PNAS RESEARCH ARTICLE | Psychological and Cognitive Sciences

Developmental changes in memory structure and precision alter the use of retrieved episodes during decisions for reward

Nora C. Harhen[a,1], Aaron M. Bornstein[b,c*], and Catherine A. Hartley[a*]

This manuscript was compiled on May 10, 2026

Most widely studied option evaluation strategies rely on knowledge accumulated across repeated experiences. But how should options be evaluated in unfamiliar environments, in which knowledge is sparse? In these situations, how do decision makers make efficient use of limited past experience to guide their choices? One possible strategy is episodic sampling, in which a decision maker retrieves a small number of past decisions from memory to estimate the value of present options. By virtue of their age, children and adolescents have less experience than adults, making episodic sampling a particularly useful strategy for them. At the same time, the effectiveness of episodic sampling derives from memory's precision and context sensitivity—properties that continue to develop into adolescence and young adulthood. This tension raises a key question: are developmental differences in episodic memory associated with differences in episodic sampling? To address this question, 106 participants, ages 8-25, completed a two-day choice task that dissociated the influence of a single episodic memory from the influence of multiple episodes sharing a common context. At all ages, single episodes biased choices. But, only adults were sensitive to the broader evoked context. Further clarifying the relationship between episodic memory and decision making, differences in memory precision predicted differences in episodic sampling, even after taking into account age, while episodic sampling, in turn, accounted for individual differences in forward planning. Together, these findings suggest that episodic memory guides decision making throughout development, but the character of its influence evolves as memory becomes more precise and richly structured.

Episodic memory | Reinforcement learning | Decision-making | Development

Significance

Decision making and memory both change dramatically across childhood and adolescence, but they have largely been studied in isolation. Here, we bridge this gap by showing that as early as middle childhood, a single fleeting event can systematically bias choices a day later. However, only older participants draw on richer episodic context to guide their choices, and this age-related difference may be tied to improvements in memory precision. Our findings suggest that developmental changes in decision making may be driven, in part, by the maturation of episodic memory.

Humans regularly make adaptive decisions in new environments despite having limited experience. This ability is especially remarkable in children and adolescents, for whom many environments are novel. Episodic memory may be a key ingredient to this "sample efficiency" (1–7). A defining feature of episodic memory is that it can be highly detailed. Within an episodic memory of a past decision, a person may encode not only what they chose and its immediate outcome but also when and where the decision occurred and its longer-term consequences. These details allow individuals to go beyond their direct experience, enabling them to vividly imagine counterfactual pasts and distant futures as they decide what to do next (8–10). A growing body of work demonstrates that adults draw on episodic memory to guide decision making (11–13), particularly in uncertain or volatile environments (14), and even when recency-based evaluation strategies should dominate (15, 16). Given their more limited experience, children and adolescents may benefit even more from episodic sampling than adults. But despite the growing evidence in adults, little is known about how episodic memory guides decision making during childhood and adolescence.

Author affiliations: [a]Department of Psychology, New York University; [b]Department of Cognitive Sciences, University of California, Irvine; [c]Center for the Neurobiology of Learning & Memory, University of California, Irvine

---

Please provide details of author contributions here.

The authors declare no competing interests.

[*] A.M.B. (Author Two) contributed equally to this work with C.A.H. (Author Three).

