Child Development, March/April 2008, Volume 79, Number 2, Pages 395–410

Judicious Imitation: Children Differentially Imitate Deterministically and Probabilistically Effective Actions

Laura E. Schulz and Catherine Hooppell
Massachusetts Institute of Technology

Adrianna C. Jenkins
Harvard University

Three studies look at whether the assumption of causal determinism (the assumption that all else being equal, causes generate effects deterministically) affects children's imitation of modeled actions. These studies show even when the frequency of an effect is matched, both preschoolers ($N = 60$; $M = 56$ months) and toddlers ($N = 48$; $M = 18$ months) imitate actions more faithfully when modeled actions are deterministically rather than probabilistically effective. A third study suggests that preschoolers' ($N = 32$; $M = 58$ months) imitation is affected not just by whether the agent's goal is satisfied but also by whether the action is a reliable means to the goal. Children's tendency to generate variable responses to probabilistically effective modeled actions could support causal learning.

Imagine that every time your Uncle Robbie makes a soufflé it rises perfectly, but when your Uncle Sam makes a soufflé, sometimes it rises and sometimes it falls. Although you might learn to cook both by observing Uncle Sam's failures and by observing Uncle Robbie's successes, if you were learning from Uncle Robbie, you would probably imitate his technique faithfully, while if you were learning from Uncle Sam, you might be inclined to vary the recipe. That is, the precision with which you imitate an observed action might be affected by your beliefs about the efficacy of the action; optimal learning might depend on knowing when to imitate and when to explore. In this study, we look at whether a similar proposition is true for young children: Do children differentially imitate deterministically and probabilistically effective actions?

Previous research on children's imitation raises a puzzle. On the one hand, children are very good at reproducing modeled actions. Indeed, in some contexts, children will faithfully copy even arbitrary, unnecessary actions. For instance, children will imitate elaborate, causally irrelevant routines to open a box even when the mechanisms that could be used to open the box directly are obvious (Horner & Whiten, 2005; Lyons, Young, & Keil, 2006; Whiten, Custance, Gomez, Teixidor, & Bard, 1996). Similarly, children will copy an actor who, for no apparent reason, activates a toy with her head, even though the children can (and often do) also activate the toy with their hands (Gergely, Bekkering, & Király, 2002; Meltzoff, 1988).

On the other hand, children will sometimes override modeled actions in order to generate their own means to inferred ends. Thus for instance, toddlers do not copy actions that fail to achieve the agent's intended outcome (Meltzoff, 1995). If an adult pulls on a barbell toy but does not pull it apart, 18-month-olds do not imitate the ''failed'' action. Instead, they act to achieve the inferred goal of the action: They pull apart the toy. Critically, children do not ''read through the goals'' of the action if the action is performed by a machine rather than a person; the physical action by itself is not sufficient to lead children to complete the causal sequence.

However, it is equally striking, and to our knowledge, less often noted, that children's only cue that the action failed to fulfill the adult's intention was the failure of the expected causal relationships (e.g., between pulling and separating). That is, the children were not given any linguistic or affective cues suggesting that the action failed to achieve the agent's goals (the actor did not say ''whoops'' or frown); the children simply saw that the actor pulled on the toy

This research was supported by a James H. Ferry Jr. Fund for Innovation in Research Education grant and a McDonnell Foundation Collaborative Initiative Causal Learning grant to L.E. S. Thank you to Elizabeth Baraff Bonawitz, Darlene Ferranti, Noah Goodman, Andrew Shtulman, and Josh Tenenbaum for helpful comments and suggestions—and particular thanks to Rebecca Saxe for help in clarifying the analysis. We are also very grateful to the Discovery Center at the Museum of Science, Boston and participating parents and children.

Correspondence concerning this article should be addressed to Laura E. Schulz, 46-4011, Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, 77 Massachusetts Avenue, Cambridge, MA 02139. Electronic mail may be sent to lschulz@mit.edu.

© 2008, Copyright the Author(s)
Journal Compilation © 2008, Society for Research in Child Development, Inc.
All rights reserved. 0009-3920/2008/7902-0010

and the toy remained intact. Thus, the only cue that the action "failed" to achieve the agent's goals was the absence of a causal relationship that the child expected (e.g., between pulling and separating). In the context of intentional action, children seem to be able to use their knowledge of the causal structure of an event to infer the intentional structure of the event.

The puzzle then concerns the role of causal knowledge in children's imitative learning. Why do children sometimes seem to suspend their own causal knowledge in order to copy modeled actions faithfully (even when there are simpler means to the end), but at other times use their causal knowledge to override modeled actions in favor of novel means to inferred ends? What predicts the fidelity with which young children reproduce modeled actions?

Gergely and colleagues have proposed that some differences in children's imitation can be explained by assuming that children respect a *principle of rational action* (Csibra & Gergely, 1998; Gergely & Csibra, 1997, 2003; Gergely, Nadasdy, Csibra, & Biro, 1995). That is, children may assume that intentional actions performed by rational agents are optimal within the constraints of the situation. If children observe actions that appear to be suboptimal (e.g., an adult using her head to activate a toy or using arbitrary routines to open a box), they may nonetheless assume that "there must be a good reason" (Gergely et al., 2002) for the agent's choice of action. In the absence of an obvious "good reason" for the action, children might assume the action has an unobserved causal relationship to the effect or they might revise their understanding of the agent's goals (e.g., they might assume the goal was to demonstrate a convention or ritual). In either case, children should imitate the modeled action faithfully.

However, situational constraints can provide an obvious "good reason" for the modeled action. As elegantly demonstrated by an extension of Meltzoff's (1988) light-activation paradigm, if the situational constraints on the actor and the child are different, the child may not imitate the modeled action. In the new paradigm, the actor again used her head to activate a toy; however, this time, the actor's hands were occupied holding a blanket. In this case, children did not imitate the head action; instead, they activated the toy with their hands (Gergely et al., 2002). The situational constraint (holding the blanket) provided an explanation for the actor's unusual action; thus, we suggest *screening-off* (Reichenbach, 1956) a causal role for the particular means used to achieve the dominant goal (activating the light). Because the child was not also imitating the ancillary goal of holding a blanket, this analysis freed the child to (a) infer the causal structure of the main event (depressing the button makes the light go on) and then (b) achieve the actor's dominant goal by novel but simpler means (using their hands).

We are sympathetic to the idea that children adopt a principle of rational action, but we note that this proposal does not resolve our original puzzle. The claim that children assume that rational agents act optimally with respect to situational constraints is prima facie incompatible with the fact that children recognize that adults' actions are sometime ineffective. As researchers have noted, "judgments about the 'rationality' of means always translate into judgments of 'efficacy'" (Gergely & Csibra, 2003, p. 290). If children assume that agents perform the most rational action available given the constraints of the situation, it is difficult to understand how children might construe modeled intentional actions as "failed" actions. Why would a child assume that an adult who activates a toy with her head (instead of her free hands) is acting optimally, but that an adult who pulls on a toy but fails to separate it is not?

It is tempting to conclude that children assume that modeled actions are optimal when the actions achieve the agent's goal and not when they fail. Note however, that this presumes that children can simply "read off" the success or failure of the agents' goal from the sequence of events. This may indeed be the case when the agent provides explicit linguistic and affective cues about whether or not her goal has been achieved (e.g., "There!" or "Whoops!"; Carpenter, Akhtar, & Tomasello, 1998). However, if the agent pulls on a toy, it is possible to infer that the agent *failed* in her goal to separate the toy, but it is also possible to infer that the agent *succeeded* in her goal to pull on the toy. Critically, if children always assume that agents act optimally, the inference that the adult succeeded should be the preferred inference. That is, under the assumption that adults always take the most rational action given the situational constraints, children should not infer that they could improve upon the observed action.

