Observing child-led exploration improves parents' causal inferences

Koeun Choi[1,2](koeun@vt.edu)
Milagros Grados[2] (milagrosgrados95@gmail.com)
Elizabeth Bonawitz[2](lbaraff@gmail.com)
[1] Department of Human Development and Family Science, Virginia Tech, VA 24061
[2] Department of Psychology, Rutgers University, Newark, NJ 07102

Abstract

Do children's flexible causal inferences promote more creative causal discovery for observing adults? Inspired by a task in which children are more likely to consider unconventional causal forms (Lucas, Bridgers, Griffiths, & Gopnik, 2014; Wente et al., 2019), we designed a new method in which child-adult pairs work together to solve a causal task and assessed the relative influence of each member of the pair on the other's causal inference. Consistent with previous research, children were better than parents at learning the unusual conjunctive relationship, suggesting that children make more flexible causal inferences than adults. Our research also revealed a surprising and new result – that observing a child explore broadly helped parents to be more flexible and open-minded in their causal learning. In contrast, a child observing an adult's exploratory interventions had no negative consequence on the child's ability to infer the correct relation. Follow-up experiments explored the degree to which this child-led bootstrapping for adults was due to the particular exploratory evidence generated by the child during play, or merely the presence of a child. Results suggest that both factors may play a role in shaping adult's causal inferences.

Keywords: causality, cognitive development, parent-child interaction

Introduction

Like scientists, children explore, discover, and learn. Those of us with the good fortune to spend ample time with these little scientists can't help but be inspired by their curiosity and reminded of our own creative and inquisitive pasts. As Gopnik (2016) has suggested, childhood may be a unique time for greater exploration, cognitive flexibility, and creativity, leading to innovation for our species driven by our youngest. Of course, much research has focused on how children learn from adults, but perhaps there are cases when adults can learn from these innovative explorers. Perhaps there are cases when the flexible minds of children lead to knowledge and learning when adults lack.

Indeed, evidence from several research studies indicated that children learn specific and abstract causal structure and sometimes do so more readily than adults (Gopnik et al., 2017; Lucas et al., 2014; Wente et al., 2019). These findings suggest that children may be more open to new possibilities and willing to consider different hypotheses than adults (Gopnik et al., 2017). Often times, children encounter and explore new information in the presence of adults who may hold contrasting ideas about the world. Here we ask, how do children and adults interact with each other to explore and come to understand the world around them? In this study, using a new method in which child-adult pairs work together to solve a causal task, we look at whether exploratory patterns differ between children and adults and the extent to which these differences have consequences for causal inference in the observers.

Young children's ability to infer abstract causal principles has been studied using the forms of overhypotheses including conjunctive and disjunctive causal relationships. An overhypothesis is a broad framework that constrains the range of hypothesis learners consider (Goodman, 1955; Griffiths & Tenenbaum, 2007). A conjunctive causal relation is a functional form in which multiple causes jointly produced an effect; a disjunctive relation is a functional form in which a single cause can bring out an outcome independently (e.g., see Cheng, 1997). These overhypotheses are not bounded to a particular context but are applicable to many other scenarios, and having these assumptions shape future learning by limiting the number of possible hypotheses that are considered.

Prior research has revealed developmental differences in inferring a certain form of overhypotheses (Lucas et al., 2014). After having the same amount of exposure to evidence that is statistically best explained by (the unconventional) conjunctive causal form, children outperformed adults by correctly generalizing the conjunctive causal relationship to new objects. While both adults and children were successful at inferring a disjunctive form, the ability to infer conjunctive forms appears to be decreased with age. When given evidence that supported a conjunctive form, adults instead maintained a disjunctive relationship (Lucas et al., 2014). The developmental differences in learning of the conjunctive (but not disjunctive) causal form suggest that young children are more flexible than adults in incorporating evidence to guide future learning (See also Gopnik et al., 2017; Wente et al., 2019; Gopnik, Griffiths, & Lucas, 2015).

In these past studies, participant's ability to infer abstract causal forms was tested by asking for judgments about the causal efficacy of each cause or to use potential causes to produce an outcome. However, in these studies, participants were not given the opportunity to explore and generate their own evidence. Thus, it remains an open question whether adult and child participants will generate different patterns of exploration when given the opportunity to test out possible causal forms.