[1] To whom correspondence should be addressed. E-mail: nora.harhen@nyu.edu

Research on developmental changes in decision making has focused largely on incremental learning. In incremental learning, an action's value is estimated by averaging over many past outcomes, thereby discarding the idiosyncratic features of each contributing experience. This form of learning strengthens across development, growing more flexible and sensitive to context (17). Episodic sampling offers an alternative decision strategy. Rather than averaging across all prior outcomes, episodic sampling estimates action values by averaging over a subset of past decisions retrieved from memory. The number of memories can vary. At one extreme, a decision maker can rely on a single memory. For instance, the decision maker, recalling a particularly delicious meal enjoyed years ago at a Japanese-Peruvian fusion restaurant in Peru, may be swayed toward trying a new ramen shop rather than returning to their beloved halal cart down the block. But, when multiple memories are retrieved, the initial memory can cue the recall of other memories that share a common context, such as time or place. In this case, recalling one meal from Peru might trigger memories of other meals from the same trip, leaving the decision maker to forgo both the ramen shop and halal cart and choose a Peruvian restaurant instead. In this way, a handful of memories from the distant past can exert an outsized influence on present choices, producing systematic departures from the predictions of incremental learning (15).

Episodic sampling should be particularly useful for children and adolescents, given their relatively limited experience. However, the very features of episodic memory that make episodic sampling efficient—context sensitivity and precision—are among the slowest to develop. Although children can recognize previously seen individual items as well as adults, they struggle to remember the contextual details that situate these items in time and place (18–25). Consequently, they have difficulty organizing experiences around overlapping contexts (26) and binding people, places, objects, and time into coherent events (27, 28). During the same developmental period, children also encode individual episodes with less precision or fidelity than adults (29–37). Accordingly, because episodic sampling ought to depend on the structure and fidelity of memory (38, 39), it should change in tandem with memory development. Here, we asked how episodic sampling varies from childhood to young adulthood. Based on prior research on the development of memory, we hypothesized that children and adolescents would use episodic memory to inform their decisions but the nature of its contribution would differ from that of adults. Specifically, we predicted that younger participants would be able to draw on information from a single episode to guide their choices but would be less likely than adults to sample multiple contextually linked episodes. To test this prediction, we recruited a large, age-continuous sample of 8- to 25-year olds to complete, over multiple days, a task designed to dissociate the influence of a reminded episode from that of its broader context (15). In addition to the main task, participants completed two more tasks. One task allowed us to test a potential mechanistic explanation for developmental differences in episodic sampling. The other allowed us to test the hypothesis that developmental differences in episodic sampling partially explain known developmental differences in planning.

Results

We analyzed the behavior of 106 participants, between the ages of 8 and 25 years (M=16.96, SD=5.09, N=54 female), who completed the two-day study. On the first day, participants completed the encoding phase. During this phase, participants made choices and encoded images of objects presented alongside the outcomes of their choices. In the task's narrative framing, participants were the captain of a pirate ship, choosing which of three crew members—Red Beard, Black Beard, or White Beard—would attempt to rob the next passing cargo ship (Fig. 1A). Each choice returned, with varying probabilities, either a pile of gold coins if the crew member was successful or a red X if they were not (Fig. 1B). Participants were told that the number of gold coins they collected across the entire study would determine their bonus pay. Each choice also returned a trial-unique image of an object alongside the outcome. On the second day of the study, these object images would serve as memory probes intended to evoke trials encountered on the previous day.

Participants completed 240 of these trials. The trials took place across six distinct "islands" or contexts, distinguished by unique background images of outdoor scenes. The payoff time series was carefully structured across the islands. On the very first trial, each crew member was assigned a target probability of yielding gold coins (60%, 30%, or 10%), which then gradually drifted over the first 10 trials before being reshuffled, with the new mapping persisting for the next 40 trials before being reshuffled again on a new island (Fig. 1B). Thus, the crew member that was most rewarded early in the island was not the same one that was most rewarded across the full duration of the island.

On the second day, participants completed the test phase consisting of 180 choice trials in which they chose between the same three crew members from the previous day. Interleaved at pseudorandom intervals were 60 recognition memory trials, on which participants indicated whether, and with what confidence, they recognized a presented object image (Fig. 1C). Participants earned additional bonus pay for correct responses. Critically, the memory probes were drawn exclusively from object images presented during the first 10 trials of each island, prior to the reshuffling of payoff probabilities. This design allowed us to dissociate the influence of a single evoked trial from that of the broader context on choice. More concretely, if the probe triggered retrieval of multiple episodes from the same context, it should increase the likelihood of choosing the option most rewarded across the context, rather than the option rewarded on the directly probed trial itself. Finally, at the end of the test phase, participants completed a source recognition memory test that assessed whether they could identify the island (context) in which the probed objects had been encountered during the encoding phase.