Here, we suggest that although children do assume that agents act rationally with respect to their goals, they do not make this assumption uncritically. We suggest that children analyze goal-directed actions in the context of their broader causal knowledge. As discussed, there is considerable evidence that, given common situational constraints, children faithfully imitate arbitrary, causally irrelevant actions (activating a toy with their heads, engaging in elaborate rituals to open a box). We suggest that this is because arbitrary actions are, by definition, actions about which children have few prior expectations. If children do not have sufficient prior causal knowledge to

evaluate the efficacy of the modeled actions, we expect that children will adopt a principle of rational action and assume the adult actions are optimal. Provided the modeled actions are not screened off by a known relationship to an ancillary goal, children should imitate such actions faithfully. Because children's tendency to imitate arbitrary actions has been well established by previous research, we will not replicate that aspect of our analysis here.

However, if children do have sufficient prior knowledge to evaluate the relationship of the modeled action to the goal, we predict that children will imitate the modeled action faithfully only if they construe the action as an optimal means to the inferred end. Here, we focus on a fundamental, context-independent criterion for the optimality of an action: whether or not the action is construed as a reliably effective means to the inferred end. (In previous studies, researchers have focused primarily on whether actions were optimal with respect to heuristics such as the ''familiarity'' or ''naturalness'' of the action; Gergely et al., 2002.) Given familiar, nonarbitrary relationships between modeled means and ends, we believe children's own causal judgments about the efficacy of the action can override the assumption that the agents' actions are optimal. If children believe the modeled action is not an effective means to the inferred end, we predict that children will innovate their own means to the inferred end rather than imitate the modeled behavior.

There is suggestive evidence that children do use their background causal knowledge to identify actions that are effective means to intended goals. For instance, when children are shown a correct and incorrect solution to a novel problem (retrieving a toy through either a blocked or an open hole in a tube), they selectively imitate the more effective action (Want & Harris, 2001). Additionally, children tend to faithfully imitate actions that enable other actions (e.g., putting a horse on a panel and then tipping the panel to make a rocking horse) but they displace or omit causally irrelevant actions (patting the horse's mane; Bauer, 1992; Bauer & Hertsgaard, 1993; Bauer & Mandler, 1989). Research also suggests that children are sensitive to hierarchies of goal-directed actions; when children fail to imitate modeled actions faithfully, they tend to err more often on subordinate goals than on dominant goals (Bekkering, Brass, Woschina, & Jacobs, 2005; Bekkering, Wohlschläger, & Gattis, 2000; Carpenter, Call, & Tomasello, 2005; Gergely et al., 2002; though see Brindley, Bird, & Heyes, 2008, for an opposing view). For example, children who infer that the actor's goal is to reach for a left or right dot rarely err in imitating the choice of dot but occasionally err in imitating the choice of the ipsilateral or contralateral hand. This suggests that children distinguish actions that are causally relevant (e.g., direction of reach) and irrelevant (choice of hand) to the dominant goal.

Similarly, we suggest that the results of Meltzoff's (1995) seminal study were predicated on children's considerable prior knowledge about both the modeled actions and outcomes. The stimuli were designed to support robust inferences about both the anticipated and the observed effects of the modeled action (e.g., that pulling would lead to separating and that the toy did not separate because the adult did not pull hard enough). The claim that children have such prior causal knowledge is substantiated by the finding that toddlers can produce the target action (pulling) even when they are shown only the initial and end state of the toy (Huang, Heyes, & Charman 2002). We suggest that such robust prior causal knowledge allows children to evaluate the efficacy of the agents' actions, both to infer the agents' goal and to identify more effective means of achieving the goal.

However, there is little direct evidence that children can use an independent causal analysis of an event to evaluate the efficacy of modeled actions. In particular, there is no evidence that children faithfully imitate actions they construe as effective but innovate their own means to ends when they construe the modeled action as ineffective. In the current set of experiments, we look at contexts in which we expect children to have strong prior causal assumptions and we look at how children's judgments of causal efficacy affect the fidelity with which they imitate modeled actions. We predict that if children construe a modeled action as an optimally effective means to the inferred goal, they will imitate the action faithfully; if they believe the observed action is not reliably effective they will be more likely to pursue their own means to the inferred end.

It is important to note that our analysis implies that children are actively interpreting the *modeled* actions and outcomes. The analysis does not depend on children's knowledge about *other* means to the inferred end. Children may know, for instance, many ways to activate a button or open a box and yet not have strong expectations (because the tasks were designed to circumvent such knowledge) about the relationship between head movements and button activation or between arbitrary rituals and box opening. In contrast, children do have prior knowledge about the relationship between pulling on an object and the object coming apart. When children have such relevant causal knowledge, we believe they to bring it to bear in analyzing the efficacy of the modeled means

to the goal. One interesting implication of this account (because it does not depend on children's knowledge of alternative means to the intended outcome) is that children should not faithfully imitate actions they construe as ineffective, *even if they do not know alternative, more effective means to the end*. That is, we predict that children will explore novel actions rather than imitate modeled actions that they believe are unreliable.

Although previous studies investigating the effect of causal knowledge on children's imitation have focused primarily on children's understanding of physical mechanisms and affordances (e.g., pull-apart toys, Huang et al., 2002; Meltzoff, 1995; rakes, Nagell, Olguin, & Tomasello, 1993; trap tubes, Want & Harris, 2002; bolts and latches, Horner & Whiten, 2005; Lyons et al., 2006; Whiten et al., 1996; balls and cups; Bauer, 1992), here, we investigate how children's imitative learning is affected by more fundamental, abstract causal beliefs. In particular, we look at whether children's imitation is affected by the assumption that physical causes generate effects deterministically.

Previous research suggests that preschoolers are causal determinists; when causes appear to act probabilistically, children infer the existence of unobserved causes (Schulz & Sommerville, 2006). If children believe that, all else being equal, physical causes always generate their effects, then they should construe deterministically effective actions as optimally effective. All else being equal, children should imitate deterministically effective actions faithfully. However, children should not accept that physical causes might generate effects probabilistically. If children observe an intervention that generates effects only some of the time, they should believe the effect could, in principle, be generated more reliably. Thus, if children observe an action that is only probabilistically effective, they should be less likely to imitate the modeled action and more likely to explore alternative actions. Note, also that if children are causal determinists, then even if they do not have any *other* prior knowledge about a modeled action (e.g., the action is pushing a novel switch), the assumption of determinism should affect their analyses of the event.

In Experiments 1 and 2, we look at whether, controlling for the frequency and salience of the action/outcome relationship, children imitate deterministically effective actions more precisely than probabilistically effective actions. Earlier research on children's belief in causal determinism focused on preschoolers, so we begin (in Experiment 1) by looking at whether 4-year-olds differentially imitate deterministically and probabilistically effective actions. In Experiment 2, we replicate the study with 18-month-olds.

Two final notes are in order. First, there has been ongoing debate (see Meltzhoff, 2007; Woodward, 2005, for reviews) over whether the ontogenesis of children's understanding of goal-directed actions involves relatively rich theory of mind inferences (that agents *want* to achieve particular outcomes) or simpler teleological inferences (that agents tend to move in straight paths, avoid obstacles, etc.; Csibra & Gergely, 1998; Gergely & Csibra, 1997, 2003; Gergely et al., 1995). Because our studies involve children 18 months old and older and there is independent evidence (e.g., Bartsch & Wellman, 1995; Repacholi & Gopnik, 1997) that children of this age attribute mental states like intentions and desires to agents, we use the terms *goal* and *intention* interchangeably throughout. (Our studies do not rely on children's understanding of representational mental states like belief.) Second, in contrast to the comparative literature (Call & Carpenter, 2001; Tomasello & Call, 1997; Whiten & Ham, 1992), the developmental literature typically uses the term ''imitation'' to cover an extensive range of responses to modeled behaviors (e.g., mimicry, outcome emulation, goal emulation, and imitation). To avoid confusion in reviewing the literature, we follow this tradition and refer primarily to ''imitation'' throughout; we distinguish contexts in which children do and do not faithfully imitate modeled actions in discussion.

