

How children make choices when information sources diverge

By

Rista C Plate

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Psychology)

at the

UNIVERSITY OF WISCONSIN-MADISON

2019

Date of final oral examination: 12/12/2019

The dissertation is approved by the following members of the Final Oral Committee:

C. Shawn Green, Associate Professor, Psychology
Chuck Kalish, Professor Emeritus, Educational Psychology
Seth Pollak (Co-Chair), Professor, Psychology
Jenny Saffran, Professor, Psychology
Kristin Shutts (Co-Chair), Professor, Psychology

Table of Contents

Acknowledgements	iii
Abstract	iv
Introduction	1
Children’s Own Observations Inform Knowledge	2
Social Input Informs Children’s Knowledge	3
Children Attend to Both their Observations and Social Input	4
Comparing Social and Nonsocial Information	5
Present research	6
Experiment 1	7
Method	8
Results	13
Discussion	16
Experiment 2	17
Method	18
Results	19
Discussion	21
Experiment 3	22
Method	23
Results	24
Discussion	26
General Discussion	27
Developmental Differences in Resolving Conflicts between Information Sources	27
Informative Value of Social Information Matters	28
Unpacking Strategy Use	28
Limitations and Future Directions	30
Conclusion	32
References	33
Tables & Figures	42

Acknowledgements

I would like to acknowledge the many who contributed to this dissertation. Thank you to my advisors, Seth Pollak and Kristin Shutts. Seth challenges me to think big and ask questions from different perspectives. Kristin challenges me to continually evolve my thinking and approach a problem from different angles to fully answer a question. Together they have taught me how to think deeply and creatively and have invaluable guided my development as a scientist. I would like to thank the members of my committee, Shawn Green, Chuck Kalish, and Jenny Saffran for their feedback and the larger roles they have played in my academic career. A special thanks to Jenny, one of my first mentors, who continues to inspire my research and encourage my ideas.

Thank you to my wonderful and supportive lab members in the Child Emotion Lab and the Social Kids Lab. And I could not have completed this project without the many research assistants who ran participants and also motivated me with their fresh ideas and insights. I would also like to thank the children and families who contributed their valuable time to participate.

Thank you to my incredible academic and emotional support system, particularly Sarah Brown, Emily Dix, Brian Leitzke, Chelsea Mitamura, Katharine Scott, Adrienne Wood, and Tina Woodard. Thank you to my family: Nan, Bill, Brian, Carolina, Tom, Paco, and Tomás for your support, encouragement, and inspiration. My parents in particular have been my well of support for every up, down, and in between. Thank you to my husband, Sam. Your support is my foundation and your patience is my strength and I am so grateful for your partnership. Finally, thank you to Will, who reminds me every day of the wonders of child development and how many exciting questions are left to explore!

Abstract

When making decisions, children often have access to multiple sources of information and must reconcile instances in which sources disagree. Three experiments investigated how children weight divergent information in a dynamic learning context. Participants (4-9-years-old) searched for rewards and saw cues that varied in how well the cue located the reward. Participants were told that the cues were from another player (Experiment 1, $N = 102$) or because of a computer glitch (Experiment 2, $N = 89$). Children did not optimally discount social cues that diverged from the rewarded locations, and younger children weighted social cues more heavily than older children when the cues were only sometimes helpful (an effect that was unique to the social, as compared to the “computer glitch,” cue). In Experiment 3 ($N = 129$), participants received social cues that unexpectedly changed in level of disagreement with the underlying reward distribution over time. All children updated how heavily they weighted the cues after the change. However, younger children were more heavily influenced by their initial experience with the cues. Taken together, these experiments provide evidence to suggest that children are skilled in detecting the usefulness of information sources. However, younger children may have more difficulty disengaging from social information in uncertain learning environments.

Keywords: statistical learning, social learning, development

How children make choices when information sources diverge

In making choices, children may need to consider multiple sources of information in order to determine the best action to take. There will be occasions when sources of information agree (e.g., another person's suggestion is consistent with your own observations) and there will be occasions when sources of information are in conflict. One challenge then, is how to consider multiple information sources when they diverge. A successful learner must be able to track the utility of divergent sources and use that information to inform appropriate inferences and choices. In the present paper, we ask how children weight conflicting sources of information in the service of making a decision. Children ages 4 to 9 years searched for rewards and received suggestions about where rewards were located. We varied the reliability of the suggestions in order to investigate the strategies children use when social information conflicts, to varying degrees, with their own observation of rewarded locations. In doing so, we are able to interrogate mechanisms of children's social learning and ask how children use information sources differently depending on features of the learning environment.

First, we review evidence that children acquire knowledge through their own observations. Next, we highlight the ways in which children use information from others and are discerning about from whom they learn. Finally, we discuss literature that has examined how children's own observations and social information interact, which serves as the foundation for the present research. The present research expands previous work by exploring the strategies children use to children make choices when information sources diverge, in particular by investigating how contextual factors—i.e., the reliability of the environment—and individual

factors—i.e., participant age—interact in a dynamic learning scenario where children have to track and weight multiple sources of information over time.

Children’s Own Observations Inform Knowledge

Children possess an arsenal of learning skills to make sense of the many complex environments they encounter. Children’s exploration of their environments leads to the discovery of new information and facilitates knowledge acquisition (Stahl & Feigenson, 2015). Although children gain knowledge through their interactions with their environments in many ways, one particularly powerful domain-general mechanism is statistical learning (Aslin, 2017). Beginning in infancy, children are sensitive to patterns of associations in their environments. For example, infants attend to the statistical regularity of sounds and recognize that certain sounds often occur together (indicating a word) and that other sounds occur less frequently together (indicating a likely word boundary; Saffran, Aslin, & Newport, 1996). Infants also rely on statistical learning to detect visual patterns (Kirkham, Slemmer, & Johnson, 2002) and use distributional information to predict outcomes, such as the likely color of a ball that has been randomly selected from a gumball-like machine (Xu & Garcia, 2008).

As children age, their ability to use statistical information influences the choices they make and the actions they take on the environment. Toddlers and preschool-age children use statistical patterns to make inferences (Denison, Bonawitz, Gopnik, & Griffiths, 2013) and choices (Waismeyer, Meltzoff, & Gopnik, 2015). Further, children’s choice strategies are not static; children update the strategies they use to obtain rewards continuously as more information about a statistical distribution or pattern is acquired (Plate, Fulvio, Shutts, Green, & Pollak, 2018). However, children’s ability to select *optimal* strategies (i.e., strategies that will garner the most rewards) depends on the statistics of the environment; when the underlying pattern of

statistical information is less clear, children spend more time using exploratory strategies at the cost of garnering rewards (Plate et al., 2018; Tummeltshammer & Kirkham, 2013).

Social Input Informs Children's Knowledge

In addition to learning from patterns in their environment, children also learn from information provided by other people (Gergely & Csibra, 2006; Tomasello, 1999). Across myriad situations, children seek and trust information provided by others (Callanan & Oakes, 1992). 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, Emde, Campos, & Klinnert, 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; Harris & Koenig, 2006; Jaswal & Kondrad, 2016; Sobel & Kushnir, 2013). Learning from other people is a skill that is universally recognized as essential for a child's success (Kline, 2015).

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, Bekkering, & Kiraly, 2002) and children favor information provided by those who appear accurate, reliable, knowledgeable, intelligent, and have access to relevant information (e.g., Buchsbaum, Gopnik, Griffiths, & Shafto, 2011; Corriveau, Meints, & Harris, 2009; Einav & Robinson, 2011; Kidd, Palmeri, & Aslin, 2013; Koenig & Harris, 2005; 2007; Kushnir, Wellman, & Gelman, 2007; Pasquini, Corriveau, Koenig, & Harris, 2007). Children's selectivity is demonstrated in the classic

selective trust paradigm in which children see one informant who is consistently accurate (e.g., correctly labels a ball ‘ball’) and another informant who is consistently inaccurate (e.g., incorrectly labels a ball ‘shoe’). After observing informants label the known objects, participants ask the former rather than the latter informant for the label of a novel object (e.g., Koenig & Harris, 2005).