Decision making is guided by individual episodic memories from childhood but is only informed by episodic context in adults. We first examined how recent and distant experience influenced choices on the second day of the study. Recent experience

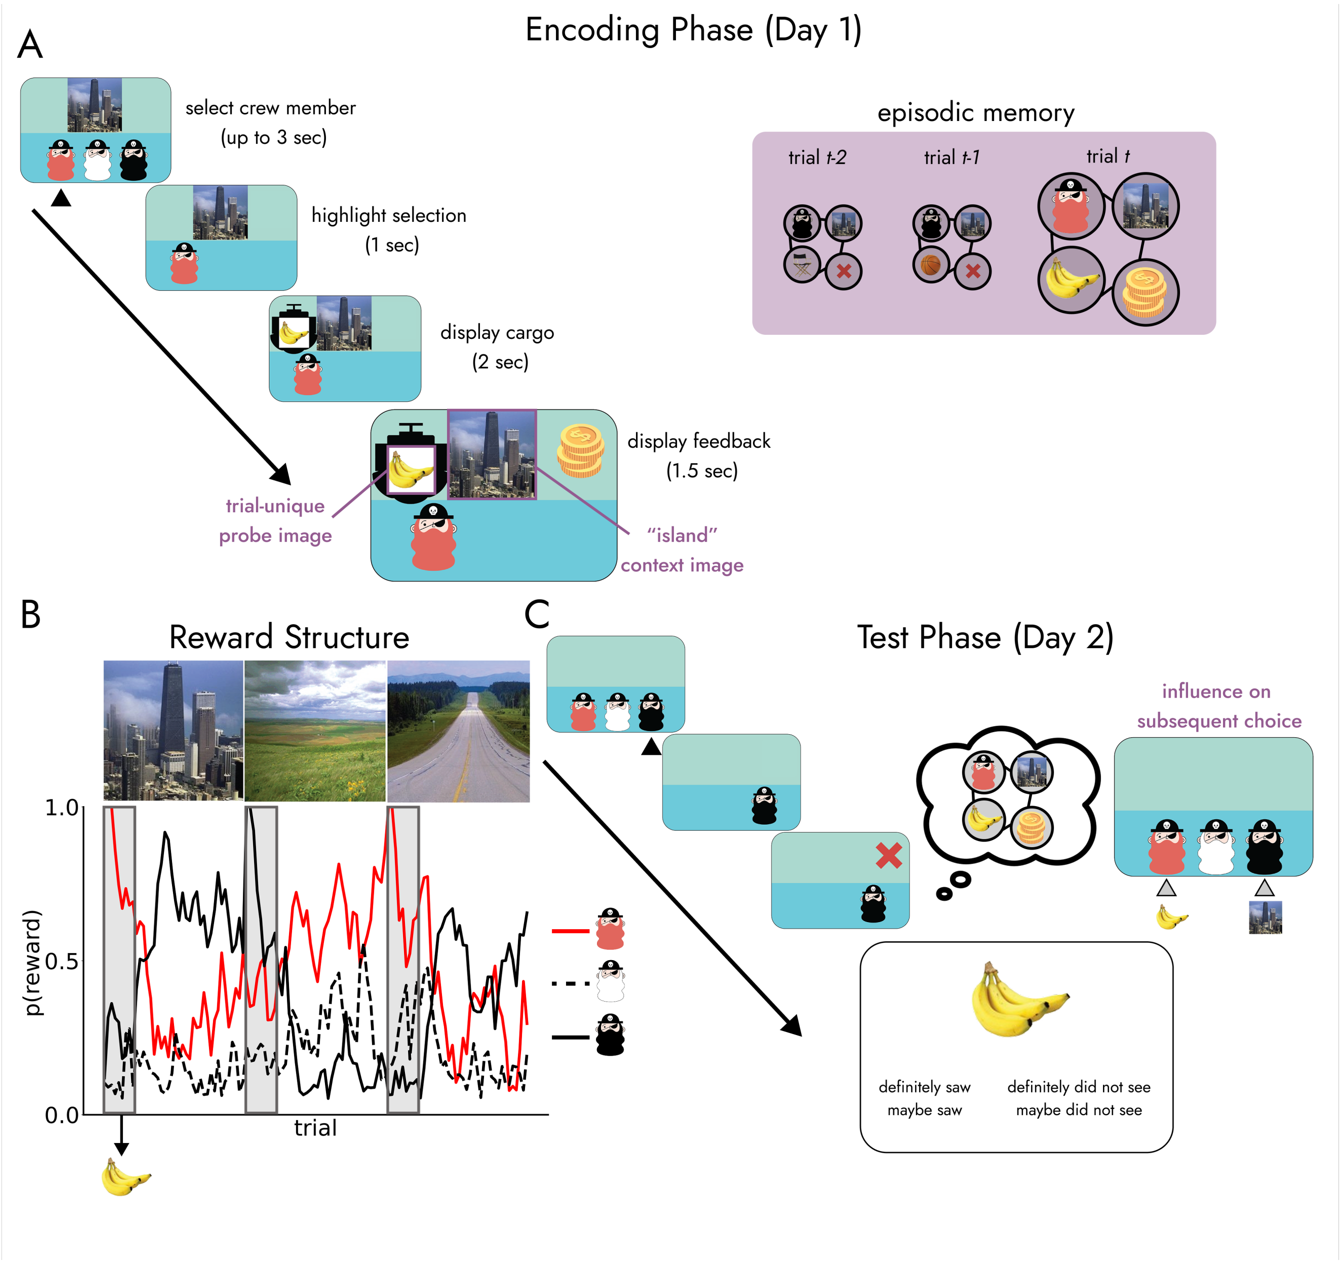

Fig. 1. Task design. Task design. Participants completed a two-session three-armed bandit task, consisting of an encoding phase and the test phase. The task was designed to isolate the effects of retrieved context on decision-making, distinct from that of the single sampled episode. **A. Day One - Encoding phase.** Participants completed 240 choice trials across six consecutive "islands." Participants experienced each island, distinguished by a unique background image, for 40 trials. On each trial, participants chose one of three crew members to rob an incoming cargo ship. After selecting a crew member, a trial-unique image of an object (the ship's "cargo") was displayed, followed by feedback indicating whether or not the crew member was successful in robbing the ship. If they were successful, a pile of gold coins was shown; if not, a red "X" was shown. Participants were instructed that their bonus payment would depend on the total amount of coins collected. We hypothesized that on each trial the trial-unique image, the context image, the chosen crew member, and the outcome of the choice were bound together into an integrated representation and stored in episodic memory. **B. Reward structure.