

The hidden structure of overimitation

Derek E. Lyons^{*†}, Andrew G. Young[‡], and Frank C. Keil^{*}

^{*}Department of Psychology, Yale University, Box 208205, New Haven, CT 06520; and [‡]Department of Psychology, University of Wisconsin, Brogden Hall, Madison, WI 53706

Edited by Susan E. Carey, Harvard University, Cambridge, MA, and approved October 18, 2007 (received for review May 11, 2007)

Young children are surprisingly judicious imitators, but there are also times when their reproduction of others' actions appears strikingly illogical. For example, children who observe an adult inefficiently operating a novel object frequently engage in what we term *overimitation*, persistently reproducing the adult's unnecessary actions. Although children readily overimitate irrelevant actions that even chimpanzees ignore, this curious effect has previously attracted little interest; it has been assumed that children overimitate not for theoretically significant reasons, but rather as a purely social exercise. In this paper, however, we challenge this view, presenting evidence that overimitation reflects a more fundamental cognitive process. We show that children who observe an adult intentionally manipulating a novel object have a strong tendency to encode all of the adult's actions as causally meaningful, implicitly revising their causal understanding of the object accordingly. This automatic causal encoding process allows children to rapidly calibrate their causal beliefs about even the most opaque physical systems, but it also carries a cost. When some of the adult's purposeful actions are unnecessary—even transparently so—children are highly prone to misencoding them as causally significant. The resulting distortions in children's causal beliefs are the true cause of overimitation, a fact that makes the effect remarkably resistant to extinction. Despite countervailing task demands, time pressure, and even direct warnings, children are frequently unable to avoid reproducing the adult's irrelevant actions because they have already incorporated them into their representation of the target object's causal structure.

causal learning | cognitive development | imitation

Much of the success of our species rests on our ability to learn from others' actions. From the simplest preverbal communication to the most complex adult expertise, a remarkable proportion of our abilities are learned by imitating those around us (e.g., refs. 1–5). Imitation is a critical part of what makes us cognitively human and generally constitutes a significant advantage over our primate relatives (6, 7). Yet for all of its usual utility, our imitative capacity also has dimensions whose benefits remain less clear. Indeed, especially in the case of young children, there are times when imitation appears to induce significant errors in reasoning.

A phenomenon that we term *overimitation* illustrates a seeming cost of our imitative prowess. Children have been observed to overimitate, or to reproduce an adult's obviously irrelevant actions (8–14), in several different contexts—even in situations where chimpanzees correctly ignored the unnecessary steps (10, 12–14). This curious contrast, however, has attracted surprisingly little interest. It has been assumed that children overimitate not for deep cognitive reasons but simply because of implicit social demands or out of imitative habit. For example, one account of overimitation emphasizes children's willingness “to copy to satisfy social motivations, to fulfill an interpersonal function of promoting shared experience with others” (ref. 15, p. 563; see also refs. 16 and 17). It is argued that this motivation for mutual social engagement causes children to approach imitation as a kind of social game, one in which they will “perform imitations of most any act modeled as a way of participating” (ref. 16, p. 7). Children, therefore, overimitate because they are more interested in the imitative interaction itself

than in the utility of the actions that they copy. Others suggest that children overimitate “because they [see] the behavior of the demonstrator as intentional, even if they did appreciate that some parts of the demonstration were causally irrelevant” (ref. 10, p. 179). That is, the intentionality of the adult's action may constitute an implicit social demand for children, leading them to infer that they are “supposed” to imitate. A final possibility is that overimitation may simply be a byproduct of habit. Overimitation may arise, in other words, because imitation “remains habitual even in a specific situation in which less fidelity would actually afford more efficiency” (ref. 14, p. 11; see also ref. 11).

These social/habitual accounts of overimitation are quite sensible, but they neglect an important alternative. Given that infants and children usually imitate selectively and rationally (18–25), might overimitation have a hidden rational structure? We hypothesized that overimitation might result from the overextension of a normally adaptive learning process, one in which children use others' actions to imitatively learn about physical causality.

Children develop in a wilderness of cultural artifacts and tools whose causal underpinnings are not just complex, but in fact often opaque to direct inspection.[§] This opacity poses a formidable challenge for children's causal learning, one that requires social catalysis to overcome (1, 2, 26, 27). We hypothesized that when children observe an adult manipulating a novel object, they may automatically (and potentially erroneously) encode all of the adult's purposeful actions as causally necessary. In other words, they may implicitly treat the adult's actions as highly reliable indicators of the object's “inner workings” or causal structure, revising their causal beliefs about the object accordingly. As adults, we recognize this learning strategy as one that we often deliberately invoke. When faced with a causally opaque device whose functions are not obvious, we frequently use others' intentional manipulations to infer causally important operations. Our proposal is that children do much the same thing, but that they do so more automatically. They treat the purposeful actions that adults direct toward novel objects as a source of privileged causal information, automatically encoding those actions[¶] as causally meaningful even when there is clear visible evidence to the contrary.

