Title

Choice Strategies in a Changing Social Learning Environment

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Plate, Rista C.

Shutts, Kristin

Cochrane, Aaron

et al.

Publication Date

2020

Choice Strategies in a Changing Social Learning Environment

Rista C. Plate
rplate@sas.upenn.edu
Department of Psychology, University of Pennsylvania, 3720 Walnut St.
Philadelphia, PA 19104 USA

Kristin Shutts, Aaron Cochrane, C. Shawn Green, & Seth D. Pollak
{kshutts, akcochrane, csgreen2, spollak} @wisc.edu
Department of Psychology, University of Wisconsin-Madison, 1202 W Johnson St.
Madison, WI 53706 USA

Abstract

One challenge that children face when learning from others is that social agents can behave in unpredictable ways. Social agents may acquire—or fail to acquire—new information that influences how they interact with the learner. Little is known about children's sensitivity to these changes or how effectively children update their own behavior in response. Participants (*N* = 129) searched for rewards while receiving suggestions from a social agent. The suggestions changed in level of reliability over time. All children updated how heavily they weighted the cues after the change. However, younger children were more influenced by their initial experience with the suggestions, indicating that younger children may have more difficulty disengaging from social information in uncertain learning environments.

Keywords: social learning; statistical learning; development

Introduction

When making choices, children often have access to multiple sources of information, including their own observations or past experiences and information from others. Information from others can be useful: Individuals may have expertise and information not directly available to the learner. However, learning from others poses a number of potential challenges. One challenge is that the learner needs to assess whether information provided from another is reliable. An additional challenge is that another's behavior may change over time. Therefore, the learner has to remain flexible and able to update their behavior as needed. In the present research, we ask whether and how children update their own choice behavior after a change in a social agent's behavior.

Social Input Informs Children's Knowledge

Across myriad situations, children seek and trust information provided by others. For example, children imitate the actions of adults (Meltzoff, 1988) and look to adults for cues about which actions are safe to perform (e.g., Sorce et al., 1985). Children's proclivity for learning from others confers many advantages. By attending to knowledgeable others, children glean information they could not acquire through direct experience, acquire information faster than the time it takes to discover it independently, and foster affiliation (Harris, 2012).

Although children look to people to learn about the world, children are discerning about when and from whom they use social information, suggesting that children are sensitive to the fact that social agents may not always provide correct information. For instance, infants tend to imitate actions *only* when the actions appear rational (Gergely et al., 2002) and children favor information provided by those who appear accurate, reliable, knowledgeable, and intelligent, as well as those who have access to relevant information (see Sobel & Kushnir, 2013 for a review).

Some research has pitted social input against children's own knowledge to determine whether children favor social input over information they have gleaned from their observations. There is some evidence to suggest that children are willing to change their behavior when social information is available (Li & Yow, 2018). For example, preschool-age children in one study abandoned their chosen label for an ambiguous object (i.e., object that was a morphed combination of two objects) in favor of a label provided by a previously accurate informant (Li & Yow, 2018). However, other evidence suggests that having the opportunity to evaluate a social agent's reliability against one's own observations can lead to reduced trust in an informant, and thus greater reliance on one's own observations when a social agent provides poor information (Bridgers et al., 2016; Ronfard & Lane, 2018; Ronfard et al., 2017). Taken together, although some evidence suggests that social input is privileged, there is growing support that children's selective trust is consistent with a rational model whereby children assess input from others against their own knowledge (Sobel & Kushnir, 2013).

The rational model perspective suggests that social learning and learning from one's direct experiences share a common mechanism that allows the learner to make inferences based both on existing knowledge and access to evidence. In other words, children are not expected to universally favor social information, but rather they are expected to attend to extant evidence in order to determine how to weight social information in any given learning scenario. Yet, previous research has not taken into account that social agents may change their minds or acquire (or fail to acquire) new knowledge over time, therefore leading to changes in their behavior. Little is known about whether and how children update their use of social information in light of changes in a social agent's behavior.

