

Research

Reward-motivated memories influence new learning across development

Alexandra O. Cohen,^{1,3,4} Camille V. Phaneuf,^{1,4} Gail M. Rosenbaum,¹
Morgan M. Glover,¹ Kristen N. Avallone,¹ Xinxu Shen,¹ and Catherine A. Hartley^{1,2}

¹Department of Psychology, New York University, New York, New York 10003, USA; ²Center for Neural Science, New York University, New York, New York 10003, USA

Previously rewarding experiences can influence choices in new situations. Past work has demonstrated that existing reward associations can either help or hinder future behaviors and that there is substantial individual variability in the transfer of value across contexts. Developmental changes in reward sensitivity may also modulate the impact of prior reward associations on later goal-directed behavior. The current study aimed to characterize how reward associations formed in the past affected learning in the present from childhood to adulthood. Participants completed a reinforcement learning paradigm using high- and low-reward stimuli from a task completed 24 h earlier, as well as novel stimuli, as choice options. We found that prior high-reward associations impeded learning across all ages. We then assessed how individual differences in the prioritization of high- versus low-reward associations in memory impacted new learning. Greater high-reward memory prioritization was associated with worse learning performance for previously high-reward relative to low-reward stimuli across age. Adolescents also showed impeded early learning regardless of individual differences in high-reward memory prioritization. Detrimental effects of previous reward on choice behavior did not persist beyond learning. These findings indicate that prior reward associations proactively interfere with future learning from childhood to adulthood and that individual differences in reward-related memory prioritization influence new learning across age.

[Supplemental material is available for this article.]

Rewarding experiences can help guide choices and actions throughout life. Studies of reinforcement learning across development indicate that the ability to use good and bad outcomes to converge on optimal decisions changes from childhood to young adulthood (Davidow et al. 2018; Nussenbaum and Hartley 2019). These studies typically require individuals to learn the values associated with novel choice options. However, in the real world, we often have past experiences that can influence how choice options are evaluated and shape how we update value associations through new learning. Across childhood and into adulthood, reward associations lead to the adaptive prioritization of information in memory (Ngo et al. 2019; Cohen et al. 2022), which may facilitate the use of reward memories to guide later behavior. Still, it remains unclear how reward-related memories impact future choices in new contexts across age.

Past research suggests that existing stimulus–reward associations can either improve or impede subsequent behaviors. For example, the presence of reward enhances response inhibition (Wang et al. 2019) and promotes reinforcement learning through increased attention to reward-related features of the task environment (Niv et al. 2015; Leong et al. 2017). Additionally, prior knowledge about stimulus reward properties has been shown to facilitate stimulus processing in incremental associative learning and is thought to provide a scaffold for incorporating new information into existing representations (Bein et al. 2019). However, previous reward associations can be detrimental to performance, particularly when they are no longer relevant to the task at hand (Krebs et al.

2011; Infanti et al. 2015). For example, past work demonstrated that previously learned but irrelevant reward associations can hamper list learning performance in a new context (Madan et al. 2012). Recent work has also shown that prior reward conditioning may hinder memory formation for previously high-reward relative to low-reward images in a nonrewarded associative learning task (Miendlarzewska et al. 2021). These impairments in learning and memory may be due to past reward memories interfering with the ability to bind new value associations to previously reward-related stimuli (Kuhl et al. 2010; Madan et al. 2012). Together, these results suggest that prior reward value can either positively or negatively influence new learning.

Extant studies examining how prior rewarding experiences influence choices across development have largely focused on the impact of reward on cognitive control. Learned reward associations have been shown to lead to better inhibitory control in children (Chevalier et al. 2014; Winter and Sheridan 2014) but either disrupt (Roper et al. 2014; Davidow et al. 2019) or improve (Insel et al. 2019) performance across adolescence. Heightened reactivity to rewarding stimuli has been observed across species during adolescence (Galván 2013; Doremus-Fitzwater and Spear 2016), but less is known about how developmental changes in reward sensitivity influence learning and memory processes. Characterizing these processes is critical for understanding how previous experience guides future behavior. Increased reward sensitivity might give rise to a particularly robust influence of reward-related

³Present address: Department of Psychology, Emory University, Atlanta, GA 30322, USA.

⁴These authors contributed equally to this work.

Corresponding authors: ali.cohen@emory.edu, cate@nyu.edu

Article is online at <http://www.learnmem.org/cgi/doi/10.1101/lm.053595.122>.

© 2022 Cohen et al. This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first 12 months after the full-issue publication date (see <http://learnmem.cshlp.org/site/misc/terms.xhtml>). After 12 months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

memories on goal-directed behaviors during this stage of development relative to childhood and adulthood. Still, whether prior reward associations facilitate or hinder new learning across age remains unclear.

Individual differences in memory for learned value associations are likely to further impact how previously rewarding experiences influence future learning. Research conducted in adults shows a high degree of individual variability in the transfer of learning to new situations. This variability has been associated with differences in brain activation during learning and with intrinsic functional connectivity of brain systems implicated in learning (Wimmer and Shohamy 2012; Gerraty et al. 2014). Individual differences in how prior reward associations modulate new learning may stem from variability in how memory for those associations affects choice behavior.

In the present study, we aimed to determine how memory for previously rewarded experiences influences new learning from childhood to adulthood. To address this question, we collected and analyzed data from 89 participants ages 8–25 yr old in a 2-d study (Fig. 1). On the first day, participants completed a reward-motivated encoding task (Fig. 1A) during which participants viewed paired associates consisting of a trial-unique object and either a face or place source image, with each category (face or place) serving as the high-reward category for half of participants. Twenty-four hours later, participants completed a memory retrieval test (Fig. 1B), from which we computed a high- versus low-reward general source memory measure for each participant. This measure served as an individual-level index of reward-related memory prioritization. The retrieval test was followed by a reinforcement learning task using source image stimuli that had been associated with high- or low-reward 24 h earlier, as well as novel stimuli, as choice options (Fig. 1C). The reinforcement learning task allowed us to assess our primary question of interest: how memory for previously rewarded experiences influences new learning across age. Finally, participants completed a test phase consisting of all possible combinations of choice stimuli (Fig. 1D). The test phase examined whether prior reward associations biased value learning, which would be indicated by a tendency to choose previously rewarded stimuli over stimuli that were equally reinforced during learning. We hypothesized that adolescents, relative to older and younger individuals, might show particularly robust influences of past reward-related memories on new learning. We fur-

ther predicted that individual differences in reward-motivated memory prioritization would modulate learning.