Under most circumstances, the inflexibility of this automatic causal encoding process would be amply compensated by its power. By deferring to adult action in this way, children would be able to

Author contributions: D.E.L., A.G.Y., and F.C.K. designed research; D.E.L. and A.G.Y. performed research; D.E.L. analyzed data; and D.E.L. and F.C.K. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

[†]To whom correspondence should be addressed. E-mail: derek.lyons@yale.edu.

[§]The causal opacity of devices like computers is obvious, but the problem extends to much simpler artifacts. As soon as tool use decouples goal-directed actions from immediately observable goal states (such as occurred when early hominids began to use tools recursively), causal learning quickly becomes an intractable inferential problem (1, 2, 26, 27).

[¶]Of course, even simple actions can be encoded in multiple ways, differing in the level of detail that is absorbed from the display (28, 29). We return to this issue in Experiment 1A, comparing the exact “style” of children's actions to that of the adult to determine the granularity with which the hypothesized causal encoding process occurs.

This article contains supporting information online at www.pnas.org/cgi/content/full/0704452104/DC1.

© 2007 by The National Academy of Sciences of the USA

rapidly calibrate their causal beliefs about even the most opaque physical systems. However, in cases where some of an adult's purposeful actions were not actually necessary, children would be expected to mis-encode them as causally meaningful, thus distorting their causal beliefs about the target object. These distortions could in turn explain why children reproduce irrelevant actions that even chimps ignore.

This strong hypothesis makes a strong prediction: If overimitation indeed arises from involuntary distortions in children's causal beliefs, then the effect should be unavoidable. When an adult's intentional manipulation of an unfamiliar object includes irrelevant components, children should overimitate even if situational factors strongly disfavor the copying of unnecessary actions. Contrastingly, if overimitation is caused by social cues or imitative habit, then it should be relatively easy to block the effect by opposing it with salient social and task demands. Here we report a series of studies that test these predictions, evaluating whether overimitation is indeed a superficial social phenomenon as previously believed, or rather a unique window onto the structure of children's causal learning.

Experiment 1A: Procedure, Results, and Discussion

We began by presenting 3- to 5-year-olds ($n = 63$, mean age 49 months) with a situation in which overimitation was strongly opposed by obvious social and task demands. These demands were instilled in an extensive initial training phase, where children were reinforced for identifying irrelevant actions performed by the experimenter as he opened familiar household objects. Immediately after training, children again saw the experimenter performing irrelevant actions while opening an object—this time a simple, causally transparent novel object. The question of interest was how children would open the novel object themselves. Would the opposing training demands cause them to ignore the adult's unnecessary actions or would an involuntary distortion in causal beliefs maintain overimitation?

Training phase stimuli were eight simple transparent containers of the sort that would be familiar to children [supporting information (SI) Fig. 6]. Participants watched the experimenter retrieve a toy dinosaur from each container using a sequence of relevant and visibly irrelevant actions. For example, the experimenter retrieved a dinosaur from a plastic jar (SI Fig. 6A) by first tapping the side of the jar with a feather and then unscrewing the lid. After each retrieval, participants were asked to identify which actions the experimenter “had to do” to get the dinosaur out, and which had been “silly” and unnecessary (see SI Movie 1). Children received detailed corrective feedback on their answers and were effusively praised when they correctly identified the irrelevant actions.

After training, participants moved immediately into the test phase of the experiment. Stimuli for the test phase were several novel “puzzle objects” (Fig. 1 and SI Figs. 7–10), each largely transparent such that the causal significance of actions performed on them was directly observable. After bringing a single puzzle object into the room, the experimenter sat next to the child (such that both had the same view of the object) and remarked: “Do you remember how those other containers had dinosaurs in them? Well, this thing [i.e., the puzzle object] has a toy turtle inside.” Just as in the training phase, he then retrieved the turtle using a short sequence of relevant and visibly irrelevant actions (Fig. 1 and Table 1; see Methods in SI Text for details). After showing the child the turtle, the experimenter reset the puzzle object outside of his/her view. He then said that he had to leave the room to check on something, telling the child “If you want to, you can get the turtle while I'm gone. You can get it out however you want.” The experimenter then left the room, remaining outside until the child retrieved the turtle. Each child was tested with two of the three puzzle objects, with pairings and presentation order counterbalanced.

