

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Young children strategically adapt to unreliable social partners

Permalink

<https://escholarship.org/uc/item/7rq378jm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Shannon, Katherine Adams

Conine-Nakano, Aneesa

Frankenhuis, Willem E.

et al.

Publication Date

2024

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Young children strategically adapt to unreliable social partners

Katherine Adams Shannon¹, Aneesa Conine-Nakano¹, Willem E. Frankenhuys^{2,3}, Michael C. Frank¹, & Hyowon Gweon¹

¹Department of Psychology, Stanford University

²Department of Evolutionary and Population Biology, University of Amsterdam

³Max Planck Institute for the Study of Crime, Security and Law, Germany

Abstract

Children learn a lot from others, but the effectiveness of their social learning depends on the reliability of others' help. How do children adapt their future learning decisions based on the past reliability of receiving help? In two experiments, 4- to 6-year-olds ($N = 60$ each) interacted with a researcher who either followed through on promised help (Reliable condition) or failed to do so (Unreliable condition). Experiment 1 was inconclusive. However, with an improved design, Experiment 2 found that children in the Unreliable condition were more likely to forego a harder but more rewarding puzzle as their next task and choose an easier, less rewarding puzzle instead compared to those in the Reliable condition. Such decisions, while seemingly maladaptive at face value, likely reflect an adaptive response to the low likelihood of receiving help. These results extend our understanding of social learning across diverse ecological contexts.

Keywords: social learning; adaptation; cognitive development

Introduction

Early learning is both active and social; children learn not only by independently exploring their environment but also by seeking help from others (Bonawitz et al., 2011; Shneidman et al., 2016; Stahl & Feigenson, 2015). Children can flexibly switch between these strategies, asking for help selectively when a learning task is too complex to solve alone (Bonawitz et al., 2018; Cluver et al., 2013; Gweon & Schulz, 2011). Yet, receiving help when requested is not a given: Teachers might promise help but fail to deliver, or parents may be unavailable to help or even unresponsive.

All children experience variation in the reliability of help across time, contexts, and social partners. Yet some children have more exposure to unreliable help than others. For example, although there is substantial heterogeneity in caregiving at all socioeconomic levels—between families living in the same community (Kuchirko & Tamis-LeMonda, 2019) and even among children within a single household (von Stumm & Latham, 2018)—children whose caregivers suffer from poverty-related stress are more likely to experience unresponsive or inconsistent caregiving (Evans et al., 1999; Evans, 2004; Frankenhuys & Amir, 2022; Pinderhughes et al., 2001; Roeters et al., 2010). Despite the significant challenge of navigating contexts that vary in the reliability of receiving help, little is known about how children's experience of reliable or unreliable help shapes their future learning decisions.

Research in cognitive development suggests that young children are 'adaptive' social learners, quickly changing their strategies depending on a particular form of reliability: in-

formativeness (i.e., whether people provide accurate and sufficient information for learning; Hembacher & Frank, 2017; Lewis & Frank, 2016; Sobel & Kushnir, 2013; Xu & Tenenbaum, 2007). For instance, preschoolers can track subtle variations in the relative accuracy between teachers (Pasquini et al., 2007), and discount prior testimony if new evidence shows that a previously informative partner has become uninformative (Corriveau & Harris, 2009; Ronfard et al., 2017; Scofield & Behrend, 2008). Furthermore, children can modulate their future learning from teachers depending on their past informativeness, showing more compensatory exploration when a teacher was underinformative (i.e., accurate but not sufficient) in the past (Gweon et al., 2014).

However, reliability or trustworthiness isn't defined solely in terms of informativeness; reliability can also mean being responsive to bids for help. Because prior research has assumed learning environments that are relatively supportive in this respect, how children adapt their behaviors to other ecological forms of social reliability—such as whether or not an adult responds as promised—remains an open question.

Kidd and colleagues (2013) took a first step in this direction using a variant of the classic delay-of-gratification 'marshmallow' task, showing that children's ability to wait for a promised reward was influenced by the researcher's past reliability. Before the marshmallow task, the researcher either showed herself to be reliable or unreliable based on whether or not she gave the child a promised reward. Children in the unreliable condition had much shorter wait times than those in the reliable condition, choosing the one guaranteed marshmallow over the promise of two marshmallows later. Typically, short wait times are taken as an index for poor self-control (Mischel & Ebbesen, 1970; Mischel et al., 1988). In the unreliable condition, however, this behavior reflects a reasonable low-risk, low-reward strategy. This study and its replications show that even a brief experimental manipulation can induce children to infer that their current social environment will likely be unreliable, and adapt their behavior accordingly (Lee & Carlson, 2015; Moffett et al., 2020).

Critically, however, it is unknown whether the reliability of help also leads to similar trade-offs in children's learning decisions. Consider, for instance, a young child who wants to learn how to ride a bike. But when she asks for help to do so, her parent responds that they are too busy. If the parent usually responds to such requests, the child might continue seeking help despite the initial rejection. However, if a series

4953

of similar bids for help all go unanswered, the child might learn something about her environment: Help may not be available when she needs it. Such an understanding may influence a host of other decisions; for instance, she may avoid challenging tasks, foregoing potential opportunities for learning. These decisions, in turn, can also lead to long-term differences in learning outcomes for children in contexts where help is more or less reliable.