One concern is whether children will be able to generate meaningful play at all. However, recent findings are supporting the claim that young children may be more competent and capable explorers than previously believed. For example, children shape their explorations to conduct inter-

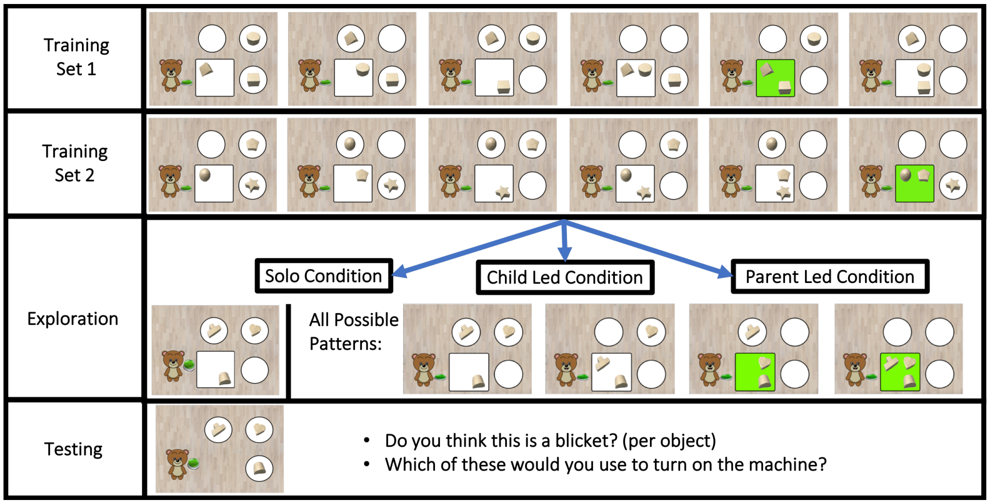

Figure 1: Stimuli and task procedure; During two sets of training, parent-child pairs watched experimenter demonstrating evidence in favor of a conjunctive causal form. Following the training, parent-child pairs were randomly assigned to solo, child-led, and parent-led conditions. In the solo condition, the parent and child each explored the set of testing objects. In the child-led group, the child explored the testing objects while parent watched. In the parent-led group, the parent explored the testing objects while the child watched. Finally, the experimenter asked the parent and child individually to judge whether each object was a blicket and to turn on the machine.

ventions to deconfound variables (Schulz & Bonawitz, 2007; Schulz, Kushnir, & Gopnik, 2007; Schulz, 2012; Schulz, Gopnik, & Glymour, 2007). Further, children plan their future exploration based on the inference about pedagogical goals of teachers based on available information (Bonawitz et al., 2011; Eaves & Shafto, 2012; Gweon, Pelton, Konopka, & Schulz, 2014). These studies provide evidence for the claim that children's exploration is guided by the evidence. However, it remains an open question whether the evidence generated during children's and adults' explorations may differently reflect beliefs going into the task.

Prior studies revealed developmental differences in the conjunctive causal inference by examining child and adult groups individually. Despite the importance of caregiver-child interaction on play and development (Weisberg, Hirsh-Pasek, & Golinkoff, 2013; Fisher, Hirsh-Pasek, Newcombe, & Golinkoff, 2013; Honomichl & Chen, 2012), little is known about the impact of observing either child's or parent's patterns of exploration on one another. Thus, we were also interested in whether observing children's broad hypothesis search would promote more creative and flexible thinking in causal learning for observing adults. Of course, observing adults' exploration may also influence children's conjunctive causal learning. For example, instructions constrain children's explorations indicating that children are sensitive to inductive biases in their explorations (Bonawitz et al., 2011). Similarly, a body of literature on guided play highlights that it is critical for adults to scaffold learning goals as well as let children direct their exploration and discovery (Weisberg et al., 2013; Fisher et al., 2013). This balance between scaffolding and letting the child take the lead could be particularly important when children and adults hold different assumptions about the world.