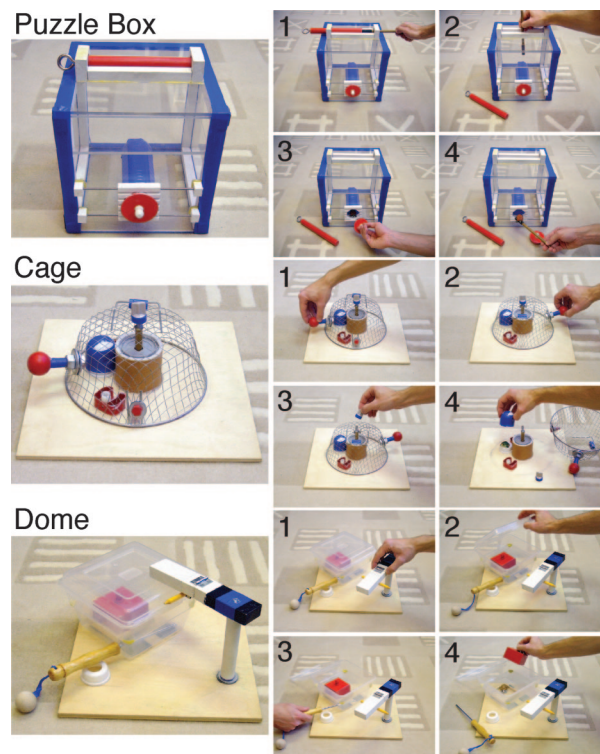


Fig. 1. The puzzle objects and examples of the corresponding experimenter action sequences (Table 1 provides text descriptions). In addition to those shown here, a second action sequence variant was also used for each object, with presentation counterbalanced across participants. The two sequences for a given object differed in the specific means that the adult used to operate each mechanism. On the Puzzle Box (based on a stimulus from ref. 10), for example, the red bolt was pushed out in one sequence and pulled out in the other. For more detail and depictions of the other action sequences, see the SI Methods and SI Figs. 7–10.

Test trial videotapes were analyzed to determine how frequently participants overimitated the experimenter's irrelevant actions. The resulting data for each puzzle object were initially segmented by age, with participants younger than 48 months being analyzed separately from older children. Within these groupings, each object's data were further subdivided by presentation order. While one significant effect of age on training outcome was detected (76% of older children received the maximum training score versus 41% of younger children [$\chi^2(1, n = 63) = 7.7, P = 0.006$], preliminary analyses showed that neither age nor order had any effect on overimitation. We thus report each object's data collapsed across these dimensions.^{||} All reported P values are two-tailed, both in this and later experiments.

The extensive training phase “taught to the test” situation in a number of important ways. Not only did it make the distinction between relevant and irrelevant actions highly salient, it also repeatedly showed participants that the experimenter was an unreliable model, one who consistently performed actions unnecessary for his goal. These factors, in combination with the praise that children received during training for identifying irrelevant actions as silly and unnecessary, created considerable situational demands opposing overimitation. Yet despite the contrary pressures, children showed a strong tendency to overimitate on all three puzzle objects (see SI Movies 2–4). Importantly, a baseline control condition established that this overimitation was not due to the puzzle

^{||}Neither age nor presentation order ever had a significant effect on overimitation, so subsequent data are all similarly collapsed.

Table 1. Experimenter action sequences for each puzzle object

Object	Panel	Sequence components
Puzzle Box*	1 (<i>Irrel</i>)	Use wand to remove red bolt by pushing from the right.
	2 (<i>Irrel</i>)	Tap wand on floor of box's empty upper compartment.
	3 (<i>Rel</i>)	Pull out round plug in center of door assembly.
	4 (<i>Rel</i>)	Use wand to remove turtle.
Cage	1 and 2 (<i>Irrel</i>)	Rotate metal basket 180° using its side handle.
	3 (<i>Rel</i>)	Unscrew locking cap on top of central spindle.
	4 (<i>Rel</i>)	Remove metal basket; get turtle from under blue/white lid.
Dome	1 (<i>Rel</i>)	Rotate white locking arm aside.
	2 (<i>Rel</i>)	Open lid of plastic box.
	3 (<i>Irrel</i>)	Pull bolt from base of plastic box using wooden handle.
	4 (<i>Rel</i>)	Get turtle from under red lid.

Panel numbers link the description to the corresponding panels in Fig. 1. The causal relevance (*Rel*) or irrelevance (*Irrel*) of each action sequence component is noted next to the panel number. See *SI Text* and *SI Figs. 7–10* for additional detail and depictions of each object's second action sequence variant.