Children's Choices Are Dynamically Influenced by the Environment

There is some evidence that children sensitively respond to changing features of the environment. Children's choice strategies are not static; children ages 4–11 update the strategies they use to obtain rewards continuously as more information about a statistical distribution or pattern is acquired (Plate et al., 2018). Additionally, children are sensitive to changes in patterns. For example, preschool-age children completed a probabilistic learning task in which one of three locations was rewarded most frequently (Starling et al., 2018). Children learned to choose the most-frequently rewarded choice more often than the other choices. When the frequencies of reward changed, children updated their choices by abandoning their previous choice and selecting the location that was presently rewarded most often. Yet, older children were more proficient at changing their behavior as compared to younger children. The studies described above provide some evidence that children continue to monitor the environment for changes that might influence their own choices.

How children respond to changing *social* information, though, is unknown. There are some reasons to think that children may have more difficulty revising their choices in response to changing social information. For instance, individuals readily apply stable traits to social agents, even on the basis of minimal information (e.g., Uleman et al., 2008) suggesting that initial interactions may obstruct future, contrasting, evidence. Additionally, children younger than 5 years have difficulty integrating information about a social agent's consistency when making predictions about future behaviors (Boseovski & Lee, 2006). Further, children may have an expectation that adults are knowledgeable (Lampinen & Smith, 1995), particularly when there are cues to suggest knowledge (e.g., pointing: Palmquist & Jaswal, 2012). Therefore, children may have more difficulty adapting to changes in a social agent's behavior because of expectations of knowledge and trait attributions.

Present Research

Participants completed a learning task adapted from Plate and colleagues (2018). In Plate et al. (2018), participants searched for a reward hidden behind multiple, differentially rewarded, locations. Children's initial choice behavior reflected the underlying statistical distribution of rewards; specifically, children adopted a probability-matching strategy of choosing each location at the rate it was rewarded. Such behavior suggested that children learned how often the reward appeared in each location. However, children changed strategies partway through the task: they transitioned to a maximizing strategy, choosing the most frequently rewarded location. Maximizing is a choice strategy that optimizes rewards because it is impossible to know where a reward would appear on any given trial.

In the present experiment, participants also searched for a reward hidden behind differentially rewarded locations. However, prior to making a selection on each trial, participants received a cue regarding the location of the reward. Participants were told that the cue was a suggestion from another player, whom they had met in the waiting room. In reality, the suggested locations were predetermined. The suggestions changed part-way through the task from being reliable (i.e., most often cueing the rewarded location) to being unreliable (i.e., most often cueing an incorrect location) or vice versa. We measured the strategies children used and the weight attributed to suggestions before and after the change. We were additionally interested in whether age might influence flexibility in response to changing social information. We expected that younger children may have difficulty changing their choice behavior when receiving suggestions that are initially reliable (and later unreliable). Given that younger children show less flexibility on a nonsocial task (Starling et al., 2018) and have an expectation that adults will be knowledgeable (Lampinen & Smith, 1995; Palmquist & Jaswal), we reasoned that they may continue to be lured by the previously reliable suggestions. Such a difference may be smaller when interacting with a previously unreliable information source. A reasonable alternative hypothesis is that, because the confederate diverges from the reward distribution for the first half of the experiment, children will learn to ignore this information source.

Method

Participants

Participants included 129 children ages 4-9-years-old (12 Hispanic or Latino; 10 African American, 4 Asian American, 9 Multiracial, 103 White, 3 chose "other" or did not report race). Nine additional participants were excluded for not completing the task (N = 6) or experimenter error (N = 3).