Importantly, avoiding challenging tasks—which might offer valuable learning opportunities—has often been characterized as unmotivated or even maladaptive (De Castella et al., 2013; Dicintio & Gee, 1999; Stipek et al., 1996). However, in an unreliable social environment where bids for help often go unanswered, it actually might be better to avoid a challenging task especially if it cannot be achieved without assistance; instead, it may be more beneficial for the learner to choose an easier task that can be achieved alone. Thus, avoiding challenging tasks, which might appear suboptimal at face value, may be contextually appropriate when a child's past experience is taken into account (Frankenhuis & Del Giudice, 2012). By considering the learner's social context, we can begin to characterize such behaviors as rational adaptations to environmental constraints, rather than deficits of the learner (Ellis et al., 2023; Frankenhuis & Nettle, 2020).

To date, there has been insufficient empirical research that directly tests a model of environmental adaptation in social learning. By taking an experimental approach to examine young children's ability to adapt their learning strategies to the reliability of their social environment, we can gain valuable insights into how children navigate the trade-offs in different environments that vary in the reliability of help. This approach can also inform a framework for exploring how past social experiences can have cascading effects on future learning through children's own adaptive decision-making.

As a first step towards these goals, we present two preregistered experiments with young children. We ask whether children can use their past experiences to infer the likelihood of receiving help in a new learning context and make adaptive decisions about whether to pursue a challenging, more rewarding task or settle for an easier, less rewarding task.

Experiment 1

Our experiment tests whether children adapt their learning decisions based on the past reliability of help. We hypothesize that children who received prior help would choose a more difficult but rewarding learning goal as their next task, because the researcher will likely respond to future requests for help if needed. By contrast, children who did not receive prior help would favor an easier but less rewarding goal, because they might fail to achieve a harder goal if future requests for help are likely to be ignored.

We focus on ages 4 to 6 because children in this age range can track and adapt their learning behaviors to account for the knowledgeability of their social partners (Gweon, 2021), and are going through an important transition into school where

they navigate learning in new social environments (Blair, 2002; Cook & Coley, 2017). We followed a sequential sampling procedure, described below, in which the sample size was not fixed in advance but determined by repeatedly evaluating the key hypothesized effect as data collection progressed, until either a predetermined evidentiary criterion or final N was met (Etikan et al., 2016). Given the relatively long procedure for data collection, this approach allowed us to stop early in the presence of sufficient evidence for or against the hypothesis, or continue to collect a full sample.

Methods

Participants Sixty children (4;1 - 6;10 years, $M_{age} = 5.50$ years, 7% Hispanic/Latino, 4% African American/Black, 37% Asian, 39% Caucasian/White, 2% Native Hawaiian/Pacific Islander, 14% Multiracial, 4% Other) were recruited at a local children's museum. Of the 53 families who reported highest parent education attainment, 81% completed graduate degrees, 4% completed some graduate training, and 15% completed college. Participants were typically developing, and heard English at least 50% of the time at home per caregiver report. Sixteen children were excluded due to child stopping early ($n = 7$), experimenter error ($n = 3$), child heard about experimental manipulation from previous participant ($n = 3$), or child did not meet the inclusion criteria ($n = 3$). Caregivers provided written consent and children verbal assent for participation. Sampling, exclusion, and analytic procedures were preregistered at https://aspredicted.org/JX3_1KQ. See Results for sequential sampling procedure to determine final sample size of $N = 60$.

Materials and procedure This between-subjects experimental design had two phases (Figure 1). In the history phase, children were presented with a project to decorate a piece of paper to display in a cup and take home. The researcher showed the child a brown crayon in a jar, explaining that normally this is the only crayon available to use. The researcher pretended to have difficulty opening the jar but succeeded, exclaiming it was a bit stuck. This sequence demonstrated the researcher's competence at opening stuck jars. The researcher then revealed another jar of cool, colorful crayons, explaining that a friend in the room next door let them borrow these to use instead. The researcher put the brown crayon away. Next, the researcher explained that she had to do some work nearby while the child colored, but the child could ring a bell on the table for any help if needed. She handed the jar of colorful crayons to the child and exited the room.

The child soon discovered that the jar of colorful crayons was also stuck closed (secretly glued shut), prompting them to ring the bell for help. One minute after leaving the room, the researcher returned. In the *reliable help* condition, the researcher secretly switched the glued jar for an identical unglued one and helped the child access the colorful crayons. In the *unreliable help* condition, the researcher explained that she was too busy to help while the child rang the bell and now the friend next door asked for the jar of colorful crayons back,