** Each crew member had a distinct, gradually changing probability of success. At the beginning of each phase, crew members were initialized with 60%, 30%, and 10% success rates. These probabilities varied across trials according to a Gaussian random walk. After the first 10 trials of each island, the probabilities were shuffled so that the most successful crew member changed. These new mappings persisted through the remainder of the island and into the next, until the next shuffle 10 trials into the new island. We show an example of how crew members' probabilities of reward might vary across three different islands. Each line represents a crew member. Background images indicate the current context for trials on the x-axis. Shaded boxes mark the first 10 trials of each island. **C. Day Two - Test phase.** The following day, participants completed 180 choice trials on a "foggy" island, during which the background image and object images were obscured. Crew member success probabilities continued to drift and were reshuffled every 40 trials, despite the background remaining unchanged. Interleaved between these 180 choice trials were 60 memory probe trials. On 50 of these trials, images from the prior day's encoding task were shown and on 10 trials, novel images were shown. The key measure of interest was the extent to which choices directly following the memory probe trials were biased by the probes and the reward information associated with them. Critically, the reward structure allowed us to isolate the contributions of a single retrieved trial from the contributions of multiple trials from a shared context. In the example displayed here, the retrieved trial should bias the participant to select Red Beard because he was rewarded on that trial while the retrieved context should bias the participant towards selecting Black Beard because he most often led to reward on the island with the cityscape background, in which the retrieved trial occurred.