In the present work, we examined the extent to which exploration patterns differed between children and adults and whether observing another's exploration shaped consequent learning. In Experiment 1, we examined children and their parents' exploration and learning following the exposure to evidence consistent with a conjunctive causal relationship. Critically, we manipulated who the actor was generating the evidence including a child-led condition, parent-led condition, and solo conditions for each group as controls. Consistent with prior research, we hypothesized that children would be more likely than adults to generalize a conjunctive causal relationship. However, we also predicted differential exploratory patterns for children and adults. We looked at whether these differences have consequences for learning in the observers. Specifically, observing parent-led exploration may restrict children's causal inference, thus resulting in more adult-like responses. On the other hand, observing child-led exploration may result in flexible learning in parents.

Experiment 1

Methods

Participants Seventy-two parent-child dyads were recruited from various settings (i.e., museum, home, and community event; Children: n = 72, 53% Female, M = 5.03, SD = 0.84, Range = 4.0-6.9 years; Parents: n = 72, 56% Mothers). The dyads were randomly assigned to the Solo (n = 24), Child-Led (n = 24), and Parent-Led (n = 24) conditions.

Additional two dyads were recruited but excluded due to not finishing the study (n = 1) or experimental errors (n = 1).

Stimuli and Apparatus Our procedure involved both real objects and interactive video stimuli. The interactive video stimuli were developed using jsPsych (De Leeuw, 2015) and displayed on a touch-screen tablet computer (10.1-in. Galaxy Tab; Samsung America, San Jose, CA). The video included images of three circles, a square, a green button, and a cartoon bear (see Figure 1). Images of objects were presented on the three circles, and each object moved to the square when tapped. The button was designed to test objects once they were placed on the square. In addition, to provide a way for participants to respond without influencing the listening other, two identical yes-no response sheets were created so participants could silently point to their response behind a barrier. The sheets included two rectangles (green and red), each includes a smiley or frowny face with "yes" or "no" written at the bottom, respectively.

Procedure Participants were tested in a quite place. The yes-no response sheets were placed in front of the participants. The experimenter asked a simple question about color (i.e., Is this white?) to both the child and the parent. If participants pointed to the wrong answer or responded verbally, the experimenter asked additional questions until children successfully responded using the yes-no response sheet.

Next, a backward blocking task was conducted (e.g. see Sobel, Tenenbaum, & Gopnik, 2004). This task was designed to acclimate the participants to an ambiguous causal reasoning task as well as familiarize participants with the instructions. The participants were introduced to a machine that detects "wugness". The experimenter explicitly stated that wugs are very rare and can not be judged solely by looking, but they possess wugness inside them. First, the experimenter placed two potential causes (Objects A and B) on the machine, which produced an outcome. Then the participants observed that the outcome occurred with the presence of only one of the causes (A). After observing these two events, the participants were asked to judge whether each object (A, B) was a wug, respectively.

Upon the completion of the backward blocking task, the experimenter introduced the tablet and stated that they would now play a completely different game (see Figure 1). The experimenter also mentioned that blickets are very rare and can not be judged just by looking, but have blicketness inside them. After a brief introduction to the features of the tablet game, the experimenter introduced an object (C) and ask if the participant thought that object as a blicket without any evidence; this allowed us to test participants' priors for the probability of an object being a blicket. Then the experimenter presented the first set of three training objects (D, E, F) which activated the machine according to a conjunctive causal rule. This was followed by another set of different training objects (D', E', F') that also provided evidence for a conjunctive rule, as in Lucas et al. (2014).

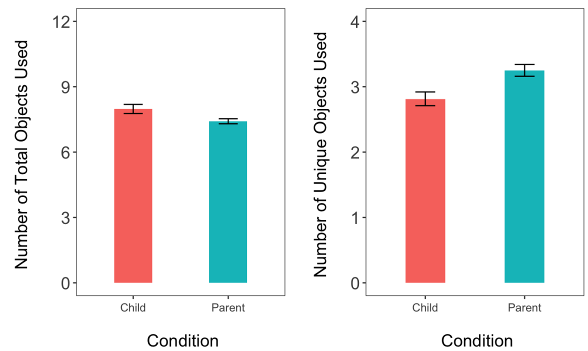

Figure 2: Differing exploratory patterns as measured by total object used (left) and unique actions (right) during the Exploration Phase. Compared to children, parents used more unique combinations of objects, but also fewer objects were tried on each trial. Error bars denote SE.