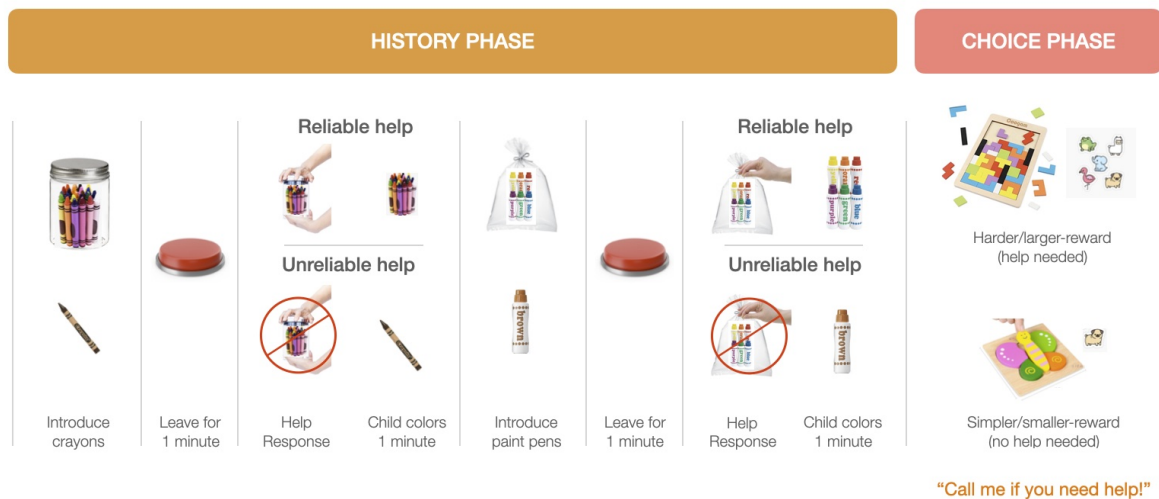


Figure 1: Schematic of the experimental design. The history phase consisted of two episodes of reliable or unreliable help, followed by the choice phase where children selected a puzzle to complete. In Experiment 1, puzzles were visible. In Experiment 2, puzzles were obscured in white boxes.

offering instead the brown crayon to the child. In both conditions, the child worked on coloring for one minute. Thus, children in both conditions tried a difficult task they could not achieve themselves (i.e., opening a stuck jar) but only those in the *reliable help* condition received help and were able to use colorful crayons; children in the *unreliable help* condition received no such help, and as a consequence, had to use the single crayon. Note that the researchers maintained a positive, encouraging demeanor throughout both conditions.

After coloring with crayons, a similar process was repeated with another set of art supplies: a bag of colorful paint pens that was secretly glued shut. As with the crayons, in the *reliable help* condition, the researcher returned and opened the bag, whereas in the *unreliable help* condition, the researcher said the bag needed to be returned and offered a single brown paint pen instead. Thus, by the end of the history phase, children had two consecutive interactions in which they asked for help, and the researcher either responded with promised help or did not respond with promised help.

The choice phase followed the history phase. The same researcher told the child that the next activity was playing with a puzzle. The researcher pointed to the simpler/smaller-reward puzzle and said, "This puzzle is really easy. Most kids can do it by themselves. If you finish this puzzle, you'll get one sticker." The researcher next pointed to the harder/larger-reward puzzle and said, "This puzzle is really hard. Most kids need help to do it. If you finish this puzzle, you'll get five stickers." Finally, the researcher prompted the child to choose a puzzle to complete: "Now listen carefully, I have to go do some more work in the room next door. But I'm really good at puzzles. So, you can call me if you need any help by pressing this bell, okay? So, you can only choose one puzzle: Which puzzle would you like to do?" The left vs. right presentation

position of harder/larger-reward and simpler/smaller-reward puzzles was counterbalanced across conditions.

After the child made the selection, the researcher 'realized' that the session was almost out of time and gave the child the simpler puzzle, explaining that they would rather stay and do this puzzle with the child rather than work more. All participants finished the simpler puzzle with help and encouragement from the researcher, and received 5 stickers upon its completion. This closing procedure allowed us to reestablish reliability for participants in the *unreliable help* condition.

Results

Data and analysis code are available at <https://osf.io/gwsxm>. Our primary dependent variable of interest was children's binary choice of either the simpler/smaller-reward puzzle or harder/larger-reward puzzle. Our key hypothesis (H1) was a negative effect of the *unreliable help* condition on choice of harder/larger-reward puzzle. We planned a logistic regression analysis in R of the form $\text{choice} \sim \text{condition}$.

We computed a Bayes factor (BF) analysis using the BF-pack package (Mulder et al., 2021) with a contingency table BF, assessing the data being more likely under the hypothesis of a condition effect (H1) relative to the null hypothesis (H0: no difference between conditions; assuming Cauchy priors). While we hypothesized finding strong evidence of a condition effect ($BF > 10$), well-recognized guidelines for reporting Bayesian analysis (Van Doorn et al., 2021) also support interpreting $BF > 5$ as moderate evidence, and $BF > 3$ as weak evidence. This analysis informed our sequential sampling procedure to determine the final sample size. We tested an initial sample of 10 children per condition, and then evaluated the BF on the hypothesis of a condition effect after each day of testing. Stopping criteria was set at a $BF > 10$ in favor