Experiment 1

We test preschoolers in four conditions. In the *deterministic* condition, children see a sliding switch activate a toy on each of the four trials. This evidence is consistent with the assumption that physical causes act deterministically, and we predict that children will think this is an optimally effective intervention. Given a chance to activate the toy, we predict that children will faithfully imitate the modeled action. In the *stochastic* condition, children again see the switch manipulated four times but the toy activates only on trials one and three. Because the stochastically effective actions violate the assumption that physical causes should generate effects deterministically, we predict that children will think that the modeled action is not optimally effective and will imitate the modeled action less precisely.

However, children might imitate the action in the stochastic condition less accurately than in the deterministic condition simply because the alternating pattern of success and failure disrupts the children's attention and impairs their ability to encode or recall the modeled action. To investigate children's ability to

remember the modeled action, we test a smaller group of children in a *memory* control condition. The memory control condition is identical to the stochastic condition with the single difference that at the end of the trials children are not asked to activate the toy; instead, they are simply asked to reproduce the experimenter's actions. If children are not distracted by the stochastic outcomes and are able to precisely recall the modeled action, then their performance in the memory control should be comparable to their performance in the deterministic condition.

Children might also differentially imitate the evidence in the deterministic and stochastic conditions because the two successful trials the stochastic condition reduce the salience of the action/outcome relationship or are less reinforcing than the four successful trials in the deterministic condition. Similarly, the reduced frequency of the effect might provide insufficient inductive evidence for the children to conclude that the observed action is actually a cause of the effect. When the action generates the outcome only twice, children might be less likely to treat the modeled action as a cause of the outcome.

To rule out the possibility that children imitate actions less faithfully when an effect occurs twice on alternating trials rather than four times on consecutive trials, we test children in a *frequency* control condition. This control condition is similar to the stochastic condition except that the children are given a deterministic explanation for the pattern of evidence: children learn that toy is "on" during the success trials and "off" during the failure trials. The evidence is thus matched for salience and "reward" value to the evidence in the stochastic condition but is consistent with the possibility that physical causes generate effects deterministically. Again, we predict that children will imitate the modeled action precisely.

Method

Participants. We recruited 60 preschoolers (mean age = 56 months; range = 48–65 months) from the Discovery Center of a metropolitan Science Museum and from urban area preschools. 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. Sixteen children were assigned to a deterministic condition, 16 to a stochastic condition, 12 to a memory control condition, and 16 to a frequency control condition. Approximately equal numbers of boys and girls participated in each condition (50% girls overall).

Materials. A freestanding toy light was used. The toy light had an on/off button in back. When the toy was turned on, it could be activated by a remote control. When activated, the toy's lights flashed. The real remote control was concealed throughout and the toy was apparently activated by a (fake) sliding switch. The sliding switch consisted of a wooden box, 25 cm × 6.9 cm × by 3.8 cm, with a center slit that ran the length of the box. A knob, 2 cm in diameter, was affixed to a screw inside the slit so that the knob could be slid along the length of the box. The top surface of the sliding switch was divided into 20 1.2 cm colored bars. Each bar was a unique color. See Figure 1.

Procedure. All children were tested individually. The experimenter placed the toy (with the button set to the "on" position) and the sliding switch on the table. She said, "I'm going to play with this toy four times and then you will get a chance to play with the toy." The experimenter slid the knob from the start position at the far end of the sliding switch (left/right position counterbalanced between children) to the middle of the sliding switch. She stopped the knob when it was centered directly over the border separating the middle two colored bars (between the zero marks in Figure 1). She said "One!" while simultaneously activating the concealed remote control. The toy activated. After approximately 3 s, the experimenter released the remote control and simultaneously returned the knob to the end of the slider. The experimenter repeated this action four times,

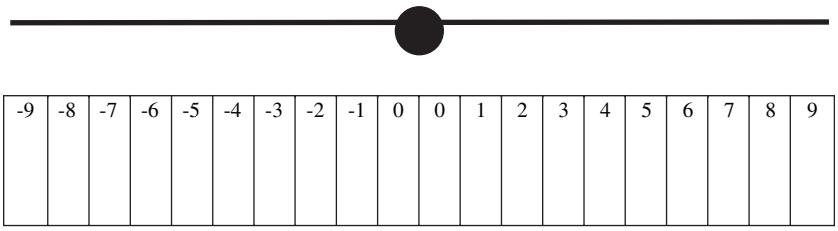

Figure 1. Schematic of the toy used in Experiment 1. (The bars are numbered to illustrate the coding scheme. The numbers below apply when the experimenter moves the knob from left to right; the numbers are reversed when the knob is moved from right to left. On the children's toy, there were no numbers and each bar was uniquely colored.)

counting aloud each time. In the deterministic condition, the experimenter activated the toy on all four trials. In the stochastic condition, the experimenter activated the toy only on Trials 1 and 3.

The evidence in the memory control condition was identical to the evidence in the stochastic condition. The frequency control condition was similar to the stochastic condition except that the toy was introduced with the button set to the "off" position. When the experimenter introduced the toy to the children, she pointed to the button in back and said, "See this? This button turns the toy on and off." On the first and third trials, she said "I'm going to turn the toy on now" and pushed the button to the "on" position before sliding the switch; on the second and fourth trials, she said "I'm going to turn the toy off now" and pushed the button to the "off" position before sliding the switch. After the fourth trial, the experimenter pushed the button back to the "on" position.

In the deterministic, stochastic, and frequency control conditions, the experimenter then said, "Now it's your turn" and passed the toy and sliding switch to the child. In the memory control condition, the experimenter first removed the toy from view. She then passed only the sliding switch to the child, saying, "Can you do exactly what I did? Can you put the knob where I put it?"

The duration of the modeled action was matched in all conditions and the experimenter did not otherwise distinguish the success and failure trials either verbally or nonverbally. Only the child's first attempt at moving the slider was coded: an "attempt" was delineated by the child either moving the slider to any position and then sliding it back to the original end or by the child moving the slider to any position and then taking her hand away. On the child's attempt, the experimenter never activated the toy. Only the first attempt was coded because (after the first attempt failed to produce an effect) children in both conditions could construe the sliding switch as stochastically effective.

Results and Discussion

Children's responses were videotaped. Children were coded as exactly imitating the modeled action if they placed the knob on the centerline (where the experimenter had put the knob – between the zero marks in Figure 1) or on the bar immediately to the left and right of the centerline. (Pilot work suggested that if no action was modeled, preschoolers acted on the toy by twirling the knob; in this study, all children's first response was to slide the knob. Thus, in all conditions, children's responses were distinct from baseline activity.) Exact imitation received a score of zero. Other responses were coded based on their deviation from the modeled action (i.e., children who stopped the knob one bar past the center bars received a score of +1; children who stopped the knob four bars before the center bars received a score of −4). The second author coded the number of the colored bar where the child placed the knob; a blind coder recoded 50% of the data. (Note, the actual sliding switch had only colored bars on it; it did not have printed numbers. The experimenter noted the color of the bar and then assigned it a number based on the criteria above.) Intercoder agreement was high (Cohen's $\kappa = .82$) disputes were resolved conservatively (i.e., the score closer to zero was used in the stochastic condition; the score further from zero was used in all other conditions).

We compared the number of children who exactly imitated the modeled action in each condition and the absolute value of children's scores in each condition. In the stochastic condition, 4 of the 16 children (25%) exactly imitated the modeled action. By comparison, in both the deterministic condition and the frequency control condition, 10 of the 16 children (62%) exactly imitated the modeled action. Similarly, in the memory control condition, 10 of the 12 children (83%) reenacted the modeled action exactly.