Results

Learning task

To investigate how previously rewarding experiences impacted new learning from childhood to adulthood, we first examined learning performance as a function of reward history and age using mixed-effects logistic regression models. We compared models containing both linear and quadratic age as interaction terms versus linear age alone to determine whether including quadratic age provided a better fit to the data. A likelihood ratio χ^2 test indicated that the learning data were significantly better fit by a quadratic age model [$\chi^2(6)=14.70$, $P=0.023$]. We report results from the mixed-effects logistic regression model that probed trial-wise optimal choice by trial number, stimulus type (previously high reward, previously low reward, and novel), continuous linear age, continuous quadratic age, and their interactions, controlling for the high-reward source image category (whether faces or places were associated with high reward during the motivated encoding task). We found a significant effect of trial number, such that optimal choice increased throughout the task [$\chi^2(1, N=89)=230.16$, $P<0.001$]. We also observed a significant effect of stimulus type, whereby optimal responding was highest for novel stimuli and lowest for previously high-reward stimuli [$\chi^2(2, N=89)=30.97$, $P<0.001$] (Fig. 2B). There was a significant effect of linear age, in which optimal choice increased with age [$\chi^2(1, N=89)=26.28$, $P<0.001$], as well as a trending effect of quadratic age [$\chi^2(1, N=89)=3.67$, $P=0.056$]. There was no significant effect of high-reward source image category [$\chi^2(1, N=89)=0.43$, $P=0.512$]. These results suggest that learning performance improved with age and throughout the task but was impeded for previously high-reward relative to low-reward stimuli.

The main effects were qualified by two-way interactions. We found a significant trial number \times stimulus type interaction, demonstrating that poorer responding early in learning for previously low-reward stimuli, and even more so for previously high-reward stimuli, were both ameliorated by the end of the task [$\chi^2(2, N=89)=9.78$, $P=0.008$]. Furthermore, significant effects of trial number \times linear age and trial number \times quadratic age indicated that, as the task progressed, younger participants did not increase their

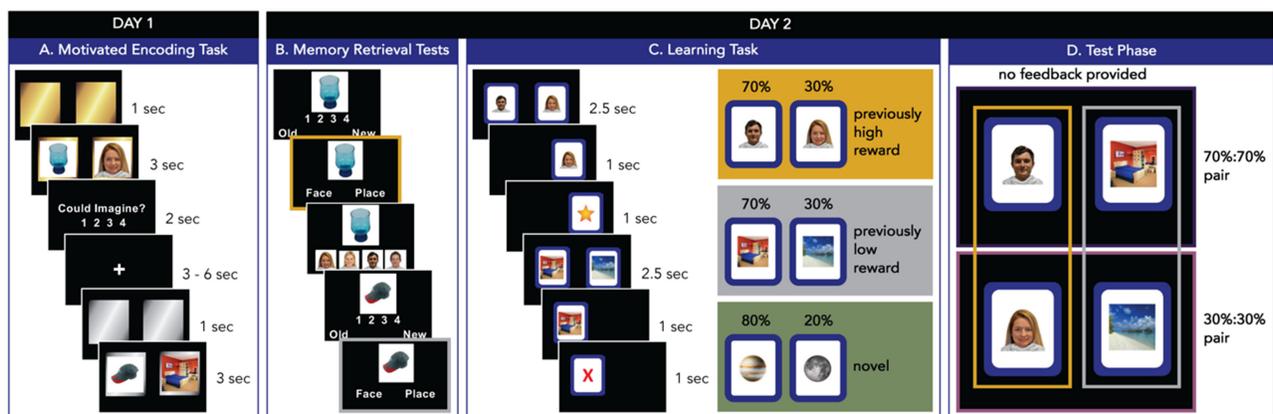


Figure 1. Experimental design. (A) On the first day of the study, participants completed two rounds of a reward-motivated encoding task with high- and low-reward trials on gold and silver backgrounds, respectively. (B) Twenty-four hours later, they completed a self-paced memory retrieval test (gold and silver outlines that visualize the reward levels of the stimuli were not depicted during the task). (C) Participants next completed a modified instructed probabilistic learning paradigm with previously high-reward, previously low-reward, and novel card deck pairs as choice options. Positive (star) or negative (X) feedback was given after each card deck choice. (D) Last, participants completed a self-paced test phase, including choices between the 70% and 30% reinforced decks from learning. Feedback was not provided during the test phase.

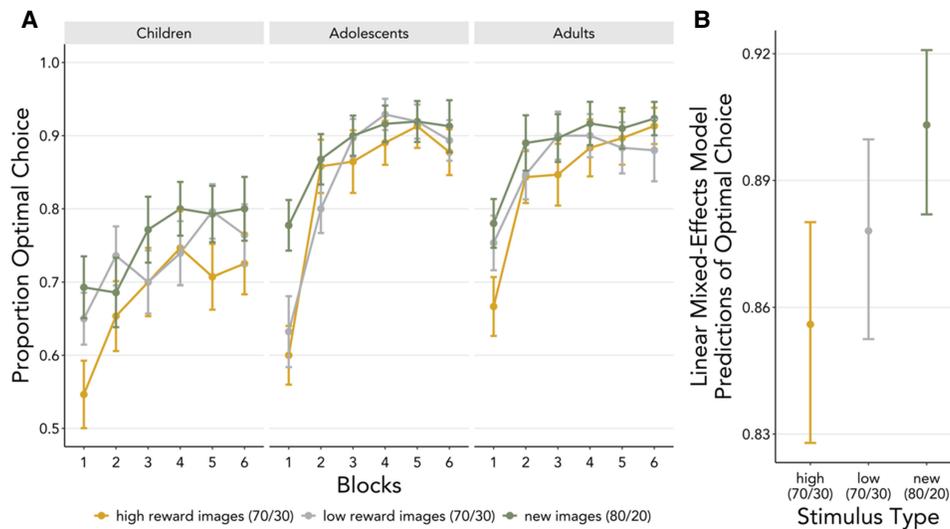


Figure 2. Prior reward associations influenced the learning of new reward contingencies. (A) Learning performance improved throughout the task and with age. Stimulus type also modulated learning, such that optimal choice was highest for novel stimuli and lowest for previously high-reward stimuli. Age is binned into child (8–12 yr), adolescent (13–17 yr), and adult (18–25 yr) groups for data visualization purposes, but age was treated as a continuous variable in all analyses. Similarly, trials are binned into 30 trial blocks (10 trials of each stimulus type) for data visualization purposes, but trial number was treated as a continuous variable in all analyses. Error bars represent standard error. (B) Proportion optimal choice was consistently diminished for previously high-reward stimuli across all ages. Points represent predictions of optimal choice from the linear mixed-effects model, including previously high-reward, previously low-reward, and novel stimuli. The vertical bars represent 95% confidence intervals around the model estimates.

rate of optimal responding as much as adolescents and older participants did [trial number \times linear age: $\chi^2(1, N=89) = 30.32, P < 0.001$; trial number \times quadratic age: $\chi^2(1, N=89) = 4.07, P = 0.044$] (trial number \times age interactions are depicted in Supplemental Fig. S1). This series of two-way interactions indicates that performance improved throughout the task, with trajectories varying by stimulus type and age (Fig. 2A; see Supplemental Fig. S2 for age visualized continuously). There were no other significant two- or three-way interactions ($P_s > 0.074$) (see Supplemental Table S1A,B for full linear mixed-effects model output). In sum, these results illustrate that prior reward associations led to initial decrements in learning (highlighted in Supplemental Fig. S3), but learning performance improved with increasing participant age as the task progressed.