refers to trials immediately preceding the current choice, while distant experience refers to trials from the previous day that have been evoked by memory probes. Following Bornstein & Norman (2017), we fit a mixed-effects logistic regression model with predictors for the previously chosen option, rewards received over the last three trials, the option chosen on the probed trial, the reward received on that trial, and the frequency with which each option was rewarded across the context in which the probed trial occurred (Table S1). The model also included age, its interactions with all predictors, and participant-level random intercepts and slopes.

This analysis revealed significant, separable effects of recent and distant experience on choice. Participants showed a tendency to repeat their last choice regardless of outcome (Fig. 2A; b=0.99, SE=0.015, p <.001). Consistent with model-free reinforcement learning, recent outcomes also influenced choice. Participants were more likely to select options rewarded over the last three trials, with more recent trials having a stronger effect (Fig. 2A; reward-1: b=0.71, SE=0.047, p <.001; reward-2: b=0.26, SE=0.026, p <.001; reward-3: b=0.18, SE=0.024, p <.001). However, when memory probes were presented, experiences from the previous day also strongly biased choice, an effect inconsistent with model-free RL. The reward on the probed trial had only a marginal effect on choice (b=0.046, SE=0.026, p =.083), but the chosen option was significantly more likely to be selected again on the current trial (Fig. 2B; b=0.079, SE=0.038, p =.036). Stronger still, rewards across the probed context exerted an effect comparable to that of a reward received only three trials earlier (b=0.17, SE=0.059, p =.0031).

Our results closely mirror those found in Bornstein & Norman (2017) despite the test phase occurring a full day after encoding rather than immediately afterward. Strikingly, these effects emerged despite relatively low overall explicit memory performance (probe recognition: mean $d' = 0.83$, SD = 0.76; source memory: mean proportion correct = 0.18, SD = 0.07; see S1.2 for additional memory performance results and analyses of relationships between memory measures). Nevertheless, both probe recognition and source memory were significantly above chance (probe recognition: one-sample t-test against 0, t(105) = 11.25, p < .001; source memory: one-sample t-test against 0.17, t(105) = 2.57, p = .012). This gap between the influence of memory probes on choice and explicit memory performance suggests that episodes and their associated contextual information are retrieved implicitly, shaping decisions even when explicit memory is weak (see S1.3 for further discussion).

Overall, our findings are consistent with a memory-sampling account in which a probe evokes a specific episode that triggers the retrieval of other contextually linked episodes. However, our findings are also consistent with an alternative account in which participants learn context-specific cached action values, as in model-free reinforcement learning. Under this account, the probe reinstates the context and its associated value, eliminating the need to retrieve individual episodes at decision time. To arbitrate between this account and our proposed sampling account, we ran two additional analyses. First, we fit a reinforcement learning model to participants' choices on the first day of the study. Importantly, at the beginning of each context, the model reset its action values and stored those of the previous context. We then used this model to extract context-specific action values and included them alongside the sampling-based context reward predictor in the regression. When both predictors were included, RL-derived action values were less predictive of choice than the sampling-based predictor (S1.4). Next, we examined whether the temporal context of the memory probe modulated its influence on choice. Under the reinforcement learning account, the probe's temporal position should not matter, as it reinstates a cached value. In contrast, the sampling-based account predicts sensitivity to temporal position, because the probed episode would trigger retrieval of temporally proximal experiences. Supporting the sampling account, the effects of both choice and reward on the probed trial were stronger when the probed trial occurred earlier in the context (S1.4.2).

