



Full Length Article

Children are more exploratory and learn more than adults in an approach-avoid task

Emily G. Liquin^{a,*}, Alison Gopnik^a

^a Department of Psychology, 2121 Berkeley Way, University of California, Berkeley, Berkeley, CA, USA



ARTICLE INFO

Keywords:

Exploration
Explore-exploit trade-offs
Development
Reinforcement learning

ABSTRACT

Intuitively, children appear to be more exploratory than adults, and this exploration seems to help children learn. However, there have been few clear tests of these ideas. We test whether exploration and learning change across development using a task that presents a “learning trap.” In this task, exploitation—maximizing immediate reward and avoiding costs—may lead the learner to draw incorrect conclusions, while exploration may lead to better learning but be more costly. In Studies 1, 2, and 3 we find that preschoolers and early school-aged children explore more than adults and learn the true structure of the environment better. Study 3 demonstrates that children explore more than adults even though they, like adults, predict that exploration will be costly, and it shows that exploration and learning are correlated. Study 4 shows that children’s and adults’ learning depends on the evidence they generate during exploration: children exposed to adult-like evidence learn like adults, and adults exposed to child-like evidence learn like children. Together, these studies support the idea that children may be more exploratory than adults, and this increased exploration influences learning.

1. Introduction

People often assume that children are more exploratory than adults, and that this contributes to their impressive learning abilities. The idea that children are active, exploratory, learners has roots in Piaget (1954) and more recent proponents in Gopnik (2012, 2020), Schulz (2012), and Xu and Kushnir (2013), among others. Recent research has found that children tend to explore when it is appropriate to do so (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Cook, Goodman, & Schulz, 2011; L. E. Schulz & Bonawitz, 2007) and take informative exploratory actions (Cook et al., 2011; Lapidow & Walker, 2020; Ruggeri, Sim, & Xu, 2017; Ruggeri, Swaboda, Sim, & Gopnik, 2019). However, the capacity for exploration does not stop in childhood: adults also explore when it is appropriate to do so (Dubey & Griffiths, 2020; Kang et al., 2009; Vogl, Pekrun, Murayama, Loderer, & Schubert, 2019) and take informative exploratory actions (Bramley, Lagnado, & Speekenbrink, 2015; Coenen, Rehder, & Gureckis, 2015; Gureckis & Markant, 2009; Rothe, Lake, & Gureckis, 2018; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). Given that adults and children alike are proficient explorers when given the opportunity, is it in fact the case that children explore *more* than adults in at least some contexts? And how might such exploration contribute to learning?

1.1. Exploration across development

There have been few direct comparisons of exploration across development. Adults have many cognitive advantages over children, so we might expect that adults would generally perform better than children on cognitive tasks, including those that involve exploration. However, we might sometimes see the opposite, counterintuitive developmental pattern—in some contexts, children might explore more than adults. In particular, in many real-world learning problems there is a trade-off between exploration and exploitation, and there is reason to believe that adults and children might resolve this trade-off differently.

The explore-exploit trade-off is well-defined in reinforcement learning problems, where agents discover which actions are rewarding and which are costly (Sutton & Barto, 1998). Agents choose actions, and they only get feedback about actions they take and not actions they avoid. As a result, solving a reinforcement learning problem requires balancing exploration—trying new actions to learn their rewards or costs—with exploitation—performing actions that you already know are rewarding (Cohen, McClure, & Yu, 2007; Hills, Todd, Lazer, Redish, & Couzin, 2015; Mehlhorn et al., 2015). For example, an agent whose goal is to eat maximally delicious pizza might initially spend some time trying a lot of different pizza toppings—that is, exploring—and this

* Corresponding author at: Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA.

E-mail addresses: emily.liquin@nyu.edu (E.G. Liquin), gopnik@berkeley.edu (A. Gopnik).

exploration will allow the agent to learn which toppings are delicious and which are not. However, exploration consumes time and energy, and comes with a risk of performing costly actions (e.g., eating a bad pizza). Therefore, after gaining some information, the agent will begin to selectively eat only demonstrably delicious pizzas—that is, exploit the information they have learned.

Many algorithms have been proposed to resolve the explore-exploit trade-off, and most share a common property: as in the pizza example, exploration is focused towards the beginning of the learning problem and decreases over time as the agent learns (Auer, 2002; Vermorel & Mohri, 2005). This is sensible, as an agent cannot properly exploit the reward structure of their environment until they have explored sufficiently. This suggests that children, who know less about the world than adults, might be biased towards exploration, while adults might be biased towards exploitation (Gopnik, 2020; Gopnik et al., 2017; Gopnik, Griffiths, & Lucas, 2015; Lucas, Bridgers, Griffiths, & Gopnik, 2014).

Related research beyond the reinforcement learning paradigm provides indirect support for this prediction: for example, children can be more flexible than adults in their hypothesis selection (Lucas et al., 2014; Seiver, Gopnik, & Goodman, 2013), language learning (Kuhl, 2004), and tool use (German & Defeyter, 2000). In addition, children have more broadly distributed attention than adults (Hanania & Smith, 2010; Plude, Enns, & Brodeur, 1994), which impacts learning and exploration (Blanco & Sloutsky, 2020; Deng & Sloutsky, 2016; Plebanek & Sloutsky, 2017). Third, adults tend to be averse to ambiguous gambles: risky gambles where the probability of reward is uncertain (Camerer & Weber, 1992). However, children as young as five display no ambiguity aversion (Li, Roberts, Huettel, & Brannon, 2017; see also Rosenbaum & Hartley, 2018), which may enable exploration in the face of risk. Finally, there is evidence that exploration decreases with aging in older adults (Chin, Payne, Fu, Morrow, & Stine-Morrow, 2015; Mata, Wilke, & Czienskowski, 2013).

While there is evidence for developmental change in reinforcement learning processes in childhood and adolescence (Nussenbaum & Hartley, 2019), researchers have only very recently begun to investigate developmental changes in exploration specifically. In a classic “multi-armed bandit task,” participants choose over a number of trials between a fixed number of options, each of which produces a fixed reward (e.g., one option provides three candies and another option provides eight candies). In this task, adults repeatedly choose the most rewarding option after a brief initial period of exploration, while children as young as 3 years old are more likely to persist in choosing both good and bad options throughout the task (Blanco & Sloutsky, 2021; Sumner et al., 2019; Sumner, Steyvers, & Sarnecka, 2019). Using a modified multi-armed bandit task, another study found that 7- to 11-year-olds showed more directed exploration (actively choosing options that resolve uncertainty) than adults (E. Schulz, Wu, Ruggeri, & Meder, 2019). However, the strategic use of directed exploration—exploring to resolve uncertainty specifically when this information will be useful for future decisions—does not emerge until later in adolescence, so adults in fact perform *more* strategic directed exploration than young adolescents (Somerville et al., 2017).

Together, these studies provide some evidence for developmental changes in exploration but their robustness across tasks and their impacts for learning remain unclear. Research directly investigating developmental change in exploration has exclusively used versions of the multiarmed bandit task, so we don’t know whether children explore more than adults in more complex settings. For example, in the classic bandit task, both good and bad options provide rewards—so there may be some opportunity cost for broad exploration, but a child has nothing to lose by continuing to explore. In the real world, learners need to contend not only with rewards, but also with costs: broader exploration might lead to a real, material cost, like eating a bad pizza or falling while climbing a tree. It remains less clear how children, particularly young children, explore in tasks that incorporate costs.

In addition, little research has investigated the connection between

explore-exploit decision making and broader learning across development. In the classic bandit task, the environment is extremely simple and there is only one distinction to learn, namely which arm is more rewarding. Real-life environments are more complex and often have abstract multi-dimensional structure, and exploration may be particularly important for learning in these settings.

In fact, earlier work (Fazio, Eiser, & Shook, 2004; Rich & Gureckis, 2018) suggests that adults under-explore in “approach-avoid” decision making tasks, and, as a result, may fail to learn the complex structure of their environment. In these tasks, participants choose to approach or avoid a single stimulus on each trial, where each stimulus is associated with some reward or cost. If the participant chooses to approach, they receive information about the stimulus and potentially any reward, but they also risk incurring a cost. If the participant chooses to avoid, they receive no information but do not incur a cost or receive a reward. In this task, adults tend to generalize from negative outcomes (e.g., they infer from one negative experience eating a banana pizza that all fruit pizzas will be bad), and so they fail to explore other related stimuli (pineapple pizzas) as a result. Because they instead avoid these stimuli, they never discover that their initial generalization was wrong. In other words, they fall into a “learning trap” (Rich & Gureckis, 2018): their initial false inference means they fail to explore, and their failure to explore means they are not exposed to the information that would correct that inference. Ultimately, this prevents them from accurately learning the structure of the environment. This mechanism has been used to explain a variety of real-life phenomena, from biases in social impression formation (Denrell, 2005) to the performance of organizations (Denrell & March, 2001; March, 1991) to clinical disorders like anxiety or phobias (Dymond, Dunsmoor, Vervliet, Roche, & Hermans, 2015). For example, one bad experience in an airplane may lead you to become anxious about airplanes and avoid flying, and so you never learn that most plane trips are fine.

This approach-avoid task provides a particularly good test of how exploration and learning are related and change across development. Some research suggests that 7- to 11-year-olds are less likely than adults to generalize from explored actions to new actions in multiarmed bandit tasks (E. Schulz et al., 2019), and there is evidence that 5- to 12-year-olds are more likely than adults to notice changes in the reward structure of a multiarmed bandit task (Sumner, Li, et al., 2019). Both these results suggest that children’s exploration might enable them to discover information that adults miss and avoid learning traps. In Studies 1, 2, and 3 of this paper, we test whether this is the case.

1.2. Exploration and learning: hypotheses for developmental change

A second aim of the paper, the focus of Studies 3 and 4, is to test different hypotheses about why children might explore and learn more than adults in these tasks. There are two different ways that exploration could change across development in these tasks, which we will call the “inferential” and “motivational” accounts.

First, children and adults might have different prior beliefs and/or make different inferences from the evidence they observe, and those beliefs or inferences might lead to different patterns of both exploration and learning. To illustrate, consider the example of exploring pizza toppings. If someone has a strong prior bias against the idea of fruit toppings, they might predict that fruit pizzas will be bad and so avoid tasting them. Moreover, if someone with such a bias does try a banana-bologna pizza and finds that it is indeed, awful, that might confirm their prior bias and lead them to generalize it to new examples. This would make them even less likely to try new fruit pizzas, such as a banana-ham pizza and a pineapple-ham pizza—the inference would make them less exploratory. As a result, because of this (lack of) exploration, the individual will fail to learn that pineapple-ham pizza is actually delicious.

In contrast, someone without this prior bias might be less likely to infer that fruit pizzas are bad from a single taste of banana-bologna pizza. Consequently, this individual would be more willing to try a

new fruit pizza, and they might eventually learn that pineapple-ham is delicious. In sum, differences in prior beliefs and inference may shape exploration, which shapes learning. Notably, a difference in exploration and learning might arise even in cases where prior beliefs are similar: two individuals might initially have no bias against fruit pizzas but draw different inferences after tasting one bad banana-bologna pizza. An individual who infers that all fruit pizzas are bad will fail to explore additional fruit pizzas, while an individual who infers that only some fruit pizzas are bad will continue to explore. Because the difference in exploration relies directly on these initial inferences (which may be indirectly informed by prior beliefs), we call this the “inferential” account.

Alternatively, children and adults might make similar inferences but differ in how much they are motivated to explore—that is, how willing they are to risk costs in order to get new information. For example, two people might agree that fruit-topped pizza is bad, predict that they will be unlikely to enjoy it, and find that view confirmed when they try banana-bologna pizza. Nevertheless, they might differ in how willing they are to continue to try new pizzas, even new fruity pizzas. We capture these differences by saying that one individual is a more adventurous eater than the other, or even by invoking the individual personality trait of “openness to experience” (McCrae & Costa, 1997). The adventurous eater will be likely to learn more about the world of food even if that means more bad food experiences as well as good ones. We call this the “motivational” account.

Children and adults might differ in either or both of these ways. Children might have different prior biases about the world than adults and therefore make different inferences from the same evidence (or they might make different inferences even with similar prior beliefs). In fact, there is evidence that this is true in some cases (e.g., Gopnik et al., 2017; Lucas et al., 2014; E. Schulz et al., 2019). In turn, this would both influence children’s and adults’ exploration and lead to differences in their learning: because children draw different inferences than adults, they would explore more and avoid learning traps. Alternatively, children and adults might make similar inferences (perhaps due to similar prior beliefs), but children might be intrinsically more inclined to explore even when anticipating costs. As a result, children and adults would get different evidence from exploration, which would influence their final conclusions. It is also possible, of course, that children and adults differ in both their inferences and their motivations.

These two accounts can be tested in two ways. First, the accounts make different predictions about adults’ and children’s inferences from initial evidence. According to the inferential account, children and adults make different initial inferences, but according to the motivational account, children and adults make similar initial inferences. We test these predictions in Study 3. Second, the accounts make different predictions about how adults and children respond to a learning trap. According to the inferential account, adults might fall into the learning trap and children might avoid it even if both groups received the same evidence. For example, adults’ prior beliefs might pull them into a learning trap even in the face of contradictory evidence (e.g., an individual who has a strong bias against fruit pizzas might resist changing their belief even after eating a pineapple-ham pizza). In contrast, according to the motivational account, whether an individual falls into the learning trap is directly dependent upon the evidence they generate during exploration. In particular, an individual who avoids relevant evidence will fall into the learning trap, while an individual who seeks relevant evidence despite potential costs will avoid the learning trap. We test these predictions in Study 4.