Children Attend to *Both* their Observations and Social Input

Some research has pitted social input against children’s own knowledge (Bridgers, Buchsbaum, Seiver, Griffiths, & Gopnik, 2016; Buchsbaum et al., 2012; Chan & Tardif, 2013; Lane, Harris, Gelman, & Wellman, 2014; Lucas et al., 2017) 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 (Bridgers et al., 2016; Ronfard & Lane, 2018; Ronfard, Lane, Wang, & Harris, 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 (e.g., Butterfield, Jenkins, Sobel & Schwertfeger, 2009; Hermes, Behne, Bich, Thielert, & Rakoczy, 2018; Hermes, Rakoczy, & Behne, 2019; Shafto, Eaves, Navarro & Perfors, 2012; Sobel, Buchanan, Butterfield & Jenkins, 2010; see Sobel & Kushnir, 2013 for a review).

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 (Sobel & Kushnir, 2013). Therefore, children are not expected to universally favor one information source (e.g., social input) over another (e.g., direct observations), but rather they are expected to attend to extant evidence to weight information sources in any given learning scenario. Yet, we understand little about the mechanisms underlying *how* children weight social information when it conflicts with their own observations. There is some evidence to suggest that children's use of social input depends on the strength of evidence they observe themselves (Bridgers et al., 2016), and that the ability to simultaneously integrate cues about a social agent's competence with task features (e.g., agent's performance on an easy vs. difficult task) might improve across early development (Leonard, Bennett-Pierre, & Gweon, 2019).

Comparing Social and Nonsocial Information

A related question to the argument above is whether *social* information is considered differently from other types of input when children are using multiple information sources to make inferences about their environments. There is some support for the idea that learners treat social information similarly to other (nonsocial) sources of input. For example, in one neuroimaging study, adult participants showed evidence of social prediction errors in similar regions in which reward prediction errors appear (Behrens, Hung, Woolrich, & Rushworth, 2009). Although the social and nonsocial information were initially encoded in different brain areas, both were integrated in the ventromedial prefrontal cortex. Others have similarly found that learners can track the associations between social cues and social rewards just as they track

the associations between cues that signal primary rewards (e.g., food; Jones et al., 2011). This research suggests that similar mechanisms underlie social and nonsocial forms of learning.

Yet, there is evidence that social learning cues are privileged under uncertainty. For infants, the reliability of a social cue influences whether that cue is treated differently from a nonsocial cue: In one study (Tummeltshammer, Wu, Sobel, & Kirkham, 2014), infants were sensitive to the reliability of both social (directed gaze) or nonsocial (arrow) cues to a rewarded location on a screen. However, when the cue pointed to a location that was never rewarded, infants only searched this new location in the social condition and not in the nonsocial condition. This finding is consistent with research suggesting that social cues can direct an infant's attention in order to make predictions, whereas nonsocial cues do not achieve the same aim (Wu & Kirkham, 2010; Wu, Tummeltshammer, Gliga, & Kirkham, 2014). Additionally, it may be that certain socially relevant learning mechanisms must be developed prior to observing differences in the use of social and nonsocial information (Sobel & Munro, 2009), again suggesting that the age of the learner might be key for discerning under what conditions social information is uniquely considered.

Present Research

In the first experiment, we seek to better understand how children weight information sources when faced with social information (“suggestions”) that diverges from the learner's own observations. We examine whether children's choices are contingent upon the reliability of the suggestions as well as children's age. In Experiment 2 we test whether children perform similarly when suggestions are not social in nature. Finally, in Experiment 3, we investigate how sensitive children are to the volatility of social information.

Participants in all experiments 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 experiments, 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 (Experiments 1 and 3) or that the cue was merely a computer glitch (Experiment 2). In reality, the cued locations were predetermined in all experiments.

Experiment 1

Experiment 1 included three between-subject conditions that varied in how often a confederate's suggestion indicated the rewarded location. We refer to these conditions in terms of "conflict," which represents the relation between the cued location and the rewarded location. The conditions included: 1) low conflict: the confederate's suggestions frequently indicated the rewarded location, 2) high conflict: the confederate's suggestions infrequently indicated the rewarded location, and 3) moderate conflict: the confederate's suggestions indicated the rewarded location on half of the trials. We first measured the types of strategies children used,

and then we calculated the weight attributed to different information sources. We reasoned that both the high- and low-conflict scenarios present little uncertainty: either the suggestions are good indicators of the underlying reward, in which case one should heavily weight those suggestions, or they are poor indicators of the reward, in which case one should ignore them. Therefore, we expected that in *both* high- and low-conflict scenarios children would use choice strategies that optimize rewards. Specifically, in the high-conflict condition we predicted that children would learn from the underlying reward distribution and use the optimal, maximizing, strategy to garner rewards. In the low-conflict condition, we predicted that children would primarily rely on the confederate's suggestions. In the moderate-conflict condition, we expected that children would have difficulty weighting information sources when the confederate's suggestions were sometimes (but not always) in conflict with the underlying rewards because there is a greater degree of uncertainty. We also asked whether these differences would be most pronounced in younger children; to test this notion, we tested participants from preschool through the early elementary school years because children as young as 4 years old can use probabilistic evidence to distinguish between a sometimes helpful and never helpful informant (Pasquini et al., 2007).

Method

The experimental paradigms, de-identified data, and analysis script for Experiments 1, 2, and 3 are available on Open Science Framework:

https://osf.io/ey4ut/?view_only=84412859abcb4189ae2e749c206d4685

Participants. The sample included 102 children ages 4-9 years (9 Hispanic or Latino, 93 not Hispanic or Latino; 6 African American, 7 Asian American, 8 Multiracial, 77 White, 4 chose “other” or did not report race or ethnicity). Children were recruited from a large Midwestern city.

Thirty-six children were in the low-conflict condition (23 male, 13 female, $M_{\text{age}} = 6.887$, $SD_{\text{age}} = 1.736$, Table 1), 34 children were in the high-conflict condition (19 male, 15 female, $M_{\text{age}} = 6.887$, $SD_{\text{age}} = 1.723$), and 32 children were in the moderate-conflict condition (17 male, 15 female, $M_{\text{age}} = 7.053$, $SD_{\text{age}} = 1.866$). There were no differences in participant age ($p = .908$) or gender ($p = .642$) by condition. A post-hoc sensitivity analysis (using the ‘pwr’ package in R) indicated that we had 85% power to detect our effect of interest (i.e., differential weighting of strategies by condition). An additional 16 participants were excluded because they quit the experiment before all trials were presented ($N = 13$) or because of experimenter error ($N = 3$). The Institutional Review Board approved the research. Parents of child participants gave informed consent; children gave verbal assent. Parents received \$20 for their time and children chose a prize after participating.

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 (adapted from Plate et al., 2018). The participant met a 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.

Once in the testing room, the participant saw a shelving unit with prizes and were told that they could win different types of prizes depending on how many coins they found. A prize from the bottom shelf (that contained prizes such as pencils and erasers) could be chosen if the

participant found a small number of coins; a prize from the middle shelf (that contained prizes such as a bouncy ball or small notebook) could be chosen if the participant found a medium number of coins; a prize from the top shelf (that contained prizes such as an art kit or Legos) could be chosen if the participant found many coins. In reality, all participants were allowed to choose a prize from any shelf after participating.

Practice phase. The participant sat in front of a 19-inch computer monitor next to the experimenter and began the practice phase. The practice phase included three example trials with coins hidden behind trees instead of rocks. A small cartoon elf appeared in the center of the screen following a fixation cross. Participants were told that whenever the elf appeared, a new coin had been hidden. Then, a circle of 8 trees appeared on the screen. The experimenter explained that a coin was hidden behind one of the trees. On the next screen, a hand pointing downward to one of the trees appeared. The experimenter described that the hand indicated the confederate's suggestion and that participants would always see the suggestion prior to making a choice themselves. The participant could choose to look for the coin under any one of the trees and saw an example of choosing a tree not suggested by the confederate. The next screen revealed that the confederate had been pointing to the correct location. Two additional example trials followed, with concurrent instructions from the experimenter and the experimenter advancing the screens. In these two additional trials, participants saw the following choice patterns and outcomes: 1) choose the same location as the confederate but the coin is in a different location and 2) choose a different location from the confederate and find the coin. Thus, participants saw multiple possible options for how to respond.