Procedure

The experimenter brought the participant to the waiting room and explained that the participant would play a computerized game where the goal was to find gold coins hidden behind rocks. The participant met an adult confederate (who, unbeknownst to the participant, was a research assistant in the lab). The majority of confederates were White and female. Participants learned that the confederate had played the game once before and would provide suggestions about which rock to choose throughout the game. The experimenter also explained that the confederate would play the game in another room and would not be able to see the participant's choices. The participant and confederate then proceeded to two separate testing rooms situated across a hallway from each other.

Example Phase The example phase allowed the participant to gain experience with the task and see multiple possibilities for choosing to follow or not follow the confederate's suggestions and possible outcomes. The participants saw the following fixed choice patterns and outcomes: 1) choose a different location from the confederate but neither finds the coin, 2) choose the same location as the confederate but the coin is in a different location, and 3) choose a different location from the confederate and find the coin. Thus, participants saw multiple possible options for how to respond. Prior to starting the testing trials, the experimenter queried the participant to ensure that the participant could verbalize: 1) that they would be looking for coins under rocks, 2) that they would be receiving suggestions regarding the location of the coins, and 3) that they could choose the same or a different location than suggested by the confederate.

Test Phase After the practice phase, an icon meant to suggest that the participant's computer was connecting to the confederate's computer appeared on the screen ("Waiting for Player 2…"). Additionally, during this time, the experimenter went to the confederate's testing room to "set up." Specifically, the experimenter entered the confederate's room and asked the confederate if they remembered how to play the game (a conversation that could be heard by the child across the hallway with both doors open) before closing the door to the confederate's room. After approximately one minute, the experimenter left the confederate's room saying "good luck" to the confederate. After 30 seconds, text that read "Player 2 ready. Please wait for experimenter." appeared on the participant's screen. The experimenter returned to the participant's room, wished the participant good luck, and instructed the participant to begin the game by clicking the mouse. The simple set-up proved to make the confederate's involvement in the game believable to participants. None of the children tested indicated doubt about the confederate and many made spontaneous comments about the confederate during the task (e.g., "Emily found the coin that time!"). There were 200 test trials, and the experiment lasted approximately 15 minutes. At the start of each test trial, eight rocks appeared on the screen with equal spacing along a horizontal line (example displays are shown in Figure 1). Before participants were allowed to respond, a pointing hand indicated the confederate's suggestion. The task was programmed such that participants were unable to respond prior to seeing the confederate's suggestion. When participants selected the correct location on a trial, a coin appeared in place of the rock they selected. When participants selected an incorrect location on a trial, a red "X" appeared in the chosen location and the coin was revealed in the correct location.

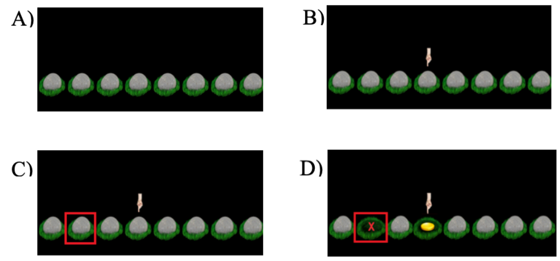

Figure 1. Progression of computerized task: A) Display prior to participant choice. B) Pointing hand indicates confederate's suggestion. C) Red box highlights participant's choice. D) Coin appears in rewarded location. Note: Figure illustrates one instance of choice behavior (i.e., choose rock not suggested by confederate) and one possible outcome (i.e., participant fails to find the coin; confederate finds the coin).

Design

From left to right, the following probabilities defined the likelihood of a coin appearing at each rock location on any given trial: 0% - 0% - 5% - 10% - 70% - 10% - 5% - 0%. To make all participants' experiences statistically equivalent, the outcomes were predetermined to ensure a perfect match to the predefined location probabilities across trial blocks (i.e., for each 100-trials, rock five was rewarded on exactly 70 trials, rocks four and six were rewarded on exactly 10 trials, etc.). Participants were not shown the probabilities; the probabilities had to be learned via experience with the task.