*Based on a stimulus from Horner and Whiten (10).

objects being too complex for children to understand on their own. Rather, when a separate group of age-matched baseline participants ($n = 62$, mean age 49 months) retrieved the toy turtles from the puzzle objects independently, i.e., without first observing the experimenter, only a small minority operated the irrelevant mechanisms. The small degree of irrelevant action production in the baseline condition was far outstripped by the extent to which experimental participants overimitated after observing the adult [Fig. 2; Puzzle Box: $\chi^2(1, n = 93) = 63.8, P < 0.001$, odds ratio = 147.0; Cage: $\chi^2(1, n = 87) = 23.3, P < 0.001$, odds ratio = 21.9; Dome: $\chi^2(1, n = 90) = 12.0, P = 0.001$, odds ratio = 5.1]. Overimitative responses were not more common in those children who had a difficult time identifying irrelevant actions during training; instead, they distributed evenly across participants. Children in both age groups who scored the highest on training—and thus received the most praise for identifying irrelevant actions as silly and unnecessary—were just as likely to overimitate as participants who found training more difficult ($\chi^2 P$ values = ns for all objects' Training Score \times Overimitation cross-tabulations).** Consistent with our hypothesis, children who found it trivially easy to identify the experimenter's unnecessary actions on the familiar training objects seemed unable to apply this ability to the equally causally transparent but unfamiliar puzzle objects.

If children are indeed automatically encoding the adult's actions as causally meaningful, what can be said about the level of detail at which this encoding occurs? As previously noted (§), even simple actions can be encoded in multiple ways, differing in their level of abstraction. The adult's irrelevant action on the Puzzle Box, for example, could potentially be encoded in a concise, high-level manner ("remove bolt") or in a more detailed fashion ("use wand to pull red bolt out from left to right"). Thus, to more precisely determine the granularity of children's causal encoding, we compared the "style" in which they operated the puzzle objects' mechanisms to the style that they saw the adult employ.

In general, these styles were very well matched, with children copying the adult's means of operation 75–94% of the time (Table 2). However, stylistic deviations did occur when the experimenter operated a mechanism in an objectively suboptimal manner. Consider, for example, the Cage. Here the adult's irrelevant action was rotating the metal basket 180° around its central axis, using either a handle on top of the basket or one on its side (*SI Fig. 9*). Using the top handle made this task needlessly difficult, because it was located much closer to the axis of rotation than the side handle. We found that 78% of the overimitators who saw the adult using the top

handle chose an objectively easier method for rotating the cage themselves, either gripping further out on the wire mesh (22%) or using the more functional side handle (56%). That is, although they overimitated the adult's inefficient use of the irrelevant mechanism, they also showed localized imitative selectivity similar to that which has been observed in other contexts (18, 19, 21–23) (see *Discussion* in *SI Text* for additional detail). A similar argument holds for the Puzzle Box. Here overimitators gravitated toward pulling the red bolt out rather than pushing it (*SI Fig. 7*), a strategy that was indeed considerably easier because of the pushing implement's short length.

This pattern of results makes a clear suggestion regarding the granularity of children's causal encoding. Specifically, children seem to process the adult's actions at a level of detail roughly corresponding to the overall state of the target object; they encode the sequence of physical state transformations that the adult performs on the object as causally meaningful but remain able to optimize the specific means by which those transformations are achieved. Although this finding weighs against the strictest formulation of our hypothesis—indicating that children do not impute causal significance down to the most fine-grained elements of the adult's actions—its overall implications support our causal-encoding account rather than diminishing it. That is, insofar as participants are demonstrably imitating in a rational framework (i.e., omitting unnecessary stylistic components of the display), we can have greater confidence that the actions that they *are* copying are construed as causally significant. Their persistent operation of the irrelevant mechanisms, despite a demonstrable concern for

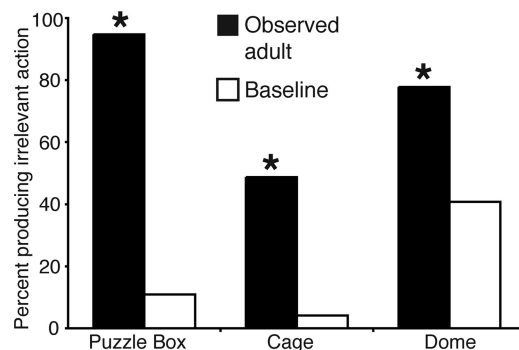


Fig. 2. Overimitation persists despite contrary task demands. Experiment 1A participants who observed the experimenter produced unnecessary actions significantly more often than baseline participants who opened the puzzle objects independently.

**Overimitation remained independent of training outcome in all subsequent experiments.

Table 2. Tendency of Experiment 1A participants to match experimenter's means of operating puzzle object mechanisms

Object: Mechanism	Experimenter's means of operation (variation 1/variation 2)	Children matching [†]	χ^2 (df = 1) [‡]
Puzzle Box: Irrelevant	Push bolt/pull bolt	33%/94%	14.6**
Puzzle Box: Relevant	Remove plug/slide frame	94%/75%	2.7
Cage: Irrelevant	Side handle/top handle	88%/22%	7.2*
Cage: Relevant	Unscrew cap/remove spindle	90%/83%	0.3
Dome: Irrelevant	Pull handle/pull ball	91%/7%	20.6**
Dome: Relevant	Rotate arm/flip up arm	90%/94%	0.3

** $P < 0.001$; * $P < 0.01$. All values are two-tailed.

[†]Of the participants who saw the experimenter use a given means of operation, the percentage that used that means themselves.