Next, the experimenter showed the participant a screen with a row of rocks and told the participant that they would look for coins under this row of rocks and could select a rock by

clicking with the computer mouse. Finally, the participant saw a screen with a big bag of coins and were encouraged to do their best to find a lot of coins. The experimenter asked participants to explain how to play the game using their own words. Prior to starting the testing trials, the experimenter ensured that participants understood: 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. Following 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.

There were two blocks of 100 test trials, separated by a break. At the start of each test trial, eight rocks appeared on the screen with equal spacing along a horizontal line (see Figure 1 A-D for experimental task). 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.

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 match to the location probabilities across trial blocks (i.e., in each 100-trial block, 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 (i.e., the underlying distribution of the confederate's choices, Figure 2) varied across three conditions; participants were randomly assigned to one of these conditions. In the low-conflict condition, the confederate suggested the correct rock 90% of the time; in the high-conflict condition, the confederate suggested the correct rock 10% of the time; in the moderate-conflict condition, the confederate suggested the correct rock 50% of the time. Participants were not told about these probabilities. All other task specifications remained constant across conditions.

Models of strategy use. Because we were interested in how children weighted strategies, we first needed to characterize the strategies participants used. In past research, participants who completed this computer task without suggestions from a confederate transitioned over time from a sub-optimal strategy (probability matching) to an optimal strategy (maximizing; Plate et al., 2018). Here, we were interested in participant strategy use in the presence of the confederate's

suggestions. Therefore, we assessed 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 *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. (For full model details, see Supplemental Materials.)

Results

Children's use of strategies. To address what approaches children across the different conditions used, we measured participants' use of the strategies described above. To do this, we determined the log likelihood of each participant's set of choices across the experiment by summing the log probabilities of the choices given the probability structure across rock locations and confederate suggestions considering each participant's past trial experience and the model being evaluated. We regressed the log likelihood on condition (moderate-conflict condition set as the referent), model (following as referent), age (mean-centered), and all possible interactions with a by-participant random intercept.

There was an interaction between model (probability matching, maximizing, following, confederate matching, and random choice) and condition (low conflict, moderate conflict, high conflict; $F(8,396) = 117.557, p < .001$; see Supplemental Materials for all pairwise comparisons). The majority of participants in the low-conflict condition were best fit by the following model (97%), the majority of participants in the high-conflict condition were best fit by the probability-matching model (85%), and the majority participants in the moderate-conflict condition were best fit by *both* the probability-matching (44%) and following (50%) models.

Weighting of social information. To better understand how participants were using each of these strategies, we examined a mixture model that included both probability matching and following because these were overwhelmingly the strategies used by participants. 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% 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 following, but occasional choices were more consistent with probability matching, this might produce a final mixture with weights of 90% following, 10% probability matching). Examining the specific weight participants attributed to the following model versus the probability-matching model thus provided critical insight into *how* children used the social and underlying reward cues. We regressed the weight attributed to the following model on condition (moderate conflict as referent), age (mean-centered), and their interaction. When comparing the weight given to the following model against that given to the probability-matching model within the mixture model that considered both, there were effects of condition ($F(2, 96) = 181.981, p < .001, \Delta R^2 = 0.776$) and age ($F(1, 96) = 9.179, p = .003, \Delta R^2 = 0.020$), which were qualified by a condition-by-age interaction ($F(2, 96) = 2.994, p = .055, \Delta R^2 =$

= 0.013, Figure 3). Specifically, participant age mattered in the moderate-conflict condition: younger participants attributed greater weight to the confederate's suggestions, with weight decreasing with age (simple slopes for relation between participant age and weight attributed to the following model: high conflict: $b = -0.009$, $p = .629$; moderate conflict: $b = -0.051$, $p = .003$; low conflict: $b = 0.005$, $p = .759$). Therefore, younger participants had more difficulty disengaging from the confederate's suggestions when their suggestions only sometimes conflicted with the underlying reward distribution. In contrast, participants in the high-conflict condition attributed little weight to the confederate's suggestions, regardless of age. As would be expected based on strategy use, participants in the low-conflict condition attributed great weight to the confederate's suggestions, again regardless of age.

Next, we compared the weight attributed to the following model against the percent of time the confederate suggested the correct rock location (i.e., 10% in the high-conflict condition, 50% in the moderate-conflict condition, 90% in the low-conflict condition). Using t-tests for each condition, there were no differences between the weight participants attributed to the confederate and the percent of time the confederate suggested the correct location (high conflict: $t(33) = -0.067$, $p = .947$; moderate conflict: $t(31) = 0.583$, $p = .564$; low conflict: $t(35) = 0.102$, $p = .920$).

Children's timing to stabilize strategies. Finally, because participants transition strategies over time (Plate et al., 2018), we used a time-evolving model to measure how many trials passed until participants reached an asymptote for a strategy (see Supplemental Materials for time-evolving model details). Timing to an asymptote could indicate whether participants are exploring multiple approaches before establishing a consistent strategy or whether they are maintaining a strategy across time. Participants were faster to stabilize on a strategy in the low-

conflict condition (as compared to both the high and moderate-conflict conditions; omnibus: $F(2, 96) = 7.286, p = .001, \Delta R^2 = 0.130$; high conflict vs. low conflict: $b = 24.776, p = .003$; moderate conflict vs. low conflict: $b = 23.878, p = .005$); however, there was no difference in time to stabilize on a strategy between the high and moderate-conflict conditions ($b = 0.898, p = .992$). Across all three conditions, participants stabilized on a strategy approximately in the first quarter of the experiment, suggesting that children readily adopted and maintained the strategies described above (trial number when asymptote reached: M high conflict = 54.273, SD high conflict = 33.377; M low conflict = 29.291, SD low conflict = 25.490; M moderate conflict = 53.123, SD moderate conflict = 32.452).

Discussion

Children were sensitive to the extent to which the confederate's suggestions aligned with the underlying reward distribution and adjusted their strategies based on this information. However, use of strategies when the underlying rewards and social suggestions were in moderate conflict was influenced by participant age. In the moderate-conflict condition, participants were split between probability-matching (choosing each location at the rate it was rewarded) and following (choosing the suggested location) strategies. Additionally, children attributed less weight to the confederate's suggestions with age. A particularly surprising finding was that participants in the high-conflict condition were not best fit by the maximizing model, despite stabilizing on a strategy early in the experiment and attributing little weight to the following model (as compared to the probability matching model). The only condition in which children demonstrated optimal choice behavior was when receiving suggestions that were in low conflict with the reward outcomes. Taken together, the results from Experiment 1 show that children monitored how reliably the suggestions indicated the underlying rewards, but that the presence of

suggestions that conflicted with the rewarded locations interfered with children's ability to enact optimal choice behavior.

To what extent was the performance of children in Experiment 1 influenced by the social nature of the presented cue (i.e., a suggestion from another person)? One account of participants' performance is that they were operating under an assumption that social agents are generally helpful. This may have led them to follow the confederate's suggestions more in the moderate-conflict condition and rely less on the underlying rewards. Another possibility, however, is that children's performance in Experiment 1 is the result of having two divergent cues (the confederate's suggestion and the rewarded location) to track. Tracking multiple cues may have led them to have difficulty discerning the best strategy more generally. Therefore, we conducted Experiment 2 to better understand how children consider cues that are not social.