1.3. The present research

In the present research, we investigate the “learning trap” cases where adults do not explore stimuli they think are costly, and so fail to learn the structure of the environment. In particular, we construct a case where the actual structure of the environment follows a complex two-

dimensional rule (e.g., choosing either black-patterned objects or spotted objects leads to a reward), and broad exploration would allow the learner to discover this rule. However, avoiding potentially risky exploration would lead the learner to fall into a learning trap and inaccurately infer a simpler one-dimensional rule (e.g., choosing black-patterned objects leads to a reward).

We test preschool and early school-aged children’s and adults’ behavior in this task to determine whether 1) children are more exploratory than adults, 2) this increased exploration allows children to learn better than adults and escape from learning traps, and 3) this increased exploration results from different inferences, different motivation to explore, or both. To preview our results, in Study 1, we demonstrate that adults fall into learning traps even in a child-friendly approach-avoid task. In Study 2, we test 4- to 7-year-olds’ behavior, finding that children are both exploratory and resistant to learning traps. In Study 3, we do not find evidence that children and adults make different initial inferences: children and adults both expect an uncertain action to be costly when it is similar to a previous costly action. However, children are more likely to risk an exploratory action even when they predict it will be costly. As a result, they get more relevant evidence than adults and learn more accurately. Finally, in Study 4, we find that children draw similar conclusions to adults when they receive the kind of evidence generated by adults, and vice-versa. Together, these findings suggest that the differences between adults and children in these tasks are more likely to be motivational than inferential. Finally, we show in Study 4 that children’s learning influences their further actions: they use the information they have gained to obtain more rewards, though they still fail to exploit as efficiently as adults.

2. Study 1

In Study 1, we attempted to replicate the basic “learning trap” results in adults, using a child-friendly causal learning task based on the method used by Rich and Gureckis (2018). Participants had to identify which objects were “zaffs”. Zaffs would activate a “zaff machine” and lead to a reward, while non-zaffs would instead incur a cost. The actual category rule was two-dimensional: for example, any object with either spots or a black pattern was a zaff, but objects with white stripes were not zaffs. Following Rich and Gureckis, we predicted that: 1) adults would increasingly avoid non-zaffs over the course of the task, but would fail to increasingly approach all zaffs, and 2) adults would learn an incorrect one-dimensional rule, such as that zaffs are spotted rather than striped, rather than the correct two-dimensional rule.

2.1. Method

2.1.1. Participants

Participants were 42 adults (age 21–62; 30 self-identified as male, 12 female) recruited from Amazon Mechanical Turk (MTurk) in June 2018. We did not collect other demographic data, but the demographic makeup of MTurk samples is reported in other research (e.g., Levay, Freese, & Druckman, 2016). The desired sample size was at least half the sample size of Study 2 (i.e., at least 32 participants); the sample size of Study 2 was based on a power analysis (see below). All participants were based in the United States and had completed at least 50 assignments on MTurk with at least a 95% approval rating. An additional 5 participants completed the study but were excluded from analyses for indicating that they had not watched all of the videos in the study. Participants were paid \$0.50, plus a bonus dependent on their performance (detailed below).

2.1.2. Materials

While this experiment was conducted online, several physical materials were used to make videos and images. Sixteen yellow wooden blocks varying along two dimensions—pattern (spots vs. stripes) and pattern color (black vs. white)—were assigned to be “zaffs” or “non-

zaffs” on the basis of a two-dimensional category rule (e.g., blocks that are white *and* striped are non-zaffs; all other blocks are zaffs). The specific rule (white/striped, black/striped, white/spotted, black/spotted) was counterbalanced across participants. Throughout the study, blocks were placed on a “zaff machine”: a laptop computer with a decorated shoebox lid placed atop the keyboard. The computer was remotely controlled with an iPhone application (Mobile Mouse Pro) or a wireless presentation remote, which the experimenter surreptitiously used to activate one of two onscreen images (a happy face or a sad face) when a block was placed on the machine. Four blue blocks varying along the same two dimensions as the yellow blocks were also created. The blue and yellow blocks were referred to as “blocks” to participants across all studies, but we refer to them as “objects” here to prevent confusion, given the common use of the word “block” to indicate sets of trials. We refer to objects with a given feature combination (e.g., white striped objects) as an “object type.”

2.1.3. Procedure

The procedure was a modified version of Experiment 1 (contingent condition) from Rich and Gureckis (2018). Participants were first told that they would be playing a game that is also played with young children, and they were encouraged to take the game seriously and try their best. They were also instructed to watch all of the videos in the experiment. After these initial instructions, participants completed the instructions phase, the approach-avoid phase, the test phase, and the generalization phase, then provided their age and gender.

In the instructions phase, participants watched a video, in which they were shown all 16 yellow objects sitting on a table. Participants were told that some of the objects were “zaffs,” which would make the screen of the “zaff machine” light up with a green happy face, and that some of the objects were “not zaffs,” which would make the screen light up with a red sad face. Participants were instructed that they would choose whether or not to put a sequence of objects on the machine, one-by-one. They would start the game with four stars, and they would gain one star for every zaff they put on the machine and lose two stars for every non-zaff they put on the machine. Any object that was not placed on the machine would have no effect on their stars. To incentivize performance, participants were told that they would be paid a bonus proportional to the number of stars they had at the end of the game. Before exiting the instructions phase, participants were asked three multiple-choice attention check questions (“What happens if you decide to put a block on the machine and it is a zaff?”; “What happens if you decide to put a block on the machine and it is not a zaff?”; and “What happens if you decide not to put a block on the machine?”) to ensure comprehension. Fourteen percent of participants answered the first question incorrectly, 7% answered the second question incorrectly, and 12% answered the third question incorrectly. Participants who answered any question incorrectly were reminded of the rules in text before proceeding.

Next, in the approach-avoid phase, participants were shown an image of one of the 16 objects and chose whether to place the object on the machine (“approach”) or to put the object away (“avoid”). If a participant chose to approach, they were shown a brief video clip of the experimenter placing the object on the machine and narrating the outcome (e.g., “It’s a zaff! So you get one more star.”). If the participant chose to avoid, they were shown the text, “Okay, we’ll put that block away,” before advancing to the next trial. This was repeated for 16 trials (one per object). Trials were organized into four trial sets, in which one object of each object type was shown. In the first trial set, the non-zaff was always presented on either the second or third trial, ensuring that participants saw both a positive and negative example in the first three trials. All other object positions were randomized. Throughout the approach-avoid phase, the number of stars a participant had and the number of trials remaining were displayed. At the end of the approach-avoid phase, participants were told how many stars they had earned and what their bonus would be at the end of the study (\$0.02 per star).

Next, in the test phase, participants were shown an image of each object type on four separate pages (white/striped, black/striped, white/spotted, black/spotted) and indicated whether they thought each object was a zaff or not a zaff. The order in which the four objects were presented was randomized.

In the generalization phase, participants were shown an image of each object type, but with a novel background color (blue rather than yellow). Participants indicated whether they thought each novel object was a zaff or not a zaff. The order in which these objects were presented was randomized.

At the end of the study, participants provided their age and gender and indicated whether they had watched all of the videos. Participants were encouraged to answer this final question honestly and were told that their response would not affect their compensation.

2.2. Results

2.2.1. Exploration

First, we investigated adults’ exploration. At the beginning of the task, participants did not know which objects were zaffs, so we expected adults to approach the first several objects. However, after receiving some initial evidence (e.g., seeing that a black striped object is a zaff and a white striped object is a non-zaff), we predicted that adults would infer an incorrect one-dimensional rule and so fall prey to a learning trap. Their subsequent decisions would be based on the incorrect one-dimensional rule they had inferred (e.g., black objects are zaffs and white objects are non-zaffs). If adults indeed fell prey to the learning trap, we would predict distinct patterns of exploration for non-zaffs and zaffs: 1) participants would be less likely to approach non-zaffs (e.g., white striped objects) over trials, as they receive additional evidence that these objects are costly, and 2) participants would fail to approach all the zaffs, even by the end of the task. After they initially erroneously infer that one of the zaffs (e.g., the white spotted object) is a non-zaff, participants would avoid this object and so would not receive corrective feedback.

To test the first prediction, we fit a generalized linear mixed-effects model with a logit link function (i.e., mixed-effects logistic regression) to adults’ approach-avoid decisions for non-zaffs. Trial set was included as a fixed effect, and random intercepts were included for participants. We compared this model to a null model with only a fixed intercept using a likelihood ratio test, to determine whether trial set made a significant contribution to the model. Supporting our prediction, we found a significant decline in the probability of approach over trial sets, $OR = 0.33$, 95% CI [0.18, 0.54], $\chi^2(1) = 24.48$, $p < .001$, indicating that participants learned to avoid non-zaffs over time (see Fig. 1a). To test the second prediction, we fit a mixed-effects regression model predicting the proportion of zaffs approached in each trial set, with trial set as a fixed effect and random intercepts for participant. This revealed no evidence for a change in the tendency to approach zaffs over the four trial sets, $b = 0.001$, 95% CI [-0.01, 0.02], $\chi^2(1) = 0.01$, $p = .92$. Across all trial sets, participants only approached an average of 59% of true zaffs (see Fig. 1a). This is consistent with the prediction that participants chose to approach or avoid on the basis of a one-dimensional rule (e.g., “all white objects are non-zaffs”), which would predict that they would only approach two of the three actual zaffs, consistently across trial sets.

2.2.2. Learning

For each participant, we classified test and generalization responses into four categories, based on their responses to each object (see Fig. 1b). Participants responded according to 1) the correct two dimensional rule (e.g., the white striped object was classified by the participant as a non-zaff; the other three objects were classified as zaffs), 2) the incorrect but appropriate one-dimensional pattern rule (e.g., striped objects were classified as non-zaffs; spotted objects were classified as zaffs), 3) the incorrect but appropriate one-dimensional color rule (e.g., white objects were classified as non-zaffs, black objects were classified as zaffs), or 4)

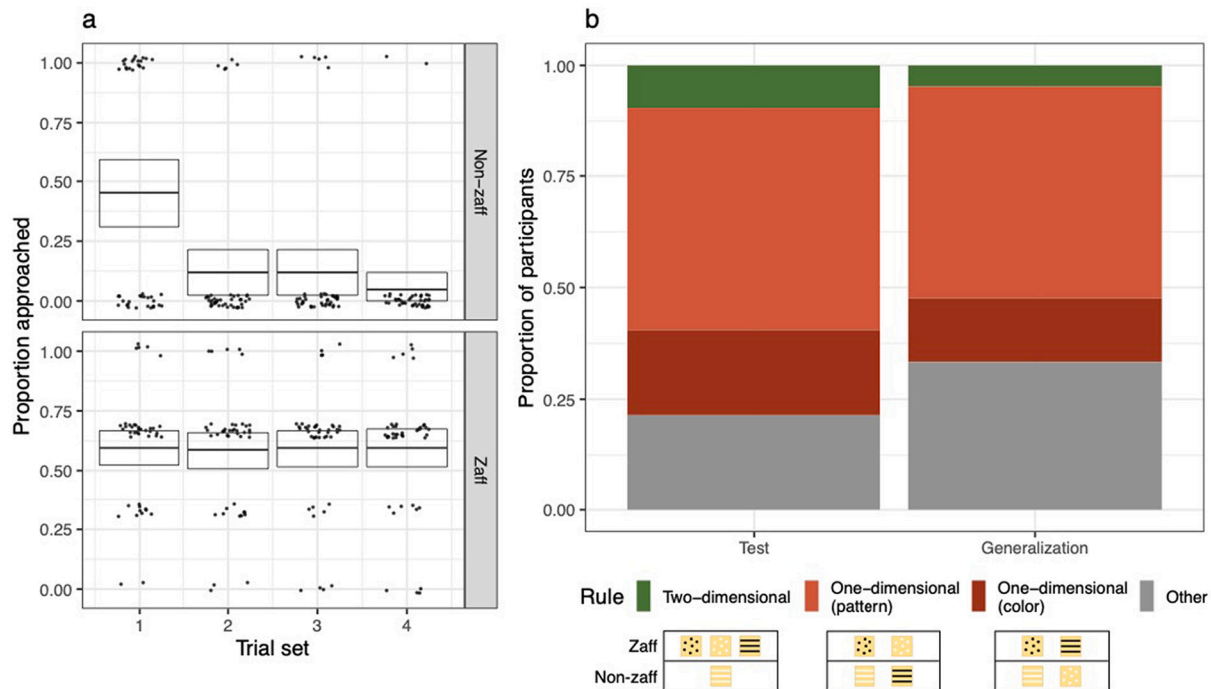


Fig. 1. Adults' exploration and learning, Study 1. A. Adults' (Study 1) approach-avoid decisions across four trial sets (each comprised of one non-zaff trial and three zaff trials). Boxplots display mean proportion approached with bootstrapped 95% confidence intervals, and jittered points represent individual participant choices (with each participant represented four times in the top panel for non-zaff trials and four times in the bottom panel for zaff trials). B. Classification of adults' learning (Study 1) at test and generalization. Each participant labelled each object as a zaff or a non-zaff; responses were classified according to whether they followed a two-dimensional or one-dimensional rule. Beneath each rule classification, the relevant pattern of test responses (where white/striped is the correct two-dimensional rule) is shown.

another response (e.g. all objects were classified as non-zaffs). In the test phase, 10% of participants perfectly followed the correct two-dimensional rule, 50% of participants perfectly followed the one-dimensional pattern rule, and 19% of participants perfectly followed the one-dimensional color rule. The remaining 21% of participants responded according to one of 7 unique additional patterns and were excluded from the following analysis. These proportions followed a similar pattern at generalization: 5% of participants perfectly followed the correct two-dimensional rule, 48% perfectly followed the one-dimensional pattern rule, 14% perfectly followed the one-dimensional color rule, and 33% responded according to some other pattern. At both test and generalization, participants responded more often with one of the incorrect one-dimensional rules (either pattern or color), than the correct two-dimensional rule, test: $\chi^2(1) = 18.94, p < .001$; generalization: $\chi^2(1) = 20.57, p < .001$, indicating that adults were more likely to fall prey to the learning trap than they were to learn the correct two-dimensional rule (see Fig. 1b).