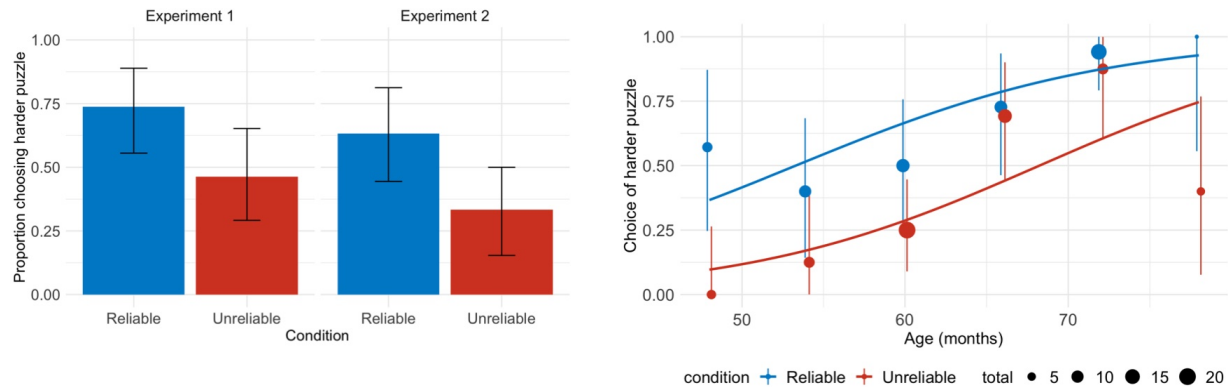


Figure 2: Proportion choosing harder puzzles for Experiments 1 and 2 (left). Choice of harder puzzle by age in pooled sample (right). Dots represent means for a given age group in 6-month bins; dot sizes represent the sample size (n) per bin. Trend lines are best-fitting logistic regression curves. Error bars are 95% bootstrapped CIs for both figures.

of the hypothesis of a condition effect ($H1$), a $BF > 3$ against the hypothesis ($H0$), or at an $n = 30/\text{condition}$ ($N = 60$ total).

Results show that children in the *reliable help* condition had about three times higher odds of choosing the harder/larger-reward puzzle ($n = 22/30$) than children in the *unreliable help* condition ($n = 16/30$; $OR = 3.14$; 95% CI [1.09, 9.64]; see Figure 2). Although the regression p -value was significant ($p = .04$), our primary evaluation of the condition effect produced a $BF = 2.18$, suggesting that the data were only about twice as likely under the hypothesis of a condition effect than that of a null hypothesis of no condition difference. This result did not reach the threshold ($BF > 3$) for weak but meaningful evidence of a condition effect (note that such divergence between frequentist and Bayesian approaches is common; see Benjamin et al. (2018) for a review). We did not reach our sampling stopping criterion of $BF > 10$, thus our sample size reflects the maximum $N = 60$.

As a planned secondary analysis, we conducted a follow-up regression including a main effect of age (in months, centered) and an interaction of age and condition. In the interaction model, *unreliable help* condition ($\beta = -1.12$, 95% CI [-2.36, 0.06], $p = .07$), age ($\beta = 0.06$, 95% CI [-0.03, 0.17], $p = .20$), and the interaction term ($\beta = 0.12$, 95% CI [-0.05, 0.34], $p = .20$) were not significantly associated with puzzle choice. The condition effect had a $BF = 1.36$. In a non-preregistered analysis dropping the interaction term, both *unreliable help* condition ($\beta = -1.27$, 95% CI [-2.51, -0.13], $p = .03$) and older age ($\beta = 0.11$, 95% CI [0.03, 0.20], $p = .01$) were significant predictors of puzzle choice; condition effect had a $BF = 2.37$.

Discussion

Experiment 1 results suggest that both children's prior experience with a more or less reliable social partner and their age may be related to their choice of future tasks. However, these results did not meet our minimum Bayes factor threshold.

Why might the negative effect of unreliable help on choice of harder/larger-reward puzzle choice be smaller than hypoth-

esized? One critical feature of Experiment 1 study design is that the puzzles were visible to the participant. Although the researcher described the puzzles as easy to do alone or hard and generally requiring help, children could also see the puzzles and make their own judgments of whether or not they would need the researcher's help to complete each puzzle. In particular, older children may have judged the presented puzzles to be generally easier to complete alone than younger children. Furthermore, because the puzzles were commercially available, some children may also have played with similar puzzles in the past and already knew how to solve them. Experiment 2 aimed to address these concerns.

Experiment 2

Experiment 2 was identical to Experiment 1, aside from how the puzzles were presented: Instead of showing the actual puzzles, the simpler/smaller-reward and harder/larger-reward puzzles were concealed in separate boxes, thus eliminating any visual cues about the puzzle design or difficulty.

By design, researchers in our study were not blind to condition. Thus, we also implemented additional procedures to minimize any unintended bias introduced by the researcher. We added a separate researcher-facing camera where the researcher's face and hands were shown in a view over the child's shoulder. A coder without knowledge of condition assignment predicted the puzzle choice of the child by watching a video clip of the researcher from the beginning of the choice phase to the moment before the selection of the puzzle. We evaluated whether the coder could systematically predict the child's puzzle choice by watching the researcher alone.