Experiment 2

In Experiment 2, we replicated the method of Experiment 1, with the exception that we provided a cue that was not associated with a social agent. Our aim was to observe whether children would behave similarly when provided with nonsocial cues that varied in the degree to which they conflicted with the underlying reward distribution. Our primary prediction was that fewer participants would continue to follow the cue in the moderate-conflict condition of Experiment 2 (compared to Experiment 1). Such a finding would be consistent with research showing that infants use social, but not nonsocial, cues in uncertain learning situations (Tummeltshammer et al., 2014). We were additionally interested in whether children would continue to probability match, instead of maximize, in the high-conflict condition. It could be that children would more readily disregard a nonsocial cue, which might facilitate more time to transition to maximizing.

Method

Participants. The sample included 89 children ages 4-9 years (5 Hispanic or Latino, 80 not Hispanic or Latino; 1 African American, 3 Asian American, 15 Multiracial, 66 White, 4 did not report race or ethnicity). Children were recruited from the same community as Experiment 1; none of the children had participated in Experiment 1. Thirty-one children received low-conflict suggestions (19 male, 12 female, $M_{\text{age}} = 7.118$, $SD_{\text{age}} = 1.779$, Table 1), 29 children received high-conflict suggestions (17 male, 12 female, $M_{\text{age}} = 6.898$, $SD_{\text{age}} = 1.671$), and 29 children received moderate-conflict suggestions (16 male, 13 female, $M_{\text{age}} = 6.986$, $SD_{\text{age}} = 1.577$). There were no differences in child age ($p = .877$) or gender ($p = .891$) by condition. An additional 17 participants were excluded because they quit the experiment before all trials were presented. The Institutional Review Board approved the research. Parents of child participants gave informed consent; children gave verbal assent. Children chose a prize after participating.

Procedure. The procedure was identical to Experiment 1 with the following exceptions. The cue signaling the suggested rock was an image of neutral-colored pixels (Figure 1, E) that were arranged to approximate the shape and size of the pointing hand cue from Experiment 1. Second, we revised the instructions of the task to remove any suggestion that the cue was social. Specifically, we told participants, “You might notice that sometimes you can see these funny shapes. They’re from the computer, and I’m not sure why they’re showing up.” The reasoning behind the change in the instructions was to emphasize the nonsocial nature of the cue. There are several difficulties in designing paradigms to contrast social and nonsocial information, particularly in statistical learning tasks. Perhaps most difficult is that there is a tendency for individuals to view patterned behavior as social in nature (Heider & Simmel, 1944). Further, children have increasing knowledge about, and interactions with, technology that might bias

them to see a computer cue as having some social characteristics. Children in kindergarten and second grade do not readily differentiate between teachers and information from the internet when seeking knowledge (Wang, Tong, & Danovitch, 2019). In order to make the present paradigm convincingly nonsocial, we had to devise an explanation for the cue that discourages social attributions.

Results

Children's use of strategies. We assessed the same five strategies as in Experiment 1. The majority of participants in the low-conflict condition were best fit by the following model (81%). The majority of participants in the high-conflict condition were best fit by the probability-matching model (83%). Contrary to Experiment 1, the majority of participants in the moderate-conflict condition were also best fit by the probability-matching model (69%; 21% fit by the following model; omnibus interaction between model and condition: $F(8,344) = 72.252, p < .001$, pairwise comparisons in Supplemental Materials).

Weighting of nonsocial information. Examining the weight attributed to the following model, there was an effect of condition such that participants attributed more weight to the cue in the low-conflict condition and less weight to the cue in the high-conflict condition compared to the moderate-conflict condition (omnibus: $F(2, 83) = 78.293, p < .001, \Delta R^2 = 0.639$); high vs. moderate: $b = -0.273, p < .001$; high vs. low: $b = -0.734, p < .001$; low vs. moderate: $b = -0.460, p < .001$). There was no effect of age and no age-by-condition interaction ($ps > .1$).

Next, we compared the weight attributed to the following model against the percent of time the cue suggested the correct rock location (i.e., 10% in the high-conflict condition, 50% in the moderate-conflict condition, 90% in the low-conflict condition). Using t-tests for each condition, participants attributed less weight to the cue compared to the percent of time the cue

indicated the correct location (effect is marginal in the low-conflict condition; high conflict: $t(28) = -3.013, p = .005$; moderate conflict: $t(28) = -3.170, p = .004$; low conflict: $t(30) = -1.951, p = .061$).

Children's timing to stabilize strategies. There were no differences in timing to stabilize on a strategy by condition or age ($ps > .3$).

Comparing Experiments 1 and 2. To better understand the differences in how the social and nonsocial cues influenced participant behavior on the task, we directly compared Experiments 1 and 2. There were no differences in participant age ($p = .795$) or gender ($p = 1$) by experiment.

Strategy use. We compared the proportion of participants using the following model in each condition (low conflict, high conflict, moderate conflict) across experiments (social, nonsocial). Exactly 0 participants used the following strategy in the high-conflict condition in both experiments. However, a greater proportion of participants used the following strategy in the moderate-conflict condition when the cue was social, compared to nonsocial ($\chi^2(1) = 4.468, p = .034$). The same pattern was marginally significant in the low-conflict condition ($\chi^2(1) = 3.281, p = .070$).

Weighting of information. To compare the weighting of strategies, we regressed the weight of the following model (within the mixture model including following and probability matching) on age (mean-centered), condition (moderate conflict as referent), experiment (nonsocial = -.5, social = .5), and all interactions. A number of effects emerged. In addition to main effects of age ($b = -0.046, F(1, 179) = 8.774, p = .003$) and condition (omnibus: $F(2, 179) = 232.446, p < .001$; high conflict vs. moderate conflict: $b = -0.360, p < .001$; high conflict vs. low conflict: $b = -0.771, p < .001$; low conflict vs. moderate conflict: $b = -0.411, p < .001$) that were

consistent with Experiment 1, there was a main effect of experiment, suggesting that more weight was attributed to the confederate's suggestions as compared to the nonsocial cue ($b = .196$, $F(1, 179) = 14.191$, $p < .001$). The main effects were qualified by a condition-by-age interaction suggesting that participants decrease the weight they attribute to the suggestions in the moderate-conflict condition with age, but age does not influence behavior in the high or low-conflict conditions (omnibus: $F(2, 179) = 4.489$, $p = .013$; simple slopes: high conflict: $b = -0.008$, $p = .615$; moderate conflict: $b = -0.046$, $p = .003$; low conflict: $b = -0.017$, $p = .239$). Finally, though not statistically significant, there was a marginal condition-by-experiment interaction ($F(2, 179) = 2.474$, $p = .087$), reflecting the pattern such that there are greater differences in weight attributed to social versus nonsocial suggestions in the moderate-conflict condition.

Timing to stabilize strategies. There was a main effect of condition such that children were faster to stabilize on a strategy in the low-conflict condition compared to the high- and moderate-conflict conditions (omnibus: $F(2, 179) = 6.008$, $p = .003$; high vs. low: $b = 16.54$, $p = .011$; moderate vs. low: $b = 18.320$, $p = .005$; no difference between the high and moderate conditions: $b = -1.779$, $p = .949$).

Discussion

Children treated the nonsocial cue similarly to the social confederate in that they were sensitive to degree of conflict between the cue and the underlying reward distribution and adjusted their behavior accordingly. This was especially true when the pattern was clear, as in both the low- and high-conflict conditions. When the pattern was less clear—as in the moderate-conflict condition—however, some differences emerged. Children attributed more weight to the suggestions when they were social (vs. nonsocial) in nature and used a strategy that relied on this

information (i.e., they followed the confederate's suggestions). When the suggestions were not tied to a social agent, children primarily focused on the underlying rewards. Therefore, the distinction between the social and nonsocial cues was contingent on the uncertainty of the environment, a finding that is consistent with previous research with infants (Tummeltshammer et al., 2014).