2.2.3. Reward

Finally, we examined adults' performance in the approach-avoid task: did adults succeed in earning rewards? If an individual perfectly exploited a two-dimensional rule across the 16 trials, they would end with 16 "stars" (4 initial stars plus 12 zaffs approached). In contrast, if an individual perfectly exploited a one-dimensional rule across the 16 trials, they would end with 12 stars (4 initial stars plus only 8 zaffs approached). On average, adults ended with 9.64 stars.

2.3. Discussion

These results are consistent with prior research (Rich & Gureckis, 2018): after very little evidence, adults fell prey to a "learning trap" that prevented them from discovering the true rule. This was reflected in both adults' exploration and their final learning. Adults were

increasingly likely to avoid non-zaffs, but they failed to approach all zaffs. Moreover, adults were far more likely to classify the objects according to a one-dimensional rule rather than a two-dimensional rule, and they generalized this rule to new objects that varied along the same dimensions. These findings confirm that this child-friendly paradigm is able to induce restricted exploration and restricted rule learning in adults.

3. Study 2

In preregistered Study 2 (preregistration: <https://aspredicted.org/cu8ia.pdf>), we used the paradigm developed in Study 1 to investigate exploration and learning in 4- to 7-year-olds. We predicted that children would generally choose to approach relatively frequently, and that children would typically discover the correct two-dimensional rule as a result. In addition, we tested whether there is developmental change in exploration within childhood, predicting that as participants grew older they would be less likely to explore, and therefore less likely to learn accurately.

3.1. Method

3.1.1. Participants

Participants were sixty-four 4- to 7-year-olds (16 from each age in years; 28 identified by their parents as male, 36 female), recruited from and tested onsite at museums and preschools in the San Francisco Bay Area, between April 2018 and January 2019. We did not collect other demographic information from participants, but museums and preschools were located in a large metropolitan area and served families resembling the diversity of the local population. The sample size was determined by power analysis, which indicated that a total sample of 48 was needed to detect an effect of age on test performance with at least 80% power, using the effect size from a pilot study ($OR = 2.23$). Because

this effect size estimate was based on only 10 pilot participants, we preregistered the target sample of 64, to ensure sufficient power. Eight additional participants were excluded from analyses because of parent interference (two children), experimenter error (five children), or refusal to continue after seeing the sad face (one child). An additional 31 participants were excluded for failing the preregistered inclusion criterion that the answers given in the test phase were fully consistent with the object information participants had seen in the approach-avoid phase. Due to the unexpectedly high rate of exclusion, we also report all analyses including these participants in Supplementary Results 1.

3.1.2. Procedure

The procedure was identical to that of Study 1, with the following exceptions. First, children played the game in-person with an experimenter, who surreptitiously controlled whether the zaff machine displayed a happy face or sad face on each trial. Instead of virtual “stars,” children started with four small stickers and gained or lost stickers based on the outcomes of their actions. Participants were not told whether or not they would get to keep these stickers at the end of the game; in fact, they chose four stickers to keep when the game had concluded. If a participant could not generate the correct response to either of the first two attention check questions, they were prompted with two alternatives (e.g., “Do you get one sticker, or do you lose two stickers?”). If the participant still could not answer or gave an incorrect answer, they were told the correct response. For the third attention check question (“What happens if you decide not to put a block on the machine?”), participants were told the correct response if they could not generate the answer or answered incorrectly. Nineteen percent of children were told the response to the first question, 5% were told the response to the second question, and 25% were told the response to the third question.

We assumed that adults in Study 1 were able to remember the specific events in the task, so we wanted to ensure that the children in this study would do so, as well. We did this in two ways. First, we gave the child participants memory aids throughout the game. Each object that was placed on the machine was moved to one of two locations based on its category: zaffs were placed on a sheet of paper with a green happy face in the bottom left corner; non-zaffs were placed on a sheet of paper with a red sad face in the bottom left corner. The left-right position of these memory aids was counterbalanced. For the test phase, one object of each object type was removed from these memory locations and presented in a random order, unless there were one or fewer objects of a given object type in the memory locations. In that case, one of the objects that had been avoided during the approach-avoid phase was used instead.

Second, to ensure our analyses included only children who actually remembered the events, we used a preregistered inclusion criterion. Children who reported a rule at test that was inconsistent with the data that they had seen during the approach-avoid phase were excluded from all analyses. In other words, a child who approached all object types at least once was required to provide correct responses for all four objects at test, whereas a child who approached only three of the four object types at least once was required to provide correct responses for those three objects at test.

3.1.3. Coding

All data was recorded during each study session by the experimenter or a trained research assistant. 78% of all sessions were video recorded, and these videos were re-coded by a trained research assistant or the first author, blind to the original recorded data. There was 98% agreement between the original data coding and the video coding. Cases of disagreement were reviewed by a trained research assistant or the first author, and the data were adjusted in favor of the video when the video provided unambiguous evidence that the original coding was incorrect.

3.2. Results

3.2.1. Exploration

We preregistered the analyses of approach/avoid decisions used in Study 1, adding age in years as a fixed effect in all models. For non-zaffs, there was no significant interaction between age and trial set, $\chi^2(1) = 0.29, p = .59$. The best-fitting model included fixed effects for both age, $OR = 0.24, 95\% CI [0.11, 0.45], \chi^2(1) = 20.69, p < .001$, and trial set, $OR = 0.46, 95\% CI [0.31, 0.65], \chi^2(1) = 21.33, p < .001$, with no interaction. That is, like adults, children avoided non-zaffs increasingly over trial sets. In addition, older children were more likely to avoid non-zaffs than younger children (see Fig. 2a). For zaffs, a model including the interaction between age and trial set outperformed a model including no interaction, $\chi^2(1) = 7.75, p = .005$. Among 6- and 7-year-olds, participants were more likely to approach zaffs over trial sets, $b = 0.05, 95\% CI [0.03, 0.07], \chi^2(1) = 18.44, p < .001$, while this effect was weaker but still significant for 4- to 5-year-olds, $b = 0.02, 95\% CI [0.005, 0.03], \chi^2(1) = 7.20, p = .007$. That is, children in both age groups became more likely to approach zaffs as the task went on (see Fig. 2a). In contrast, adults in Study 1 increasingly avoided non-zaffs but did not increasingly approach zaffs. In addition, children approached 84% of zaffs across all trial sets, a numerically larger proportion than the 67% that would be expected if participants chose to approach or avoid on the basis of a one-dimensional rule.

3.2.2. Learning

We preregistered the classification of test and generalization responses used in Study 1. In the test phase, 70% of participants perfectly followed the correct two-dimensional rule, 17% of participants perfectly followed an appropriate though incorrect one-dimensional pattern rule, and 11% of participants perfectly followed an appropriate though incorrect one-dimensional color rule. A single child provided responses inconsistent with these three response patterns (labelling as non-zaffs both the true non-zaff and the object sharing neither color nor pattern with the non-zaff). Children’s performance was less clear at generalization: 25% responded according to the correct two-dimensional rule, 27% responded according to the incorrect one-dimensional pattern rule, and 13% responded according to the incorrect one-dimensional color rule, while the remaining 36% of participants responding according to one of nine unique additional patterns (percentages do not add up to 100 due to rounding). We fit a logistic regression model predicting whether participants responded according to a one-dimensional rule (based on pattern or color) or two-dimensional rule as a function of age in years. In the test phase, there was a significant effect of age, $OR = 1.89, 95\% CI [1.12, 3.40], \chi^2(1) = 5.79, p = .02$, with increasing age corresponding to an increase in using a one-dimensional rule. In the generalization phase, there was no evidence for an effect of age, $OR = 1.27, 95\% CI [0.72, 2.30], \chi^2(1) = 0.71, p = .40$. In other words, children did not generalize a rule consistently from the objects they had seen to new and different objects, and thus there was no evidence for an age difference within children at generalization. (see Fig. 2b).

3.2.3. Reward

Finally, we investigated children’s performance in the approach-avoid task. On average, children ended with 9.41 stickers. An exploratory regression model revealed that the number of stickers earned increased with age, $b = 0.54, 95\% CI [0.15, 0.92], t = 2.77, p = .007$. Thus, even though younger children were more likely to learn a two-dimensional rule than older children, they did not successfully exploit this rule to earn more rewards.

3.3. Discussion

In Study 2, we investigated children’s exploration and learning in an approach-avoid decision making task. Children frequently chose to approach rather than avoid, and younger children explored more than

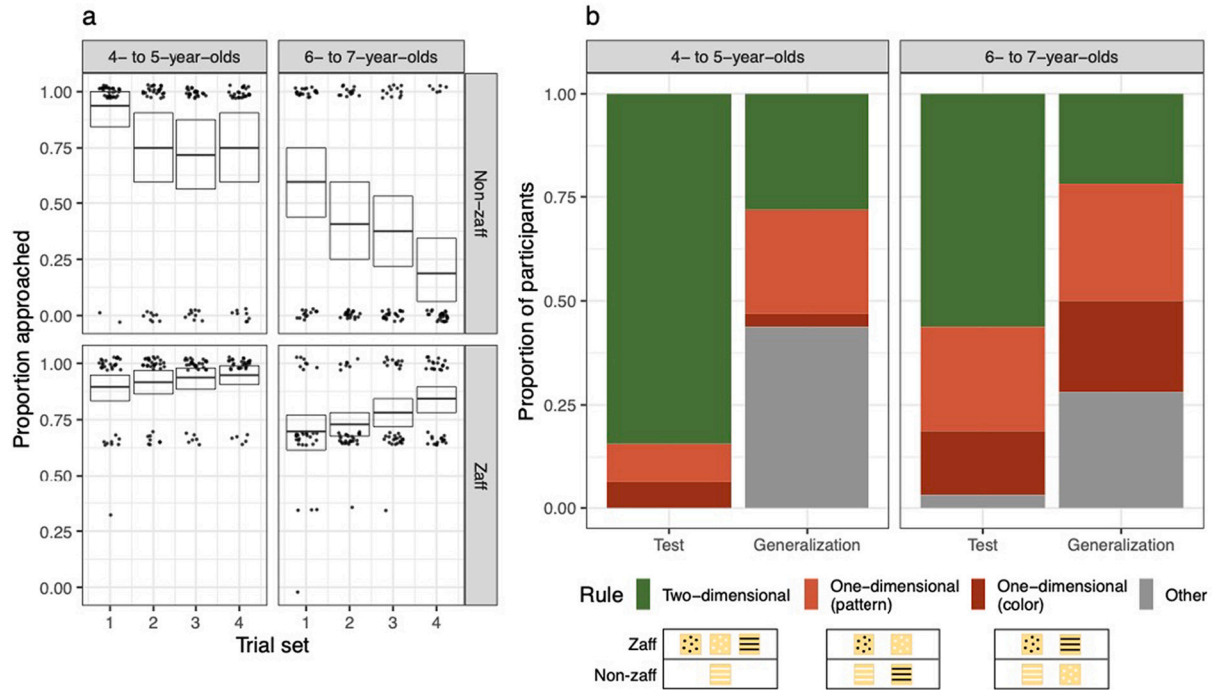


Fig. 2. Children’s exploration and learning, Study 2. A. Children’s (Study 2) approach-avoid decisions across four trial sets (each comprised of one non-zaff trial and three zaff trials). Boxplots display mean proportion approached with bootstrapped 95% confidence intervals, and jittered points represent individual participant choices (with each participant represented four times in one top panel for non-zaff trials and four times in one bottom panel for zaff trials). B. Classification of children’s learning (Study 2) at test and generalization. Each participant labelled each object as a zaff or a non-zaff; responses were classified according to whether they followed a two-dimensional or one-dimensional rule. Beneath each rule classification, the relevant pattern of test responses (where white/striped is the correct two-dimensional rule) is shown.

older ones. Moreover, children were more likely to later classify objects according to the correct two-dimensional rule than an incorrect one-dimensional rule, and again younger children were more accurate than older ones. These findings from Studies 1 and 2 support our predictions: children are more exploratory than adults, and this broader exploration can lead them to learn regularities in the environment that adults miss.

However, in Study 2, as part of our preregistered inclusion criterion, we only analyzed data from children who learned a rule that was consistent with the data they observed, to ensure that children had attended to and remembered the data. Many children were excluded as a result, and this makes it difficult to compare children’s performance in Study 2 to adults’ performance in Study 1. In addition, when these participants were included in analyses, the difference between younger and older children in test phase performance was not significant, although the other results were similar (see Supplementary Results 1). In Study 3, we remove this restriction to conduct a stronger test of developmental differences in exploration and learning.