Having replicated the results from Bornstein & Norman (2017), we next examined how the effects of recent and distant experience varied with age. Echoing prior work showing that reward sensitivity increases from childhood into adolescence (40–42), the effects of the most recent choice and its outcome were stronger in older participants (last choice x age: b=0.28, SE=0.015, p <.001; reward-1 x age: b=0.10, SE=0.047, p =.035). However, rewards from two and three trials back showed no age-related differences (reward-2 x age: b=0.0042, SE=0.026, p =.87; reward-3 x age: b=0.018, SE=0.024, p =.47). Turning to the influence of distant experience, we found distinct age-related patterns for the effects of individual episodes versus broader contexts. The choice and outcome on the probed trial had similarly strong effects across the age range (Fig. 2C; probed trial choice x age: b=-0.040, SE=0.038, p =.29; probed reward trial x age: b=-0.0089, SE=0.026, p =.74). Even the youngest participants' choices were biased by a single episode from the previous day. In contrast, the influence of the probed context was stronger in older participants (Fig. 2D; probed context x age: b=0.14, SE=0.059, p =.02). Despite comparable recognition memory across age groups (b=0.052, SE=0.075, p =.49), adults relied more on outcomes across the probed context than did children or adolescents. These findings suggest that, from middle childhood onward, decision makers rely on individual episodes to guide choice, but only in adolescence and young adulthood do they draw on multiple contextually linked episodes.

Memory precision mediates the relationship between age and context-guided episodic sampling. Given the age-related differences in the extent to which choices are biased by an evoked context, we next asked why this variation arises. Specifically, we examined whether memory precision, assessed in a separate task, might account for this variation. Memory precision reflects the dimensionality and resolution of mnemonic representations–the number of perceptual, spatial, temporal, and contextual features on which retrieval can be anchored (43). If younger participants encode episodes with less precision, as prior work suggests (29–37), they may have fewer retrieval anchors, reducing their ability to rapidly reinstate and integrate reward information across an entire context before making a choice.

To test this hypothesis, we invited participants to complete the Mnemonic Similarity Task in a separate session;