The learning task design was modified from that of an instructed probabilistic selection task (Doll et al. 2009; Decker et al. 2015) in order to assess any potential biases in value-based learning induced by the presence of previously rewarded stimuli. Consistent with past developmental work (Decker et al. 2015), the novel stimulus pair was reinforced at a different reward rate (80%:20%) than the previous high- and low-reward pairs (70%:30%). These reward contingencies were used so that a “prior reward bias on learning” measure could be computed, similar to a measure of instruction bias used in this earlier research. However, the use of these contingencies could have resulted in a main effect of stimulus type (which included previously high-reward, previously low-reward, and novel images as levels of the stimulus type variable) that was driven by differences in the difficulty of learning from pairs with dissimilar reinforcement schedules. To mitigate this concern, we directly examined the effect of the level of prior reward (i.e., high vs. low) by conducting a mixed-effects logistic regression without the novel stimulus trials. In this model, we still observed a main effect of stimulus type [$\chi^2(1, N=89) = 8.13, P = 0.004$], indicating that the influence of stimulus type on learning was not solely driven by novel stimuli. We also continued to see main effects of trial number [$\chi^2(1, N=89) = 208.63, P < 0.001$] and linear age [$\chi^2(1, N=89) = 24.35, P < 0.001$]. Additionally, the interactions between trial number and age persisted [trial number \times linear age: $\chi^2(1, N=89) =$

30.53, $P < 0.001$; trial number \times quadratic age: $\chi^2(1, N=89) = 8.48, P = 0.004$]. However, the trial number \times stimulus type interaction was not statistically significant [$\chi^2(1, N=89) = 0.92, P = 0.337$], suggesting that differences in the learning trajectories for each stimulus type were driven by the previously reward-associated versus novel pairs. There were no other significant main effects or two- or three-way interactions ($P_s > 0.120$) (see Supplemental Table S1C,D for full linear mixed-effects model output). Taken together, these findings indicate that overall learning performance improved with increasing age and that prior high-reward associations stymied learning. Analyses using learning reaction times, rather than learning performance, as the response variable produced a similar pattern of results and further support our conclusions (see Supplemental Fig. S4; Supplemental Table S2A–D).

Individual differences in learning

The extent to which high- versus low-reward information was differentially prioritized in memory across participants on day 1 may have impacted the degree to which previously acquired reward associations influenced day 2 learning. To test this possibility, we computed a participant-specific high- versus low-reward general source memory measure, derived from the memory retrieval test (Fig. 1B), as an index of high-reward memory prioritization. This measure has a participant-specific denominator that accounts for overall differences in memory performance (for more details, see the Materials and Methods). Although our prior work indicated that all ages showed reward-motivated memory enhancement for specific associations between trial-unique objects and the faces or places they were paired with (see Cohen et al. 2022), the high- versus low-reward general source memory measure calculated here reflects the extent to which reward associations formed during the motivated encoding task were maintained at the level of the rewarded source images (faces or places), even when participants lacked memory for the specific associations (for extensive characterization of the memory retrieval data from this task, see Cohen et al. 2022). We chose to examine the trials in which the general

source image category was correct but the specific source image was incorrect because we wanted the individual difference measure to assess memory at the same level of representation that was used in the reinforcement learning task (i.e., the reward category [face or place] level). If the degree to which prior reward associations influence new learning depends on individual variability in the prioritization of reward-related information in memory, we would expect our general reward source memory measure to modulate the acquisition of new reward contingencies for high-reward relative to low-reward stimuli.

As a first step, we investigated how general reward source memory related to age in our developmental sample. We found a quadratic relationship between age and general reward source memory [$\beta = -0.025$, $t(85) = -2.871$, $P = 0.005$] such that, on average, adolescents displayed relatively equivalent high- and low-reward general source memory, while younger and older participants demonstrated better low-reward general source memory (see the inverted U-shape in Fig. 3A). We also saw a main effect of the high-reward source image category control variable, such that individuals for whom faces were the high-reward image