4. Study 3

In preregistered Study 3 (preregistration: <https://aspredicted.org/yu7mf.pdf>), we attempt to replicate the results of Studies 1–2, testing whether there are age-related changes in exploration and learning. We equate all procedures and exclusion criteria across ages, allowing us to compare children and adults directly. In addition to testing whether children explore more than adults and ultimately learn more accurately, we also test whether exploration and subsequent learning are associated. In other words, are the participants who explore more, regardless of age, also more likely to learn the correct two-dimensional rule?

In addition, we test the inferential account by examining whether children draw different initial inferences than adults. Such inferences are likely to inform exploration: for example, if an adult observes that a

white striped object is not a zaff and infers that all white objects are not zaffs, the adult is likely to avoid a white spotted object. Children might draw a different inference from this early evidence, inferring that a white spotted object is a zaff (or at least could be a zaff) and therefore choosing to explore it. If this is the case, age-related changes in exploration could be explained by age-related changes in initial inferences.

Notably, these inferences might change across development for several reasons. One possibility is that children have different prior beliefs than adults do. For example, adults might have a strong initial belief that color information is likely to be relevant for category membership, while children might be more open to a variety of category rules. If this is the case, the “learning trap” to which adults are susceptible might be better characterized as a “learned trap,” in that it reflects prior beliefs. Another possibility is that adults and children use different hypothesis selection processes (perhaps despite similar prior beliefs). For example, children might be more likely to favor unlikely hypotheses than adults due to “high temperature search” through the hypothesis space (see Gopnik et al., 2015). We do not differentiate between these possibilities in the present study. Instead, we focus on whether children’s and adults’ initial inferences are, in fact, different.

To test this, we ask children (ages 4–7 years) and adults to complete the same task as in Studies 1 and 2, with one exception: for the first four trials, participants were asked to guess the category of each object (zaff or non-zaff) after they chose to approach or avoid the object but before they saw the outcome. Because these are the first four trials of the task, each guess is made in the absence of any direct evidence about the particular object type the participant is considering. As a result, these guesses reflect the initial inferences participants draw about each object type, before they have any direct evidence about whether that object type is a zaff. For example, a participant might see that a white striped object is a non-zaff and a black striped object is a zaff, and then be asked to guess whether a white spotted object is a zaff. Although the participant has no evidence about white spotted objects, they could use

evidence about the earlier objects to inform their guesses about the current one—for example, they might infer that the white spotted object is a non-zaff, on the basis of its shared color with the white striped object. Age-related differences in these initial guesses would support the hypothesis that adults and children differ in their initial inferences, which might in turn explain the differences in their exploration and subsequent rule learning. For example, individuals who infer that the white spotted object is a zaff might be more willing to try it on the machine than individuals who infer that the white spotted object is not a zaff.

4.1. Method

4.1.1. Participants

Participants were sixty-four 4- to 7-year-olds (16 from each age in years; 35 identified by their parents as male, 29 female) recruited from and tested onsite at museums in the San Francisco Bay Area, and 45 adults (age 23–72; 28 self-identified as male, 17 female) recruited from MTurk. We did not collect other demographic data, but information about the typical demographic makeup of Amazon Mechanical Turk samples is reported in other research (e.g., Levay, Freese, & Druckman, 2016). The museums from which we recruited child participants served families resembling the diversity of the local population. These sample sizes were chosen to match those of Studies 1 and 2 (the target sample size for adults was 32 participants, but fewer participants were excluded from analyses than anticipated). All adult participants were based in the United States and had completed at least 50 assignments on Mturk with at least a 95% approval rating. All participants were tested in January 2019.

An additional eight children were tested but excluded from analyses because of parent interference (one child), experimenter error (five children), zaff machine malfunction (one child), and refusal to make predictions on the first four trials (one child). An additional three adult participants were excluded for indicating that they had not watched all of the videos in the online study. As compensation, adult participants received \$0.50 plus a performance bonus as in Study 1, while child participants received all of the stickers they had earned by the end of the approach-avoid phase.

4.1.2. Procedure

The procedure was identical to that of Study 1 (for adult participants) and Study 2 (for child participants), with the following exceptions. First, children were told that they would get to keep the stickers they earned during the approach-avoid phase, more closely matching the child procedure to the adult procedure. Second, as a measure of initial inferences, all participants were asked to guess whether each of the first four objects was a zaff or not a zaff, after they chose whether or not to put the object on the machine but before they observed the outcome. In other words, on each of the first four trials, participants (1) were presented with an object, (2) decided whether to approach or avoid, (3) guessed whether it was a zaff or not a zaff, then (4) if they decided to approach, observed the outcome of placing the object on the machine. After completing these four steps for a single object, participants proceeded to the next trial. Finally, the second object in the first trial set was always a non-zaff, and the third and fourth objects each matched the non-zaff on a single dimension (pattern or color).

As in the previous studies, most children and adults answered the attention check questions with ease. Among adults, 2% answered the first question incorrectly, 4% answered the second question incorrectly, and 2% answered the third question incorrectly. Among children, 9% answered the first question incorrectly (i.e., were told the correct response after failing to produce the correct response or identify it from two alternatives), 5% answered the second question incorrectly, and 27% answered the third question incorrectly. As in Studies 1–2, those who were unable to generate a correct response were reminded of the rules before proceeding.

4.1.3. Coding

All child data was recorded during each study session by the experimenter or a trained research assistant. 77% of all sessions were video recorded, and these videos were re-coded by a trained research assistant blind to the original data. There was 94% agreement between the original data coding and the video coding. Cases of disagreement were reviewed by the first author, and the data were adjusted in favor of the video when the video provided unambiguous evidence that the original coding was incorrect. All adult data was recorded automatically as participants completed the study online.

4.2. Results

4.2.1. Exploration and learning

First, we compared children's and adults' exploration and learning. We tested whether participants' tendency to approach non-zaffs and zaffs was predicted by age group (4- to 5-year-olds, i.e., "younger children"; 6- to 7-year-olds, i.e., "older children"; adults) and trial set. For non-zaffs, the interaction between age group and trial set was not significant, $\chi^2(2) = 0.84$, $p = .66$. The best-fitting model included fixed effects for both trial set, OR = 0.37, 95% CI [0.26, 0.52], $\chi^2(1) = 46.67$, $p < .001$, and age group, younger vs. adult: OR = 174.66, 95% CI [29.71, 1026.69], Wald $z = 5.71$, $p < .001$, older vs. adult: OR = 6.31, 95% CI [1.59, 24.96], Wald $z = 2.63$, $p = .009$ (overall age effect: $\chi^2(2) = 53.90$, $p < .001$). In sum, approach of non-zaffs decreased across trial sets within all age groups, but children generally approached more non-zaffs than adults (see Fig. 3a). For zaffs, a model including the interaction between age group and trial set outperformed a model including no interaction, $\chi^2(2) = 12.26$, $p = .002$. Fitting separate regression models within each age group revealed that 4- and 5-year-olds were more likely to approach zaffs over time, $b = 0.02$, 95% CI = [0.01, 0.04], $\chi^2(1) = 7.96$, $p = .005$, and this was also the case for 6- to 7-year-olds, $b = 0.02$, 95% CI = [0.002, 0.03], $\chi^2(1) = 5.09$, $p = .02$. However, there was no evidence for change over trial sets in adults' approach of zaffs, $b = -0.01$, 95% CI = [-0.02, 0.003], $\chi^2(1) = 2.02$, $p = .16$. In other words, children became more likely to approach zaffs over trial sets, while adults did not (see Fig. 3a). Across all trial sets, 4- to 5-year-olds approached 89% of zaffs, 5- to 6-year-olds approached 77% of zaffs, and adults approached 70% of zaffs. These results suggest that children explore more than adults. In particular, children and adults both increasingly avoided non-zaffs over trial sets, but children approached non-zaffs more frequently than adults overall. In addition, children increasingly approached zaffs over the course of the task, while adults did not.

To test age-related differences in rule learning, we classified participants based on their responses at test and generalization, as preregistered (see Fig. 3b). For both test and generalization responses, we fit a logistic regression model predicting whether participants learned a one- or two-dimensional rule as a function of age group. At test, there was a significant effect of age, $\chi^2(2) = 17.33$, $p < .001$, with both younger children, OR = 10.37, 95% CI = [3.32, 36.71], Wald $z = 3.85$, $p < .001$, and older children, OR = 3.43, 95% CI = [1.09, 11.62], Wald $z = 2.07$, $p = .04$, reporting the correct two-dimensional rule with higher probability than adults. At generalization, the overall effect of age was not significant, $\chi^2(2) = 3.50$, $p = .17$: the difference between younger children and adults was not significant, OR = 3.72, 95% CI = [0.93, 15.63], Wald $z = 1.86$, $p = .06$, nor was the difference between older children and adults, OR = 1.82, 95% CI = [0.45, 7.44], Wald $z = 0.86$, $p = .39$. In sum, children were more likely to report the correct category rule than adults at test, but not at generalization (see Fig. 3b).

We also conducted exploratory analyses replicating Study 2: do exploration and learning change within childhood? Like Study 2, we found no evidence for an interaction between age (in years) and trial set in predicting approach of non-zaffs, $\chi^2(1) = 2.92$, $p = .09$. In a model excluding the interaction, we found evidence for a main effect of age, OR = 0.10, 95% CI [0.01, 0.22], $\chi^2(1) = 34.30$, $p < .001$, and trial set,