As predicted, children were less likely to imitate the modeled action exactly in the stochastic condition than in the deterministic condition, $\chi^2(1, N = 32) = 4.57$, $p < .05$. The differential imitation cannot be explained by a difference in children's encoding or recall for the modeled action or effects due to the reduced frequency of the outcome because children were also less likely to imitate the modeled action exactly in the stochastic condition than the memory, $\chi^2(1, N = 28) = 9.33$, $p < .01$, and the frequency, $\chi^2(1, N = 32) = 4.57$, $p < .05$, control conditions. There was no difference in children's tendency to perform the exact modeled action in the deterministic condition and the frequency control, $\chi^2(1, N = 32) = 0$, *ns*, the deterministic condition and the memory control, or the memory control and the frequency control, $\chi^2(1, N = 28) = 1.46$, *ns* in both cases.

Children received a mean score of 2.72 in the stochastic condition compared with a mean score of 1.75 in the deterministic condition, .75 in the memory control, and .41 in the frequency control. Children's scores were significantly higher in the stochastic condition than in the deterministic condition (one-tailed Mann-Whitney $U = 176$, $N = 32$, $p < .05$), the memory control (one-tailed Mann-Whitney $U = 42.5$, $N = 28$, $p < .005$), and the frequency control (one-tailed Mann-Whitney $U = 264.5$, $N = 32$,

$p < .001$). There was no significant difference between children's scores in the deterministic condition and the frequency control (two-tailed Mann-Whitney $U = 131.5$, $N = 32$, *ns*), the deterministic condition and the memory control (two-tailed Mann-Whitney $U = 77$, $N = 28$, *ns*), or the memory control and the frequency control (two-tailed Mann-Whitney $U = 81.5$, $N = 28$, *ns*). In all conditions, children who did not imitate the modeled action exactly were just as likely to stop the knob prematurely as to overshoot the modeled action (46% of nonimitators received scores < 0; 54% received scores > 0; $N = 26$, *ns* by binomial test).

Overall, the results of Experiment 1 suggest that children differentially imitate deterministically and stochastically effective actions. The results are consistent with the idea that children expect causes to generate effects reliably; they thus construe deterministically but not stochastically effective actions as optimally effective and precisely imitate those actions. Children imitate stochastically effective actions with less fidelity.

From this experiment, it is not clear whether children exhibit more variable responses to stochastically effective actions because they are less motivated to copy actions that are only effective some of the time or because they are more motivated to explore alternative, potentially more reliable actions. Further research might clarify the precise motivation behind children's differential imitation. In this article, we offer a computational level account (i.e., addressing goals and logic of the behavior) of how children's causal knowledge affects their imitation rather than an account at the level of the representational algorithm (i.e., how this logic might be implemented; Marr, 1982). It is important to note, therefore, that both motivations result in an equivalent, adaptive outcome. Whether children are motivated by dissatisfaction with the modeled action or by interest in exploring other actions, their tendency to generate variable responses to probabilistically effective actions will increase the probability that they discover more reliable means to the end.

Importantly, the control conditions rule out several relatively less interesting explanations for children's differential imitation. One concern was that children might have more difficulty encoding or recalling actions when effects occur stochastically than when they occur deterministically. That is, the pattern of alternating success and failure might disrupt children's attention to the action on the slider (they might instead attend more to the toy) and degrade children's memory for the modeled action. Critically, however, children's accurate reproduction of the modeled actions in the memory control condition suggests that children do not have difficulty remembering the action when they observe probabilistic evidence. Although children in the memory control observed the very same evidence as children in the stochastic condition, they performed as accurately as children in the deterministic condition; children's ability to recall the precise action was not impaired by the probabilistic evidence.

Another concern was that the reduced frequency of the effect might impair children's understanding of the relationship between the action and the outcome. The two reinforced trials might be less salient or less reinforcing of the action/outcome relationship than the four successful trials or might provide weaker inductive evidence that the switch was a genuine cause of the effect. However, the evidence from the frequency control condition suggests that merely reducing the frequency of the effect does not impair children's ability to recognize that the sliding switch activates the toy. In the frequency control condition, the two success trials (alternating with two failed trials) sufficed for children to reproduce the modeled action precisely. Additionally, as with the memory control condition, children's success in the frequency control condition mitigates against the possibility that the alternating pattern of success and failure might have degraded children's memory for the modeled action. Overall, the results are consistent with the hypothesis that children expect physical causes to generate effects deterministically and thus faithfully imitate deterministically effective but not probabilistically effective actions.

Experiment 2

Experiment 1 suggests that the imitative behavior of preschoolers is affected by how reliably actions generate outcomes; however, much of the research on children's imitative learning has focused on much younger children. In Experiment 2, we replicate the procedure of Experiment 1 with 18-month-olds. Because we could not be certain that 18-month-olds would distinguish the verbal instructions in the test conditions from those in the memory control condition (and because success in the frequency control condition suggests that the modeled action is indeed recalled), we eliminated the memory control condition from Experiment 2.

Method

Participants. Forty-eight toddlers ($M = 18$ months; range = 15–21 months) were tested at the Discovery

Center of a metropolitan Science Museum. 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. Sixteen children were assigned to a deterministic condition, 16 to a stochastic condition, and 16 to a frequency control condition. Approximately equal numbers of boys and girls participated in each condition (45% girls overall). Three toddlers in the deterministic condition, 3 in the stochastic condition, and 4 in the frequency control condition were dropped from the study and replaced due to the experimenter's accidental activation of the buzzer prematurely (1 child in each condition) or the child's unwillingness to touch the toy (the remaining 5 children).

Materials. The toy light from Experiment 1 was not used in Experiment 2. Instead, a concealed buzzer was used to create the illusion that the slider itself made noise. Additionally, the sliding switch from Experiment 1 was modified to make it less distracting for toddlers. The top surface of the toy was divided into five (rather than 20) uniquely colored 5 cm regions. Each region was further subdivided into four 1.2 cm regions by black lines (identical to the regions in Experiment 1). Instead of the knob, a 2-cm-long cylindrical bead was threaded onto a wooden dowel that ran the length of the switch so that the bead could be moved along on top of the center slit. In the frequency control condition, a separate button with "on" and "off" positions marked was also used.

Procedure. The procedure in the deterministic and stochastic conditions was identical to the procedure in Experiment 1, except that in each condition, the experimenter activated the switch by sliding the center of the bead to the center of the switch while activating the concealed buzzer. After approximately 4 s, she slid the bead back to the end of the switch and simultaneously released the buzzer. Because the sliding switch itself produced the effect (the buzzing noise) and there was no separate toy, we referred to the sliding switch as the "toy" throughout.

In the frequency control, the experimenter introduced the on/off button to the child before the first trial. She said, "See this? This makes my toy go! Now I'm going to turn the toy on!" The experimenter pushed the button into the "on" position and said "On!" After the first trial, the experimenter said, "Now I'm going to turn the toy off!" and pushed the button into the "off" position, saying "Off!" The button was similarly turned "on" for the third trial and "off" for the fourth and final trial. After the fourth trial, the experimenter turned the button "on" as before and then passed the toy to the child, saying, "Now you get to play with the toy."

Results and Discussion

Children's responses were coded as in Experiment 1. Intercoder agreement was high (Cohen's $\kappa = .88$); disputes were resolved conservatively. Unsurprisingly, as a group, the toddlers' imitative responses were less exact than the preschoolers'. Using the criteria set for preschoolers, toddlers rarely produced exact imitations of the modeled action; across all conditions, 23% of the toddlers scored 0 compared with 57% of the preschoolers, $\chi^2(1, N = 108) = 4.93$, $p < .05$. We therefore used more liberal criteria for coding imitative behavior in the toddlers. Responses in the middle of the sliding switch (i.e., with an absolute value between 0 and 4) were considered approximate imitation; responses in the far two quartiles of the sliding switch (i.e., with an absolute value between 5 and 9) were considered imprecise imitation. (Again, responses in all conditions were different from baseline responding: If no action was modeled, toddlers' modal response was to lift the entire switch; in this study, all children's first response was to slide the knob.)