The choices of the confederate varied across two conditions: reliable$\rightarrow$unreliable (i.e., the confederate suggested the correct rock 90% of the time for the first 100 trials and 10% of the time for the second 100 trials, $N = 63$, 31 male, 32 female, $M_{age} = 6.938$, $SD_{age} = 1.691$) and unreliable$\rightarrow$reliable (i.e., the confederate suggested the correct rock 10% of the time for the first 100 trials and 90% of the time for the second 100 trials, $N = 66$, 37 male, 29 female, $M_{age} = 7.068$, $SD_{age} = 1.630$; no difference in participant age ($p = .531$) or gender ($p = .328$) by condition). There was no break or any other change in the experiment that would draw participants' attention to the change in reliability.

Models of Strategy Use

We first characterized the strategies participants used by assessing the extent to which individual participant choices were best captured by one of five different possible models of choice behavior. In brief, the first model was a *probability-matching model*. Here participants were expected to choose each option in proportion to the probability that each location had been observed to be correct up to the current trial. The second model was a *maximizing model*. Under this model, participants were expected to choose the option that had been observed to have the highest probability of reward up to the current trial. The third model was a *confederate-matching model,* in which the participant's distribution of choices was expected to be the same as the confederate's distribution of choices. The fourth model was a *confederate-following model,* in which participants were expected to choose the option suggested by the confederate (this model can be thought of as maximizing on the confederate's suggestions). The final model was a *random choice model*, in which there was an equal and constant probability of the participant selecting each of the eight options. We also included a "lapse rate" of 0.01 added to all locations that otherwise would have had a probability of zero, with probabilities at all other locations being proportionally adjusted such that the total probability at all locations would sum to one. This small offset allowed all log likelihoods to be evaluated without a single unexpected choice causing the probability of a model to immediately fall to zero and is standard practice in behavioral model fitting (e.g., Kattner et al., 2017, Klein, 2001; Wichmann & Hill, 2001). We tested various lapse rates (e.g., .02, .005) and inferences from models were not changed. The log likelihood was computed for each trial of each participant.

For any given model, the fit of the model was calculated as the summed log likelihood of the predicted probabilities at the chosen locations. For the purposes of optimization, this value was then multiplied by -1 to find the negative log likelihood. These values were then summed within participant, resulting in negative log likelihoods for each model for each task half for each participant.

Results

Children's Use of Strategies

Using a direct model selection approach of choosing the best-fitting model based on the highest log likelihood, we found that prior to the switch in confederate reliability, the majority of participants who received reliable suggestions were best fit by the confederate-following model (92%) while participants who received unreliable suggestions were best fit by the probability-matching model (79%; see Table 1 for percent of participants best fit by all models). Post-switch, participants changed their strategy use. The majority of participants who initially received reliable suggestions, but now received unreliable suggestions, were best fit by the probability-matching model (68%). Participants who initially received unreliable suggestions, but now received reliable suggestions, were best fit by the confederate-following model (80%). Regardless of whether participants received reliable suggestions in the first half or second half of the experiment, there was no difference in the proportion of participants who were best fit by the confederate-following model ($X^2(1) = 2.798$, $p = .094$; similarly there was no difference in proportion of participants receiving unreliable suggestions fit by the probability-matching model in the first or second half of the experiment, $X^2(1) = 1.340$, $p = .247$). Given the continuous nature of the fit likelihoods and task half as a variable of interest, we used a Linear Mixed Effect Model (using lme4; Bates et al., 2015) as a unified method to present a statistical test of model fits. We regressed summed log likelihood fit on condition (unreliable→reliable = -.5, reliable→unreliable = .5), model (confederate-following as referent given our primary interest in this model), task half (first half = -.5, second half = .5), and all possible interactions. We included by-participant random slopes for model and task half. The three-way interaction was significant ($F(4, 508) = 731.166$, $p < .001$) in accordance with the pattern described above (see Table 2 for mean log likelihood by each condition and task half). For more details, see analysis code and de-identified dataset that are available on Open Science Framework: https://osf.io/ey4ut/?view_only=84412859abcb4189ae2e749c206d4685.