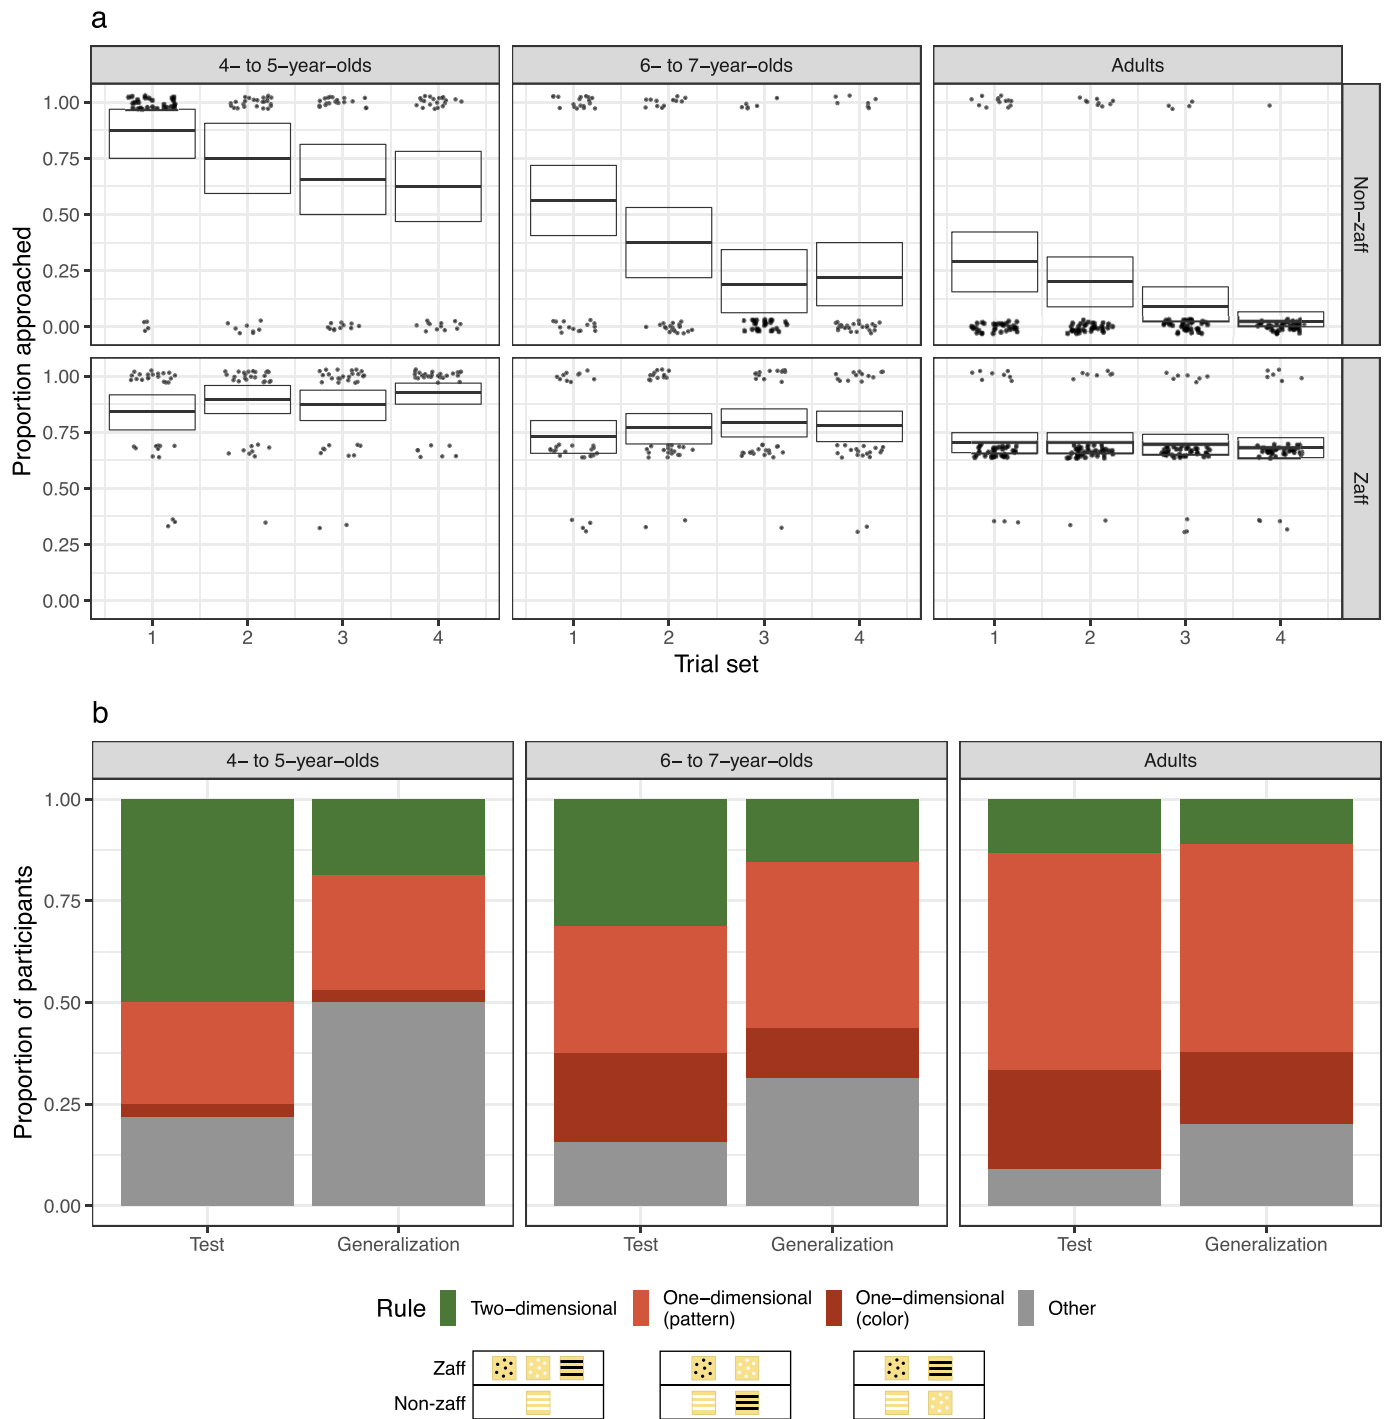


Fig. 3. Children’s and adults’ exploration and learning, Study 3. **a.** Participants’ approach-avoid decisions across four trial sets (each comprised of one non-zaff trial and three zaff trials). Boxplots display mean proportion approached with bootstrapped 95% confidence intervals, and jittered points represent individual participant choices. **b.** Classification of learning at test and generalization. Each participant labelled each object as a zaff or a non-zaff; responses were classified according to whether they followed a two-dimensional or one-dimensional rule. Beneath each rule classification, the relevant pattern of test responses (where white/striped is the correct two-dimensional rule) is shown.

OR = 0.34, 95% CI [0.17, 0.52], $\chi^2(1) = 31.04, p < .001$. As found in Study 2, children increasingly avoided non-zaffs across trial sets and with increasing age. Unlike Study 2, there was no evidence for a significant interaction between age and trial set in predicting approach of zaffs, $\chi^2(1) = 0.40, p = .53$. In a model excluding the interaction, we found evidence for a main effect of age, $b = -0.07, 95\% \text{ CI } [-0.10, -0.03], \chi^2(1) = 12.48, p < .001$, and trial set, $b = 0.02, 95\% \text{ CI } [0.01, 0.03], \chi^2(1) = 12.93, p < .001$. In other words, younger children were

more likely to approach zaffs than older children, and children were increasingly likely to approach zaffs over the course of the task. Finally, at test, there was a significant effect of age on rule learning, OR = 0.47, 95% CI [0.25, 0.81], $\chi^2(1) = 7.67, p = .006$, and at generalization, the effect of age was not significant, OR = 0.60, 95% CI [0.27, 1.20], $\chi^2(1) = 2.10, p = .15$. With increasing age, children became less likely to respond according to a two-dimensional rule at test.

4.2.2. Reward

Although children were more likely to learn the correct two-dimensional rule than adults, it is an open question whether children capitalized on their learning to gain rewards. While accurate learning of the two-dimensional rule would enable an individual to earn up to 16 stickers/stars in the approach-avoid task, this would require both (1) very rapid learning (this assumes that none of the non-zaffs are approached), and (2) the ability and motivation to exploit one's learning to seek rewards and avoid costs. In fact, children did not outperform adults: on average, 4- to 5-year-olds earned 8.75 stickers, 6- to 7-year-olds earned 10.53 stickers, and adults earned 11.16 stars. In a model predicting rewards received, there was a significant effect of age group, $F(2, 106) = 13.12, p < .001$. While the reward earned by 6- to 7-year-olds was not significantly different from that earned by adults, $b = -0.62, 95\% \text{ CI} [-1.57, 0.32], t(106) = -1.31, p = .19$, 4- to 5-year-olds earned significantly fewer rewards relative to adults, $b = -2.41, 95\% \text{ CI} [-3.35, -1.46], t(106) = -5.05, p < .001$, despite being more likely to learn the correct two-dimensional rule. This suggests that young children in particular might learn more slowly than adults and/or lack the ability or motivation to exploit the information they gain during exploration, at least in the context of this task.

4.2.3. Differences in early inferences

Next, we tested whether age-related differences in exploration may be driven by different patterns of early inference. In particular, we focused on the predictions participants made on the third and fourth trials, after encountering the first zaff on trial one and the first non-zaff on trial two. On these early trials, participants could make a one-dimensional inference, assuming that a single feature distinguished the zaffs and non-zaffs, or they could make the correct two-dimensional inference that more than one feature was involved. At this point, the data they had obtained on the first two trials would not distinguish between these two possibilities.

It should be noted that not all participants approached the first two objects. On the first trial, 88% of 4- to 5-year-olds, 88% of 6- to 7-year-olds, and 93% of adults chose to approach. On the second trial, 88% of 4- to 5-year-olds, 56% of 6- to 7-year-olds, and 29% of adults chose to approach. Therefore, children's and adults' predictions were based on different evidence even on these early trials of the task. In Study 4, we further investigate children's and adults' inferences with matched evidence. However, importantly, neither common pattern of exploration on the first two trials—approaching both objects, or approaching the first and avoiding the second—differentiates between a one- and two-dimensional rule. Again, if children explore more than adults because they have different initial beliefs or hypothesis selection processes, we would expect children and adults to make different predictions on these early trials.

In fact, there was no evidence for an age-related difference in predictions on these trials between younger children and adults, $b = 1.31, 95\% \text{ CI} [0.69, 2.51], \text{Wald } z = 0.82, p = .41$, or older children and adults, $b = 1.08, 95\% \text{ CI} [0.57, 2.07], \text{Wald } z = 0.25, p = .81$ (overall age effect: $\chi^2(2) = 0.69, p = .71$). Inspection of the data showed that participants across ages were more likely to predict that an object was a zaff on trials one and three and to predict that an object was a non-zaff on trials two

and four,¹ so trial was treated as a categorical variable for the following exploratory analysis. In a regression model predicting responses across all four prediction trials, there was no significant age by trial interaction, $\chi^2(6) = 12.12, p = .06$, and there was no significant effect of age controlling for the effect of trial, $\chi^2(2) = 3.28, p = .19$ (see Supplementary Results 2 for a preregistered comparison of predictions against chance). These analyses provide no evidence that children and adults differ in their initial inferences (Fig. 4a).

However, despite similar expectations about each object, there were differences in the association between expectations and approach-avoid decisions across ages (Fig. 4b). In an exploratory analysis, we fit a mixed-effects logistic regression model to approach decisions across the first four trials. There was a significant interaction between age group and prediction, $\chi^2(2) = 15.18, p < .001$. The model coefficients, translated from log odds ratios to probabilities, revealed that younger children, older children, and adults almost always approached objects they expected to be zaffs (with 99.0%, 98.6%, and 99.7% predicted probabilities, respectively). However, 4- to 5-year-old children also approached objects they expected to be non-zaffs with 75.7% predicted probability, 6- to 7-year-olds did so with 28.5% predicted probability, and adults did so with only 6.3% predicted probability, and these differences were significant. Note also that in the attention checks, 95% of the children correctly answered that placing a non-zaff on the machine would result in losing two stickers, demonstrating that they understood that non-zaffs were costly. In other words, adults consistently avoided objects they expected to be non-zaffs, whereas children, particularly younger children, frequently approached objects that they expected to be non-zaffs, even though they understood that this action would be costly.

Together, these results suggest that children were more motivated than adults to explore despite the fact that they drew similar inferences about what the objects would do. Children were more willing than adults to try an object that they thought was a non-zaff, and would therefore lead to a cost. They were willing to lose stickers to gain information.

4.2.4. Is exploration associated with rule learning?

Finally, we tested the association between exploration and learning the correct category rule. We fit a model predicting test performance (one- or two-dimensional rule, classified as preregistered using the response patterns described previously), with age and number of approached objects as predictors. There was no significant interaction between age and number of approached objects, $\chi^2(2) = 2.93, p = .23$, nor was there evidence for an independent influence of age group controlling for the number of approached objects, $\chi^2(2) = 0.61, p = .74$. In the best-fitting model, including only number of approached objects as a predictor, there was a strong positive association between this variable and test performance, $\text{OR} = 2.16, 95\% \text{ CI} [1.70, 2.97], \chi^2(1) = 66.42, p < .001$. The more objects a participant of any age approached, the more likely that participant was to learn the two-dimensional rule. In an additional exploratory analysis, we fit a model predicting rule learning at test with four predictors that captured the number of times a given participant placed each of the four object types on the machine. The four object types were coded as the non-zaff, the zaff that did not match the

¹ One concern on the basis of this data is that most children and adults merely alternate in their responses, regardless of which objects they encounter. This does not appear to be the case: only 41% of 4- to 5-year-olds, 28% of 5- to 6-year-olds, and 51% of adults alternated between guessing zaff and non-zaff across the four prediction trials. In addition, only 13% of 4- to 5-year-olds and 34% of 5- to 6-year-olds alternated between guessing zaff and non-zaff across the four test trials (this data is not available for adults because the randomized order of objects in the test phase was not recorded). Together, this suggests that children and adults do not use a simple alternation strategy when predicting objects' category labels.

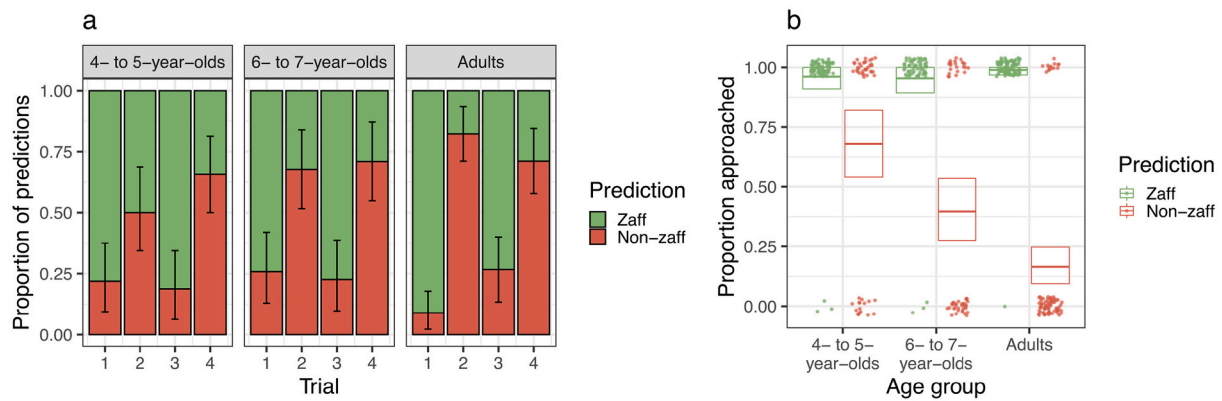


Fig. 4. Predictions and approach decisions. a. Younger children's, older children's, and adults' predictions, with bootstrapped 95% confidence intervals, in the first four trials of Study 3. Trials one, three, and four were always zaffs, and trial two was a non-zaff. b. Younger children's, older children's, and adults' decisions to approach as a function of their predictions in the first four trials. Boxplots display mean proportion approached with bootstrapped 95% confidence intervals, and jittered points represent individual participant choices (with each participant represented four times, corresponding to the four prediction trials).

non-zaff in either color or pattern, the zaff that matched the non-zaff in pattern, and the zaff that matched the non-zaff in color. Approaching the first two types of objects would not allow one to differentiate between the one-dimensional and two-dimensional rules. However, approaching the last two types of objects would generate the critical data that was needed to escape the learning trap. In fact, only approach of those critical objects predicted rule learning; pattern-match zaff: OR = 4.48, 95% CI [2.41, 10.89], $\chi^2(1) = 32.63$, $p < .001$, color-match zaff: OR = 4.03, 95% CI [2.08, 10.56], $\chi^2(1) = 23.49$, $p < .001$, other $ps > .05$.