[‡]For each mechanism, tests whether children matched one means of operation significantly more than the other.

efficiency, thus argues that they have indeed encoded the adult's use of those mechanisms as causally necessary.

In summary, the findings from this experiment seriously challenge the view that overimitation occurs for superficial social reasons. The data are instead consistent with our hypothesis, arguing that overimitation may be driven by observationally induced distortions in children's causal beliefs.

Experiment 1B: Procedure, Results, and Discussion

Perhaps though, despite the contrary task demands, children still assumed that they were supposed to copy the experimenter. Children may also have been reluctant to contradict an adult through their actions, despite noticing unnecessary steps. Both of these views see overimitation as situational; they predict that children will stop reproducing irrelevant actions when removed from the unusual social context of the experiment. Our hypothesis makes a different prediction. If overimitation is mediated by distortions in children's causal beliefs, then it should persist even after the experiment has ended; children should continue to overimitate even when manipulating the puzzle objects as part of a practical real-world task. We used a surreptitious follow-up experiment to test this prediction.

After completing Experiment 1A, each child was told that the study was over and given a congratulatory prize. While this was occurring, an assistant carried the puzzle objects back into the room, explaining that they were there for a new participant due to arrive shortly. After the assistant left, the experimenter began gathering his notes and preparing to lead the child back to his or her classroom. Suddenly though, he froze as though he had just remembered something important. He told the child that he was worried about whether his assistant had done her job correctly. She had previously forgotten to put the toy turtles back into the puzzle objects between participants; had she forgotten again this time?




Fig. 3. Overimitation persists beyond the boundaries of the experiment. The apparent conclusion of the study in Experiment 1B did not significantly change overimitation levels for the Cage and Dome. Overimitation on the Puzzle Box was attenuated but remained four times more frequent than in the baseline condition.

Explaining that he needed to rush to prepare for the next participant, the experimenter asked the child to help by checking to see whether the turtles were indeed back in the puzzle objects. The child was then left unobserved while the experimenter busied himself with his other tasks (SI Movie 5 illustrates this procedure).

From the child's perspective, the experiment has ended and they are simply being asked to help a busy adult by gathering pragmatic information. Moreover, the experimenter's rush and worry naturally emphasizes economy of action, creating a strong impetus for the child to obtain the desired information as quickly as possible. Children should thus express their most efficient real-world causal theories of the puzzle objects, stripped of any artifice introduced by being part of an overt experiment.

Despite the considerable contrary pressure, overimitation remained robust. Indeed, for two of the three puzzle objects, frequency of overimitation did not decline from Experiment 1A levels [Fig. 3; Cage: McNemar Test $\chi^2(1, n = 33) = 1.5, P = \text{ns}$; Dome: McNemar Test $\chi^2(1, n = 30) = 3.3, P = \text{ns}$]. Overimitation on the Puzzle Box, although reduced [McNemar Test $\chi^2(1, n = 30) = 13.1, P < 0.001$], remained substantial, with participants operating the irrelevant mechanism four times as frequently as was observed in the baseline condition [$\chi^2(1, n = 85) = 13.1, P < 0.001$, odds ratio = 6.9]. Importantly, although children observed the adult act on each object just once, this surreptitious follow-up occurred only after the full Experiment 1A procedure was completed (during which time participants interacted with one to two additional puzzle objects). Thus, the memory load of this task alone presents a formidable obstacle to reproducing irrelevant actions—further reinforcing the significance of the high overimitation rates.

Experiment 2A: Procedure, Results, and Discussion

Children's robust reproduction of irrelevant actions—even when they believe the experiment has ended and after a considerable intervening delay—supports our contention that overimitation is mediated by distortions in underlying causal beliefs. However, we can test our hypothesis in an even more stringent manner. If children are actually encoding the adult's irrelevant actions as causally functional, then they should continue to overimitate even when directly instructed to copy only necessary actions. They should be unable, in other words, to avoid the irrelevant steps even when consciously attempting to do so. Experiment 2A tested this prediction.

Three- to five-year-olds who had not taken part in the prior studies ($n = 29$, mean age 50 months) underwent training as in Experiment 1A and were then tested with either the Puzzle Box or the Dome object.^{††} While introducing this object, the experimenter

^{††}Experiments 2A and 2B were run concurrently with the same participants. Because the instructions for Experiment 2A could have biased future responses, the Experiment 2B puzzle object was always presented first.