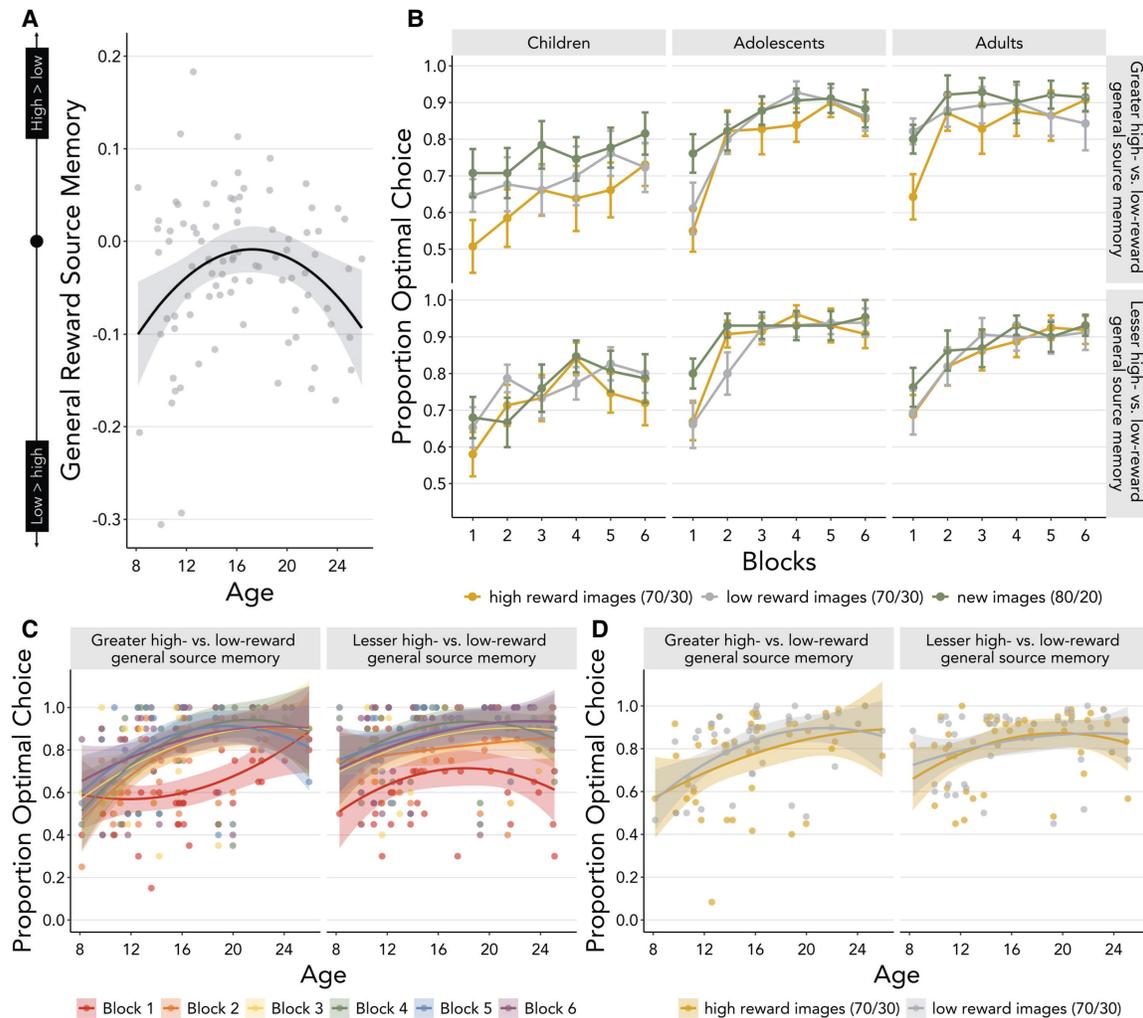


Figure 3. Reward source memory prioritization changed with age and influenced how new reward contingencies were learned throughout the task. (A) High- versus low-reward general source memory was greatest for adolescents, relative to younger and older participants. (B) Learning trajectories for previously high- and low-reward stimuli in younger and older participants varied with high- versus low-reward general source memory. Age is binned into child (8–12 yr), adolescent (13–17 yr), and adult (18–25 yr) groups for data visualization purposes, but age was treated as a continuous variable in all analyses. Similarly, trials are binned into 30 trial blocks (10 trials of each stimulus type) for data visualization purposes, but trial number was treated as a continuous variable in all analyses. Novel stimuli (green) are depicted here for illustrative comparison, but only previously high- and low-reward images were included in the corresponding analysis. Error bars represent standard error. (C) Visualization of block-wise learning performance across age, according to high- versus low-reward general source memory. While early learning was hindered across all ages for individuals with relatively less high-reward memory prioritization, adolescents that demonstrated greater high-reward memory prioritization also showed early learning deficits. Trials are binned into blocks of 20 trials (10 trials of each stimulus type) for data visualization purposes, but trial number was treated as a continuous variable in all analyses. (D) Visualization of learning performance for previously high- and low-reward stimuli across age, according to high- versus low-reward general source memory. Across age, individuals demonstrating greater high-reward memory prioritization showed worse learning for previously high-reward relative to low-reward stimuli, but similar learning for both stimulus types was observed for those demonstrating relatively less high-reward memory prioritization. A median split on general reward source memory (median general reward source memory = -0.02) is shown as “greater” or “lesser” for data visualization purposes in B–D, but general reward source memory was treated as a continuous variable in all analyses. Shading represents a 95% confidence interval around fitted lines in A, C, and D.

category demonstrated greater high-reward general source memory than individuals for whom places were the high-reward image category [$\beta = -0.042$, $t(85) = -2.507$, $P = 0.014$].

Next, to assess how prior reward level and individual differences in reward memory prioritization impacted future choice, we included high- versus low-reward general source memory as a predictor of interest in our learning model, which excluded novel stimuli. We continued to see significant main effects of trial number, stimulus type (previously high reward and previously low reward), and linear age on optimal responding [trial number: $\chi^2(1, N = 89) = 211.80$, $P < 0.001$; stimulus type: $\chi^2(1, N = 89) = 9.91$, $P = 0.002$; linear age: $\chi^2(1, N = 89) = 24.11$, $P < 0.001$]. As in the learning model, there was no significant effect of high-reward source image category [$\chi^2(1, N = 89) = 0.16$, $P = 0.691$], and again, we found significant two-way interactions between trial number \times linear age and trial number \times quadratic age, whereby performance improved throughout the task and across development, particularly for adolescents [trial number \times linear age: $\chi^2(1, N = 89) = 15.48$, $P < 0.001$; trial number \times quadratic age: $\chi^2(1, N = 89) = 8.73$, $P = 0.003$]. We also observed a significant interaction between trial number and high- versus low-reward general source memory, such that increased high-reward general source memory impaired optimal responding, especially early in the task [$\chi^2(1, N = 89) = 7.08$, $P = 0.008$]. Additionally, there was a trending interaction between stimulus type and quadratic age, with younger and older participants differentially responding to previously high- and low-reward stimuli [$\chi^2(1, N = 89) = 3.50$, $P = 0.061$].

We also found two three-way interactions. There was a significant effect of trial number \times general reward source memory \times quadratic age on learning, such that early learning was hindered for individuals with relatively less high-reward memory prioritization across all ages and for adolescents that demonstrated greater high-reward memory prioritization [$\chi^2(1, N = 89) = 5.32$, $P = 0.021$] (Fig. 3C). There was also a significant effect of stimulus type \times general reward source memory \times quadratic age on learning. Individuals with greater high-reward memory prioritization showed worse learning for previously high-reward relative to low-reward stimuli across all ages, while those demonstrating relatively less high-reward memory prioritization showed similar learning for both stimulus types [$\chi^2(1, N = 89) = 6.20$, $P = 0.013$] (Fig. 3D). There were no other significant main effects or two-, three-, or four-way interactions ($P_s > 0.118$) (see Supplemental Table S3A,B for full linear mixed-effects model output). These results suggest that the extent to which high- versus low-reward associations were prioritized in memory affected learning for previously rewarded stimuli across age and across the task. Learning of new reward contingencies for stimuli with existing high- and low-reward associations was modulated by the degree to which individuals preferentially maintained past high-reward experiences in memory. Additionally, adolescents showed early learning decrements for previously rewarded stimuli irrespective of the extent to which they prioritized high- over low-reward-associated experiences in memory.

Test phase

The learned value of stimuli can be revealed by participants' decisions in a test

phase, where choices are made between all possible pairings of stimuli in the absence of feedback. In order to assess whether the influence of reward history on choices persisted beyond new learning, we considered test phase decisions for stimuli that were reinforced at the same rate (e.g., 70%) during learning but were associated with different levels of prior reward (i.e., high or low reward) during the motivated encoding task. Specifically, we examined the proportion of previously high-reward choices made by continuous age, controlling for the high-reward source image category (faces or places) and using separate linear regressions for the 70%:70% and 30%:30% pairings (see Supplemental Table S4A,B for full regression model outputs). Neither regression revealed significant effects of age [70%:70% pairing: $\beta = -0.031$, $t(86) = -0.825$, $P = 0.412$; 30%:30% pairing: $\beta = -0.014$, $t(86) = -0.404$, $P = 0.687$] or high-reward source image category [70%:70% pairing: $\beta = 0.108$, $t(86) = 1.458$, $P = 0.148$; 30%:30% pairing: $\beta = 0.090$, $t(86) = 1.325$, $P = 0.189$] (Fig. 4). Because we did not observe significant differences in these analyses of the test phase data, we did not compute a "prior reward bias on learning" measure or fit reinforcement learning models to test phase choices to characterize how participants integrated feedback into their stimulus value estimates during the learning task, as has been done in previous work using similar paradigms (Doll et al. 2009; Decker et al. 2015). Reinforcement learning models fit to participants' learning task data did not indicate a clear best-fitting model, and the models were not recoverable; therefore, we do not report these analyses here. The test phase results indicate that prior reward value did not continue to bias choices following new learning.