4.3. Discussion

In Study 3, we replicated the main results of Studies 1–2: 4- to 5-year-olds and 6- to 7-year-olds were more likely than adults to approach both zaffs and non-zaffs. Moreover, children in both age groups were more likely than adults to respond according to the correct two-dimensional rule at test, indicating that they did not fall prey to the learning trap. However, there was no evidence that children were more likely than adults to respond according to the correct two-dimensional rule at generalization. This was consistent with the mixed and weak results on the generalization task in Study 2: children did not appear to generalize the rule they had learned to new objects that were a different color.

Study 3 also tested whether children and adults draw different inferences at an early stage of learning, as one possible explanation for the age-related changes in exploration. However, there was no evidence that children and adults made different inferences about the objects in the third and fourth trials of the approach-avoid task—the objects that could disambiguate the one- and two-dimensional rules. Instead, both younger and older children were more willing to approach objects than adults, even when they predicted that they would incur a cost by doing so.

Finally, we demonstrated in Study 3 that exploration was associated with learning. The more objects a participant approached in the approach-avoid phase (especially objects that matched the non-zaff in pattern or color), the more likely the participant was to respond according to the correct two-dimensional rule at test. Moreover, age group did not predict test performance above and beyond the number of objects approached.

All this suggests that the differences in learning between adults and children are more due to motivational than inferential differences, and that they are the result of the different patterns of data that children and adults generate. However, it is also possible that children and adults explore differently and also, independently, infer different rules in the test phase. Perhaps young children simply prefer two-dimensional rules or adults simply prefer one-dimensional rules, explaining both their patterns of exploration (more objects approached by children than

adults) and rule learning (responding according to a two-dimensional rule by children and a one-dimensional rule by adults). The similarity in the early inferences of children and adults weighs against this idea, but the clearest test of whether inference or motivation explains the developmental differences in exploration and subsequent learning would be to give children the kind of data generated by adults, give adults the kind of data generated by children, and see which rules they learn.

5. Study 4

In preregistered Study 4a (preregistration: <https://aspredicted.org/4wb49.pdf>) and 4b (preregistration: <https://aspredicted.org/xe52g.pdf>), we tested experimentally whether 4- to 5-year-olds' and adults' rule learning is affected by the evidence they observe.² That is, if children observed evidence similar to the evidence an adult would generate during the approach-avoid task (i.e., if they only tested three rather than four object types on the machine), would they be more likely to learn the one-dimensional rule? Similarly, if adults observed evidence similar to that a 4- to 5-year-old would generate during the approach-avoid task (i.e., if they tested all four object types on the machine), would they be more likely to learn the two-dimensional rule? This would strongly support the idea that differences in rule learning are due to the fact that children and adults generate different kinds of evidence, rather than because children prefer different rules than adults or make different inferences from the same evidence.

Participants either observed “child-like evidence”—testing all four object types on the machine—or “adult-like evidence”—testing three of the four object types on the machine. In the child-like evidence condition, the observations were consistent only with a two-dimensional rule; in the adult-like evidence condition, the observations were ambiguous between a one-dimensional and a two-dimensional rule. In the test phase, participants judged whether each object in a new set of four was a zaff, with a single trial for each object type.

In the previous studies, adults typically received evidence about two or three object types across the full approach-avoid phase (i.e., they approached two or three object types at least once across the four trial sets). In Study 1, adults received evidence about 2.5 object types on average, while in Study 3, adults received evidence about 2.69 object types on average. For the purpose of this study, we define “adult-like

² Studies 4a and 4b were conducted and preregistered separately, with Study 4a (children) preceding Study 4b (adults). As a result, we conduct all analyses within each age group. We also report analyses preregistered as part of Study 4b, comparing adults' data from Study 4b to children's data from Study 4a.

evidence” as approaching *three* out of the four objects, providing participants with the maximal amount of evidence that is consistent with both a two-dimensional rule and a one-dimensional rule.

Study 4 also examined whether children have the capacity to exploit the information they learn, and how their exploitation compares to that of adults. In reinforcement learning tasks, early exploration is useful because it leads to more effective subsequent exploitation. However, 4- to 5-year-olds earned fewer rewards than adults in Study 3, so it is unclear whether children can exploit the accurate rule they learn to gain reward. In Study 4, after assessing what participants learned from adult-like or child-like evidence, we both told participants and demonstrated that the one-dimensional rule was correct (in the adult-like evidence condition) or that the two-dimensional rule was correct (in the child-like evidence condition), to ensure that all participants learned a particular rule. Then, participants completed four approach-avoid trials as in Studies 1–3. If children and adults have the capacity to exploit the information they learned, we would expect their approach-avoid decisions to differ depending on the rule they had learned.

5.1. Method

5.1.1. Participants

Participants in Study 4a were sixty-four 4- to 5-year-olds (32 from each age in years; 31 identified by their parents as male, 33 female) recruited from and tested onsite at museums and preschools in the San Francisco Bay Area, in January 2020. As in the previous studies, additional demographic information was not collected. Participants were randomly assigned to the child-like evidence condition or the adult-like evidence condition, with equal numbers of 4-year-olds and 5-year-olds assigned to each condition. The sample size was chosen so that the number of participants in each condition matched the number of 4- to 5-year-olds in Studies 2 and 3. An additional four children were tested but excluded from analyses because their date of birth did not match the age group in which they were tested.

Participants in Study 4b were 70 adults (age 19–71; 27 identified as male, 41 female, and 2 other) recruited from Prolific in June 2021. More information about the demographic characteristics on participants on Prolific is reported in prior research (Peer, Brandimarte, Samat, & Acquisti, 2017). In order to participate, participants were required to reside in the US and to have completed at least 50 previous studies on Prolific with a minimum 95% approval rate. Participants were randomly assigned to the child-like evidence condition or the adult-like evidence condition, with equal numbers of adults assigned to each condition. The sample size was chosen so that the number of participants matched the number of 4- to 5-year-olds in Study 4a. An additional two adults were tested but excluded from analyses because they indicated that they did not watch all of the videos in the study. Adults were also required to pass an attention check to be included in our analyses, but all adults successfully passed the attention check.

As compensation, adult participants received \$1.00 plus a performance bonus, while child participants received the stickers they had earned by the end of the approach-avoid phase.

5.1.2. Procedure

We begin by describing the procedure used with children (Study 4a), who were recruited and tested in person. Following this description, we note all departures from this procedure with adults (Study 4b), who were recruited and tested online.

First, in the instructions phase, children were shown a box with objects inside (all but three of which were concealed—the first three presented in the learning phase) and were told that some of the objects were “zaffs” and some of the objects were “not zaffs.” As in Studies 1–3, children were told the effect of each type of object when placed on the “zaff machine.” Children were given four small stickers to start and were told they would gain one star for every zaff placed on the machine and lose two stars for every non-zaff placed on the machine. Participants

were asked two attention check questions: “What happens if you put a block on the machine and it is a zaff?” and “What happens if you put a block on the machine and it is not a zaff?” Among children, 21% were told the answer to the first question after providing a wrong answer or providing no response, and 6% were told the answer to the second question after providing a wrong answer or providing no response. Among adults, 1% were told the answer to the first question after providing a wrong answer, and 4% were told the answer to the second question after providing a wrong answer.

In the learning phase, children were given one object at a time to test on the zaff machine. Unlike in Studies 1–3, children were not given the choice of approaching or avoiding each object. In the child-like evidence condition, children first saw one object of each type (white stripes, white spots, black stripes, black spots), then a second set of objects of each type. In the adult-like evidence condition, children first saw two sets of only three of the four objects, where the unseen object shared one dimension (color or pattern) with the non-zaff. In the first set, participants always observed the zaff that shared neither color nor pattern with the non-zaff first, and the non-zaff second. All remaining objects were presented in a random order within each set.

Next, in the test phase, children saw a new set of all four object types and were asked the identity (zaff or non-zaff) of each object. The order in which these objects were presented was randomized.

In the teaching phase, children were provided with feedback on their responses in the test phase. In the child-like evidence condition, children were told that in fact, the previously tested zaffs were all zaffs, while the previously tested non-zaff was a non-zaff (two-dimensional rule). In the adult-like evidence condition, children were told that in fact, the previously tested zaffs were zaffs, while the previously tested non-zaff and the previously untested object were both non-zaffs (one-dimensional rule). In both conditions, children were also provided the relevant general rule: in the child-like evidence condition, for example, “So blocks with a white striped pattern are not zaffs, and all the other blocks are zaffs,” and in the adult-like evidence condition, for example “So blocks with a striped pattern are not zaffs, and all the other blocks are zaffs.” Then, the experimenter placed each object from the test phase on the machine (with zaffs followed by non-zaffs) to demonstrate.

In the second instructions phase, children were told that they would be shown more objects, but this time they could choose whether to put each object on the machine (approach) or to not put it on the machine (avoid). Instructions regarding the reward associated with zaffs and non-zaffs were repeated. All children were given four new stickers to start, and they were told they would keep the stickers they had remaining at the end of this final task.

Finally, in the approach-avoid phase, children were presented with a final set of four objects, one at a time. The non-zaffs were always first or second to be presented in the child-like evidence condition; or first and third or second and fourth in the adult-like evidence condition. The position of the remaining objects was randomized. The procedure in this phase was otherwise identical to the approach-avoid phase in Study 2.

The adult procedure was nearly identical, with the following exceptions. Similar to Study 1, adults received all instruction through videos narrated by the primary experimenter (and first author). In the instructions phase, all objects were concealed inside the bin. As in Studies 1 and 3, adult participants were not provided with memory aids throughout the task. Participants earned “stars” instead of stickers, and the stars earned in the approach-avoid phase were converted to a monetary bonus at the end of the study (two cents per star). Finally, the four objects were presented in a randomized order during the approach-avoid phase.

5.1.3. Coding

All child data was recorded during each study session by the experimenter. 70% of all sessions were video recorded, and these videos were re-coded by a trained research assistant blind to the original data. There was 96% agreement between the original data coding and the video

coding. Cases of disagreement were reviewed by the first author, and the data were adjusted in favor of the video when the video provided unambiguous evidence that the original coding was incorrect. All adult data was recorded automatically as participants completed the study online.

5.2. Study 4a results: Children

5.2.1. Does the evidence children receive affect learning?

In the child-like evidence condition, 59% of children responded according to the correct two-dimensional rule, classifying only the true non-zaff as a non-zaff. An additional 9% perfectly followed a one-dimensional pattern rule, classifying as non-zaffs only the true non-zaff and the object that matched it in pattern. An additional 9% perfectly followed a one-dimensional color rule, classifying as non-zaffs only the true non-zaff and the object that matched it in color. The remaining 22% of participants responded according to one of seven unique additional patterns (percentages do not add up to 100 due to rounding). In the adult-like evidence condition, only 16% of children responded according to the correct two-dimensional rule. Instead, 34% perfectly followed a one-dimensional pattern rule, and 28% perfectly followed a one-dimensional color rule. The final 22% of participants responded according to one of six unique additional patterns. As predicted, children in the adult-like evidence condition were significantly more likely to respond according to a one-dimensional rule (based on either pattern or color) over a two-dimensional rule, relative to children in the child-like evidence condition (Fig. 5a), $OR = 12.67$, 95% CI [3.55, 53.65], $\chi^2(1) = 16.66$, $p < .001$. Notably, the evidence children received in the adult-like evidence condition was ambiguous between a one-dimensional and a two-dimensional rule, but children learned a one-dimensional rule more frequently than a two-dimensional rule (Fig. 5a).

Children's performance in the child-like evidence condition was not significantly different from children's performance in Study 3 (see Supplementary Results 3), providing no evidence that the procedural differences across studies impacted performance. Additionally, children's performance in the adult-like evidence condition was not significantly different from adults' performance in Study 1 (see Supplementary Results 3). This suggests that children's learning depends on the evidence they observed—if children had explored like adults, they would have learned inaccurately like adults.

5.2.2. Do children "exploit"?

Next, we tested whether children can exploit the information they learn. Children avoided significantly more objects in the adult-like evidence condition, in which two objects were non-zaffs ($M = 0.94$), compared to the child-like evidence condition, in which only one object was a non-zaff ($M = 0.34$), $b = 0.59$, 95% CI [0.23, 0.96], $t(62) = 3.25$, $p = .002$. Across conditions, children avoided 0% of zaffs and 43% of non-zaffs (Fig. 5b; see Supplementary Results 4 for further analyses). Performance in the child-like evidence condition was not significantly different from the last four trials in Study 2 (see Supplementary Results 4), providing no evidence that the procedural differences across studies affected performance. Children exploited less than was optimal, because they continued to approach non-zaffs, but their exploitation was sensitive to the structure of the environment and the rule they had learned.

5.3. Study 4b results: Adults

5.3.1. Does the evidence adults receive affect learning?

In the child-like evidence condition, 43% of adults responded according to the correct two-dimensional rule, 6% perfectly followed a one-dimensional pattern rule, 23% perfectly followed a one-dimensional color rule, and the remaining 29% of participants responded according to one of seven unique additional patterns (percentages do not add up to 100 due to rounding). In the adult-like evidence condition, only 6% of adults responded according to the correct two-dimensional rule. Instead, 54% perfectly followed a one-dimensional pattern rule, and 23% perfectly followed a one-dimensional color rule. The final 17% of participants responded according to one of four unique additional patterns. As predicted, adults in the adult-like evidence condition were significantly more likely to respond according to a one-dimensional rule (based on either pattern or color) over a two-dimensional rule, relative to adults in the child-like evidence condition (Fig. 5a), $OR = 20.25$, 95% CI [4.66, 144.29], $\chi^2(1) = 19.07$, $p < .001$.