In the stochastic condition, 5 of the 16 children (31%) approximately imitated the modeled action, whereas 11 of the 16 (69%) imitated imprecisely. By comparison, in both the deterministic condition and the frequency control condition, 12 of the 16 children (75%) approximated the modeled action and only 4 of the 16 (25%) imitated imprecisely. Children were significantly less likely to approximate the modeled action in the stochastic condition than in the deterministic condition, $\chi^2(1, N = 32) = 6.14$, $p < .025$, or the frequency control condition, $\chi^2(1, N = 32) = 6.14$, $p < .025$. There was no difference between children's responses in the deterministic condition and the frequency control, $\chi^2(1, N = 32) = 0$, *ns*.

Children received a mean score of 7 in the stochastic condition compared with a mean score of 3.5 in the deterministic condition and a score of 2.5 in the frequency control. There was a trend for children's scores to be higher in the stochastic condition than in the deterministic condition (one-tailed Mann-Whitney $U = 170$, $N = 32$, $p = .056$) and the frequency control (one-tailed Mann-Whitney $U = 166$, $N = 32$, $p = .079$). There was no significant difference between children's scores in the deterministic condition and the frequency control (two-tailed Mann-Whitney $U = 131.5$, $N = 32$, *ns*). Toddlers who approximated the modeled action were just as likely to undershoot as to overshoot the zero point (64% scored in the -4 to $-.5$

range; 36% scored in the +5 to +4 range; $N = 17$, *ns* by binomial test). By contrast, toddlers who imitated imprecisely were more likely to overshoot than undershoot the modeled action across all conditions (94% scored in the +5 to +9 range; 6% scored in the −5 to −9 range; $N = 16$, $p < .001$ by binomial test). These results corroborate the results in Experiment 1 and suggest that even 18-month-olds imitate deterministically effective actions more precisely than probabilistically effective ones.

Experiment 3

Experiments 1 and 2 suggest that children differentially imitate deterministically effective and probabilistically effective actions. In both experiments, children's accurate reproduction of the modeled action in the control conditions suggests that children do not have difficulty encoding actions associated with alternating patterns of successes and failures nor do children have difficulty inferring that an action is causally effective from the evidence of only two successful, nonconsecutive, trials. Rather, the results are consistent with the hypothesis that children expect physical causes to act deterministically and thus construe deterministically but not probabilistically effective actions as effective means to intended outcomes.

However, there are two distinct ways in which the assumption of causal determinism and the observed evidence might affect children's imitative learning: (a) by affecting children's inferences about whether the agent's goal is satisfied or (b) by affecting children's inferences about whether the modeled action is a reliable means of achieving the intended goal. The distinction between these two accounts is quite subtle. Indeed, we believe that in everyday reasoning, the inferences may be closely related: Children might infer agents' goals by inferring that agents intend their actions to generate expected effects (as when children use the knowledge that pulling causes separating to infer that an actor who pulls on a toy intends to separate it; Meltzoff, 1995). Nonetheless, there is a difference between imitating an action because the action satisfies the agent's goal and imitating an action because the action is perceived as an effective means to the goal. The former is based on a contingent fact about the action (whether or not it happens to achieve the intended outcome); the latter is based on a more stable causal inference (i.e., whether the action is a deterministically effective means of achieving the intended outcome).

In the preceding experiments, the satisfaction of the agent's goals and the efficacy of the actions were conflated: Much as children might construe a ''failed'' pulling event as a failure to achieve the agent's goals, children might construe probabilistically effective actions as failures to achieve agent's intentions. A different but related possibility is that children might not faithfully imitate the probabilistically effective actions because the agents' intentions are unclear: Given the probabilistic outcomes, children might not know what the agents' intentions actually were.

Thus, children might not faithfully imitate stochastically effective actions either because they do not believe such actions satisfy the agents' goals or because they believe such actions are ineffective means to the goals. One way to dissociate these accounts is by having the agent make her intentions transparent. If agents provide explicit (affective and linguistic) information about the desired outcome of their actions, then judgments about agents' goal satisfaction can be separated from judgments of causal efficacy.

Imagine for instance that Sally and Jane both slide a switch but Sally wants the toy to activate and Jane does not. If the toy activates whenever either Sally or Jane slide the switch to the middle (the deterministic condition), then Sally will always be satisfied with the outcome of her actions and Jane will never be satisfied with the outcome of her actions. Additionally, the switch will deterministically achieve Sally's intended outcome (activating the toy) and reliably fail to achieve Jane's intended outcome (not activating the toy). By contrast, if the toy activates when Sally slides the switch but not when Jane does (the stochastic condition), then both Sally and Jane will be satisfied with the outcome of their actions. However, although the action, as a contingent fact, always satisfies the agent who performs the action, the action itself is only stochastically effective with respect to both intended outcomes: it neither reliably activates nor fails to activate the toy.

Suppose the child is asked to adopt Sally's goal and make the toy go. There are three possibilities for how children might imitate the modeled action (see Table 1). One hypothesis is that children fail to imitate faithfully whenever a modeled action fails to fulfill an agent's goal. In the stochastic condition, both agent's goals are always satisfied; in the deterministic condition, one agent's goal is never satisfied. This suggests that the results of Experiments 1 and 2 should reverse in this experiment: children should imitate the modeled action more faithfully in the stochastic condition than in the deterministic condition. A second possibility is that children selectively attend to whether the action satisfies the goals of the agent who shares their

Table 1
Predictions for Experiment 3 Under Three Different Hypotheses About What Might Increase the Fidelity of Children's Imitation

	Deterministic	Stochastic	Prediction
Action satisfies the goals of all agents who perform the action	No	Yes	Children will imitate more faithfully in Stochastic condition than in Deterministic condition
Action satisfies the goals of the agent who shares the child's goals	Yes	Yes	Children will imitate faithfully in both conditions
Action is reliable means of achieving the shared goal	Yes	No	Children will imitate more faithfully in Deterministic than in Stochastic condition

own goals. Because both Sally and the children want to make the toy turn on, children might attend only to whether or not the action satisfies Sally's goals. If this is the case, then the children should not show differential imitation in this experiment; Sally is equally satisfied with the outcome of her actions in both conditions. The third possibility is that children attend to whether or not the action is a reliable means of achieving the shared goal. If so, the results of Experiments 1 and 2 should replicate in Experiment 3. Children should faithfully imitate the action in the deterministic condition (because the action deterministically turns on the toy) but not in the stochastic condition (because the action only sometimes turns on the toy).

In Experiment 3, we look at children's imitation and predict that children attend not only to whether actions happen to satisfy an agent's goal but also to whether actions are reliably effective means to the intended outcome. We predict that even when agents are as or more satisfied with the outcome of their actions in the stochastic condition than in the deterministic condition, children will imitate the modeled action more faithfully in the deterministic condition than in the stochastic condition.

Method

Participants. Thirty-two preschoolers (M = 58 months; range = 46–69 months) were tested at the Discovery Center of a metropolitan Science Museum. 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. Sixteen children were assigned to a deterministic condition and 16 to a stochastic condition. Approximately equal numbers of boys and girls participated in each condition (47% girls overall). One child in the stochastic condition declined to touch the sliding switch and was replaced.

Materials. The toy light from Experiment 1 was used. Although this study involved preschoolers, the sliding switch from Experiment 2 was used due to mechanical problems with the sliding switch in Experiment 1. In addition, two stuffed animal puppets were used: a horse and a dog.