Experiment 3

When the conflict between social and reward information was variable, as in the moderate-conflict condition of Experiment 1, younger children attributed more weight to the social information. Yet, there may be more than one way in which a social information presents uncertainty, for example social information can be volatile. In Experiment 3 we tested how children weighted strategies when suggestions changed from being reliable (low conflict) with the underlying rewards to being unreliable (high conflict) with the underlying rewards or vice versa. In general, we expected that children would be able to flexibly update their behavior in this changing social environment, particularly given how quickly they stabilized on a strategy in Experiment 1. We were interested in whether age might influence flexibility. In Experiment 1, younger children attributed more weight to the confederate than older children when the confederate's suggestions only sometimes indicated the correct reward location. We reasoned that younger children may also attribute more weight to a confederate who is reliable early in the experiment even if that reliability changes over time. We predicted that younger children would attribute more weight to the confederate's suggestions if the confederate had previously reliably indicated the rewarded location, even if that confederate currently provided suggestions that were in conflict with the underlying rewards. Such a difference may not be present 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. Experiment 2 suggests that children will monitor information sources, even when they have no reason to expect that they will be useful. Therefore, we predicted that, even if participants did not use the confederate in the first half of the experiment, they would be sensitive to the change in conflict and update their choices accordingly.

Method

Participants for Experiment 3 included 129 children ages 4-9 (12 Hispanic or Latino, 117 not Hispanic or Latino; 10 African American, 4 Asian American, 9 Multiracial, 103 White, 3 chose “other” or did not report race) who did not participate in Experiments 1 or 2. Nine additional participants were excluded for not completing the task ($N = 6$) or experimenter error ($N = 3$).

The design, procedure, and data analysis approach were the same as Experiment 1, with the following exceptions: First, in Experiment 3, the conditions included low→high conflict (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, $M_{\text{age}} = 6.938$, $SD_{\text{age}} = 1.691$, Table 1) and high→low conflict (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, $M_{\text{age}} = 7.068$, $SD_{\text{age}} = 1.630$; 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 participant’s attention to this change. As in Experiments 1 and 2, we examined participant strategy use and weight attributed to confederate’s suggestions. Because the confederate changed in accuracy halfway through the task, there were fewer trials (100 instead of 200), therefore

impairing our ability to assess time to stabilize on a strategy (see Supplemental Materials for full explanation).

Results

Children's use of strategies. We assessed the same five strategies as in Experiment 1. We regressed the log likelihood fit on condition (high→low conflict = -.5, low→high conflict = .4), model (following as referent), 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 ($F(4, 508) = 731.166, p < .001$). Prior to the switch in confederate reliability, the majority of participants who received low-conflict suggestions were best fit by the following model (92%), participants who received high-conflict suggestions were best fit by the probability-matching model (79%). Post-switch, participants changed their strategy use. The majority of participants who initially received low-conflict suggestions, but now received high-conflict suggestions, were best fit by the probability-matching model (68%). Participants who initially received high-conflict suggestions, but now received low-conflict suggestions were best fit by the following model (80%). Regardless of whether participants received low-conflict 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 following model ($p = .094$; similarly no difference in proportion of participants receiving high-conflict suggestions fit by the probability matching model in the first or second half of the experiment, $p = .247$).

Weighting of social information with an unexpected change. We next examined the weight attributed to the following model and tested whether weight was influenced by initial experience (i.e., receiving low-conflict suggestions initially or receiving high-conflict suggestions initially) and age. We regressed weight attributed to the following model on

condition (high→low conflict = -.5, low→high conflict = .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 following model in the low→high condition). This effect was qualified by a condition-by-half interaction, $b = -1.20$, $F(1,125) = 485.318$, $p < .001$ (Figure 4). Prior to the switch, participants who received low-conflict suggestions weighted the following model more heavily than participants who received high-conflict suggestions (M low conflict = .838 (SD = .241) vs. M high conflict = .140 (SD = .126), $b = 0.698$, $t = 20.571$). Following a change in the confederate's behavior, participants who received high-conflict suggestions (from a confederate who had previously provided low-conflict suggestions) attributed less weight to the following model compared to those who received low-conflict suggestions (from a confederate who had previously diverged from the underlying reward distribution, M low conflict = .756 (SD = .243) vs. M high conflict = .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* accuracy, rather than continuing to be influenced by the confederate's previous record of accuracy. No other effects or interactions were significant ($ps > .1$).

Likelihood of following the confederate's suggestion differs by age. The interaction between condition and age in the above model assessing weight attributed to the following model was not significant ($p = .186$). However, we wondered whether age differences might emerge if 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 5). For participants in the low→high condition, likelihood of agreeing with the confederate decreased with age ($b = -0.058$, $p < .001$) whereas in the high→low 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), following the same patterns as described when assessing weight attributed to the 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).

Discussion

Whether encountering a low-conflict learning environment following a period of high conflict 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 decreases with age). Therefore, younger participants again appeared to be more influenced by the uncertainty—as was also seen in the moderate confederate condition of Experiment 1—and had difficulty adapting to the changing environment as compared to older children.

General Discussion

The aim of the present research was to investigate how children weight multiple, divergent, sources of information when they vary in level of conflict. By creating a task that included statistical information regarding the underlying reward distribution and suggestions from a confederate, we were able to measure what strategies children used, how children weighted reward and social information, how long it took to stabilize on a strategy, and developmental changes. We found that children were sensitive to the level of conflict with the reward information but failed to use optimal strategies when cues did not align with the underlying rewards. Task characteristics (i.e., reliability of cues, whether the cues were social or nonsocial, whether cues changed over time) and participant characteristics (i.e., participant age) influenced how children weighted information sources.

Developmental Differences in Resolving Conflicts between Information Sources

In Experiments 1 and 3 children had to reconcile conflicts between a social agent's suggestion and the participants' observations of the underlying reward distribution. In both experiments, participant age influenced the pattern of choice behavior. Younger children appear to be especially sensitive to the initial behavior of a social agent and have relatively more difficulty updating their behavior. This finding stands in contrast to other 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; Lucas, Bridgers, Griffiths, & Gopnik, 2014). 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 old 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 uncertain environments when social information conflicts with another source of, also uncertain, information. The idea that young children have particular difficulty resolving these conflicts is consistent with previous research in which approximately half of the 4-year-old participants continued to endorse a block as being a causal mechanism when endorsed by an ostensibly knowledgeable informant, even after seeing probabilistic evidence that countered the informant's selection (Bridgers et al., 2016).

Informative Value of Social Information Matters

Social information, like language or the rules of causality, can follow statistical regularities. The regularity of social information may be important because more regular, or predictable, information (be it in conflict with other sources of information or consistent with other sources of information) is more informative than irregular, or unpredictable, social information. Therefore, children faced with processing social information—as in the present experiment—are tasked with both discerning whether the suggestions are accurate or inaccurate *and* whether the information is consistent. That children adjusted their behavior when the suggestions were clear (either low or high conflict)—even after an unexpected change in their utility—is consistent with findings that social information is not uniformly privileged or discounted. Rather, children pay attention to the utility of social information and compare that utility to that of their first-hand knowledge (Bridgers et al., 2016; Sobel & Kushnir, 2013). However, this adjustment process is influenced by the regularity of social information—and the age of the learner—and made more difficult when social information is unpredictable.

Unpacking Strategy Use

One unexpected finding was that children who received high- and moderate-conflict suggestions from the confederate or the nonsocial cue did not transition to maximizing (unlike children in Plate et al., 2018, who were tested in the absence of a confederate). Participants' failure to maximize was particularly surprising when the confederate's suggestions were frequently in conflict with the underlying rewards because the results showed that participants were able to disengage from the confederate's suggestions and rely on the reward distribution (probability matching). However, previous research has also shown reduced rates of maximizing when the statistics of the task are made more difficult to distinguish (Plate et al., 2018).