Table 1: Proportion of participants fit best by each model. No participants were best fit by the confederate-matching model.

	Random	Matching	Maximiz-ing	Following
reliable→unreliable: pre-switch	4.76	3.17	0	92.06
reliable→unreliable: post-switch	14.29	68.25	4.76	12.70
unreliable→reliable: pre-switch	19.70	78.79	1.52	0
unreliable→reliable: post-switch	0	15.15	4.55	80.30

Table 2: Mean log likelihood by experiment condition and task half (standard error in parentheses). The random model fits all blocks with the same value, -207.94.

	Matching	Maximizing	Con. Match.	Following
reliable→unreliable: pre-switch	-132.02 (3.97)	-187.77 (5.47)	-304.24 (1.89)	-72.94 (11.82)
reliable→unreliable: post-switch	-150.02 (7.85)	-217.12 (13.11)	-254.3 (2.18)	-303.60 (11.68)
unreliable→reliable: pre-switch	-166.28 (5.93)	-244.88 (9.84)	-253.91 (1.67)	-358.88 (6.26)
unreliable→reliable: post-switch	-110.10 (2.99)	-153.40 (6.02)	-308.99 (1.40)	-73.87 (8.64)

Weighting of Social Information

To further examine how participants were using each of these strategies, we used a mixture model that included both probability matching and confederate following. We utilized only these models because, as demonstrated in Table 1, the overwhelming majority of participants were best fit by either one or the other of these models. The decision to use the mixture model was intended to reduce the dimensionality of our hypothesis tests to these two most-common strategies. One key aspect of the mixture model was that the mixture itself was not constrained to be a perfect average of the given two models at hand (e.g., 50% probability matching, 50% confederate following). Instead, the fitting procedure involved finding the best-fitting weighted average of the two models (e.g., if a participant's choices were largely consistent with confederate following, but occasional choices were more consistent with probability matching, this might produce a final mixture with weights of 90% confederate following, 10% probability matching). Examining the specific weight participants attributed to the confederate-following model versus the probability-matching model thus provided critical insight into *how* children used the social and underlying reward cues.

The weighting calculation was implemented using the *optim* function in R. In the previous example the full predicted probability of one model was used to calculate the likelihood. When optimizing weights between two models, each model's prediction was multiplied by a weight w between 0 and 1, while the other model's prediction was multiplied by 1-w. Thus, the predictions of the two models were combined using one free parameter, and these combined predictions were used in the following equation where r is a single choice on a trial, c is the total set of choices (for a given participant), and M is the model being evaluated:

$$\ell(c|M_i, M_j) = \sum_{t=1}^{n} \log\left(p(r_t|M_i)w + p(r_t|M_j)(1-w)\right)$$

We used a linear mixed-effects model to regress weight attributed to the confederate-following model on condition (unreliable→reliable = -.5, reliable→unreliable = .5), task half (first half = -.5, second half = .5), age (mean-centered), and all possible interactions. We included a by-participant random slope for task half. There was a main effect of condition (b = 0.097, F(1, 125) = 12.568, p < .001; greater weight was attributed to the confederate-following model in the reliable→unreliable condition). This effect was qualified by a condition-by-half interaction, b = -1.20, F(1,125) = 485.318, p < .001 (Figure 2). Prior to the switch, participants who received reliable suggestions weighted the confederate-following model more heavily than participants who received unreliable suggestions (M reliable = .838 (SD = .241) vs. M unreliable = .140 (SD = .126), b = 0.698, t = 20.571). After a change in the confederate's behavior, participants who received unreliable suggestions (from a confederate who had previously provided reliable suggestions) attributed less weight to the confederate-following model compared to those who received reliable suggestions (from a confederate who had previously offered unreliable suggestions (M reliable = .756 (SD = .243) vs. M unreliable = .255 (SD = .247), b = -0.505, t = -13.088). Therefore, participants updated the weight they attributed to the confederate's suggestions versus the underlying reward distribution based on the confederate's *current* pattern of behavior, rather than continuing to be influenced by the confederate's previous pattern of behavior. No other effects or interactions were significant (ps > .1).