Discussion

The present study examined how memory for existing reward associations influences new learning from childhood to adulthood. Although learning improved throughout the task and with increasing age, we observed that prior high-reward associations impeded learning across all participants. We also found that individual

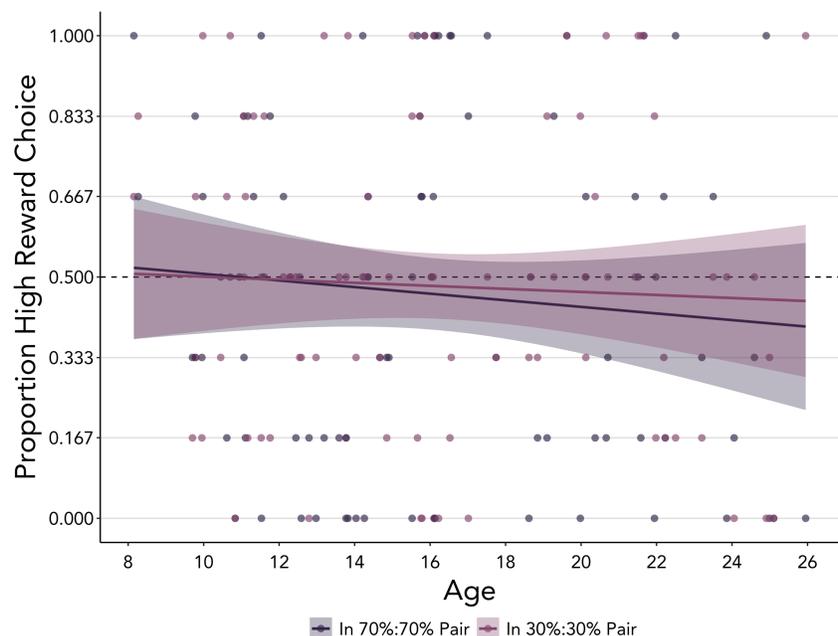


Figure 4. Prior reward value did not persist following new learning across age. Shading represents 95% confidence intervals for predictions from each linear regression.

variability in the degree to which high-reward associations were prioritized over low-reward associations in memory was related to differences in learning performance. Specifically, we saw that greater high-reward general source memory prioritization was associated with worse learning for previously high-reward relative to low-reward stimuli across age and that adolescents showed hindered early learning regardless of stimulus type and whether high- or low-reward associations had been more strongly prioritized in memory. Influences of past rewards on learning were ameliorated by the end of the task and did not bias later test phase choices. Taken together, these results suggest that prior reward associations, as well as individual differences in how reward associations are prioritized in memory, modulate how individuals learn from childhood to adulthood.

Across all ages, previous high-reward associations led to decrements in learning. Consistent with past research conducted in adults examining how prior rewards influence new memory formation (Madan et al. 2012; Miendlarzewska et al. 2021), we found that existing high-reward associations proactively interfered with learning the values of stimuli in a new reinforcement learning task across all participants. Although previous work has demonstrated age-related differences in susceptibility to proactive interference in different tasks (Loosli et al. 2014), our findings illustrate that reward-related proactive interference can be observed in learning from childhood to adulthood. This aligns with research demonstrating that children, adolescents, and adults all show value-related memory enhancements (Cohen et al. 2019, 2022; Ngo et al. 2019) and suggests that reward-related memories influence subsequent choice across age.

We observed a nonlinear age effect in our individual difference measure examining general reward source memory. This measure assessed memory at the same level of representation that was used in the reinforcement learning task (i.e., the reward category [face or place] level) of the present study. In recent work, we showed that reward-motivated memory enhancements consist of both age-invariant and nonlinear age-related patterns (Cohen et al. 2022). We found that while specific source memory for high-reward associations was enhanced across all ages, general source memory was better for low-reward relative to high-reward associations in younger and older participants. A similar pattern of results has been reported in a separate study that examined general source memory after 24 h for the outcomes of choices observed during a value-based learning task (Katzman and Hartley 2020). We suggest that these converging findings may be due to younger and older individuals encoding high- and low-reward associations at different levels of specificity as a function of reward level (i.e., encoding high-reward associations with greater specificity and low-reward associations with less specificity), while adolescents may prioritize remembering specific high-reward associations and not “hold on” to more general representations of low-reward associations. While, on average, younger and older individuals tended to show better low-reward general source memory, we observed that across all ages, some participants show better high-reward general source memory while others show better low-reward general source memory. In the present study, we aimed to leverage this individual variability in memory to characterize how individual differences in reward-motivated general source memory prioritization influenced new learning.

Prior research has suggested a high degree of individual variability in the transfer of learned reward associations to new stimuli and tasks in adults (Wimmer and Shohamy 2012; Gerraty et al. 2014). Our results examining value transfer across time and context suggest that individual variation in reward-related memory prioritization differentially impacts new learning in several different ways across age. We found that greater high-reward general source memory prioritization was associated with impeded learn-

ing for previously high-reward relative to low-reward stimuli across age. This supports the conclusion that prior high-reward associations interfere with new learning by demonstrating that individuals with greater high-reward general source memory prioritization show worse learning for previously high-reward stimuli. We also found that early learning was hindered across individuals of all ages who showed greater low-reward memory prioritization, but that adolescents also showed relatively worse early learning regardless of individual differences in memory prioritization and the level (high or low) of prior reward associations. Studies in rodents examining reward learning and subsequent extinction of reward associations have shown that adolescents, relative to adults, persevere more in response to previously rewarded cues (Sturman et al. 2010; Meyer and Bucci 2016). Similarly, in the present study, adolescent humans may show proactive interference of prior reward associations particularly in the early phase of new learning. Further research is needed to characterize how changes in the reward properties of cues impact learning dynamics across development. Taken together, these findings indicate that the extent to which high-reward associations were prioritized in memory impacted learning across age.

To assess whether prior reward associations may have biased value-based learning, we examined participants' choices during a test phase. Specifically, we examined choices between options that were equivalently reinforced during learning but had different reward histories. Across all ages, we found no significant differences in participants' choices. The lack of differentiation between previously high- and low-reward stimuli in participants' test phase choices suggests that any initial detrimental effects of prior reward value on decision-making were ameliorated over the course of learning. Previous research investigating the influence of instruction on learning has demonstrated robust instruction biases that distort value learning in adults, whereas children and adolescents do not show such biases, instead estimating stimulus values based on experience (Doll et al. 2009; Decker et al. 2015). The results of the present study suggest that prior rewarding experiences influence the trajectory of learning but do not significantly distort the ultimate outcome. Together, these findings illustrate that distinct types of prior knowledge (e.g., instructed vs. directly experienced) may differentially impact new learning across age.

The current findings highlight several unresolved questions about the cognitive mechanisms through which reward associations influence reinforcement learning in a new context. Attentional mechanisms likely contributed to how prior reward value modulated new learning; however, we did not specifically assess attention in this study. Future work that measures attention throughout learning may provide important insights into how reward memories modulate attention, which in turn dynamically guides later learning. Additionally, established value associations are a specific type of prior knowledge, which has been shown to help or hinder new knowledge acquisition across age (Shing and Brod 2016). Congruence between prior knowledge and new information has been proposed as an important determinant of how existing knowledge impacts learning (Brod 2021). Information that is incongruent with prior knowledge is typically less well remembered than congruent information (Schulman 1974; Bein et al. 2015; Brod and Shing 2019). Although the experimental design in the present study, which involved choosing between two previously rewarded options, may resemble real-world decision scenarios, this manipulation led to mismatches in the congruence of prior and current reward value for one of the two choice options in each pair (e.g., a previously high-reward stimulus currently reinforced 30% of the time). Future work that eliminates such congruency mismatches is needed to assess whether prior reward associations consistently interfere with new learning. Taken together, our results highlight that the opposing effects of prior

knowledge on learning may be further clarified by also considering the temporal dynamics of learning, individual variation in memory, and potential individual-level and age-related variability in the cognitive mechanisms that facilitate the integration of previous experiences with new learning.