Following the two training trials, the experimenter introduced the new set of three testing objects (G, H, I) that the participants would have the opportunity to explore and test themselves. During this phase, one object (I) was permanently attached to the square. This was designed so that the evidence generated during the free intervention phase would maintain ambiguity. At the beginning of the exploration phase, the participants were told that the object (I) was stuck on the machine and that they can test the object (I) by itself or with the other objects. There were four exploration trials, and participants could choose one of the four possible options (I, GI, HI, GHI) in each trial. Our critical between-subjects design varied who controlled the interventions. In the solo condition, both parent and child had their own tablets, and could not see the screens or exploratory choices of the other. In the Child-Led condition, the child made all intervention choices while the parent watched. In the Parent-Led condition, the parent made all intervention choices while the child watched. Participants were given four intervention trials (a trial was counted once the participant depressed the test button); the intervention choices were recorded automatically with the tablet software. Next, the experimenter asked the parent and child to judge whether each object was a blicket. Lastly, as our critical test measure, the experimenter asked the parent and child individually to generate the effect using the objects ("Which of these objects would you use to turn on the machine?"). We coded whether two or more objects were used.

Results

We first assessed what kinds of interventions children and adults performed during the exploration phase. Results revealed that overall, children were more likely than adults to explore objects jointly, Welch $t(90.5) = 3.19$, p = .001 (Figure 2, left), suggesting that children may have been more amenable to the conjunctive rule as early as the exploration phase. However, the quality of children's interventions was not strictly better than adults: parents tried more deconfounding causal explorations by testing more unique combinations of objects, $t(72.9) = -2.36$, $p = .021$ (Figure 2, right). Within each age group (child and adult, respectively), there was no significant difference between solo and joint groups, $p > .250$.

Critically, we explored whether any particular group was more successful at generating the correct response in the final test phase. Overall, and replicating previous findings, children performed better on average than the adults, $\chi^2(1) = 20.39$, $p < .001$. We conducted a logistic regression to predict the probability of selecting one or more objects to activate the machine as a function of condition. As shown in Figure 3, the parents in the Child-Led group were more likely to use multiple objects to activate the machine than those in the Parent-Led group, $b = 1.44$, $p = .022$, suggesting that observing evidence generated by children helped parents to be more flexible and exploratory in their own causal inferences. The probability of choosing multiple objects as blickets in the Parent-Solo condition did not differ from that in the Child-Led condition $b = 0.84$, $p = .152$, or the Parent-Led condition, $b = 0.59$, $p = .353$.

Consistent with previous research, children were better than parents at learning the unusual conjunctive relationship, suggesting that children make more flexible causal inferences than adults. Our research also revealed a surprising result – that merely observing a child's exploratory behavior may suffice to help parents to be more flexible and open-minded in their causal learning. Two possible explanations exist for this result. One possibility is simply that watching a child interact with the toy (regardless of the patterns of exploration) was sufficient to get adults in a childlike frame of reference, opening their mind to a broader set of hypotheses. Another possibility is that the particular evidence generated by children (which differed from adults) was critical in helping adults infer the conjunctive form. In Experiment 2, we explored these possibilities by having adults view a child actor perform interventions for all conditions. Critically we varied the particular interventions presented, yolking the evidence to the specific interventions attempted in the Child-Solo, Child-Led, and Parent-Led conditions of Experiment 1.

Experiment 2

Methods

Adults participants were recruited from Amazon Mechanical Turk. The final sample of 72 participants were randomly assigned to the Child-Solo-Exploration-Data ($n = 24$), Child-Led-Exploration-Data ($n = 24$), and Parent-Led-Exploration-Data ($n = 24$) conditions. An additional 10 participants were excluded from analysis because of the failure to pass an attention check question. Participants were paid \$.75 for completing the 6-8 minute survey.

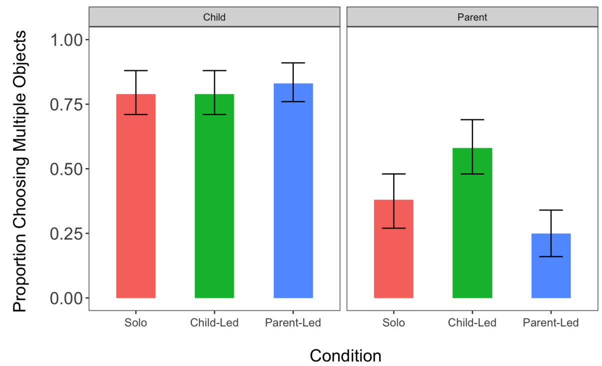

Figure 3: Proportion of participants who selected multiple objects on the final test trial by condition (Solo, Child-Led, Parent-Led). Children correctly attempted multiple objects to turn on the machine regardless of condition. However, adults were only more likely to test multiple objects in the Child-Led condition. Error bars denote *SE*.