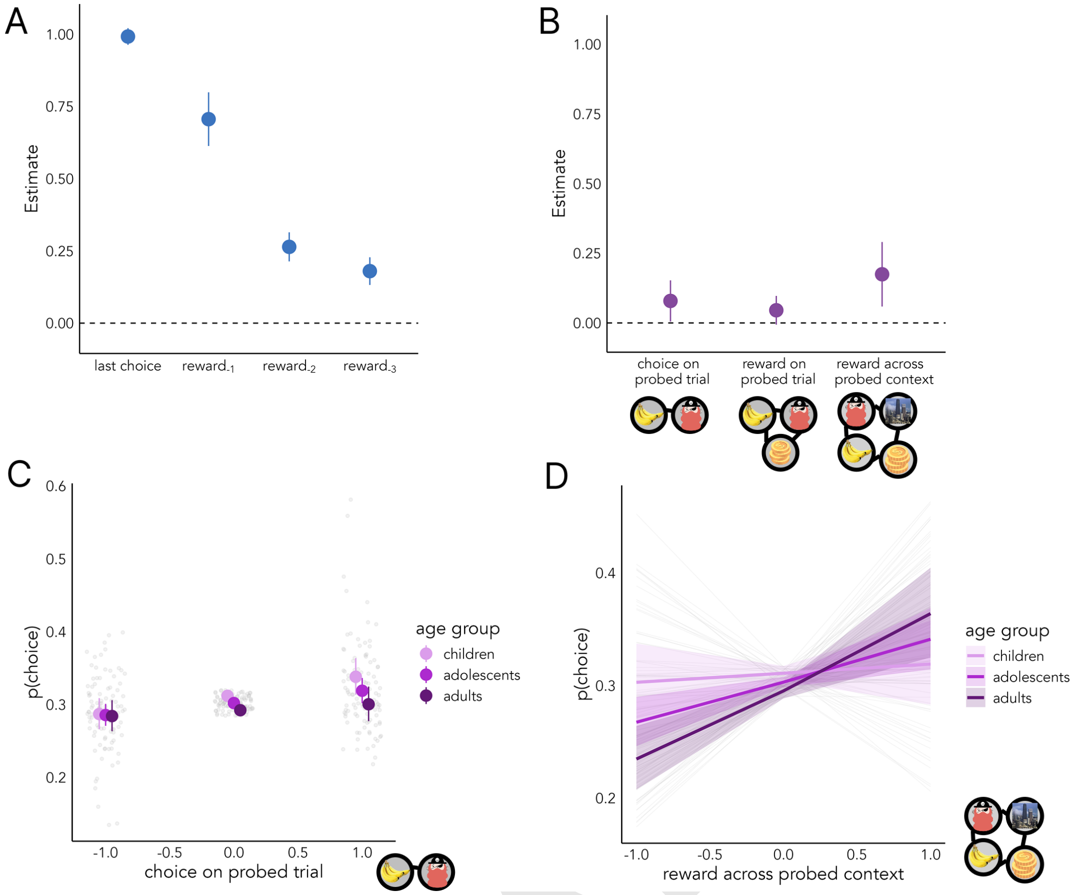

Fig. 2. Episodic memory guides choice in middle childhood and continues to strengthen into young adulthood. A-B. Episodic memory guides choice in middle childhood and continues to strengthen into young adulthood. Estimates of β weights from a mixed-effects logistic regression are shown with medians and 95% confidence intervals. **A. Recent experience shapes choice.** Participants tended to repeat their selection on the last trial, regardless of its outcome ("last choice"), and to choose options rewarded up to three trials earlier (reward-1, reward-2, reward-3). **B. Episodic memories of distant experience shapes choice.** On trials directly following a memory probe (the ship's "cargo"), participants favored the option chosen on the trial evoked by the probe ("choice on probe trial"), while the evoked choice's evoked outcome did not affect their choice ("reward on the probed trial"). Participants also favored options that were frequently rewarded across the context ("island") associated with the evoked trial. **C. The influence of a single evoked trial is stable from middle childhood to young adulthood.** Markers show the estimated probability of choosing an option when it was selected on the evoked trial ("same"), was not selected ("different"), and when no trial was recently evoked ("irrelevant"). Purple markers denote group averages, and gray markers denote individual participants. Choice on the evoked trial significantly influenced subsequent choices, whereas reward on that trial did not. Error bars indicate 95% confidence intervals. While age is plotted discretely (children: 8–12.99 years; adolescents: 13–17.99; adults: 18–25), it was included as a continuous variable in the statistical model. **D. The influence of an evoked context strengthens with age.** Lines show the estimated probability of choosing an option as a function of how frequently it was rewarded across the context ("island") evoked by the probe. Purple lines denote group averages, and gray lines denote individual participants. Shading indicates 95% confidence intervals.