Several explanations are plausible for participants' failure to maximize in the present research. First, perhaps participants simply did not have enough trials to capture maximizing in this task. Given that children typically probability match for a period of time prior to maximizing, it is possible that the addition of a second source of information required a prolonged period of probability matching prior to transitioning to maximizing. Second, it is possible that children considered alternative strategies beyond probability matching. One such strategy could have involved taking into account how "off" the confederate's suggestions were. As shown in Figure 2, the confederate's error had a narrow distribution, and the confederate never selected a location more than two rocks away from the correct rock. This decision was made to avoid the confederate frequently choosing the never-rewarded rock locations, however some participants made spontaneous comments about the confederate typically choosing a rock that was 1-2 rocks away from the correct location. While there are not any deterministic trials (i.e., trials in which one could know the correct rock with 100% certainty based on the confederate's suggestion), it is possible that participants were trying to discern such a pattern. Future research, that includes trial types that systematically vary on determinism, is needed to

better understand children's approaches. Finally, it is possible that the confederate's behavior established a norm for exploration. There is evidence to suggest that 2-year-olds will try more actions on a novel object after having seen an agent try many (as compared to fewer) actions (Hoicka, Powell, Knight, & Norwood, 2017). Perhaps children continued to probability match (instead of maximize) because the confederate appeared to continue sampling options, as opposed to choosing a single option.

Limitations and Future Directions

We observed both similarities and differences in comparing the social and nonsocial cues. Differences between social and nonsocial conditions can be particularly difficult to interpret because of individuals' inclinations to attribute agency to nonsocial cues (e.g., Heider & Simmel, 1944; Wang et al., 2019). One possibility is that children privileged social information over nonsocial information. This view is in line with research suggesting that social cues can direct attention and action above and beyond nonsocial cues (Tummeltshammer, Gliga, & Kirkham, 2014; Tummeltshammer et al., 2014; Wu & Kirkham, 2010). Another possibility is that the differences were driven by the difference in instructions. In other words, because participants were told that the cue was incidental, perhaps they did not have any reason to think that the cue would help them find the reward. Despite this confound, however, participants in the low-conflict condition across both experiments used the cue. Nevertheless, children could have been more willing to discount the nonsocial cue after having been told that it was not an intended aspect of the game.

There is a body of research on pointing gestures that suggests that pointing is a particularly strong cue for preschool age children and that these children may have a bias to search behind pointed locations (Couillard & Woodward, 1999; Palmquist, Burns, & Jaswal,

2012; Palmquist & Jaswal, 2012). However, young children are able to distinguish between accurate and inaccurate pointers (Palmquist & Jaswal, 2015). Additionally, evidence suggests that preschoolers are lured by pointers because they have difficulty overcoming a bias that these individuals are knowledgeable (as opposed to the alternate hypothesis that young children have difficulty inhibiting a response to the visual point display; Palmquist, Kondrad, & Norris, 2018). Therefore, the difference in response to a point versus other visual display may be reflective of being drawn to social information, rather than a consequence of the particular image.

Considering these issues, the social versus nonsocial contrast could cut both ways: Following the cue in Experiment 1 could be too alluring because it is pointer-based and children are given instructions that might prime them to use the cue, therefore inflating the differences between Experiment 1 and Experiment 2. Alternatively, Experiment 2 could be unintentionally social because the cue is not entirely random, therefore diminishing the differences between Experiment 1 and Experiment 2. This complex nature of cues is precisely why investigations are needed to make headway in contrasting the mechanisms that underlie social learning and those that underlie nonsocial learning. We believe that the present research highlights an important early step in this area of research. Further, manipulations of the social cue could add to our insights regarding how individuals weight social input. For example, systematically varying the child's interactions with the confederate or the characteristics of the confederate (e.g. demographic characteristics or socio-emotional characteristics such as a friendly confederate vs. unfriendly confederate) could bear on how children make initial choices to follow a confederate's suggestions as well as why, and for how long, children persist in following a confederate's suggestions.

Another question for future research is how the confederate's distribution of responses might influence children's weighting of social information. As mentioned above, children may pay attention to the amount of error in the confederate's responses; therefore, future research may consider varying the amount and type of error in a confederate's response. 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 weight information when two sources diverge. Children were sensitive to the reward structure and the statistics of the cues, and they flexibly updated their behavior according to the current information available. However, children had difficulty adjusting behavior when social input was inconsistent with reward cues, and younger children over-relied on the social information in this instance. Younger children were also more strongly influenced by their initial interactions with the confederate, an effect that decreased with age. Through these experiments, we shed light on children's integration of multiple cues and elucidate the helpful and hindering power of social information. Taken together, this research makes progress in exploring the dynamic interplay between the multiple sources of information children must weight when making choices.

References

- Aslin, R. N. (2017). Statistical learning: a powerful mechanism that operates by mere exposure. *Wiley Interdisciplinary Reviews: Cognitive Science*, 8(1-2). doi:10.1002/wcs.1373
- Behrens, T. E., Hunt, L. T., Woolrich, M. W., & Rushworth, M. F. (2008). Associative learning of social value. *Nature*, 456(7219), 245–249. doi: 10.1038/nature07538
- Bonawitz, E., & Shafto, P. (2016). Computational models of development, social influences. *Current Opinion in Behavioral Sciences*, 7, 95-100. doi: 10.1016/j.cobeha.2015.12.008
- 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. doi: 10.1016/j.cognition.2010.10.001
- 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. doi: 10.1037/a0039830
- Buchsbaum, D., Bridgers, S., Whalen, A., Seiver, E., Griffiths, T., & Gopnik, A., (2012). Do I know that you know what you know? Modeling testimony in causal inference. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 34(34).
- Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children’s imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120(3), 331-340. doi: 10.1016/j.cognition.2010.12.001

- Butterfield, J., Jenkins, O. C., Sobel, D. M., & Schwertfeger, J. (2009). Modeling aspects of theory of mind with Markov random fields. *International Journal of Social Robotics, 1*(1), 41-51. doi: 10.1007/s12369-008-0003-1
- Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development, 7*(2), 213-233. doi:10.1016/0885-2014(92)90012-G
- Chan, C. C., & Tardif, T. (2013). Knowing better: the role of prior knowledge and culture in trust in testimony. *Developmental Psychology, 49*(3), 591. doi:10.1037/a0031336
- 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. doi:10.1348/026151008x310229
- Couillard, N. L., & Woodward, A. L. (1999). Children's comprehension of deceptive points. *British Journal of Developmental Psychology, 17*(4), 515-521. doi: 10.1348/026151099165447
- Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition, 126*(2), 285-300. doi: 10.1016/j.cognition.2012.10.010
- 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. doi:10.1177/0956797611416998
- Gergely, G., Bekkering, H., & Király, I. (2002). Developmental psychology: Rational imitation in preverbal infants. *Nature, 415*(6873), 755-755. doi:10.1038/415755a

- 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
- 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
- Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Harvard University Press.
- Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child development*, *77*(3), 505-524. doi:10.1111/j.1467-8624.2006.00886.x
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American journal of psychology*, *57*(2), 243-259. doi: 10.2307/1416950
- Hermes, J., Behne, T., Bich, A. E., Thielert, C., & Rakoczy, H. (2018). Children's selective trust decisions: Rational competence and limiting performance factors. *Developmental science*, *21*(2), e12527. doi: 10.1111/desc.12527
- Hermes, J., Rakoczy, H., & Behne, T. (2019). Making sense of conflicting information: A touchscreen paradigm to measure young children's selective trust. *Infant and Child Development*, *28*(1), e2119. doi: 10.1002/icd.2119
- Hoicka, E., Powell, S., Knight, J., & Norwood, M. (2018). Two-year-olds can socially learn to think divergently. *British Journal of Developmental Psychology*, *36*(1), 22-36. doi: 10.1111/bjdp.12199

- Jaswal, V. K. (2010). Believing what you're told: Young children's trust in unexpected testimony about the physical world. *Cognitive psychology*, *61*(3), 248-272. doi: 10.1016/j.cogpsych.2010.06.002
- Jaswal, V. K., Croft, A. C., Setia, A. R., & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science*, *21*(10), 1541-1547. doi:10.1177/0956797610383438
- Jaswal, V. K., & Kondrad, R. L. (2016). Why children are not always epistemically vigilant: Cognitive limits and social considerations. *Child Development Perspectives*, *10*(4), 240-244. doi:10.1111/cdep.12187
- Jones, R. M., Somerville, L. H., Li, J., Ruberry, E. J., Libby, V., Glover, G., ... Casey, B. J. (2011). Behavioral and Neural Properties of Social Reinforcement Learning. *Journal of Neuroscience*, *31*(37), 13039–13045. doi: 10.1523/JNEUROSCI.2972-11.2011
- 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
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*(2), B35-B42. doi:10.1016/S0010-0277(02)00004-5
- Kline, M. A. (2015). How to learn about teaching: An evolutionary framework for the study of teaching behavior in humans and other animals. *Behavioral and Brain sciences*, *38*. doi: 10.1017/S0140525X14000090
- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child development*, *76*(6), 1261-1277. doi:10.1111/j.1467-8624.2005.00849.x