Procedure. The experimenter introduced the child to each of the two animals (''This is Sally and this is Jane'' for girls; ''This is Paul and this is Joe'' for boys). Using 'Sally' as a hand puppet (particular puppet counterbalanced across children), the experimenter had Sally slide the knob to the middle of the sliding switch, as in Experiment 1. When the knob reached the middle, the experimenter (through the puppet) said ''One!'' and simultaneously activated the concealed remote control. The toy activated. After approximately 3 s, the experimenter released the remote control and simultaneously had Sally return the knob to the end of the slider. The experimenter had Sally say, ''I like the lights! I wanted the toy to turn on!'' (We had the agent express her intention after performing the action rather than beforehand because we were concerned that the reverse order might leave the impression that the agent could ''magically'' control the outcome.) The experimenter then removed Sally from view and repeated the procedure using ''Jane'' as the puppet. In the deterministic condition, the light again activated; in the stochastic condition, it did not. In either case, Jane said, ''I don't like the lights! I didn't want the toy to turn on!'' The entire procedure was then repeated. After the fourth trial, the children were asked, ''Can you show me which puppet wanted the toy to turn on?'' ''Can you show me which puppet didn't want the toy to turn on?'' (order of questions counterbalanced between children). The experimenter then passed the sliding switch and toy to the child and said, ''Can you make the toy turn on?''

Results and Discussion

All the children correctly identified which puppet did and did not want the toy to turn on. Although

children were using the sliding switch from Experiment 2 (hence, there were only 5 colored regions), there were still 20 marked bars on the switch; thus, responses were coded as in Experiment 1. Reponses were coded by the first author; a blind coder recoded 90% of the data. Intercoder agreement was high (Cohen's $\kappa = .95$); disputes were resolved conservatively.

We compared the number of children who exactly imitated the modeled action in each condition and the absolute value of children's scores in each condition. In the stochastic condition, 5 of the 16 children (31%) exactly imitated the modeled action. By comparison, in the deterministic condition, 11 of the 16 children (69%) exactly imitated the modeled action. Although the outcomes in the stochastic condition always fulfilled the agents' goals, children were less likely to imitate the modeled action exactly in the stochastic condition than in the deterministic condition, $\chi^2(1, N = 32) = 4.5, p < .05$. Children received a mean score of 4.94 in the stochastic condition compared with a mean score of .78 in the deterministic condition. Children's scores were significantly higher in the stochastic condition than in the deterministic condition (one-tailed Mann-Whitney $U = 190.5$, $N = 32$, $p < .01$). Children who did not imitate the modeled action exactly were more likely to overshoot the modeled action than to stop the knob prematurely (12% of nonimitators received scores < 0; 88% received scores > 0; $N = 16$, $p < .005$ by binomial test).

These results are consistent with the hypothesis that in imitating goal-directed actions, children attend not only to whether the actions fulfill the actor's goal but also to whether the actions reliably achieve the intended outcome. Even though the outcomes in the stochastic condition always fulfilled the agents' goals and even though the children were asked to adopt Sally's goal of activating the toy and Sally achieved her goal in both conditions, children imitated the deterministically effective actions more faithfully than the stochastically effective actions. This suggests that children's assumption of determinism affects children's inferences about effective means to inferred ends, not just children's inferences about whether or not the agent's goal is fulfilled.

Arguably, children's responses might have been affected by the change in task instructions from "It's your turn" (in Experiments 1 and 2) to "Can you make the toy go?" We needed to establish a specific goal for the children because if the children could, in principle, have adopted the goal of "not activating the toy," the predicted results (i.e., in Table 1) would differ and there would be no uniform way to interpret the data. However, we believe the instructions were functionally equivalent to the instructions in Experiments 1 and 2 in that all children spontaneously seemed to adopt the goal of turning on the toy. In all three experiments, children almost always expressed dissatisfaction when their attempts failed to make the toy go. Indeed, to avoid frustrating the children, we let them "try again" after the last trial and allowed them to activate the toy. Because, in the absence of task instructions, children spontaneously try to activate the toy, we think it is unlikely that the task instructions had a significant effect on children's responses.

Note that this experiment does not rule out the possibility that children's imitative responses were influenced by whether the actions happened to satisfy the goals of the agents performing the actions. The incongruence, for instance, between Jane's goal and the outcomes on the deterministic condition (and the congruence between her goal and the outcomes in the stochastic condition) might have highlighted the fact that the actions in the deterministic condition always turned on the toy and the actions in the stochastic condition sometimes did not. This experiment does suggest, however, that children's imitative learning does not depend only on whether or not the outcomes of actions satisfy the agents who perform them; children also attend to whether or not actions are deterministically effective means to inferred ends.

General Discussion

Overall, our results suggest that children's imitative learning is affected by the probability with which observed actions generate outcomes. Both preschoolers and toddlers faithfully imitate deterministically effective actions, consistent with the hypothesis that children expect physical causes to generate effects deterministically and, all else being equal, construe deterministically effective actions as reliable means to inferred ends. However, although children seem to have no difficulty remembering and encoding probabilistically effective actions, children are less precise in their imitation of actions that generate effects probabilistically. Indeed, even when a probabilistically effective action explicitly satisfies the goals of the agent who performs the action, children seem to infer that the action is not a reliable means to the intended outcome and do not imitate the action faithfully.

More broadly, this research suggests that children bring their broader causal knowledge to bear on their analysis of modeled actions. When children have prior knowledge about the relationship between modeled actions and outcomes (whether derived

from abstract assumptions like determinism or, as in previous studies, familiarity with particular physical mechanisms and affordances; e.g., Meltzoff, 1995), they can use this knowledge to evaluate the extent to which the modeled actions are effective with respect to the intended outcomes. If children construe a modeled action as an effective means to an inferred end, they imitate it faithfully; otherwise they generate their own means to inferred ends.

Our proposal contrasts with previous work on causal knowledge and imitative learning in several respects. First, some researchers have suggested that children's imitation depends on their knowledge of how actions are "designed to bring about [a] goal" (Tomasello, 1996, p. 323) and is characterized by "an understanding of both the behavior's goal and its strategy for achieving that goal" (Tomasello, 1996, p. 324). By contrast, we suspect that children frequently imitate modeled actions precisely because they do not understand the relationship between the modeled action and the goal. That is, we agree with other researchers (Gergely & Csibra, 2006) that, when modeled actions are arbitrary, or cognitively opaque, children respect a principle of rational action. In the absence of knowledge to the contrary, children assume adults act optimally with respect to their goals given the situational constraints; thus, children faithfully imitate modeled actions when they do not understand them.

Nonetheless, in at least some respects, we share Tomasello's (1996) perspective that children's own understanding of the relationship between the means and the goal is relevant to their imitative learning. In particular, we believe that when children have sufficient causal knowledge to evaluate the relationship between modeled actions and a goal, children will imitate the modeled actions faithfully only if they construe the modeled actions as an effective means to the goal. If they believe the modeled actions are not reliably effective, they will generate more variable, exploratory behavior.

One implication of this account is that toddlers and preschoolers can use their own knowledge of a domain to assess the rationality of adult actions and, if necessary, to override information provided by an adult. This prediction is consistent with recent research (Jaswal & Neely, 2006; Koenig, Clement, & Harris, 2004; Koenig & Harris, 2006) suggesting that children can use their independent knowledge of a domain to evaluate the reliability of adults as sources of information. If, for instance, an adult routinely provides incorrect labels for familiar referents (e.g., calling a pencil a shoe), preschoolers ignore the adult's novel label for a novel referent. However, if the adult generally provides correct labels for familiar referents, children accept the adult's novel label for the novel referent (Jaswal & Neely, 2006). We suspect that the children would also be likely to accept the adult's novel label for the novel referent had the adult previously offered only novel labels for novel referents. That is, analogously with our account, we suggest that if the children are either assured of the adults' accuracy (because it is consistent with their own knowledge) or ignorant of the adults' accuracy (because they don't know enough to judge themselves), children assume that the adults' information is reliable and worth learning. However, if the child knows enough to infer that the adult is unreliable, children will override the assumption that the modeled behavior is informative. Our study thus adds to recent evidence suggesting that even very young children evaluate the reliability of adult behavior with respect to their own understanding of a domain.