In summary, we found that past rewarding experiences lead to impairments in future learning that are modulated by individual variability in reward memory prioritization and age. Despite differences in learning trajectories for stimuli with different reward histories, individuals of all ages update reward associations and learn new stimulus values. Further research is needed to characterize how various types of motivationally salient prior experiences, such as punishments, influence how individuals learn. Insights from this line of work have the potential to inform our understanding of the interplay between learning and memory processes (Hartley et al. 2021) and how variation in individuals' past experiences affect value-based learning.

Materials and Methods

Participants

Eighty-nine child, adolescent, and adult participants ages 8–25 yr old (mean age = 16.16 yr, standard deviation age = 4.67 yr, 45 females) from the New York City community completed this 2-d study. We aimed to evenly distribute participant recruitment across age groups and sex (children [8–12 yr old]: $N=28$ participants, mean age = 11.04 yr, standard deviation age = 1.24 yr, 14 females; adolescents [13–17 yr old]: $N=31$ participants, mean age = 15.35 yr, standard deviation age = 1.19 yr, 16 females; adults [18–25 yr old]: $N=30$ participants, mean age = 21.78 yr, standard deviation age = 2.13 yr, 15 females). While age groups guided recruitment and are defined here to characterize the sample, all analyses treated age as a continuous variable. The age and sex distribution is visualized in Supplemental Figure S5.

The first day of testing occurred in the fMRI scanner (neuroimaging results reported in Cohen et al. 2022), while the second day of testing took place outside the scanner. Individuals diagnosed with a learning, anxiety, and/or mood disorder, or consumers of medications that influence the nervous system, were ineligible to participate. The target sample size was consistent with prior studies detecting a significant age-related change in goal-directed behavior (e.g., Decker et al. 2015; Insel et al. 2017, 2019; Insel and Somerville 2018; Katzman and Hartley 2020; Raab et al. 2022). In accordance with the National Institutes of Health Policy on Reporting Race and Ethnicity Data, our final sample self-identified as 25% Asian, 12% Black/African American, 38% Caucasian/White, 24% mixed race, and <2% Native American. Sixteen percent self-identified as Hispanic.

Participants were paid \$75 for completing the scanning and behavioral testing sessions. In addition, all participants had the opportunity to earn up to \$21 in bonus payments by the end of the second experiment day. This study was approved by the Institutional Review Board at New York University. For minors, informed parental permission and child assent were acquired prior to participation. Informed consent was obtained for adult participants.

Procedures

We examined the influence of prior reward associations on new learning through a 2-d study (Fig. 1). On the first day, participants completed a reward-motivated encoding task. Twenty-four hours later, they underwent a memory retrieval test, followed by a learning task with previously high-reward, previously low-reward, and novel stimuli as choice options, and then a test phase. All tasks were programmed in Epyriment (Krause and Lindemann 2014) using Pygame v1.9.4 and Python v2.7.15. Stimuli were created using images from RADIATE (Conley et al. 2018), the Chicago Face Database (Ma et al. 2015), Harvard's Konkle laboratory (Konkle et al. 2010), and Massachusetts Institute of Technology's Places

Scene Recognition (<https://dspace.mit.edu/handle/1721.1/96941?show=full>) databases.

Reward-motivated encoding task

Participants completed a motivated encoding paradigm (Duncan et al. 2014; Murty et al. 2017) during day 1 of the study (Fig. 1A). On each trial, participants first saw two squares, either gold or silver, that indicated remembering the upcoming pair of images would help them win a big \$15 bonus (high reward on gold square backgrounds) or a small \$1 bonus (low reward on silver square backgrounds). A trial-unique picture of an object was overlaid on the left, and one of eight repeated source images was overlaid on the right. Source images consisted of four faces (two women and two men) and four places (two indoor and two outdoor). Participants were pseudorandomly assigned faces or places as their high-reward source image category based on subject ID number. Participants were instructed to create a story involving both images to promote deep encoding. They then rated how well they could imagine the story that they told themselves, on a scale from 1 (very easy to imagine) to 4 (very hard to imagine). There were two encoding phases, each consisting of 64 trials (32 high reward and 32 low reward). For detailed day 1 methods, see Cohen et al. (2022).

Memory retrieval test

Participants completed a memory retrieval test (Fig. 1B) 24 h after the reward-motivated encoding task. They were presented with all 128 objects observed during encoding and 128 new objects included as lure images for 256 trials total. On each trial, participants saw an object and indicated whether the object was definitely old, maybe old, maybe new, or definitely new. If they endorsed the object as "old," they then indicated whether that object was previously paired with a face or place. Participants then selected a specific source image. If the participant endorsed the object as "new," they were not asked further questions about that object. The memory retrieval test was self-paced.

In order to examine how individual variation in long-term memory for previously high- versus low-reward associations modulated future choice behavior, we used the retrieval data to compute a high- versus low-reward general source memory measure for each participant. To do so, we found the number of trials in which the high-reward source image category (e.g., faces) was correctly identified, but the specific source image (e.g., a particular woman) was not. This isolated a general associative memory measure corresponding to the source image category. Next, we found the number of trials in which the low-reward source image category (e.g., places) was correctly identified, but the specific source image (e.g., a particular bedroom) was not. Because source memory was only queried for items that participants identified as old, each measure was divided by the total number of items correctly identified as maybe/definitely old by each participant. Thus, the denominator for high- and low-reward general source memory was the same within each participant and largely accounted for overall differences in memory performance. Finally, low-reward general source memory was subtracted from high-reward general source memory. The resulting measure quantified the degree to which an individual demonstrated better general memory for the previously high-reward relative to low-reward associations, thus providing a participant-specific metric of high- versus low-reward general source memory corresponding to the images used in the learning task (described below). This metric allowed us to investigate how reward-related memory prioritization impacted new learning.

Learning task

Participants completed a probabilistic selection task (modified from Doll et al. 2009; Decker et al. 2015) after the memory retrieval test. This task was adapted to examine the influence of reward-motivated experiences on subsequent learning and decision-making (Fig. 1C). Participants chose between six unique card decks, grouped into three pairs. They saw one pair of card decks

at each choice, with one deck on the left side of the screen and the other deck on the right side. Card decks were distinguishable by the image on the front of each deck. One pair of card decks had two randomly selected high-reward source images from the motivated encoding task on the front (one woman and one man), the second pair of card decks had two randomly selected low-reward source images on the front (one indoor and one outdoor), and the third pair of card decks had two novel images (planets) on the front.

Every card deck had either a star or an X on the back, which participants could reveal by selecting one deck at each choice. Participants were instructed to determine which card decks were good and which were bad so that they could collect as many cards with stars on the back as possible in order to win bonus money. To make selections, participants were told to press 1 for the deck on the left and to press 2 for the deck on the right. Each pair of card decks was displayed on the screen for 2.5 sec, and the participant's choice and subsequent feedback were each displayed for 1 sec. Additionally, there were 1-sec fixations between trials.

All stimuli were probabilistically reinforced, with an optimal choice in each pair. In line with instructed probabilistic selection task paradigms (Doll et al. 2009; Decker et al. 2015), pairs of card decks with previously rewarded source images on the front included one stimulus that was positively reinforced 70% of the time while the other stimulus was positively reinforced 30% of the time. For the pair of decks with novel images on the front, one stimulus was positively reinforced 80% of the time while the other stimulus was positively reinforced 20% of the time. These reward contingencies were selected so that a prior reward bias could be computed, similar to a measure of instruction bias used in earlier work (Decker et al. 2015). Participants saw each pair of decks 60 times, pseudorandomized in 10-trial blocks, for a total of 180 learning trials.