Fig. 4. Overimitation is not blocked by direct contrary instruction. Explicitly warning Experiment 2A participants to ignore any unnecessary actions performed by the experimenter failed to diminish overimitation.

gave children an explicit warning, explained in terms of the simplest training item (SI Fig. 6.4). Taking out the jar and feather, the experimenter re-performed the relevant and irrelevant actions used during training, and reminded the child that using the feather had been a “silly extra thing” that hadn’t helped to get the dinosaur. The child was then warned to watch very closely for similarly unnecessary actions: “I want you to watch really carefully, because when I open this [puzzle object], I might do something that’s silly and extra, just like the feather.” The child was firmly instructed to ignore any such silly actions and to do only what was necessary when retrieving the turtle for him or herself. After opening the puzzle object as in Experiment 1A, the experimenter reminded the child once again of the instructions (“Remember, don’t do anything silly and extra, okay? Only do the things you *have* to do”) and then left him or her unobserved to retrieve the turtle.

Directly warning participants to ignore unnecessary actions failed to attenuate overimitation. Despite deliberately monitoring for irrelevant steps, children continued to overimitate as frequently as they did in Experiment 1A [Fig. 4; Puzzle Box: $\chi^2(1, n = 53) = 2.7$, $P = ns$; Dome: $\chi^2(1, n = 50) = 0.2$, $P = ns$]. Again, this continued overimitation cannot be explained by positing that the puzzle objects were too complex for children to understand; age-matched participants in the baseline condition, who did not observe the experimenter, had no difficulty determining the minimal set of actions needed to retrieve the turtles. Thus, participants in the present experiment failed in their deliberate attempts to identify the adult’s irrelevant actions despite the demonstrated causal transparency of the puzzle objects. These data support our hypothesis, arguing that children can’t help but perceive the adult’s purposeful behavior as causally meaningful. Children are largely unable to circumvent overimitation, even when directly instructed to do so, because the adult’s irrelevant actions have already been absorbed into their representation of the puzzle object’s causal structure.

Experiment 2B: Procedure, Results, and Discussion

The robustness of overimitation has so far stood up well to the strong predictions of our theory. It is important, however, to address a final alternative possibility. Namely, the very persistence with which children overimitate—even under circumstances that should strongly promote efficient action—might be interpreted as supporting prior views of overimitation as a kind of social game (15, 16). Perhaps the phenomenon has little to do with causal reasoning at all, hinging instead on simple curiosity or on an innate motivation to copy others’ actions. Because both of these possibilities place overimitation outside the domain of causal reasoning altogether, they both predict that overimitation should persist regardless of how brazenly the adult’s irrelevant actions flaunt basic causal principles. Contrastingly, if overimitation instead arises from observationally induced distortions in causal beliefs, then we would expect that the effect might have some kind of causal boundary conditions.




Fig. 5. Overimitation is subject to contact constraints. Overimitation was significantly more frequent on the connected form of the Igloo than on the disconnected form. Overimitation on the disconnected form failed to exceed the background level of irrelevant action production observed in the baseline condition.

One such boundary condition might be the contact principle: the rule that mechanical interactions cannot occur at a distance. Even 3-month-olds are sensitive to this regularity, reacting with surprise when inanimate objects appear to interact without touching (30). The contact principle is thus thought to be part of human “core knowledge,” a set of innate expectations that structure our earliest interpretation of events (30, 31). We therefore predicted that children would not encode irrelevant actions as causally necessary if doing so obliged them to encode a violation of the contact principle; we expected that the foundational status of the principle would instead block overimitation.^{##} Conversely, if overimitation is driven by simple curiosity or by an intrinsic motivation to copy, then an implied violation of the contact principle should have no effect.

We tested this prediction using a new puzzle object consisting of two spatially separated halves and a removable connector (the Igloo). After undergoing training as in Experiment 1A, participants ($n = 29$, mean age 50 months)^{††} watched the experimenter retrieve a toy turtle from this object by performing an irrelevant action on one of its halves and a relevant action on the other (SI Fig. 11). However, whereas one group of children saw the object’s halves joined by the connector (such that the relevant and irrelevant actions occurred on the same continuous object), the other group saw the halves presented with no connector (such that the relevant and irrelevant actions occurred on two distinct objects). Although the adult’s actions were identical in both cases, we predicted that only children in the connected condition would show a significant degree of overimitation. In the disconnected condition, where encoding the irrelevant action as causally meaningful would imply a violation of the contact principle, we predicted that overimitation would be blocked.