Other recent evidence is also congruent with our analysis. Studies suggest, for instance, that children will faithfully imitate modeled actions if they do not know the intended effect of the action (e.g., making a smiley face) but will generate their own means to the end once they do understand the goal (Williamson & Markman, 2006; Williamson, Meltzhoff, & Markman, in press). That is, when actions appear to be arbitrary, children imitate faithfully; however, when children have sufficient knowledge to evaluate the relationship between the means and the ends, they are not committed to the modeled action and can generate their own, more efficient, means to the end.

In accordance with recent suggestions, we find it plausible that pedagogical cues (e.g., calling the child's name, making eye contact, ostensive pointing) increase the probability that children will faithfully imitate modeled actions (e.g., Gergely & Csibra, 2003, 2006). Critically, however, our experiments show that even in pedagogically rich contexts, children do not always faithfully learn from adult actions. Although all conditions in our studies provided equivalent pedagogical cues, children did not always seem to construe the adult's intentional actions as reliably effective and thus did not always imitate the modeled action faithfully. We suggest that the assumption of rational action, even in pedagogical contexts, can be overridden by children's independent causal analysis of the event.

Our study also raises several questions. First, as noted (in the Discussion of Experiment 1), it is not clear whether children's variable imitation of probabilistically effective actions results from children's relatively low motivation to copy actions that only work some of the time or from their relatively high

motivation to look for more reliable ways to generate the effect. Future research might dissociate these motivations. However, as discussed, children's tendency to produce varied responses to stochastic causality might be advantageous for causal learning, regardless of how it is implemented. Varying their own actions from those of a stochastically effective model will increase the probability that children discover unobserved variables (i.e., the factors differentiating the successful and failed trials) and more effective ways of producing the intended outcome.

Our study was also limited in that the range of children's responses was deliberately constrained to facilitate coding. The sliding switch moved only along a single track and children were allowed only a single attempt to generate the effect. It is important, therefore, to note that in discussing differential imitation, we are focusing on relatively small changes in children's responses. Even when children did not faithfully imitate the modeled action with regard to the end point of the action, they did imitate the modeled action insofar as they manipulated the switch and ran it along the track. Importantly, this action was distinct from baseline responding (in which children tended to twirl the knob or pick up the entire switch). Thus, one possibility is that even if children believe a modeled action is not reliable, they may anchor on the modeled behavior and assume that exploring a range of behaviors around the action may be useful. In future research, a paradigm that allows children less restricted opportunities for exploration might provide further insights into the nature of both imitative and nonimitative behaviors.

Note, also, that when we discuss the ''variability'' of children's responses to probabilistically effective actions, we refer to children's deviation from the modeled action rather than their absolute variability (e.g., in the sense of the range of the switch used). In these studies, for ease of coding, we designed the switch so the affordances encouraged children to explore the full range of the switch when the modeled action was not reliably effective. However, the claim that children will exhibit more variable behavior when a modeled action is probabilistically than deterministically effective does not predict that children will necessarily exhibit less precise behavior than the behavior modeled. In principle, when the modeled action was stochastically effective, children might have gripped the knob more closely, moved it more slowly, and inspected more carefully to ensure the knob was precisely at the center (while merely copying the action as we performed it when the action was deterministically effective). Note that had children taken such pains only in the stochastic condition, this would still constitute evidence for differential imitation: children's actions would be more variable (in the sense of deviating more from the modeled action) in the stochastic condition than in the deterministic condition even if they were less variable in an absolute sense (e.g., in using a more narrow range of the switch).

We also note that even in the stochastic condition success and failure trials alternated reliably. A more strict definition of stochasticity might require a nonfixed, trial-by-trial contingency between successive states (J. Cutting, personal communication, May 12, 2006). However, for our purposes, it was sufficient that in the deterministic condition, the intervention on the observed cause fully accounted for the evidence; in the stochastic condition, the observed cause did not. We note, however, that even the deterministic condition provided only weak evidence (four trials) that the action really was deterministically effective. Given previous research suggesting that children have an initial inductive bias toward assuming that physical causes act deterministically (Schulz & Sommerville, 2006), this small amount of data presumably sufficed for children to be relatively confident that the observed action always produced the effect. However, it would be interesting to know how varying the frequency and conditional probability of the events in both conditions might affect children's imitative learning.

Although in these studies we look exclusively at children's assumption of physical causal determinism, we do not mean to suggest that children are only determinists about physical causality. It is possible that children (and adults) might be no more willing to accept genuinely indeterminate events in the psychological domain than the physical domain. Future research might investigate the extent to which the assumption of determinism holds across domains. Additionally, as we have noted elsewhere (Schulz & Sommerville, 2006), a belief in determinism might be best characterized as a belief in both of the following propositions (a) that causes generate effects deterministically (as studied here) and (b) that all events have causes. The first proposition implies that children should infer unobserved causes whenever effects appear to occur stochastically (i.e., whenever the probability of an effect given the known causes is less than 100%); the second implies that children should infer unobserved causes whenever effects appear to occur spontaneously (i.e., whenever the probability of an effect in the absence of known causes is more than 0%). Other researchers have captured a similar distinction, noting that we may assume both that the complete set of candidate causes would be sufficient

to generate an event and that a sufficient cause of an event is necessary (Gergely & Watson, 1996; Watson, 1979, 1985, 1994). Our experiment suggests that children do not imitate faithfully when the observed causes are not sufficient to generate the effect. It would be interesting to know whether children's imitative learning would also be affected if the known causes were not necessary: that is, if the effect sometimes appeared to occur spontaneously.

Our findings may also be interestingly related to proposals suggesting that even infants engage in exploratory behavior to estimate the extent to which their own actions (e.g., babbling, cooing) are both sufficient and necessary causes of maternal responses (the contingency maximization hypothesis; Gergely & Watson, 1996). It seems possible that children might generate exploratory actions until they can establish that either their own actions or some other event fully predicts maternal responses. It would be interesting to know whether infants' exploratory behavior in social domains is affected by the assumption of determinism.

Critically, although we focus on the relationship between causal knowledge and imitative learning, we do not mean to imply that causal inference underlies all instances of imitation. Much research on imitation has focused on what has been called *the correspondence problem* (Brass & Heyes, 2005; Heyes, 2001): the problem of how perception of a motor action is mapped onto performance of a comparable action. Recently, ''mirror neurons'' (Rizzolatti, Fogassi, & Gallese, 2001) have attracted considerable attention as a candidate mechanism for solving the correspondence problem by enabling such supramodal representations. Arguably a wide range of imitative behaviors, including neonatal imitation (Meltzoff & Moore, 1977), unconscious imitation of the mannerisms of conversational partners (Chartrand & Bargh, 1999), and what researchers have called empathic or altercentric imitation (as when parents feeding an infant unconsciously open their mouths when the infant does; Braten, 1988), might rely primarily on the activation of such common representations. In such cases, causal knowledge might be irrelevant. However, as noted, children do not always ''mirror'' observed actions. Previous research has demonstrated that children's imitation of goal-directed action is sensitive to children's inferences about the intentional structure of the event (Carpenter et al., 1998; Meltzoff, 1995, 2007). Here, we suggest that it is also sensitive to children's inferences about the causal structure of the modeled event.

Similarly, we do not mean to suggest that children's belief in determinism is only or even chiefly manifest in imitative learning. Children respond differentially to probabilistic and deterministic evidence in a wide range of contexts (Kushnir & Gopnik, 2005; Schulz & Sommerville, 2006). However, we believe that imitative learning is important both as a method for exploring children's causal beliefs and as a topic of study in its own right. That is, imitation paradigms can help us understand how children represent the causal structure of events—and understanding children's beliefs about the causal structure of events can help us understand how and why children imitate modeled actions the way they do. Thus, although we expect that children's fundamental assumptions about physical causal relations affect their behavior quite broadly, the relationship between causal knowledge and imitative learning is a fruitful area for investigation.