Test phase

Participants were presented with all 15 possible pairs of card decks (three original and 12 new) during the test phase that immediately followed learning (Fig. 1D). Prior to beginning the test phase, participants were instructed to choose which of the two decks was best based on what they had learned. Participants were told that there would be no feedback associated with their choices. Each pair of card decks appeared six times in a random order (90 total trials). One-second fixations were displayed between trials. The test phase was self-paced, and no time limits were imposed.

Analyses

Statistical analyses

Data processing and analyses were conducted in R version 3.6.2 (R Core Team 2019). Linear mixed-effects models were run using the “glmer” function in the “lme4” package (version 1.1-23) (<https://cran.r-project.org/web/packages/lme4/index.html>) for analyzing the learning task data. For the glmer learning models, we first compared models containing both linear and quadratic age as interaction terms versus linear age alone to determine whether including quadratic age provided a better fit to the data. A likelihood ratio χ^2 test indicated that the learning data were significantly better fit by a quadratic age model [$\chi^2(6) = 14.70, P = 0.023$]. Next, we fit the maximal glmer model and iteratively removed the random slopes when the model failed to converge until arriving at a reduced, converging model (Bates et al. 2015). Through exploratory data visualization, we examined whether the individual variability in the learning data could be meaningfully explained by within-participant differences in learning for each stimulus type. We did not find this to be true, so the final glmer learning models only included random intercepts for each participant to account for repeated measurements within participants and the individual variability in learning that they explain.

To disentangle the influences of prior reward presence and level on learning, we first ran our learning model with all three stimulus types (previously high reward, previously low reward, and novel) included, and then ran our learning model excluding

the novel stimulus trials to directly compare the previously rewarded choice options. When analyzing individual differences in learning, the learning model including all three stimulus types and high-versus low-reward general source memory as a predictor did not converge. Thus, here we only share the results from the model excluding novel stimuli, which directly assesses learning differences as a function of prior reward level. Statistics were reported from analysis of deviance (type III Wald χ^2 tests using Satterthwaite approximations for degrees of freedom) performed on glmer learning models and linear regressions performed on high- versus low-reward general source memory and test phase choices using the “lm” function. Both age and trial number were treated as continuous variables in all analyses and z-scored across all participants.

Data availability

Data and the code used to produce the results and figures are available at https://github.com/hartleylabnyu/reward_memories_influence_learning.

Acknowledgments

We thank Betsy Deza, Anastasia Filimontseva, Hannah Walker, and James Wyngaarden for help with data collection, and the Hartley Laboratory for providing thoughtful feedback on the analyses. This work is supported by a Klingenstein-Simons Fellowship Award in Neuroscience, a Jacobs Foundation Research Fellowship, and National Science Foundation CAREER Award 1654393 to C.A.H.; a Department of Defense National Defense Science and Engineering Graduate Fellowship to C.V.P.; and National Science Foundation Social, Behavioral, and Economic Sciences Postdoctoral Research Fellowship grant number 1714321 and National Institute on Drug Abuse grant number K01DA053438 to A.O.C.

References

- Bates D, Kliegl R, Vasishth S, Baayen H. 2015. Parsimonious mixed models. arXiv doi:10.48550/arXiv.1506.04967
- Bein O, Livneh N, Reggev N, Gilead M, Goshen-Gottstein Y, Maril A. 2015. Delineating the effect of semantic congruency on episodic memory: the role of integration and relatedness. *PLoS One* **10**: e0115624. doi:10.1371/journal.pone.0115624
- Bein O, Trzewik M, Maril A. 2019. The role of prior knowledge in incremental associative learning: an empirical and computational approach. *J Mem Lang* **107**: 1–24. doi:10.1016/j.jml.2019.03.006
- Brod G. 2021. Toward an understanding of when prior knowledge helps or hinders learning. *NPJ Sci Learn* **6**: 24. doi:10.1038/s41539-021-00103-w
- Brod G, Shing YL. 2019. A boon and a bane: comparing the effects of prior knowledge on memory across the lifespan. *Dev Psychol* **55**: 1326–1337. doi:10.1037/dev0000712
- Chevalier N, Chatham CH, Munakata Y. 2014. The practice of going helps children to stop: the importance of context monitoring in inhibitory control. *J Exp Psychol Gen* **143**: 959–965. doi:10.1037/a0035868
- Cohen AO, Matese NG, Filimontseva A, Shen X, Shi TC, Livne E, Hartley CA. 2019. Aversive learning strengthens episodic memory in both adolescents and adults. *Learn Mem* **26**: 272–279. doi:10.1101/lm.048413.118
- Cohen AO, Glover MM, Shen X, Phaneuf CV, Avallone KN, Davachi L, Hartley CA. 2022. Reward enhances memory via age-varying online and offline neural mechanisms across development. *J Neurosci* **42**: 6424–6434. doi:10.1523/JNEUROSCI.1820-21.2022
- Conley MI, Dellarco DV, Rubien-Thomas E, Cohen AO, Cervera A, Tottenham N, Casey BJ. 2018. The racially diverse affective expression (RADIATE) face stimulus set. *Psychiatry Res* **270**: 1059–1067. doi:10.1016/j.psychres.2018.04.066
- Davidow JY, Insel C, Somerville LH. 2018. Adolescent development of value-guided goal pursuit. *Trends Cogn Sci* **22**: 725–736. doi:10.1016/j.tics.2018.05.003
- Davidow JY, Sheridan MA, Van Dijk KRA, Santillana RM, Snyder J, Vidal Bustamante CM, Rosen BR, Somerville LH. 2019. Development of prefrontal cortical connectivity and the enduring effect of learned value on cognitive control. *J Cogn Neurosci* **31**: 64–77. http://dx.doi.org/10.1162/jocn_a_01331.
- Decker JH, Lourenco FS, Doll BB, Hartley CA. 2015. Experiential reward learning outweighs instruction prior to adulthood. *Cogn Affect Behav Neurosci* **15**: 310–320. doi:10.3758/s13415-014-0332-5