This is in fact exactly what we observed. Overimitation was much more frequent for the connected form of the Igloo than for the disconnected form [Fig. 5; $\chi^2(1, n = 29) = 4.2$, $P = 0.04$, odds ratio = 5.3]. Indeed, overimitation on the disconnected form failed to exceed the background level of irrelevant action production observed when a separate, age-matched group of baseline participants ($n = 25$, mean age 49 months) operated the object without observing the adult [$\chi^2(1, n = 25) = 2.3$, $P = ns$; see SI Text for additional discussion]. These data support our characterization of overimitation as properly a causal reasoning phenomenon and not simply the product of curiosity or of an indiscriminate social motivation to reproduce others’ actions.

^{##}One might ask why contact principle violations would be expected to block overimitation when violations of the similarly foundational “efficiency principle” (32) (i.e., the adult operating mechanisms in a suboptimal way) did not diminish the effect (Experiment 1A). We return to this matter in the Discussion section of SI Text, arguing that contact principle violations undermine the causal plausibility of the target object in a way that efficiency violations do not.

Conclusion

For humans, surrounded as we are by causally opaque tools and artifacts, causal learning is an uphill battle (1, 2, 26, 27). Adults overcome this causal opacity in part through deliberate social inference: When making sense of a complex object, we use the intentional manipulations of more knowledgeable individuals to infer causally important operations. Children, it seems, do much the same thing. However, as this work has now shown, they do so in a surprisingly automatic way—one that leads to the phenomenon of overimitation.

Children who observe an adult manipulating an unfamiliar object show a strong tendency to encode all of the adult's purposeful actions as causally meaningful, revising their causal beliefs about the object accordingly. Although generally a powerful learning strategy, the apparent automaticity of this causal encoding process carries a cost. When some of the adult's purposeful actions are unnecessary—even transparently so—children are highly prone to mis-encoding them as causally significant. The resulting distortions in children's causal beliefs are the true cause of overimitation, not implicit social demands (10, 15, 16) or imitative habit (11, 14) as previously believed. This deeper cause makes overimitation remarkably resistant to extinction. Despite countervailing task demands, time pressure, and even direct warnings, children are frequently unable to avoid reproducing the adult's irrelevant actions because they have already incorporated them into their representation of the target object's causal structure.

The revised conception of overimitation presented here suggests many possibilities for future work. As we elaborate in the *General Discussion* in the *SI Text*, one such avenue would be to further investigate constraints on overimitation. That is, in addition to the contact constraint already identified, what kinds of boundaries and preconditions apply to children's automatic causal encoding? Data from other imitation studies, for example, suggest that an adult's actions may need to be more than simply intentional for children to encode them as causally meaningful; the qualities of being both unconstrained (i.e., not determined or limited by external factors; see ref. 23) and potentially communicative or pedagogical in nature (see refs. 2 and 18) may also be prerequisites. Overimitation may also be bounded by developmental factors. In particular, our theory predicts that overimitation may actually increase from infancy to early childhood as socially derived inferences begin to play a larger role in causal learning. Preliminary evidence from related tasks is consistent with this prediction (18).

We will close this paper as we began, by observing that imitation is a remarkably potent learning strategy. Indeed, as we have now seen, it can at times be too potent for the integrity of children's causal knowledge. All of which recommends caution the next time you idly fidget with a complex device. You never know who might be watching.

Materials and Methods

Training Phase. Training began with a plastic jar containing a toy dinosaur (SI Fig. 6A). Explaining that he was going to retrieve the

dinosaur, the experimenter proceeded to do so by unscrewing the jar's lid. He then reinforced the meaning of “have to” for the child by pointing out, “I *have* to take the lid off to get the dinosaur. If I don't take it off, he won't come out, see?” Next the experimenter told the child that he was going to get the dinosaur out in a different way; this time he tapped the jar with a feather before unscrewing the lid. The child was then asked whether each of these actions—the feather tap and the unscrewing of the lid—had been necessary (e.g., “Did I have to tap the jar with the feather to get the dinosaur?”). Correct responses were praised and reiterated (“That's right! I *didn't* have to use the feather to get the dinosaur out. The feather was extra.”); wrong responses were corrected verbally and with an accompanying demonstration (“Well actually, I *can* get the dinosaur out without using the feather, see?”).

For each of the remaining training objects, the child saw the experimenter retrieve the dinosaur in just one way, using either (i) one relevant and one irrelevant action, or (ii) in one case, two relevant actions (SI Fig. 6D; this item served as an attentional control). Children were questioned after each object as above. Corrective feedback was withheld on the final two objects, and children were assigned a training score between 0 and 2 on the basis of how many of these final items they responded correctly to.

Baseline Condition. Baseline participants underwent the same initial training procedure used in the experimental conditions. The testing procedure was also similar, the critical difference being that baseline participants did not see the experimenter open the puzzle objects. Instead, they were asked to find the toy turtle in each object while the experimenter was out of the room and then to show the experimenter how to retrieve it. This allowed us to evaluate children's baseline level of causal understanding for each puzzle object, i.e., how frequently they would operate the irrelevant mechanisms when opening the objects independently. Baseline participants were tested with three puzzle objects in counterbalanced order.