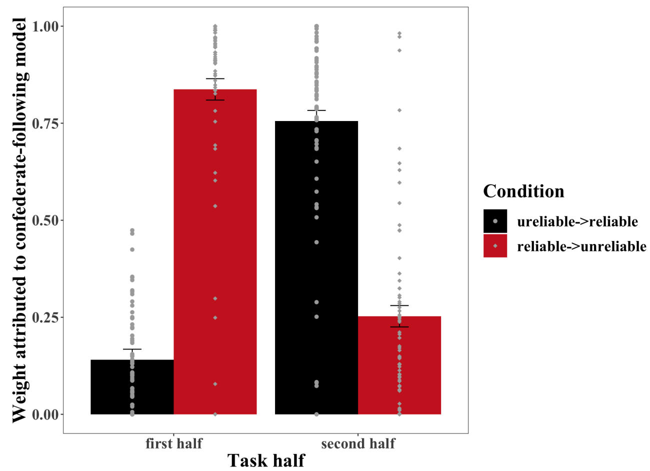

Figure 2. Weight attributed to the confederate-following model by condition and task half: Model predictions and participant-level data. Error bars represent standard error of the estimates. Points are individual participants' weights.

Likelihood of Following Suggestions

In order to better understand children's use of the suggestions, we simply examined the participant's likelihood to agree with the confederate's suggestions. We used a logistic regression model to regress whether the participant chose the same rock that was suggested by the confederate (0 = no, 1 = yes) on condition, age, task half, and all possible interactions. We included a by-participant random slope for task half and included a by-trial random intercept. The interaction between condition and age was significant (b = -0.126, $X^2(1)$ = 41.430, p < .001; odds ratio (OR) = 0.882; Figure 3). For participants in the reliable→unreliable condition, likelihood of agreeing with the confederate decreased with age (b = -0.058, p < .001) whereas in the unreliable→reliable condition, the likelihood of agreeing with the confederate increased with age (b = 0.068, p < .001). In sum, younger participants were more strongly influenced by the confederate's initial behavior.

In addition to the condition-by-age interaction, there was a main effect of condition (b = 0.373, $X^2(1)$ = 122.024, p < .001, OR = 1.452) and condition-by-half interaction (b = -5.805, $X^2(1)$ = 587.391, p < .001, OR = -.003), consistent with the patterns as described when assessing weight attributed to the confederate-following model. There was also a main effect of half (b = 0.237, $X^2(1)$ = 3.879, p = .049, OR = 1.267; greater likelihood of agreement in the second half of the experiment) and a marginal condition-by-age-by-half interaction (b = -0.245, $X^2(1)$ = 2.880, p = .090, OR = -0.783).

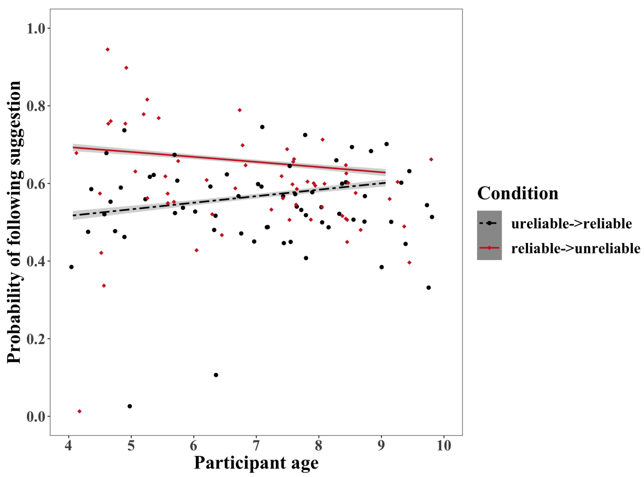

Figure 3. Likelihood of agreeing with the confederate's suggestion by age and condition: Model predictions and participant-level data. Note: Lines are point estimates from logistic mixed-effects models with the interaction between condition and participant age, and lower-order effects. Error bands represent standard error of the point estimates. Points are individual participants' proportion of choices that followed the confederate's suggestion.