5.3.2. Do adults "exploit"?

Next, we tested whether adults, like children, exploit the information they learn. Adults avoided significantly more objects in the adult-like evidence condition, in which two objects were non-zaffs ($M = 1.86$), compared to the child-like evidence condition, in which only one object was a non-zaff ($M = 1.31$), $b = 0.54$, 95% CI [0.24, 0.85], $t(68) = 3.55$, $p < .001$. Across conditions, adults avoided 7% of zaffs and 93% of non-

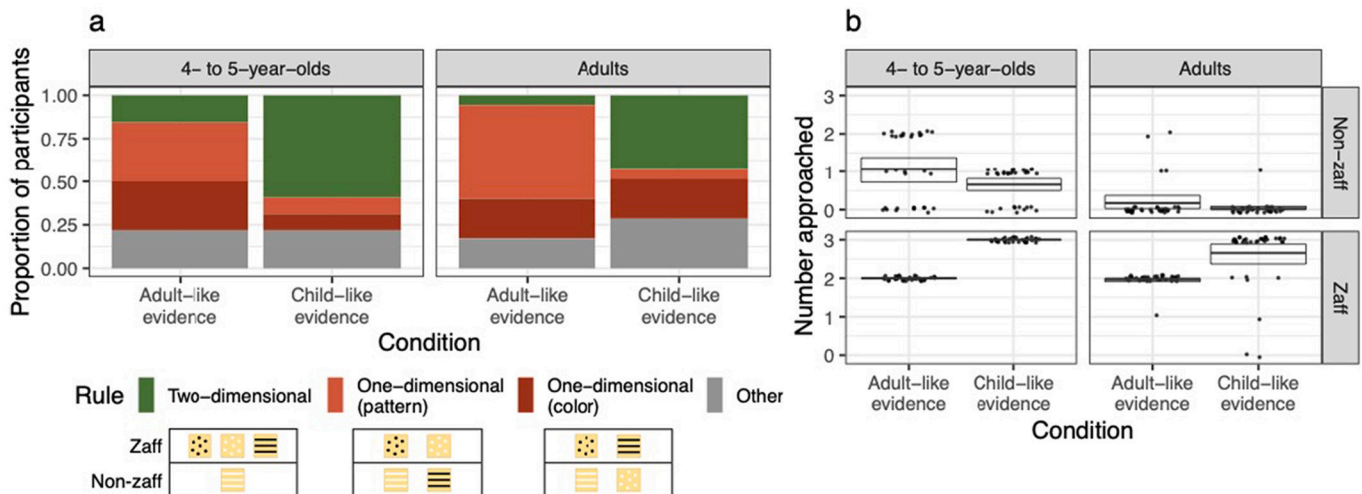


Fig. 5. Learning and approach-avoid decisions by condition, Study 4. a. Classification of learned rules by 4- and 5-year-olds and adults in Study 4, as a function of whether they received adult-like evidence or child-like evidence. b. Approach-avoid decisions of 4- to 5-year-olds and adults in Study 4, after learning a one-dimensional rule based on adult-like evidence or a two-dimensional rule based on child-like evidence. Boxplots display mean number of objects approached with 95% bootstrap confidence intervals, and jittered points represent individual participant choices (with each participant represented once in the left panel for non-zaff trials and once in the right panel for zaff trials).

zaffs (Fig. 5b; see Supplementary Results 4 for further analyses).

5.3.3. Age-related changes

Finally, we compared adults' learning and exploitation to children's. First, we fit a logistic regression model predicting whether participants responded according to a one-dimensional rule or two-dimensional rule as a function of age group (adults vs. children), condition (adult-like evidence vs. child-like evidence), and their interaction. There was no evidence for a significant interaction between age group and condition, $\chi^2(1) = 0.19, p = .66$. In other words, the evidence did not support the possibility that adults' and children's learning was differentially affected by received evidence. In addition, we conducted an exploratory analysis testing whether responding at test was associated with age when averaged across the two conditions. If children are biased towards learning two-dimensional rules or if adults are biased towards learning one-dimensional rules (even holding evidence constant), we would expect age group to predict rule learning. In fact, there was no evidence for a significant effect of age, $OR = 0.50, 95\% CI [0.22, 1.10], \chi^2(1) = 2.98, p = .08$. Thus, the evidence did not support the hypothesis that either age group was biased towards learning a particular rule.

Next, we investigated age-related differences in exploitation. We fit a linear regression model predicting the number of objects a participant avoided during the approach-avoid phase, with age group, condition, and their interaction as predictors. There was no evidence for an interaction between age group and condition, $F(1, 130) = 0.05, p = .83$. Again, we conducted an exploratory analysis testing whether the number of avoided objects was associated with age when averaged across the two conditions. Indeed, there was a significant effect of age: children avoided fewer objects than did adults, $b = -0.95, 95\% CI [-1.20, -0.69], t(132) = -7.41, p < .001$. Thus, though both adults and children were selective in primarily avoiding non-zaffs, adults avoided more objects than children across experimental conditions.

Finally, we conducted exploratory analyses examining the rewards received by children and adults during the approach-avoid phase. Children earned an average of 3.88 stickers in the adult-like evidence condition and 5.69 stickers in the child-like evidence condition. Adults earned an average of 5.63 stickers in the adult-like evidence condition and 6.60 stickers in the child-like evidence condition. Mirroring the results above, there was no evidence for an interaction between age group and condition, $F(1,130) = 3.75, p = .055$. However, there was a significant effect of age on reward when averaging across the two conditions, $b = -1.33, 95\% CI [-1.83, -0.84], t(132) = -5.35, p < .001$. Thus, children earned fewer rewards than adults, suggesting a lower capacity or motivation to exploit.

5.4. Discussion

In Study 4, we addressed three questions. First, does the information children are exposed to during exploration cause them to avoid the learning trap, or do children simply have an initial preference for the correct two-dimensional rule? Second, does the information adults are exposed to during exploration cause them to fall into the learning trap, or do adults simply have an initial preference for the incorrect one-dimensional rule? And finally, to what extent are children capable of exploiting the information they've learned, relative to adults?

Answering the first question, we found that children who were exposed to adult-like evidence (i.e., information that is ambiguous between a one- and two-dimensional rule) were more likely to learn a one-dimensional rule than were children who were exposed to child-like evidence (i.e., evidence that is only consistent with a two-dimensional rule). That is, young children are not generally biased to infer the correct two-dimensional rule—instead, they only learn this rule when the evidence supports it. In contrast, when young children receive evidence that is ambiguous, they tend to learn a one-dimensional rule, like adults.

Answering the second question, we found similar results in adults. Adults who were exposed to child-like evidence (i.e., full information

about all four object types) were more likely to learn the correct two-dimensional rule than a one-dimensional rule. Moreover, there was no evidence that adults were more likely to respond according to a one-dimensional rule than children, averaged across evidence conditions. Therefore, adults' susceptibility to the learning trap appears to be a result of the evidence they generate during exploration. This is consistent with prior research (Rich & Gureckis, 2018). Taken together, the answers to these first two questions provide support for the motivational account—that children are more exploratory than adults because they are more motivated to gain information when facing costs—rather than the inferential account.

As mentioned earlier, many adults in Studies 1 and 3 in fact approached only two object types and avoided the other two (both the non-zaff and one zaff) across trial sets. It remains an open question whether children would also infer a one-dimensional rule if they only saw evidence about two objects, rather than three; further research is needed to test whether this is the case.

Answering the final question, we found that children did exploit the information they learned, though imperfectly. Children were more likely to avoid non-zaffs than zaffs, and children who learned a one-dimensional rule (which classified two objects as non-zaffs) avoided more objects than children who learned a two-dimensional rule (which classified only one object as a non-zaff). However, adults avoided more objects than children and earned more rewards. In sum, children do exploit the information they have gathered to moderate their approach-avoid decision making, though adults do so more reliably.

6. General discussion

In four studies, we tested how exploration changes across the lifespan and whether changes in exploration are related to learning. Replicating prior work (Rich & Gureckis, 2018), we found that adults under-explore following negative outcomes (Study 1), which prevents them from learning the true structure of the environment. We also found that children explore more, allowing them to reach final conclusions that are more accurate than those reached by adults (Studies 2–3). Finally, these developmental changes in rule learning seem to be caused primarily by differences in the evidence that adults and children generate during exploration, rather than differences in how adults and children make inferences from that evidence (Studies 3–4). In particular, we found no evidence that children made different predictions than adults in the initial trials of the approach-avoid task (Study 3), indicating that children's broad exploration is unlikely to stem from differences in their initial inferences or prior biases. Instead, children were more likely than adults to test an object on the machine even when they predicted that the object was a non-zaff and therefore costly. In addition, when 4- to 5-year-olds were presented with the evidence adults would have generated during exploration, they inferred the same rule that adults inferred, and when adults were presented with the evidence 4- to 5-year-olds would have generated during exploration, they inferred the same rule children inferred (Study 4).

Together, these results provide evidence that children explore more than adults and that children's exploration can have learning benefits. Children and adults do not appear to explore and learn differently in these tasks because they make different inferences, have different initial beliefs, have different procedures for hypothesis selection, or prefer different types of rules, although this may indeed be true in other contexts. Instead, the difference seems to be motivational. Unlike adults, children are motivated to explore despite the costs. As a result, they generate more extensive data and learn from that data accurately.

These findings leave open several important questions about why and how children explore more than adults. In the reinforcement learning literature, there is a distinction between "directed exploration" and "random exploration" (Gershman, 2018, 2019; Wilson, Geana, White, Ludvig, & Cohen, 2014). In directed exploration, the learner chooses to explore just those options that will be more informative,

whereas in random exploration, the learner simply chooses to explore more options overall. One possibility is that children in the present research were more uncertain about the structure of the task than adults and tried to reduce that uncertainty by exploring. Even though their initial predictions appear to be similar, children's predictions in Study 3 might represent their "best guess," while adults might be strongly committed to their predictions. If children are more uncertain than adults during reinforcement learning, their broader exploration might reflect "directed exploration": exploration specifically aimed towards reducing uncertainty. Consistent with this possibility, E. Schulz et al. (2019) found that changes in exploration between 7 and 11 years and adulthood can be explained in part by a developmental decrease in directed exploration.

Another possibility is that children explore in a more random way. They might be more impulsive, acting in a way that will generate effects in the world regardless of their informativeness for a particular task. This is unlikely to be the full story, however. Earlier studies have shown that children's exploration is rational and reflects their understanding (Bonawitz et al., 2012; Cook et al., 2011; Legare, 2012; Mills, Sands, Rowles, & Campbell, 2019; L. E. Schulz & Bonawitz, 2007; Sim & Xu, 2017; Stahl & Feigenson, 2015), and developmental changes in exploration between 7 and 11 years and adulthood are not explained by changes in random exploration (E. Schulz et al., 2019). Several features of the current results also suggest that children's exploration is not purely random or impulsive. Children did not act indiscriminately to generate effects—in other words, they did not approach all objects. Instead, children systematically adjusted their approach-avoid behavior across trials (Studies 2–3), increasingly avoiding the non-zaffs but not the zaffs. In Study 4, they approached objects differently depending on the rules they had learned. Together, this suggests that children's exploration is not solely the result of noise or impulsiveness (see also Blanco and Sloutsky, 2020, 2021; Sumner, Steyvers, et al., 2019).

On the other hand, children still often approached non-zaffs, despite what they had learned. Therefore, children's desire to generate effects in the world might play a role in explaining their broad exploration—and perhaps their ability to avoid the learning trap. Moreover, a general motivation to intervene in the world and bring about effects might play a role in the wide-ranging causal learning that is a feature of early childhood, given the importance of intervention in causal learning and understanding. Determining the balance between directed and random exploration in children, as well as charting developmental changes in that balance, are particularly important directions for future research (see E. Schulz et al., 2019).

Finally, it remains an open question why children failed to maximally exploit their learning, especially in Studies 3–4 where children were made aware that they could keep any stickers they received during the approach-avoid task. In particular, adults earned more reward than children despite less accurate learning. Note that children clearly were sensitive to the reward value of the stickers, increasingly avoiding the objects that would lead to a loss and increasingly approaching those that would lead to a gain. Nevertheless, they failed to optimally exploit their knowledge. One possible, though rather unlikely, explanation is that older children value stickers more than younger children, and adults place still more value on small monetary rewards. If younger children do not value the stickers we used as rewards as much as older children, they would have less incentive to exploit any learned information. A second more plausible possibility is that, as noted above, children experience placing an object on the machine and generating an effect as intrinsically rewarding, enough to partially offset the material cost of approaching a non-zaff. Of course, this in itself would suggest that children are intrinsically motivated to explore. If this is the case, children's exploitation might be improved by reducing the intrinsic reward of approaching and testing an object, or by increasing the cost associated with non-zaffs.

6.1. Limitations

There are several limitations to this work. First, while we found fairly robust age differences in learning between children and adults at test (when participants classified the same objects that they had seen during exploration), age differences were less clear at generalization (when participants classified a new set of objects that varied along the same dimensions). This suggests that there are indeed differences in the way that children and adults make inferences and generalize from data (see E. Schulz et al., 2019). However, future research is needed to determine when children and adults are willing to extend learned rules to new stimuli.