- Koenig, M. A., & Harris, P. L. (2007). The basis of epistemic trust: Reliable testimony or reliable sources?. *Episteme*, 4(3), 264-284. doi:10.1353/epi.0.0017
- Kushnir, T., Wellman, H. M., & Gelman, S. A. (2008). The role of preschoolers' social understanding in evaluating the informativeness of causal interventions. *Cognition*, 107(3), 1084-1092. doi: 10.1016/j.cognition.2007.10.004
- Lane, J. D., Harris, P. L., Gelman, S. A., & Wellman, H. M. (2014). More than meets the eye: Young children's trust in claims that defy their perceptions. *Developmental Psychology*, 50(3), 865. doi:10.1037/a0034291
- Leonard, J., Bennett-Pierre, G., & Gweon, H. (2019). Who is better? Preschoolers infer relative competence based on efficiency of process and quality of outcome. *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- 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. doi: 10.1016/j.jecp.2018.03.008
- 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: 10.1016/j.cognition.2013.12.010
- Lucas, A. J., Burdett, E. R., Burgess, V., Wood, L. A., McGuigan, N., Harris, P. L., & Whiten, A. (2017). The Development of Selective Copying: Children's Learning From an Expert Versus Their Mother. *Child Development*, 88(6), 2026-2042. doi:10.1111/cdev.12711

- 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. doi:10.1037/0012-1649.24.4.470
- Palmquist, C.M., Burns, H.E., & Jaswal, V.K. (2012). Pointing disrupts preschoolers' ability to discriminate between knowledgeable and ignorant informants. *Cognitive Development*, 27(1), 54-63. doi: 10.1016/j.cogdev.2011.07.002
- Palmquist, C.M. & Jaswal, V.K. (2012). Preschoolers expect pointers (even ignorant ones) to be knowledgeable. *Psychological Science*, 23(3), 230-231. doi: 10.1177/0956797611427043.
- Palmquist, C.M. & Jaswal, V.K. (2015). Preschoolers' inferences about pointers and labelers: The modality matters. *Cognitive Development*, 35, 178 - 185. doi: 10.1016/j.cogdev.2015.06.003.
- Palmquist, C.M., Kondrad, R.L., & Norris, M.N. (2018). Follow my point? Preschoolers' expectations of veridicality disrupt their understanding of deceptive points. *Cognitive Development*, 48, 190 - 202. doi: 10.1016/j.cogdev.2018.08.009.
- Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, 43(5), 1216. doi:10.1037/0012-1649.43.5.1216
- 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. doi:10.1111/cdev.12718

- 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.
doi:10.1111/cdev.12720
- 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
- Sabbagh, M. A., & Shafman, D. (2009). How children block learning from ignorant speakers. *Cognition*, 112(3), 415-422. doi: 10.1016/j.cognition.2009.06.005
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928. doi:10.1126/science.274.5294.1926
- 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
- Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15(3), 436-447.
doi: 10.1111/j.1467-7687.2012.01135.x
- Sobel, D. M., Buchanan, D. W., Butterfield, J., & Jenkins, O. C. (2010). Interactions between causal models, theories, and social cognitive development. *Neural Networks*, 23(8-9), 1060-1071. doi: 10.1016/j.neunet.2010.06.004
- 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.
doi:10.1037/a0034191

- Sobel, D. M., & Munro, S. E. (2009). Domain generality and specificity in children's causal inference about ambiguous data. *Developmental Psychology*, *45*(2), 511. doi: 10.1037/a0014944
- 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. doi:10.1037/0012-1649.21.1.195
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91-94. doi: 10.1126/science.aaa3799
- Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology*, *28*(1), 509-529. doi: 10.1146/annurev.anthro.28.1.509
- Wu, R., Tummeltshammer, K. S., Gliga, T., & Kirkham, N. Z. (2014). Ostensive signals support learning from novel attention cues during infancy. *Frontiers in Psychology*, *5*, 251. doi: 10.3389/fpsyg.2014.00251
- Tummeltshammer, K. S., & Kirkham, N. Z. (2013). Learning to look: probabilistic variation and noise guide infants' eye movements. *Developmental science*, *16*(5), 760-771. doi: 10.1111/desc.12064
- Tummeltshammer, K. S., Wu, R., Sobel, D. M., & Kirkham, N. Z. (2014). Infants track the reliability of potential informants. *Psychological Science*, *25*(9), 1730-1738. doi: 10.1177/0956797614540178
- Waismeyer, A., Meltzoff, A. N., & Gopnik, A. (2015). Causal learning from probabilistic events in 24-month-olds: an action measure. *Developmental Science*, *18*(1), 175-182. doi: 10.1111/desc.12208

Wang, F., Tong, Y., & Danovitch, J. (2019). Who do I believe? Children's epistemic trust in internet, teacher, and peer informants. *Cognitive Development, 50*, 248-260. doi:

10.1016/j.cogdev.2019.05.006

Wu, R., & Kirkham, N. Z. (2010). No two cues are alike: Depth of learning during infancy is dependent on what orients attention. *Journal of Experimental Child Psychology, 107*(2),

118-136. doi: 10.1016/j.jecp.2010.04.014

Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences, 105*(13), 5012-5015. doi:10.1073/pnas.0704450105

Table 1

Participant Characteristics

	N	N male	N female	Mean Age	SD Age
Experiment 1					
high conflict	34	19	15	6.887	1.723
moderate conflict	32	17	15	7.053	1.866
low conflict	36	23	13	6.887	1.736
Experiment 2					
high conflict	29	17	12	6.898	1.671
moderate conflict	29	16	13	6.986	1.577
low conflict	31	19	12	7.118	1.779
Experiment 3					
high->low conflict	66	37	29	7.068	1.63
low->high conflict	63	31	32	6.938	1.691

Figures

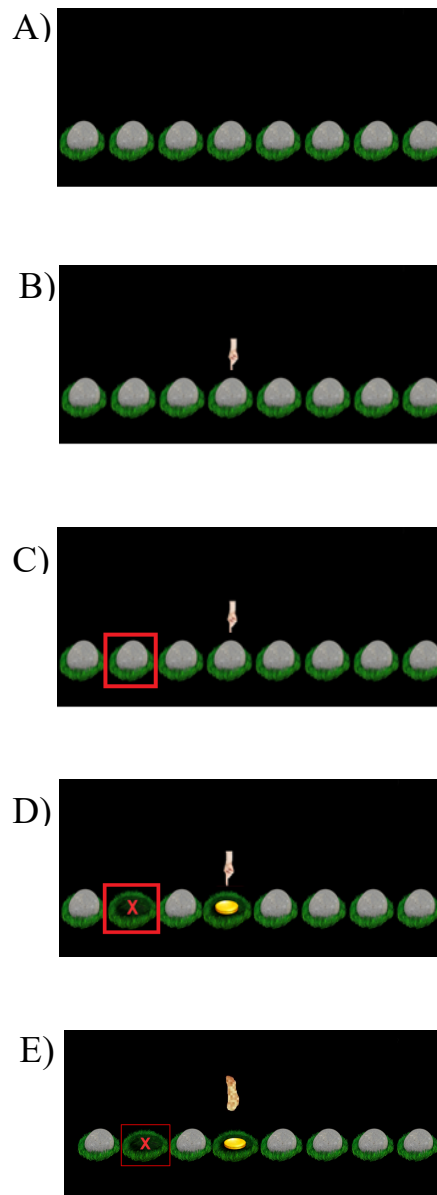


Figure 1. Progression of computerized task: A) Display prior to participant choice, B) Pointing hand indicates confederate's suggestion, C) Red box highlights participant's choice, D) Coin appears in rewarded location. E) Illustrates the nonsocial cue. Note: Figure illustrates one instance of choice 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).