Procedure

The stimuli and procedure was the same as that of Experiment 1 except the following differences. First, the data collection was conducted online; thus, the training phase was introduced with a series of screenshots of tablet games, and the participants were required to click a button to proceed. Second, for the exploration phase, the participants saw a video of a preschooler trying to figure out which objects are blickets using the tablet game. The same child actor generated different exploratory patterns, which was organized to match the patterns of exploration data generated from the three conditions (Child-Solo, Child-Led, Parent-Led) in Experiment 1. The preschooler was described to the participants as being a randomly selected example of a child exploring the toy.

Results

We used a logistic regression model as a function of condition (Child-Solo-Exploration-Data, Child-Led-Exploration-Data, Parent-Led-Exploration-Data) to predict the probability of using two or more objects to turn on the machine. There was a significant difference between the Child-Solo-Exploration-Data (67%) condition and the Parent-Led-Exploration-Data (38%) condition such that the group of participants who observed child-solo-exploration data showed a higher probability of using multiple objects to activate the machine compared to those who observed parent-led-exploration data, $b = 1.20$, $p = .046$. Unexpectedly, there was also a marginally significant difference between the Child-Solo-Exploration-Data (67%) and Child-Led-Exploration-Data condition (42%), $b = 1.02$, $p = .085$.

These results revealed that observing children's exploratory patterns based on broad, exploratory hypotheses supported adults' learning of an unconventional abstract causal form (at least in cases when data from the Child-Solo-Exploration-Data condition were observed). Experiment 2 provides additional support for the idea that observing child-generated exploratory patterns increases the flexibility of adult's causal reasoning. It remains unclear whether the child-generated exploratory evidence alone would be sufficient to promote adults' causal reasoning, or whether it is this evidence *in conjunction* with a child-directed play that helps adults. Further, Experiment 2 was conducted via an online survey platform; thus, the findings may be limited in their generalization to adults in a live setting. To explore this further, in Experiment 3, we used an in-lab setting to conduct child-yoked interventions, but performed by adults. We focus our attention on the two critically different yolked-data conditions: Child-Led-Exploration-Data and Parent-Led-Exploration-Data.

Experiment 3

Methods

Forty-eight undergraduate students (M = 20.60, SD = 3.13, range: 18-31 years) were randomly assigned to the Child-Led-Exploration-Data (n = 24) or Parent-Led-Exploration-Data (n = 24) conditions. An additional 5 participants were excluded from analysis because of experimental errors in generating the data from the yolked trials.

Procedure

The stimuli and procedure was the same as to that of Experiment 1 except the following differences. First, the data collection was conducted in the lab. Second, each participant was paired with a confederate who secretly worked with the lab but who was introduced as another naive participant. During the exploration phase, the participant observed the confederate exploring the testing objects, as if choices were "in the moment" decisions. Instead, however, the confederate completed the four exploration trials as yolked to the data generated from the two conditions (Child-Led, Parent-Led) in Experiment 1.

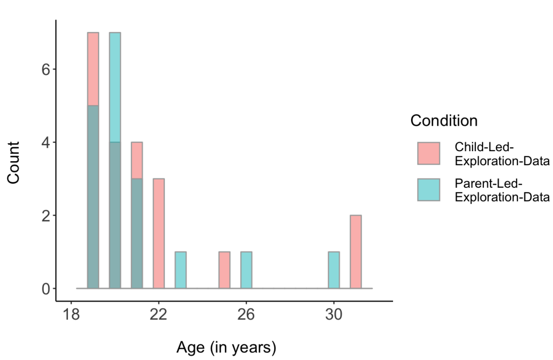

Figure 4: Histogram of age of participants for each condition (Child-Led-Exploration-Data, Parent-Led-Exploration-Data) in the sample. Age was ranged from 18 to 31 years.