Discussion

The aim of the present research was to investigate how children respond to an unexpected change in a social agent's behavior. Whether encountering a reliable confederate after a period of unreliability or vice versa, participants readily adjusted strategies to match the changing environment. However, younger participants continued to be influenced by the confederate's initial behavior as revealed by the difference in young children's likelihood of agreeing with the confederate's suggestions by condition (an effect that decreased with age). Therefore, younger participants appeared to be more influenced by the confederate's behavior prior to an unexpected change.

Developmental Differences in Using Social Information

In the present research children had to update their own behavior in response to changes in another's behavior. Participant age influenced the pattern of choice behavior. Younger children were especially sensitive to the initial behavior of a social agent and had relatively more difficulty updating their behavior. That younger children were more influenced by their initial observations is consistent with research showing that school-age children will more readily abandon a rule than preschool-age children (Sanders, 1971) and more recent research indicating that children (particularly younger children) do not shift behavioral predictions or choice strategies as readily as adults (Boseovski & Lee, 2006; Starling et al., 2018). However, the present findings are paradoxical in light of a body of research suggesting that younger children are particularly flexible in their ability to update beliefs in light of extant evidence (e.g., Gopnik et al., 2017). Additionally, children as young as 4 years of age learn as well as older children on the present task without social information (Plate et al., 2018). Children as young as 4 years of age are also able to distinguish between the reliability of confederates who differ in accuracy that is probabilistic (e.g., distinguish between an informant who is accurate 25% of the time and one who is accurate 100% of the time; Pasquini et al., 2007). The developmental differences observed here may represent a special case of interference in unpredictable environments, particularly when needing to track multiple sources of probabilistic information.

Limitations and Future Directions

The present research illustrated one type of uncertainty when considering social information, namely that social agents might change their behaviors unpredictably. However, social agents can be unpredictable in many other ways, including changing back and forth between behaviors frequently. Additionally, there may be clues that help learners anticipate changes in another's behavior. More research is needed to better understand the nuance in how children respond to various types of unpredictability in social information.

Another question for future research is which, if any, patterns of confederate behavior would result in the participant entirely ignoring the confederate. Presumably, in the current research, participants continued to monitor the confederate's choices as evidenced by their effective updating of their own choice strategies. Children may pay attention to both the amount, and type (e.g., whether the confederate was close to or far from the rewarded location), of error in the confederate's responses; therefore, future research may consider varying these features of the confederate's suggestions. Research on selective trust suggests that children are responsive to both accuracy and inaccuracy (Corriveau et al., 2009), and younger children are particularly sensitive to inaccuracy (Pasquini et al., 2007). Additionally, there is evidence that children are sensitive to an informant's error magnitude and use error magnitude as a cue for generalizing an informant's credibility (Einav & Robinson, 2011). In the present experiment, it is unclear if or how children's behavior would differ based on different error magnitudes from a confederate.

Conclusion

Overall, this research provides evidence regarding how children respond to changing social information. Children were sensitive to both changes from reliable to unreliable suggestions and vice versa. However, younger children were more strongly influenced by their initial interactions with the confederate, an effect that decreased with age. In sum, we highlight how children remain nimble to adjust to a dynamic social environment.