Methods

Participants Sixty children (4;0 - 6;10 years, $M_{age} = 5.36$ years, 4% Hispanic/Latino, 2% African American/Black, 52% Asian, 31% Caucasian/White, 15% Multiracial) were recruited and assessed at a local children's museum. Of the 50 families who reported highest parent education attainment,

78% completed graduate degrees, 2% completed some graduate training, 16% completed college, 2% completed some college, and 2% completed high school. Participants were typically developing. All participants heard English at least 25% of the time; 94% of participants heard English at least 50% of the time per caregiver report. All participants belonged to separate households. Eight children were excluded due to experimenter error ($n = 4$), child stopping early ($n = 3$), or child did not meet the inclusion criteria ($n = 1$). Sampling, exclusion, and analytic procedures were preregistered at https://aspredicted.org/WSS_KQ6. The sequential sampling procedure as described in Experiment 1 was implemented, leading to our final sample size of $N = 60$.

Materials and procedure Materials and procedure were identical to Experiment 1, with the exception that the puzzles were hidden in a box. The harder/larger-reward puzzle was hidden in a medium, unmarked white box, and the simpler/smaller-reward puzzle was hidden in a small, unmarked white box. The sizes differed to offer an additional cue to help children distinguish between the boxes as the researcher described each puzzle within.

Results

Except where noted, all sampling and analytic procedures were the same as in Experiment 1.¹

In this second version of the experiment with puzzles obscured ($N = 60$, see Figure 2), children in the *reliable help* condition were more likely to choose the harder/larger-reward puzzle ($n = 19/30$) than children in the *unreliable help* condition ($n = 10/30$; $OR = 3.45$; 95% CI [1.22, 10.35], $p = .02$). This analysis produced a $BF = 3.5$ in favor of the hypothesis of a condition effect over the null hypothesis, which we interpret as weak but meaningful evidence of an effect (Van Doorn et al., 2021) although it indicates less support than our hypothesis of a strong condition effect ($BF > 10$).

In our planned secondary analysis investigating the effects of age on puzzle choice, we find no significant interaction of age and condition ($\beta = -0.05$, 95% CI [-0.24, 0.12], $p = .54$). *Unreliable help* condition ($\beta = -1.51$, 95% CI [-2.86, -0.29], $p = .02$) and age ($\beta = 0.17$, 95% CI [0.06, 0.33], $p = .01$) were both significant predictors of puzzle choice, and the condition effect had a $BF = 3.86$. In a non-preregistered analysis dropping the interaction term, both *unreliable help* condition ($p = .02$) and age ($p < .001$) remained significant main effects, and the condition effect had a $BF = 4.35$.

Importantly, we consider it unlikely that the researcher's behavior affected the children's puzzle choice because a blind coder could not reliably predict the child's choice of the puzzle given the researcher-facing video. The unweighted/weighted Cohen's Kappa value of 0.19 (95% CI [-0.061, 0.45]) suggests little systematic agreement between

the blind coder and the child's actual choice of puzzle.

Pooled results

Finally, in a post-hoc, non-preregistered analysis, we pooled participants across Experiment 1 and 2 to increase power to detect an age-by-condition interaction, adding experiment as a fixed effect (see Figure 2, right). In the pooled data ($N = 120$), we again find no significant interaction of age and condition ($\beta = 0.03$, 95% CI [-0.09, 0.15], $p = .61$) nor experiment ($\beta = -0.48$, 95% CI [-1.35, 0.35], $p = .26$). *Unreliable help* condition ($\beta = -1.39$, 95% CI [-2.28, -0.55], $p = .002$) and age ($\beta = 0.11$, 95% CI [0.04, 0.20], $p = .004$) were both significant predictors of puzzle choice. The condition effect had a $BF = 26.30$, an expected increase given the doubled sample size. Both *unreliable help* condition ($p = .002$) and age ($p < .001$) remained significant when the interaction term was removed, and the condition effect had a $BF = 28.18$.

General Discussion

The current project examined environmental adaptation in children's social learning behavior. We asked whether 4- to 6-year-olds tailored their choice of future learning goal based on the past reliability of receiving help. Experiment 1 was inconclusive; we changed the design in Experiment 2 to reduce unintended variation in children's judgment of task difficulty and, in turn, need for help. Results show weak but meaningful evidence that children who experienced unreliable help were more likely to forego a challenging but more rewarding task and choose an easier but less rewarding task instead, compared to those who received help when needed.

Why did children make these different learning decisions depending on the reliability of help? In the *reliable help* condition, the harder/larger-reward puzzle was a reasonable choice because children could harness a responsive partner who could provide help if needed. Within an explore/exploit framework (Frankenhuis & Gopnik, 2023), the harder/larger-reward puzzle choice may be viewed as the child exploring their skills ('Can I solve this to earn 5 stickers?') with little risk of failure if the task was too difficult. In the *unreliable help* condition, however, that would have been a risky choice; given that their bids for help are likely to go unanswered, children played it safe by attempting the simpler puzzle on their own to get only 1 sticker than gain nothing by trying (and ultimately failing to solve) a very difficult task (i.e., choosing the simpler puzzle exploited the child's known skills; 'I can solve this to earn 1 sticker'). In other words, what counts as a more adaptive choice differed depending on the reliability of help, and children made choices that reflected this pattern.

