causal Explanations of Animate and Inanimate Motion. *Child Development* 67(5), 1970-1987.
- Gottfried, G. M., & Gelman, S. A. (2005). Developing domain-specific causal-explanatory frameworks: The role of insides and immanence. *Cognitive Development*, 20(5), 137-158.
- Guthrie, S. (1993). *Faces in the Clouds*. Oxford University Press.

- Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, 332(6037), 1524-1524.
- Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not So Innocent: Toddlers' Inferences About Costs and Culpability. *Psychological Science*, 26(5), 633-640.
- Johnson, S. P., Fernandes, K. J., Frank, M. C., Kirkham, N., Marcus, G., Rabagliati, H., & Slemmer, J. A. (2009). Abstract rule learning for visual sequences in 8-and 11-month-olds. *Infancy*, 14(1), 2-18.
- Kirkham, N. Z., Slemmer, J. A., Richardson, D. C., & Johnson, S. P. (2007). Location, location, location: Development of spatiotemporal sequence learning in infancy. *Child Development*, 78(5), 1559-1571.
- Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, 21(8), 1134-1140.
- Kuzmak, S. D., & Gelman, R. (1986). Young children's understanding of random phenomena. *Child Development*, 559-566.
- Laplace, P.-S. (1951). *A philosophical essay on probabilities* (Truscott, F. W. & Emory, F. L. Trans.). New York: Dover. (Original work published 1814.)
- Leslie, A. M. (1984). Infant perception of a manual pick-up event. *British Journal of Developmental Psychology*, 2, 19-32.
- Lewkowicz, D. J. (2004a). Perception of serial order in infants. *Developmental Science*, 7(2), 175-184.
- Lewkowicz, D. J. (2004b). Serial order processing in human infants and the role of multisensory redundancy. *Cognitive Processing*, 5(2), 113-122.

- Luo, Y., & Baillargeon, R. (2005). Can a self-propelled box have a goal? Psychological reasoning in 5-month-old infants. *Psychological Science, 16*, 601–608.
- Luo, Y., Kaufman, L., & Baillargeon, R. (2009). Young infants' reasoning about physical events involving self- and nonself- propelled objects. *Cognitive Psychology, 58*, 441–486.
- Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences, 104*(50), 19751-19756.
- Ma, L., & Xu, F. (2013). Preverbal infants infer intentional agents from the perception of regularity. *Developmental Psychology, 49*(7), 1330.
- Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science, 283*(5398), 77-80.
- Markson, L., & Spelke, E. S. (2006). Infants' rapid learning about self-propelled objects. *Infancy, 9*, 45–71.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*, Henry Holt and Co. Inc., New York, NY.
- Meltzoff, A. N., Waismeyer, A., & Gopnik, A. (2012). Learning about causes from people: observational causal learning in 24-month-old infants. *Developmental Psychology, 48*(5), 1215.
- Muentener, P., Bonawitz, E., Horowitz, A., & Schulz, L. (2012). Mind the gap: Investigating toddlers' sensitivity to contact relations in predictive events. *PLoS ONE, 7*(4).
- Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology, 61*, 63–86.
- Muentener, P., & Schulz, L. (2014). Toddlers infer unobserved causes for spontaneous

- events. *Frontiers in Psychology*, 5, 1-9.
- Newman, G., Keil, F., Kuhlmeier, V., & Wynn, K., (2010). Sensitivity to Design: Early Understandings of the Link Between Agents and Order. *Proceedings of the National Academy of Sciences*, 107, 17140–17145.
- Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs?. *Science*, 308(5719), 255-258.
- Piaget, J. (1951). *The child's conception of physical causality*. London: Routledge & Kegan Paul.
- Piaget, J., & Inhelder, B. (1975). *The origin of the idea of chance in children*. New York: Norton.
- Premack, D. (1990). Do infants have a theory of self-propelled objects? *Cognition*, 36, 1–16.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Saxe, R., Tenenbaum, J., & Carey, S. (2005). Secret agents: 10 and 12-month-olds infer an unseen cause of the motion of an inanimate object. *Psychological Science*, 16(12), 995–1001.
- Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: infants identify the causal agent in an unseen causal interaction. *Developmental Psychology*, 43(1), 149–158.
- 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. E., & Sommerville, J. (2006). God does not play dice: Causal determinism in children's inferences about unobserved causes. *Child Development*, 77(2), 427–442.

- Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42*, 1103-1115.
- Song, H. J., Onishi, K. H., Baillargeon, R., & Fisher, C. (2008). Can an agent's false belief be corrected by an appropriate communication? Psychological reasoning in 18-month-old infants. *Cognition, 109*(3), 295-315.
- Southgate, V., Senju, A., & Csibra, G. (2007). Action anticipation through attribution of false belief by 2-year-olds. *Psychological Science, 18*(7), 587-592.
- Spelke, E. S., Phillips, A., & Woodward, A. L. (1995). Infants' knowledge of object motion and human action. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 44-78). New York: Clarendon Press/Oxford University Press.
- Surian, L., Caldi, S., & Sperber, D. (2007). Attribution of beliefs by 13-month-old infants. *Psychological Science, 18*(7), 580-586.
- Téglás, E., Girotto, V., Gonzalez, M., & Bonatti, L. L. (2007). Intuitions of probabilities shape expectations about the future at 12 months and beyond. *Proceedings of the National Academy of Sciences, 104*(48), 19156-19159.
- Téglás, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L. L. (2011). Pure reasoning in 12-month-old infants as probabilistic inference. *Science, 332*(6033), 1054-1059.
- Vredenburg, C., Kushnir, T., & Casasola, M. (2014). Pedagogical cues encourage toddlers' transmission of recently demonstrated functions to unfamiliar adults. *Developmental Science*.

- Walker, C. M., & Gopnik, A. (2014). Toddlers infer higher-order relational principles in causal learning. *Psychological Science, 25*(1), 161-169.
- Wang, S., Kaufman, L., & Baillargeon, R. (2003). Should all stationary objects move when hit? Developments in infants' causal and statistical expectations about collision events. *Infant Behavior and Development (special issue), 26*, 529-568.
- Wellman, H. M. (1990). *The child's theory of mind*. Cambridge, MA: MIT Press.
- Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology, 43*(1), 337-375.
- Wellman, H. M., & Wooley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition, 35*, 245-275.
- Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition, 13*(1), 103-128.
- Xu, F., & Denison, S. (2009). Statistical inference and sensitivity to sampling in 11-month-old infants. *Cognition, 112*(1), 97-104.
- Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences, 105*(13), 5012-5015.
- Yu, Y., & Kushnir, T. (2014). Social context effects in 2-and 4-year-olds' selective versus faithful imitation. *Developmental Psychology, 50*(3), 922.


Figure Captions

Fig. 1. Procedure of the experiments.

Fig. 2. Toddlers' looking times to the test trials in Experiments 2-4. * $p < .05$, ** $p < .01$. Error bars indicate 95% confidence intervals of means based on 10000 bootstrap samples.


Initial Display

A box with two handles.



Familiarization Trial

First the hand presses the red handle and the lollipop pops up; then the hand presses the green handle and the cake pops up.



Red handle – Lollipop

This sequence is repeated three times.




Green handle – Cake

Switch Trial

Deterministic Condition
(Switch the order.)



Green handle – Cake




Red handle – Lollipop

Probabilistic Condition
(Switch the relation.)



Red handle – Cake




Green handle – Lollipop


EXP 1 (Manipulation check)

Test Trial


When the hand presses the red handle, a human hand (EXP 2), a novel agent (EXP 3) or a non-agent (EXP 4) emerged from inside the box.



EXP 2



EXP 3



EXP 4