- Doll BB, Jacobs WJ, Sanfey AG, Frank MJ. 2009. Instructional control of reinforcement learning: a behavioral and neurocomputational investigation. *Brain Res* **1299**: 74–94. doi:10.1016/j.brainres.2009.07.007
- Doremus-Fitzwater TL, Spear LP. 2016. Reward-centricity and attenuated aversions: an adolescent phenotype emerging from studies in laboratory animals. *Neurosci Biobehav Rev* **70**: 121–134. doi:10.1016/j.neubiorev.2016.08.015
- Duncan K, Tomparay A, Davachi L. 2014. Associative encoding and retrieval are predicted by functional connectivity in distinct hippocampal area CA1 pathways. *J Neurosci* **34**: 11188–11198. doi:10.1523/JNEUROSCI.0521-14.2014
- Galván A. 2013. The teenage brain: sensitivity to rewards. *Curr Dir Psychol Sci* **22**: 88–93. doi:10.1177/0963721413480859
- Gerraty RT, Davidow JY, Wimmer GE, Kahn I, Shohamy D. 2014. Transfer of learning relates to intrinsic connectivity between hippocampus, ventromedial prefrontal cortex, and large-scale networks. *J Neurosci* **34**: 11297–11303. doi:10.1523/JNEUROSCI.0185-14.2014
- Hartley CA, Nussenbaum K, Cohen AO. 2021. Interactive development of adaptive learning and memory. *Annu. Rev. Psychol* **3**: 59–85. doi:10.1146/annurev-devpsych-050620-030227
- Infanti E, Hickey C, Turatto M. 2015. Reward associations impact both iconic and visual working memory. *Vision Res* **107**: 22–29. doi:10.1016/j.visres.2014.11.008
- Insel C, Somerville LH. 2018. Asymmetric neural tracking of gain and loss magnitude during adolescence. *Soc Cogn Affect Neurosci* **13**: 785–796. doi:10.1093/scan/nsy058
- Insel C, Kastman EK, Glenn CR, Somerville LH. 2017. Development of corticostriatal connectivity constrains goal-directed behavior during adolescence. *Nat Commun* **8**: 1605. doi:10.1038/s41467-017-01369-8
- Insel C, Charifson M, Somerville LH. 2019. Neurodevelopmental shifts in learned value transfer on cognitive control during adolescence. *Dev Cogn Neurosci* **40**: 100730. doi:10.1016/j.dcn.2019.100730
- Katzman PL, Hartley CA. 2020. The value of choice facilitates subsequent memory across development. *Cognition* **199**: 104239. doi:10.1016/j.cognition.2020.104239
- Konkle T, Brady TF, Alvarez GA, Oliva A. 2010. Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *J Exp Psychol Gen* **139**: 558–578. doi:10.1037/a0019165
- Krause F, Lindemann O. 2014. Expyriment: a Python library for cognitive and neuroscientific experiments. *Behav Res Methods* **46**: 416–428. doi:10.3758/s13428-013-0390-6
- Krebs RM, Boehler CN, Egner T, Woldorff MG. 2011. The neural underpinnings of how reward associations can both guide and misguide attention. *J Neurosci* **31**: 9752–9759. doi:10.1523/JNEUROSCI.0732-11.2011
- Kuhl BA, Shah AT, DuBrow S, Wagner AD. 2010. Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *Nat Neurosci* **13**: 501–506. doi:10.1038/nn.2498
- Leong YC, Radulescu A, Daniel R, DeWoskin V, Niv Y. 2017. Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron* **93**: 451–463. doi:10.1016/j.neuron.2016.12.040
- Loosli SV, Rahm B, Unterrainer JM, Weiller C, Kaller CP. 2014. Developmental change in proactive interference across the life span: evidence from two working memory tasks. *Dev Psychol* **50**: 1060–1072. doi:10.1037/a0035231
- Ma DS, Correll J, Wittenbrink B. 2015. The Chicago face database: a free stimulus set of faces and norming data. *Behav Res Methods* **47**: 1122–1135. doi:10.3758/s13428-014-0532-5
- Madan CR, Fujiwara E, Gerson BC, Caplan JB. 2012. High reward makes items easier to remember, but harder to bind to a new temporal context. *Front Integr Neurosci* **6**: 61. doi:10.3389/fnint.2012.00061
- Meyer HC, Bucci DJ. 2016. Age differences in appetitive Pavlovian conditioning and extinction in rats. *Physiol Behav* **167**: 354–362. doi:10.1016/j.physbeh.2016.10.004
- Miendlarzewska EA, Aberg KC, Bavelier D, Schwartz S. 2021. Prior reward conditioning dampens hippocampal and striatal responses during an associative memory task. *J Cogn Neurosci* **33**: 402–421. doi:10.1162/jocn_a.01660
- Murty VP, Tomparay A, Adcock RA, Davachi L. 2017. Selectivity in postencoding connectivity with high-level visual cortex is associated with reward-motivated memory. *J Neurosci* **37**: 537–545. doi:10.1523/JNEUROSCI.4032-15.2016
- Ngo CT, Newcombe NS, Olson IR. 2019. Gain-loss framing enhances mnemonic discrimination in preschoolers. *Child Dev* **90**: 1569–1578. doi:10.1111/cdev.13297
- Niv Y, Daniel R, Geana A, Gershman SJ, Leong YC, Radulescu A, Wilson RC. 2015. Reinforcement learning in multidimensional environments relies on attention mechanisms. *J Neurosci* **35**: 8145–8157. doi:10.1523/JNEUROSCI.2978-14.2015
- Nussenbaum K, Hartley CA. 2019. Reinforcement learning across development: what insights can we draw from a decade of research? *Dev Cogn Neurosci* **40**: 100733. doi:10.1016/j.dcn.2019.100733
- Raab H, Foord C, Ligneul R, Hartley CA. 2022. Developmental shifts in computations used to detect environmental controllability. *PLoS Comput Biol* **18**: e1010120. doi:10.1371/journal.pcbi.1010120
- R Core Team. 2019. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Roper ZJJ, Vecera SP, Vaidya JG. 2014. Value-driven attentional capture in adolescence. *Psychol Sci* **25**: 1987–1993. doi:10.1177/0956797614545654
- Schulman AI. 1974. Memory for words recently classified. *Mem Cognit* **2**: 47–52. doi:10.3758/BF03197491
- Shing YL, Brod G. 2016. Effects of prior knowledge on memory: implications for education. *Mind Brain Educ* **10**: 153–161. doi:10.1111/mbe.12110
- Sturman DA, Mandell DR, Moghaddam B. 2010. Adolescents exhibit behavioral differences from adults during instrumental learning and extinction. *Behav Neurosci* **124**: 16–25. doi:10.1037/a0018463
- Wang Y, Braver TS, Yin S, Hu X, Wang X, Chen A. 2019. Reward improves response inhibition by enhancing attentional capture. *Soc Cogn Affect Neurosci* **14**: 35–45. doi:10.1093/scan/nsy111
- Wimmer GE, Shohamy D. 2012. Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science* **338**: 270–273. doi:10.1126/science.1223252
- Winter W, Sheridan M. 2014. Previous reward decreases errors of commission on later ‘No-Go’ trials in children 4 to 12 years of age: evidence for a context monitoring account. *Dev Sci* **17**: 797–807. doi:10.1111/desc.12168

Received March 30, 2022; accepted in revised form September 14, 2022.



Reward-motivated memories influence new learning across development

Alexandra O. Cohen, Camille V. Phaneuf, Gail M. Rosenbaum, et al.

Learn. Mem. 2022, **29**:

Access the most recent version at doi:[10.1101/lm.053595.122](https://doi.org/10.1101/lm.053595.122)

**Supplemental
Material**

<http://learnmem.cshlp.org/content/suppl/2022/10/13/29.11.421.DC1>

References

This article cites 45 articles, 8 of which can be accessed free at:
<http://learnmem.cshlp.org/content/29/11/421.full.html#ref-list-1>

**Creative
Commons
License**

This article is distributed exclusively by Cold Spring Harbor Laboratory Press for the first 12 months after the full-issue publication date (see <http://learnmem.cshlp.org/site/misc/terms.xhtml>). After 12 months, it is available under a Creative Commons License (Attribution-NonCommercial 4.0 International), as described at <http://creativecommons.org/licenses/by-nc/4.0/>.

**Email Alerting
Service**

Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#).