The decision to avoid a challenging task (and therefore forego an opportunity to gain a higher reward) may, at face value, seem like a maladaptive behavior. However, in a context where help is not a given, it may be adaptive behavior; in the context of our experiment, these choices can be explained as an inference informed by utility-based social reasoning (see Baker et al., 2017; Jara-Ettinger et al., 2016 for utility-based reasoning as a computational framework for

¹Due to unequal distribution of age across conditions, we continued data collection despite reaching our sequential sampling stopping criterion of $BF(H1) = 19.17$ at $n = 28$. Age was better distributed across condition for the subsequent BF analysis at $n = 34$, and the criterion of $BF(H1) > 10$ was no longer met.

social cognition), wherein children consider the behavior of their social partners when reasoning about the costs and benefits of a given learning decision. Such reasoning may not necessarily involve explicit awareness of their own utilities, but nonetheless inform children's implicit decision-making processes. Although children around this age can reason about another learner's expected utilities to decide what to teach (Bridgers et al., 2018), little work has directly investigated how utility-based social reasoning may account for learning adaptation in contexts where support is unavailable.

Extending this work to children's own learning decisions, children's adaptations to past reliability can be explained as decisions that maximize their own expected utilities (i.e., difference between costs and benefits) that take into account the difficulty and reward of the puzzle and the inferred likelihood of receiving help based on past experience. If children can make accurate predictions about the responsiveness of their social partners, they can make utility-maximizing decisions for themselves in their choice of future learning tasks.

However, children's predictions about the availability of help may not always be accurate. In particular, children who might face variable reliability—either due to unpredictability within an environment or large discrepancies across different environments (e.g., home vs. school)—might miscalibrate their anticipation of help and make suboptimal learning decisions. Such cases highlight the double-edged nature of these adaptive responses. For instance, preemptively avoiding tasks that require help, though well-reasoned, could lead to missed opportunities in new contexts where help is available; conversely, a child who is used to receiving help may lose valuable opportunities for self-guided discovery by waiting for help that is ultimately undelivered. By demonstrating how children adapt to past reliability of help, the current work provides empirical grounds for further studying the consequences of such miscalibrated adaptations.

In a similar vein, children may also be miscalibrated in their estimates of future task difficulty. In the current study, the two puzzles were explicitly marked as difficult and easy, and it was implied that the difficult puzzle may be *too* difficult for children to solve on their own. Yet, in real-world contexts, tasks are rarely marked in terms of whether children can do it on their own. Thus, there is rarely a single 'best' strategy in learning; both help-seeking and exploration are important and useful, and the effectiveness of children's learning decisions may be further modulated by their persistence as well as their ability to estimate the cost of exploration or the difficulty of new tasks (Gweon & Schulz, 2019; Gweon et al., 2017; Leonard et al., 2017, 2020). In ongoing work, we are examining how children's past experiences of unreliable help influence their exploration and discovery.

While we found a main effect of age, suggesting that older children were more likely to choose the harder task, we did not find an age-by-condition interaction. Children across ages 4 to 6 appeared to adapt similarly to the reliability of help on their choice of learning goals. This result highlights the nu-

anced nature of how children may balance factors like task demands, their own skill, and the availability of help when setting future goals; older children may have been motivated to try the harder task despite unreliable help because they are generally more confident in their abilities. An understanding of how diverse experiences, samples, and contexts shape environmental adaptation is critical to pursue in future research.

Similar to the prior work by Kidd et al. (2013), we found evidence of adaptation to a relatively 'light' dose of unreliability. Children modified their choice of future learning goal after only two brief interactions with a new adult who was either a reliable or unreliable helper. The effectiveness of the manipulation is noteworthy because much care was taken to ensure that the researchers, in both conditions, maintained a warm and kind demeanor. The unreliable helper also gave the excuse of being 'too busy' (rather than refusing) to ensure that they were not perceived as antagonistic. Thus, our manipulation of reliability was not only effective but also ecologically valid in ways that ensured the well-being of our participants.

One related consideration for interpreting our results is the socioeconomic background of our sample. Our participants were recruited in a science museum located in an affluent neighborhood. It is notable that we found an adaptive response to unreliable help even in these children who likely have strong priors that adults are generally dependable (Stamos et al., 2019). Yet, it is also possible that the experience of unreliable help was particularly salient for these children, contributing to the effectiveness of such a light dose of unreliability. In future work, we plan to extend these findings by recruiting participants from diverse backgrounds, including those from more or less supportive contexts.

By understanding how children's expectations for future help might shape their choice of what to learn, we can recast behaviors previously considered maladaptive as an adaptive response to the social environment. From this point of view, the educational implications of this direction of work are clear. Educators (and social workers) who may at first view a help-avoidant child as unmotivated or antisocial may instead consider how to establish trust and aim to help children show flexibility in their social learning strategies, such as relying more or less on others' help depending on the context (Landry et al., 2017; Sheridan, 2008).