Overall, our results suggest that young children's imitative behavior is remarkably sophisticated. Children analyze goal-directed actions, not just with respect to physical affordances and visible mechanisms but also with respect to more abstract assumptions, including the assumption of causal determinism. Children's evaluation of the efficacy of the modeled action seems to affect their decision about when to faithfully imitate and when to innovate with respect to a modeled action. Such differential imitation of deterministically and stochastically effective actions could provide very young children with an adaptive mechanism for causal learning.

References

Bartsch, K., & Wellman, H. M. (1995). *Children talk about the mind*. New York: Oxford University Press.

Bauer, P. J. (1992). Holding it all together: How enabling relations facilitate young children's event recall. *Cognitive Development*, *7*, 1–28.

Bauer, P. J., & Hertsgaard, L. A. (1993). Increasing steps in recall of events: Factors facilitating immediate and long-term memory in 13.5- and 16.5-month-old children. *Child Development*, *64*, 1204–1223.

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*, 197–206.

Bekkering, H., Brass, M., Woschina, S., & Jacobs, A. M. (2005). Goal-directed imitation in patients with Ideomotor Apraxia. *Cognitive Neuropsychology, 22*, 419–432.

Bekkering, H., Wohlschläger, A., & Gattis, M. (2000). Imitation is goal-directed. *Quarterly Journal of Experimental Psychology, 53A*, 153–164.

Brass, M., & Heyes, C. (2005). Imitation: Is cognitive neuroscience solving the correspondence problem? *Trends in Cognitive Science*, *9*, 489–495.

Braten, S. (1988). Dialogic mind: The infant and the adult in protoconversation. In M. Carvallo (Ed.), *Nature, cognition and system* (Vol. 1, pp. 187–205), Dordrecht, Netherlands: Kluwer Academic.

Brindley, R., Bird, G., & Heyes, C. M. (2008). *General processes, rather than "goals," explain imitation errors*. Manuscript under review.

Call, J., & Carpenter, M. (2001). Do chimpanzees and children know what they have seen? *Animal Cognition*, *4*, 207–220.

Carpenter, M., Akhtar, N., & Tomasello, M. (1998). 14- through 18-month-old infants differentially imitate intentional and accidental actions. *Infant Behavior and Development*, *21*, 315–330.

Carpenter, M., Call, J., & Tomasello, M. (2005). Twelve- and 18-month-olds copy actions in terms of goals. *Developmental Science*, *8*, F13–F20.

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology, 76*, 893–910.

Csibra, G., & Gergely, G. (1998). The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, *1*, 255–259.

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

Gergely, G., & Csibra, G. (1997). Teleological reasoning in infancy: The infant's naïve theory of rational action. A reply to Premack & Premack. *Cognition*, *63*, 227–233.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in Cognitive Sciences*, *7*, 287–292.

Gergely, G., & Csibra, G. (2006). Sylvia's recipe: The role of imitation and pedagogy in the transmission of human culture. In: N. J. Enfield & S. C. Levinson (Eds.), *Roots of human sociality: Culture, cognition, and human interaction* (pp. 229–255). Oxford, UK: Berg.

Gergely, G., Nadasdy, Z., Csibra, G., & Biro, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*, 165–193.

Gergely, G., & Watson, J. S. (1996). The social biofeedback theory of parental affect- mirroring: The development of emotional self-awareness and self-control in infancy. *International Journal of Psychoanalysis*, *77*, 1181–1212.

Heyes, C. M. (2001). Causes and consequences of imitation. *Trends in Cognitive Sciences*, *5*, 253–261.

Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal Cognition*, *8*(3), 164–181.

Huang, C.-T., Heyes, C. M., & Charman, T. (2002). Infants' behavioral re-enactment of "failed attempts": Exploring the roles of emulation learning, stimulus enhancement and understanding of intentions. *Developmental Psychology, 38*, 840–885.

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

Koenig, M. A., Clement, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, *15*, 694–698.

Koenig, M. A., & Harris, P. L. (2006). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*, 1261–1277.

Kushnir, T., & Gopnik, A. (2005). Children infer causal strength from probabilities and interventions. *Psychological Science*, *16*, 678–683.

Lyons, D., Young, A., & Keil, F. (2006, May). *Observing an adult model can disrupt children's detection of causally irrelevant actions*. Poster presented at the annual convention of the Association for Psychological Science, New York.

Marr, D. (1982). *Vision: A computational investigation in the human representation of visual information*. San Francisco: Freeman.

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

Meltzoff, A. N. (1995). What infant memory tells us about infantile amnesia: Long-term recall and deferred imitation. *Journal of Experimental Child Psychology, 59*, 497–515.

Meltzoff, A. N. (2007). "Like me": A foundation for social cognition. *Developmental Science*, *10*, 126–134.

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, *198*, 75–78.

Nagell, K., Olguin, R. S., & Tomasello, M. (1993). Processes of social learning in the tool use of chimpanzees (Pan troglodytes) and human children (Homo sapiens). *Journal of Comparative Psychology, 107*(2), 176–186.

Reichenbach, H. (1956). *The direction of time*. Berkeley: University of California Press.

Repacholi, B., & Gopnik, A. (1997). Early understanding of desires: Evidence from 14 and 18-month-olds. *Developmental Psychology, 33*(1), 12–21.

Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, *2*, 661–670.

Schulz, L. E., & Sommerville, J. (2006). God does not play dice: Causal determinism in children's inferences about unobserved causes. *Child Development*, *77*, 427–442.

Tomasello, M. (1996). Do apes ape? In J. Galef & C. Heyes (Eds.), *Social learning in animals: The roots of culture* (pp. 319–346). New York: Academic Press.

Tomasello, M., & Call, J. (1997). *Primate cognition*. New York: Oxford University Press.

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

Want, S. C., & Harris, P. L. (2002). How do children ape? Applying concepts from the study of non-human primates to the developmental study of "imitation" in children. *Developmental Science*, *5*(1), 1–13.

Watson, J. S. (1979). Perception of contingency as a determinant of social responsiveness. In E. Thoman (Ed.), *Origins of the infant's social responsiveness* (pp. 33–64). Hillsdale, NJ: Erlbaum.

Watson, J. S. (1985). Contingency perception in early social development. In T. M. Field & N. A. Fox (Eds.), *Social perception in infants* (pp. 157–176). Norwood, NJ: Ablex.

Watson, J. S. (1994). Detection of the self: The perfect algorithm. In S. T. Parker, R. W. Mitchell, & M. L. Boccia (Eds.), *Self-awareness in animals and humans* (pp. 131–148). New York: Cambridge University Press.

Whiten, A., Custance, D. M., Gomez, J. C., Teixidor, P., & Bard, K. A. (1996). Imitative learning of artificial fruit processing in children (Homo sapiens) and chimpanzees (Pan troglodytes). *Journal of Comparative Psychology, 110*(1), 3–14.

Whiten, A., & Ham, R. (1992). On the nature and evolution of imitation in the animal kingdom: Reappraisal of a century of research. In P. J. B. Slater, J. S. Rosenblatt, C. Beer, & M. Milinski (Eds.), *Advances in the study of behavior* (Vol. 21, pp. 239–283). New York: Academic Press.

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

Williamson, R. A., Meltzoff, A. N, & Markman, E. M. (in press). Prior experiences and perceived efficacy influence 3-year-olds' imitation. *Developmental Psychology.*

Woodward, A. L. (2005). The infant origins of intentional understanding. In R. V. Kail (Ed.), *Advances in child development and behavior* (Vol. 33, pp. 229–262). Oxford, UK: Elsevier.