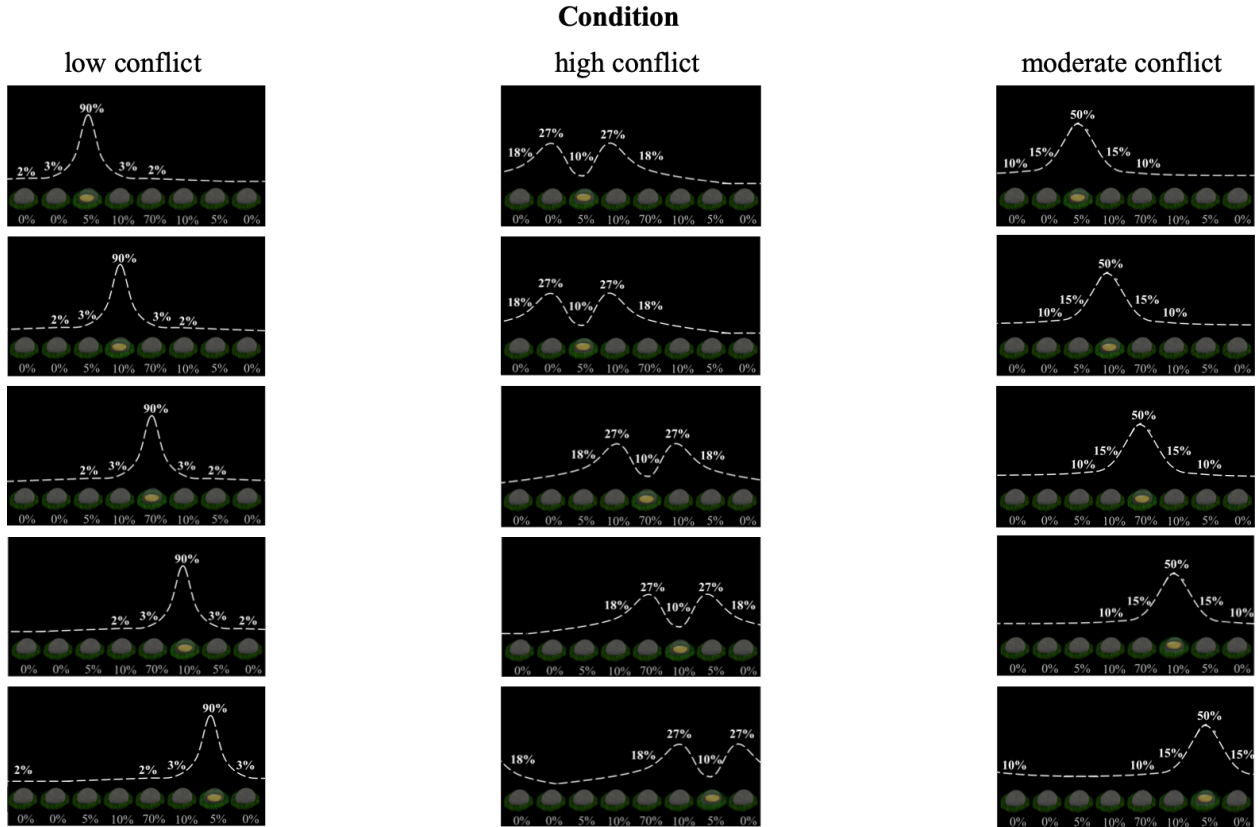


Figure 2. Likelihood of the cue indicating the rewarded location by condition (low-conflict condition on the left, high-conflict condition in the middle, moderate-conflict condition on the right) indicated by dotted lines. Likelihood of reward appearing behind each rock depicted below row of rocks (the same across all conditions). Note: The cue distribution was centered on the rewarded location on each trial and therefore was not static (in contrast to the reward distribution). As illustrated in the images on the bottom row, if the reward appeared under the seventh rock, the distribution “wrapped around.”

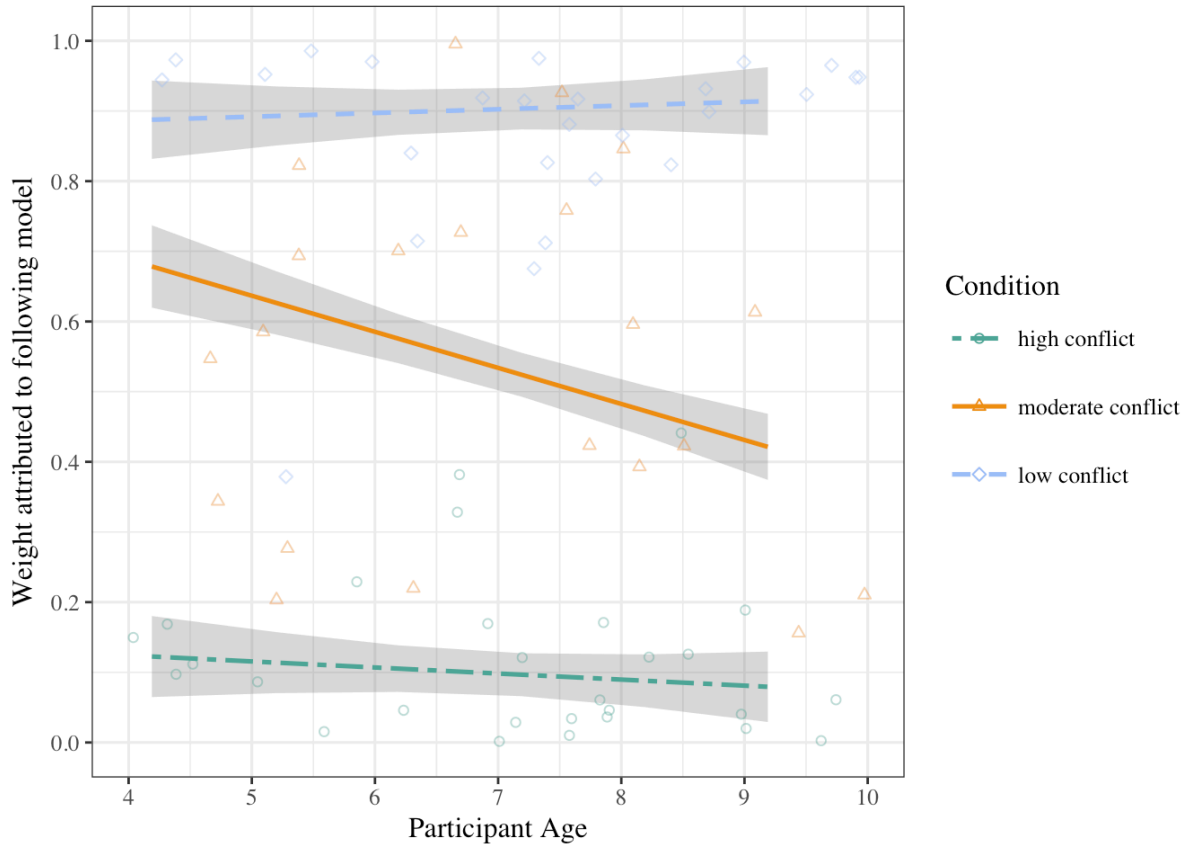


Figure 3. Weight attributed to the confederate's suggestion by age and condition for Experiment 1: Model predictions and participant-level data. Note: Lines are point estimates from linear model with the interaction between condition and participant age, and lower-order effects. Error bands represent standard error of the estimates. Points are individual participants' weights.




Figure 4. Weight attributed to the following model by condition and age for Experiment 3:

Model predictions and participant-level data. Error bars represent standard error of the estimates.

Points are individual participants' weights.




Figure 5. Likelihood of agreeing with the confederate's suggestion by age and condition in Experiment 3: 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 by age.