Acknowledgments

This research was supported by the National Institute of Mental Health through grant R01MH61285 to S.D.P. and in part by a core grant to the Waisman Center from the National Institute of Child Health and Human Development (U54 HD090256). R.C.P. was supported by a National Science Foundation Graduate Research Fellowship (DGE-1256259) and the Richard L. and Jeanette A. Hoffman Wisconsin Distinguished Graduate Fellowship. We thank the families who participated in this study and the research assistants who helped conduct the research.

Correspondence concerning this article should be addressed to Rista C. Plate, rplate@sas.upenn.edu

References

Bates D, Mächler M, Bolker B, Walker S (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software, 67*(1), 1–48.

Boseovski, J. J., & Lee, K. (2006). Children's use of frequency information for trait categorization and behavioral prediction. *Developmental psychology*, *42*(3), 500.

Bridgers, S., Buchsbaum, D., Seiver, E., Griffiths, T. L., & Gopnik, A. (2016). Children’s causal inferences from conflicting testimony and observations. *Developmental psychology*, *52*(1), 9.

Corriveau, K. H., Meints, K., & Harris, P. L. (2009). Early tracking of informant accuracy and inaccuracy. *British Journal of Developmental Psychology*, *27*(2), 331-342.

Einav, S., & Robinson, E. J. (2011). When being right is not enough: Four-year-olds distinguish knowledgeable informants from merely accurate informants. *Psychological science*, *22*(10), 1250-1253.

Gergely, G., Bekkering, H., & Király, I. (2002). Developmental psychology: Rational imitation in preverbal infants. *Nature*, *415*(6873), 755-755.

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.

Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Harvard University Press.

Kattner, F., Cochrane, A., & Green, C. S. (2017). Trial-dependent psychometric functions accounting for perceptual learning in 2-AFC discrimination tasks. *Journal of vision*, *17*(11), 3-3.

Klein, S. A. (2001). Measuring, estimating, and understanding the psychometric function: A commentary. *Perception & psychophysics, 63*(8), 1421-1455.

Lampinen, J. M., & Smith, V. L. (1995). The incredible (and sometimes incredulous) child witness: Child eyewitnesses' sensitivity to source credibility cues. *Journal of Applied Psychology*, *80*(5), 621.

Li, X., & Yow, W. Q. (2018). Willingness to revise own testimony: 3-and 4-year-olds’ selective trust in unexpected testimony from accurate and inaccurate informants. *Journal of experimental child psychology*, *173*, 1-15.

Meltzoff, A. N. (1988). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental Psychology*, *24*(4), 470.

Palmquist, C.M. & Jaswal, V.K. (2012). Preschoolers expect pointers (even ignorant ones) to be knowledgeable. *Psychological Science*, *23*(3), 230-231.

Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, *43*(5), 1216.

Plate, R. C., Fulvio, J. M., Shutts, K., Green, C. S., & Pollak, S. D. (2018). Probability Learning: Changes in Behavior Across Time and Development. *Child development*, *89*(1), 205-218.

Ronfard, S., & Lane, J. D. (2018). Preschoolers continually adjust their epistemic trust based on an informant's ongoing accuracy. Child development, 89(2), 414-429.

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.

Sanders, B. (1971). Factors affecting reversal and nonreversal shifts in rats and children. Journal of Comparative and Physiological Psychology, 74(2), 192.

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.

Sorce, J. F., Emde, R. N., Campos, J. J., & Klinnert, M. D. (1985). Maternal emotional signaling: Its effect on the visual cliff behavior of 1-year-olds. *Developmental Psychology*, *21*(1), 195-200.

Starling, S. J., Reeder, P. A., & Aslin, R. N. (2018). Probability learning in an uncertain world: How children adjust to changing contingencies. *Cognitive development*, *48*, 105-116.

Uleman, J. S., Adil Saribay, S., & Gonzalez, C. M. (2008). Spontaneous inferences, implicit impressions, and implicit theories. *Annu. Rev. Psychol.*, *59*, 329-360.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & psychophysics, 63*(8), 1293-1313.