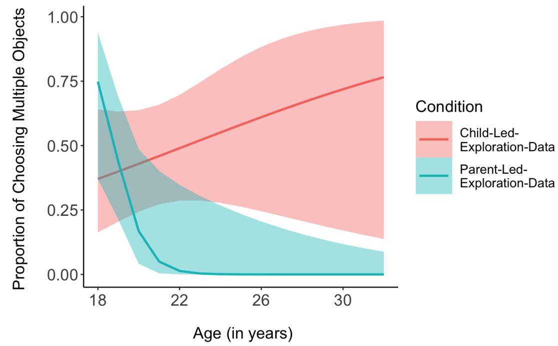

Figure 5: Proportion of participants who selected multiple objects for each condition (Child-Led-Exploration-Data, Parent-Led-Exploration-Data). Adults in the Child-Led-Exploration-Data condition were more likely to use two or more objects to turn on the machine than those in the Parent-Led-Exploration-Data condition. Shading indicates 95% confidence intervals.

Results

Comparing overall performance in terms of endorsement of two or more casual blocks between Child-Led-Exploration-Data and Parent-Led-Exploration-Data conditions revealed no overall differences, $\chi^2(1) = 0.35$, $p = .555$. This result is surprising, given the condition differences observed in Experiment 2. One possible explanation for this difference is that indeed the presence of the child generating the particular interventions was required to help adults consider the unlikely conjunctive form.

However, we also noticed that the age of the participants in our lab sample (ranging from 18-31 years; see Figure 4) significantly differed from the parents in Experiment 1 (ranging in mid-thirties to forties), and so we performed an unplanned exploratory analysis using a logistic regression model with age as a continuous variable and condition (Child-Led-Exploration-Data, Parent-Led-Exploration-Data) to predict the probability of using two or more objects to turn on the machine. In fact, there was a significant interaction between age and condition such that the group of participants who observed child-led-exploration data showed a higher probability of using multiple objects to activate the machine compared to those who observed parent-led-exploration data with increasing age, $b = 1.46$, $p = .018$ (see Figure 5). The pattern stayed the same for a narrower age range (18-24 years).

These results suggest that while child-yoked interventions may assist adults with causal form inferences, age may moderate this effect. Specifically, observing child-generated evidence was particularly helpful for the older participants of our sample.

Discussion

Consistent with previous research showing that adults are less likely to generate a conjunctive causal form than children, children were better than parents at generalizing the unusual conjunctive form to their exploration and learning (Gopnik et al., 2017; Lucas et al., 2014; Wente et al., 2019). By examining exploration, we revealed that parents tried more deconfounding explorations than children. In contrast, children performed more interventions that involved multiple blocks, suggesting that children were engaging in hypothesis confirmation consistent with having inferred the conjunctive form from the previous training trails.

Strikingly, parents in the Child-Led group were more likely to generalize the conjunctive relationship than those in the Parent-Led group. Child-yoked interventions performed by either a child or adult similarly improved causal form inferences, suggesting that observing evidence generated by children may help adults to be more flexible in their own causal inferences.

In our study, young children generalized the unconventional conjunctive relationship to their exploration and learning regardless of whether the free play period was led by an adult or not. Of course, if children had already inferred the correct causal form from the initial training trials, than any intervention observed would continue to confirm children's overhypothesis because we designed the toy to produce outcomes consistent with the conjunctive form. In contrast, if adults had not yet inferred the correct form prior to the exploration phase, then observing their children repeatedly use multiple blocks to activate the machine may have been sufficient to raise the salience of this alternative hypothesis and facilitate learning.

The results from Experiments 2 and 3 suggest that adults' causal inferences can benefit from observing child-yoked explorations, especially when those exploratory patterns were generated by a child than an adult. However, as these two experiments were conducted in different environments (in-person vs. online), the contextual factor may have contributed to the differences. Thus, an important next step would be to test the effect of the age of the model who demonstrates child-led exploratory patterns in the same setting. Future work will examine whether watching a video of an adult demonstrator performing child-yoked interventions similarly improves adults' causal form inferences, controlling for the familiarity of the adults to the demonstrator. Further, future studies could explore the characteristics of adult observers such as age and experience working with young children.