Most significantly, our study only addresses a limited population: US 4- to 7-year-olds and adults. There is evidence for changes in exploration both later in childhood (E. Schulz et al., 2019), in adolescence (Nussenbaum et al., 2020; Somerville et al., 2017), and in older adulthood (Chin et al., 2015; Mata et al., 2013). There is also some evidence that children across cultures may be more flexible learners than adults in causal learning tasks (Wente et al., 2019). However, children's tendency to make risky decisions (e.g., explore when it could be costly) varies across cultures (Amir et al., 2020) and contexts (Tottenham, Shapiro, Flannery, Caldera, & Sullivan, 2019) and depends on early experience (Humphreys et al., 2015). Thus, it will be important for future work to study the connection between exploration and learning in broader populations and contexts.

6.2. Conclusion

Although many questions remain to be answered, the present research supports the claim that children can be more exploratory than adults, explained at least in part by a stronger motivation to explore. Furthermore, we show that this exploratory behavior can make children better learners, particularly in situations where the environment is complex and there are costs and risks. The protected period of immaturity we call childhood may reflect a division of labor that helps human beings to balance exploration and exploitation.

Data availability statement

All data, analysis scripts, and experimental materials are available at <https://osf.io/m3h6u/>.

Declaration of Competing Interest

None.

Acknowledgements

We would like to thank Teresa Garcia, Lena Gavenas, Cortland Jell, Rachel Marcus, Prnsi Patel, Isha Shah, Kaitlan Tseng, and Sarah Yang for assistance with data collection and processing. Some of the results reported here were presented at the 2019 meeting of the Society for Research in Child Development and the 2018 Cognitive Science Society pre-conference workshop on "Understanding Exploration-Exploitation Trade-offs". We are grateful to these audiences for their discussion and feedback. This work was supported by an NSF Graduate Research Fellowship to EGL [grant numbers DGE-1752814 and DGE-1656466], and by grants from the Bezos Family Foundation, the Templeton Foundation and the DARPA Machine Common Sense program to AG. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funders. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2021.104940>.

References

- Amir, D., Jordan, M. R., McAuliffe, K., Valeggia, C. R., Sugiyama, L. S., Bribiescas, R. G., ... Dunham, Y. (2020). The developmental origins of risk and time preferences across diverse societies. *Journal of Experimental Psychology: General*, 149(4), 650–661. <https://doi.org/10.1037/xge0000675>
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov), 397–422.
- Blanco, N. J., & Sloutsky, V. M. (2020). Attentional mechanisms drive systematic exploration in young children. *Cognition*, 202, 104327. <https://doi.org/10.1016/j.cognition.2020.104327>
- Blanco, N. J., & Sloutsky, V. (2021). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*, 24(2), e13026. <https://doi.org/10.1111/desc.13026>
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, 64(4), 215–234. <https://doi.org/10.1016/j.cogpsych.2011.12.002>
- Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Conservative forgetful scholars: How people learn causal structure through sequences of interventions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(3), 708–731. <https://doi.org/10.1037/xlm0000061>
- Camerer, C., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4), 325–370. <https://doi.org/10.1007/BF00122575>
- Chin, J., Payne, B. R., Fu, W.-T., Morrow, D. G., & Stine-Morrow, E. A. (2015). Information foraging across the life span: Search and switch in unknown patches. *Topics in Cognitive Science*, 7(3), 428–450. <https://doi.org/10.1111/tops.12147>
- Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, 79, 102–133. <https://doi.org/10.1016/j.cogpsych.2015.02.004>
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(1481), 933–942. <https://doi.org/10.1098/rstb.2007.2098>
- Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3), 341–349. <https://doi.org/10.1016/j.cognition.2011.03.003>
- Deng, W. S., & Sloutsky, V. M. (2016). Selective attention, diffused attention, and the development of categorization. *Cognitive Psychology*, 91, 24–62. <https://doi.org/10.1016/j.cogpsych.2016.09.002>
- Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review*, 112(4), 951–978. <https://doi.org/10.1037/0033-295X.112.4.951>
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12(5), 523–538. <https://doi.org/10.1287/orsc.12.5.523.10092>
- Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127(3), 455–476. <https://doi.org/10.1037/rev0000175>
- Dymond, S., Dunsmoor, J. E., Vervliet, B., Roche, B., & Hermans, D. (2015). Fear generalization in humans: Systematic review and implications for anxiety disorder research. *Behavior Therapy*, 46(5), 561–582. <https://doi.org/10.1016/j.beth.2014.10.001>
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87(3), 293–311. <https://doi.org/10.1037/0022-3514.87.3.293>
- German, T. P., & Defeyter, M. A. (2000). Immunity to functional fixedness in young children. *Psychonomic Bulletin & Review*, 7(4), 707–712. <https://doi.org/10.3758/BF03213010>
- Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, 173, 34–42. <https://doi.org/10.1016/j.cognition.2017.12.014>
- Gershman, S. J. (2019). Uncertainty and exploration. *Decision*, 6(3), 277–286. <https://doi.org/10.1037/dec0000101>
- Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102), 1623–1627. <https://doi.org/10.1126/science.1223416>
- Gopnik, A. (2020). Childhood as a solution to explore-exploit tensions. *Philosophical Transactions of the Royal Society B*, 375, 20190502. <https://doi.org/10.1098/rstb.2019.0502>
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24(2), 87–92. <https://doi.org/10.1177/0963721414556653>
- Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wentz, 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. <https://doi.org/10.1073/pnas.1700811114>
- Gureckis, T. M., & Markant, D. B. (2009). Active learning strategies in a spatial concept learning game. In N. A. Taatgen, & H. van Rijn (Eds.), *Proceedings of the 31st annual conference of the cognitive science society* (pp. 3145–3150). Cognitive Science Society.
- Hanania, R., & Smith, L. B. (2010). Selective attention and attention switching: Towards a unified developmental approach. *Developmental Science*, 13(4), 622–635. <https://doi.org/10.1111/j.1467-7687.2009.00921.x>
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19(1), 46–54. <https://doi.org/10.1016/j.tics.2014.10.004>
- Humphreys, K. L., Lee, S. S., Telzer, E. H., Gabard-Durnam, L. J., Goff, B., Flannery, J., & Tottenham, N. (2015). Exploration-exploitation strategy is dependent on early experience. *Developmental Psychobiology*, 57(3), 313–321. <https://doi.org/10.1002/dev.21293>
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, 20(8), 963–973. <https://doi.org/10.1111/j.1467-9280.2009.02402.x>
- Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*, 5(11), 831–843. <https://doi.org/10.1038/nrn1533>
- Lapidow, E., & Walker, C. M. (2020). Informative experimentation in intuitive science: Children select and learn from their own causal interventions. *Cognition*, 201, 104315. <https://doi.org/10.1016/j.cognition.2020.104315>
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, 83(1), 173–185. <https://doi.org/10.1111/j.1467-8624.2011.01691.x>
- Levy, K. E., Freese, J., & Druckman, J. N. (2016). The demographic and political composition of mechanical Turk samples. *SAGE Open*, 6(1). <https://doi.org/10.1177/2158244016636433>
- Li, R., Roberts, R. C., Huettel, S. A., & Brannon, E. M. (2017). Five-year-olds do not show ambiguity aversion in a risk and ambiguity task with physical objects. *Journal of Experimental Child Psychology*, 159, 319–326. <https://doi.org/10.1016/j.jecp.2017.02.013>
- 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. <https://doi.org/10.1016/j.cognition.2013.12.010>
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87. <https://doi.org/10.1287/orsc.2.1.71>
- Mata, R., Wilke, A., & Czienskowski, U. (2013). Foraging across the life span: Is there a reduction in exploration with aging? *Frontiers in Neuroscience*, 7, 53. <https://doi.org/10.3389/fnins.2013.00053>
- McCrae, R. R., & Costa, P. T. (1997). Conceptions and correlates of openness to experience. In R. Hogan, J. Johnson, & S. Briggs (Eds.), *Handbook of personality psychology* (pp. 825–847). Academic Press. <https://doi.org/10.1016/B978-012134645-4/50032-9>
- Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2(3), 191–215. <https://doi.org/10.1037/dec0000033>
- Mills, C. M., Sands, K. R., Rowles, S. P., & Campbell, I. L. (2019). “I want to know more!”: Children are sensitive to explanation quality when exploring new information. *Cognitive Science*, 43(1), Article e12706. <https://doi.org/10.1111/cogs.12706>
- Nussenbaum, K., Cohen, A. O., Davis, Z. J., Halpern, D. J., Gureckis, T. M., & Hartley, C. A. (2020). Causal information-seeking strategies change across childhood and adolescence. *Cognitive Science*, 44(9), Article e12888. <https://doi.org/10.1111/cogs.12888>
- Nussenbaum, K., & Hartley, C. A. (2019). Reinforcement learning across development: What insights can we draw from a decade of research? *Developmental Cognitive Neuroscience*, 40, 100733. <https://doi.org/10.1016/j.dcn.2019.100733>
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163. <https://doi.org/10.1016/j.jesp.2017.01.006>
- Piaget, J. (1954). *The construction of reality in the child*. Basic Books.
- Plebanek, D. J., & Sloutsky, V. M. (2017). Costs of selective attention: When children notice what adults miss. *Psychological Science*, 28(6), 723–732. <https://doi.org/10.1177/0956797617711664>
- Plude, D. J., Enns, J. T., & Brodeur, D. (1994). The development of selective attention: A life-span overview. *Acta Psychologica*, 86(2–3), 227–272. [https://doi.org/10.1016/0001-6918\(94\)90004-3](https://doi.org/10.1016/0001-6918(94)90004-3)
- Rich, A. S., & Gureckis, T. M. (2018). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General*, 147(11), 1553–1570. <https://doi.org/10.1037/xge0000466>
- Rosenbaum, G. M., & Hartley, C. A. (2018). Developmental perspectives on risky and impulsive choice. *Philosophical Transactions of the Royal Society B*, 374, 20180133. <https://doi.org/10.1098/rstb.2018.0133>
- Rothe, A., Lake, B. M., & Gureckis, T. M. (2018). Do people ask good questions? *Computational Brain & Behavior*, 1(1), 69–89. <https://doi.org/10.1007/s42113-018-0005-5>
- Ruggeri, A., Sim, Z. L., & Xu, F. (2017). “Why is Toma late to school again?” preschoolers identify the most informative questions. *Developmental Psychology*, 53(9), 1620–1632. <https://doi.org/10.1037/dev0000340>
- Ruggeri, A., Swaboda, N., Sim, Z. L., & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, 193, 104013. <https://doi.org/10.1016/j.cognition.2019.104013>

- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, 30(11), 1561–1572. <https://doi.org/10.1177/0956797619863663>
- Schulz, L. E. (2012). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, 16(7), 382–389. <https://doi.org/10.1016/j.tics.2012.06.004>
- Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, 43(4), 1045–1050. <https://doi.org/10.1037/0012-1649.43.4.1045>
- Seiver, E., Gopnik, A., & Goodman, N. D. (2013). Did she jump because she was the big sister or because the trampoline was safe? Causal inference and the development of social attribution. *Child Development*, 84(2), 443–454. <https://doi.org/10.1111/j.1467-8624.2012.01865.x>
- Sim, Z. L., & Xu, F. (2017). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, 35(4), 596–608. <https://doi.org/10.1111/bjdp.12198>
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology: General*, 146(2), 155–164. <https://doi.org/10.1037/xge0000250>
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94. <https://doi.org/10.1126/science.aaa3799>
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453–489. [https://doi.org/10.1016/S0364-0213\(03\)00010-7](https://doi.org/10.1016/S0364-0213(03)00010-7)
- Sumner, E., Li, A. X., Perfors, A., Hayes, B., Navarro, D., & Sarnecka, B. W. (2019). *The exploration advantage: Children's instinct to explore allows them to find information that adults miss*. PsyArXiv. <https://doi.org/10.31234/osf.io/h437v>.
- Sumner, E., Steyvers, M., & Sarnecka, B. W. (2019). It's not the treasure, it's the hunt: Children are more explorative on an explore/exploit task than adults. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st annual conference of the cognitive science society* (pp. 2891–2897). Cognitive Science Society.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1). MIT press.
- Tottenham, N., Shapiro, M., Flannery, J., Caldera, C., & Sullivan, R. M. (2019). Parental presence switches avoidance to attraction learning in children. *Nature Human Behavior*, 3(10), 1070–1077. <https://doi.org/10.1038/s41562-019-0656-9>
- Vermorel, J., & Mohri, M. (2005). Multi-armed bandit algorithms and empirical evaluation. *Machine Learning: ECML, 2005*, 437–448.
- Vogl, E., Pekrun, R., Murayama, K., Loderer, K., & Schubert, S. (2019). Surprise, curiosity, and confusion promote knowledge exploration: Evidence for robust effects of epistemic emotions. *Frontiers in Psychology*, 10, 2474. <https://doi.org/10.3389/fpsyg.2019.02474>
- Wente, A. O., Kimura, K., Walker, C. M., Banerjee, N., Flecha, M. F., MacDonald, B., ... Gopnik, A. (2019). Causal learning across culture and socioeconomic status. *Child Development*, 90(3), 859–875. <https://doi.org/10.1111/cdev.12943>
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6), 2074–2081. <https://doi.org/10.1037/a0038199>
- Xu, F., & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, 22(1), 28–32. <https://doi.org/10.1177/0963721412469396>