Looking forward, an important direction for this research is to consider what it means to provide supportive learning environments, especially given the variability in children's experience and expectations of reliability. By encouraging flexibility in both help-seeking and exploration, we can help children maximize their learning potential.

Acknowledgments

This work was supported by funding from NICHD 1F32HD105705-01 (KAS); the Dutch Research Council (V1.Vidi.195.130) and the James S. McDonnell Foundation (<https://doi.org/10.37717/220020502>) (WEF); McDonnell Scholars Award and NSF BCS-2019567 (HG).

References

- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 1–10. doi: 10.1038/s41562-017-0064
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., ... others (2018). Redefine statistical significance. *Nature human behaviour*, 2(1), 6–10.
- Blair, C. (2002). School readiness: Integrating cognition and emotion in a neurobiological conceptualization of children's functioning at school entry. *American psychologist*, 57(2), 111.
- Bonawitz, E., Bass, I., & Lapidow, E. (2018). Choosing to Learn: Evidence Evaluation for Active Learning and Teaching in Early Childhood. In M. M. Saylor & P. A. Ganea (Eds.), *Active Learning from Infancy to Childhood* (pp. 213–231). Cham: Springer International Publishing. doi: 10.1007/978-3-319-77182-3_12
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322–330. (00278) doi: 10.1016/j.cognition.2010.10.001
- Bridgers, S. E. C., Jara-Ettinger, J., & Gweon, H. (2018). *Young children consider the expected utility of others' learning to decide what to teach* (preprint). PsyArXiv. doi: 10.31234/osf.io/mwr6a
- Cluver, A., Heyman, G., & Carver, L. J. (2013). Young children selectively seek help when solving problems. *Journal of Experimental Child Psychology*, 115(3), 570–578. doi: 10.1016/j.jecp.2012.12.011
- Cook, K. D., & Coley, R. L. (2017). School transition practices and children's social and academic adjustment in kindergarten. *Journal of Educational Psychology*, 109(2), 166.
- Corriveau, K., & Harris, P. L. (2009). Choosing your informant: weighing familiarity and recent accuracy. *Developmental Science*, 12(3), 426–437. doi: 10.1111/j.1467-7687.2008.00792.x
- De Castella, K., Byrne, D., & Covington, M. (2013). Unmotivated or motivated to fail? a cross-cultural study of achievement motivation, fear of failure, and student disengagement. *Journal of educational psychology*, 105(3), 861.
- Dicintio, M. J., & Gee, S. (1999). Control is the key: Unlocking the motivation of at-risk students. *Psychology in the Schools*, 36(3), 231–237.
- Ellis, B. J., Abrams, L. S., Masten, A. S., Sternberg, R. J., Tottenham, N., & Frankenhuis, W. E. (2023). *The hidden talents framework: Implications for science, policy, and practice*. Cambridge University Press.
- Etikan, I., Alkassim, R., & Abubakar, S. (2016). Comparison of snowball sampling and sequential sampling technique. *Biometrics and Biostatistics International Journal*, 3(1), 55.
- Evans, G. W. (2004). The Environment of Childhood Poverty. *American Psychologist*, 59(2), 77–92. doi: 10.1037/0003-066X.59.2.77
- Evans, G. W., Maxwell, L. E., & Hart, B. (1999). Parental language and verbal responsiveness to children in crowded homes. *Developmental psychology*, 35(4), 1020.
- Frankenhuis, W. E., & Amir, D. (2022). What is the expected human childhood? insights from evolutionary anthropology. *Development and Psychopathology*, 34(2), 473–497.
- Frankenhuis, W. E., & Del Giudice, M. (2012). When do adaptive developmental mechanisms yield maladaptive outcomes? *Developmental Psychology*, 48(3), 628–642. doi: 10.1037/a0025629
- Frankenhuis, W. E., & Gopnik, A. (2023). Early adversity and the development of explore–exploit tradeoffs. *Trends in Cognitive Sciences*, 27(7), 616–630. doi: 10.1016/j.tics.2023.04.001
- Frankenhuis, W. E., & Nettle, D. (2020). The strengths of people in poverty. *Current Directions in Psychological Science*, 29(1), 16–21.
- Gweon, H. (2021). Inferential social learning: Cognitive foundations of human social learning and teaching. *Trends in Cognitive Sciences*, 25(10), 896–910.
- Gweon, H., Asaba, M., & Bennett-Pierre, G. (2017). Reverse-engineering the process: Adults' and preschoolers' ability to infer the difficulty of novel tasks. In *Cogsci*.
- 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. (00049) doi: 10.1016/j.cognition.2014.04.013
- Gweon, H., & Schulz, L. (2011). 16-Month-Olds Rationally Infer Causes of Failed Actions. *Science*, 332(6037), 1524–1524. doi: 10.1126/science.1204493
- Gweon, H., & Schulz, L. (2019). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child development*, 90(1), e148–e164.
- Hembacher, E., & Frank, M. C. (2017). Children's social referencing reflects sensitivity to graded uncertainty. In *Proceedings of the 38th annual conference of the cognitive science society* (pp. 495–500).
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The Naïve Utility Calculus: Computational Principles Underlying Commonsense Psychology. *Trends in Cognitive Sciences*, 20(8), 589–604. doi: 10.1016/j.tics.2016.05.011
- Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental

- reliability. *Cognition*, 126(1), 109–114. doi: 10.1016/j.cognition.2012.08.004
- Kuchirko, Y. A., & Tamis-LeMonda, C. S. (2019). The cultural context of infant development: Variability, specificity, and universality. In *Advances in child development and behavior* (Vol. 57, pp. 27–63). Elsevier.
- Landry, S. H., Zucker, T. A., Williams, J. M., Merz, E. C., Guttentag, C. L., & Taylor, H. B. (2017). Improving school readiness of high-risk preschoolers: Combining high quality instructional strategies with responsive training for teachers and parents. *Early Childhood Research Quarterly*, 40, 38–51.
- Lee, W. S. C., & Carlson, S. M. (2015). Knowing When to Be “Rational”: Flexible Economic Decision Making and Executive Function in Preschool Children. *Child Development*, 86(5), 1434–1448. (eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cdev.12401>) doi: 10.1111/cdev.12401
- Leonard, J. A., Garcia, A., & Schulz, L. E. (2020). How adults’ actions, outcomes, and testimony affect preschoolers’ persistence. *Child development*, 91(4), 1254–1271.
- Leonard, J. A., Lee, Y., & Schulz, L. E. (2017). Infants make more attempts to achieve a goal when they see adults persist. *Science*, 357(6357), 1290–1294.
- Lewis, M. L., & Frank, M. C. (2016). Understanding the effect of social context on learning: A replication of Xu and Tenenbaum (2007b). *Journal of Experimental Psychology: General*, 145(9), e72–e80. (00006) doi: 10.1037/xge0000203
- Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of personality and social psychology*, 16(2), 329.
- Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. *Journal of personality and social psychology*, 54(4), 687.
- Moffett, L., Flannagan, C., & Shah, P. (2020). The influence of environmental reliability in the marshmallow task: An extension study. *Journal of Experimental Child Psychology*, 194, 104821. doi: 10.1016/j.jecp.2020.104821
- Mulder, J., Williams, D. R., Gu, X., Tomarken, A., Böing-Messing, F., Olsson-Collentine, A., ... van Lissa, C. (2021). Bfpack: Flexible bayes factor testing of scientific theories in R. *Journal of Statistical Software*, 100(18), 1–63. doi: 10.18637/jss.v100.i18
- Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, 43(5), 1216–1226. doi: 10.1037/0012-1649.43.5.1216
- Pinderhughes, E. E., Nix, R., Foster, E. M., Jones, D., & Group, C. P. P. R. (2001). Parenting in context: Impact of neighborhood poverty, residential stability, public services, social networks, and danger on parental behaviors. *Journal of marriage and family*, 63(4), 941–953.
- Roeters, A., Van Der Lippe, T., & Kluwer, E. S. (2010). Work characteristics and parent-child relationship quality: the mediating role of temporal involvement. *Journal of Marriage and family*, 72(5), 1317–1328.
- Ronfard, S., Lane, J. D., Wang, M., & Harris, P. L. (2017). The impact of counter-perceptual testimony on children’s categorization after a delay. *Journal of Experimental Child Psychology*, 163, 151–158. doi: 10.1016/j.jecp.2017.06.006
- Scofield, J., & Behrend, D. A. (2008). Learning words from reliable and unreliable speakers. *Cognitive Development*, 23(2), 278–290. doi: 10.1016/j.cogdev.2008.01.003
- Sheridan, S. M. (2008). *Getting ready: Promoting school readiness through a relationship-based partnership model* Unpublished doctoral dissertation. University of Nebraska-Lincoln.
- Shneidman, L., Gweon, H., Schulz, L. E., & Woodward, A. L. (2016). Learning From Others and Spontaneous Exploration: A Cross-Cultural Investigation. *Child Development*, 87(3), 723–735. (00006) doi: 10.1111/cdev.12502
- Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, 120(4), 779–797. doi: 10.1037/a0034191
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants’ learning and exploration. *Science*, 348(6230), 91–94. (00065)
- Stamos, A., Altsitsiadis, E., & Dewitte, S. (2019). Investigating the effect of childhood socioeconomic background on interpersonal trust: Lower childhood socioeconomic status predicts lower levels of trust. *Personality and Individual Differences*, 145, 19–25.
- Stipek, D. J., et al. (1996). Motivation and instruction. *Handbook of educational psychology*, 1, 85–113.
- Van Doorn, J., Van Den Bergh, D., Böhm, U., Dablander, F., Derks, K., Draws, T., ... others (2021). The jasp guidelines for conducting and reporting a bayesian analysis. *Psychonomic Bulletin & Review*, 28, 813–826.
- von Stumm, S., & Latham, R. M. (2018). Early life experiences: Meaningful differences within and between families. *Infant Behavior and Development*, 53, 56–63.
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, 114(2), 245–272. (00564) doi: 10.1037/0033-295X.114.2.245