Coding of Data. Trials were videotaped by using a camera unobtrusively positioned behind participants, above their line of sight. Two independent coders, blind to the experimenter's actions, then analyzed the tapes to determine whether and how participants had operated the puzzle objects' relevant and irrelevant mechanisms. Cohen's κ values were uniformly high (SI Table 3; mean $\kappa = 0.934$; $P < 0.01$ in every case), indicating reliable inter-rater agreement in each experiment.

We thank the children and preschools who made this work possible; Jennifer Barnes, Diana Damrosch, Jennifer Lin, Neha Mahajan, Deanna Simeone, and Deena Weisberg for their assistance; and Paul Bloom and Laurie Santos for their helpful comments. D.E.L. was supported by a National Defense Science and Engineering Graduate Fellowship. F.C.K. was supported by Yale University and National Institutes of Health Grant R37 HD023922.

- Gergely G, Csibra G (2005) *Interact Stud* 6:463–481.
- Gergely G, Csibra G (2006) in *Roots of Human Sociality: Culture, Cognition, and Human Interaction*, eds Enfield NJ, Levinson SC (Berg Publishers, Oxford), pp 229–255.
- Heyes CM, Foster CL (2002) *Q J Exp Psychol* 55A:593–607.
- Meltzoff AN, Gopnik A (1993) in *Understanding Other Minds: Perspectives from Autism*, eds Baron-Cohen S, Tager-Flusberg H, Cohen DJ (Oxford Univ Press, Oxford), pp 335–366.
- Tomasello M (1999) *The Cultural Origins of Human Cognition* (Harvard Univ Press, Cambridge, MA).
- Tennie C, Call J, Tomasello M (2006) *Ethology* 112:1159–1169.
- Tomasello M, Call J (1997) *Primate Cognition* (Oxford Univ Press, Oxford).
- Call J, Carpenter M, Tomasello M (2005) *Anim Cognit* 8:151–163.
- Carpenter M, Call J, Tomasello M (2002) *Child Dev* 73:1431–1441.
- Horner V, Whiten A (2005) *Anim Cognit* 8:164–181.
- McGuigan N, Whiten A, Flynn E, Horner V (2007) *Cognit Dev* 22:353–364.
- Nagell K, Olguin K, Tomasello M (1993) *J Comp Psychol* 107:174–186.
- Want SC, Harris PL (2002) *Dev Sci* 5:1–41.
- Whiten A, Custance DM, Gomez J-C, Teixidor P, Bard KA (1996) *J Comp Psychol* 110:3–14.
- Nielsen M (2006) *Dev Psychol* 42:555–565.
- Uzgiris IC (1981) *Int J Behav Dev* 4:1–12.
- Tomasello M, Carpenter M, Call J, Behne T, Moll H (2004) *Behav Brain Sci* 28:675–691.
- Brugger A, Lariviere LA, Mumme DL, Bushnell EW (2007) *Child Dev* 78:806–824.
- Bekkering H, Wohlschläger A, Gattis M (2000) *Q J Exp Psychol* 53A:153–164.
- Carpenter M, Akhtar N, Tomasello M (1998) *Infant Behav Dev* 21:315–330.
- Carpenter M, Call J, Tomasello M (2005) *Dev Sci* 8:F13–F20.
- Gleissner B, Meltzoff AN, Bekkering H (2000) *Dev Sci* 3:405–414.
- Gergely G, Bekkering H, Király I (2002) *Nature* 415:755.
- Meltzoff AN (1995) *Dev Psychol* 31:838–850.
- Schwier C, Van Maanen C, Carpenter M, Tomasello M (2006) *Infancy* 10:303–311.
- Csibra G, Gergely G (2006) in *Processes of Change in Brain and Cognitive Development. Attention and Performance, XXI*, eds Munakata Y, Johnson MH (Oxford Univ Press, Oxford), pp 249–274.
- Gergely G, Egyed K, Király I (2007) *Dev Sci* 10:139–146.
- Csibra G (2007) in *Sensorimotor Foundations of Higher Cognition: Attention and Performance, XXII*, eds Haggard P, Rosetti Y, Kawato M (Oxford Univ Press, Oxford), pp 427–451.
- Lyons DE (2008) in *Mirror Neuron Systems: The Role of Mirroring Processes in Social Cognition*, ed Pineda J (Humana, Totowa, NJ), in press.
- Spelke E (1994) *Cognition* 50:431–445.
- Spelke ES, Breinlinger K, Macomber J, Jacobson K (1992) *Psychol Rev* 99:605–632.
- Gergely G, Csibra G (2003) *Trends Cognit Sci* 8:396–403.