The current findings show the importance of observing other's exploration when beliefs are in conflict with each other. Adults at least may be able to recognize the relationship between attempted interventions and considered hypotheses, raising awareness of hypotheses that were not previously considered. Such an account is consistent with the Wisdom of the Crowds (Vul & Pashler, 2008) or the adage that two heads are better than one. However, our results go one step further, suggesting that even observing the exploratory actions of another may help bootstrap inference to the best explanation.

More broadly, these findings support the importance of adult-child play, but with a surprising twist. Adults may benefit from play with children (rather than the other way around, as is often considered in the literature). Such work suggests the importance of giving children opportunities to lead their own exploration and discovery.

Parent-child joint play occurs in numerous settings involving both concrete and digital materials. In light of the pervasive interactive technology in young childrens everyday lives, it is important to understand how these tools can be used not only to transmit information but also support active exploration and discovery. This line of research can help us understand the ways in which parents and children conjunctively learn about the world.

Acknowledgments

We thank staff, parents, and children of Newark Museum for their participation and Cierra Clark, Kiabeth Guazhco, and Micah Connolly for data collection. This research has been supported by NSF-CNS-1623486 (EB), Jacobs Foundation Fellowship (EB), and NSF-SES 1627971 (EB).

References

Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. E. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, *120*(3), 322–330. doi: 10.1016/j.cognition.2010.10.001

Cheng, P. W. (1997). From covariation to causation: a causal power theory. *Psychological Review*, *104*(2), 367. doi: 0033-295 X/97/S3.00

De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, *47*(1), 1–12. doi: 10.3758/s13428-014-0458-y

Eaves, B. S., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in Child Development and Behavior*, *43*, 295–319. doi: 10.1016/B978-0-12-397919-3.00011-3

Fisher, K. R., Hirsh-Pasek, K., Newcombe, N., & Golinkoff, R. M. (2013). Taking shape: Supporting preschoolers' acquisition of geometric knowledge through guided play. *Child Development*, *84*(6), 1872–1878. doi: 10.1111/cdev.12091

Goodman, N. (1955). *Fact, fiction, & forecast*. Cambridge, MA, US: Harvard University Press.

Gopnik, A. (2016). *The gardener and the carpenter: What the new science of child development tells us about the relationship between parents and children*. New York, NY: Farrar, Straus, and Giroux.

Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, *24*(2), 87–92. doi: 10.1177/0963721414556653

Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., ... Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, *114*(30), 7892–7899. doi: 10.1073/pnas.1700811114

Griffiths, T. L., & Tenenbaum, J. B. (2007). From mere coincidences to meaningful discoveries. *Cognition*, *103*(2), 180–226. doi: 0.1016/j.cognition.2006.03.004

Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335–341. doi: 10.1016/j.cognition.2014.04.013

Honomichl, R. D., & Chen, Z. (2012). The role of guidance in children's discovery learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *3*(6), 615–622. doi: 10.1002/wcs.1199

Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, *131*(2), 284–299. doi: 10.1177/0963721414556653

Schulz, L. E. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, *16*(7), 382–389. doi: 10.1016/j.tics.2012.06.004

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045. doi: 10.1037/0012-1649.43.4.1045

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, *10*(3), 322–332. doi: 10.1111/j.1467-7687.2007.00587.x

Schulz, L. E., Kushnir, T., & Gopnik, A. (2007). Learning from doing: Intervention and causal inference. *Causal learning: Psychology, Philosophy, and Computation*, 67–85. doi: 10.1093/acprof:oso/9780195176803.003.0006

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and bayesian reasoning in preschoolers. *Cognitive Science*, *28*(3), 303–333. doi: $10.1207/s15516709cog2803_1$

Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, *19*(7), 645–647. doi: 10.1111/j.1467-9280.2008.02136.x

Weisberg, D. S., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Guided play: Where curricular goals meet a playful pedagogy. *Mind, Brain, and Education*, *7*(2), 104–112. doi: 10.1111/mbe.12015

Wente, A. O., Kimura, K., Walker, C. M., Banerjee, N., Fernández Flecha, M., MacDonald, B., ... Gopnik, A. (2019). Causal learning across culture and socioeconomic status. *Child Development*, *90*(3), 859-875. doi: 10.1111/cdev.12943