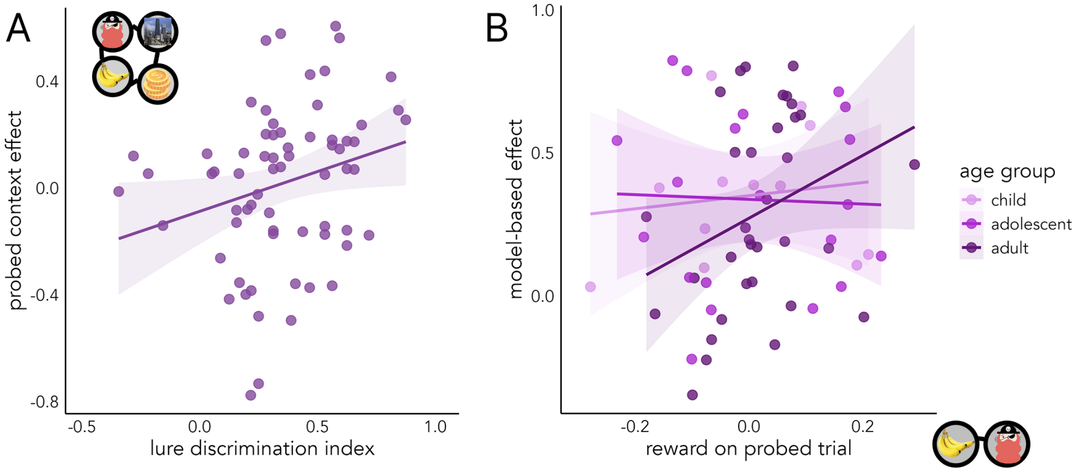

Fig. 3. Memory precision and forward planning track the influence of episodic memory on decision making. Markers indicate individual participants and shaded bands show 95% confidence intervals. **A.** Across participants, greater memory precision (higher lure discrimination index) predicted a stronger influence of the evoked context's reward history on choice. This was the case even when taking age into account. **B.** A greater influence of the reward received on the probed trial was associated with greater forward planning (model-based effect), most strongly in adults.

67 of them returned (see S2 for demographic information). The Mnemonic Similarity Task is a recognition memory task in which participants distinguish previously seen objects from perceptually similar lures and novel foils (44). As our measure of memory precision, we used the lure discrimination index (LDI), the difference between the rate of “similar” responses to lures versus foils. Consistent with prior findings, LDI was higher in older participants (b = 0.064, SE = 0.030, p = .037), indicating greater memory precision.

We next tested for associations between memory precision and each of the distant-experience predictors (i.e., probed trial choice, probed trial reward, and rewards across the probed context), extracting participant-level random slopes from the regression model (Tables S5,S6,S7). Because the Mnemonic Similarity Task measures the precision with which individual objects are encoded, one might expect memory precision to be related to the influence of individual episodes on choice. Contrary to this, memory precision was not significantly associated with sensitivity to the probed trial’s choice or outcome (probed trial choice: b = -0.23, SE = 0.14, p = .096; probed trial reward: b = 0.046, SE = 0.060, p = .45). Strikingly, however, there was a significant association with sensitivity to the probed context (Fig. 3A; b = 0.30, SE = 0.14, p = .039); individuals with greater memory precision were more likely to select an option that was rewarded across the entire probed context. Initially, this may seem surprising given that prior work has found memory precision for individual items to be unrelated to item–context binding in children and young adults (45–47). However, memory precision for individual items may nonetheless track precision for contextual details without necessarily reflecting the ability to bind items to context.

Given that age predicted memory precision and memory precision predicted sensitivity to context, we next conducted a mediation analysis with age as the predictor, context sensitivity as the outcome, and memory precision (LDI) as the mediator. Because the effects of probed-trial reward and choice did not vary with age in our main sample, we did not conduct a mediation analysis for these predictors. Supporting our hypothesis, memory precision fully mediated the relationship between age and sensitivity to the probed context (b = 0.019, SE = 0.00029, p = .048). Thus, the general precision of mnemonic representations is related to their retrieval during decision making.

Use of individual episodic memories is positively associated with forward planning in adults. Having found age-related differences in episodic sampling, we next asked whether these differences might explain developmental variation in planning (48–50). We were motivated by proposals that hippocampal forward replay, the sequential reactivation of previously experienced states, supports planning over a mental model of the environment (51, 52). To address this question, we had the same subset of participants who completed the Mnemonic Similarity Task also complete a two-stage decision making task (48, 53), in which first-stage choices probabilistically lead to one of two second-stage states, where participants choose between two options that probabilistically yield reward. This task dissociates reliance on a mental model of the environment (model-based control) from reliance on cached action values (model-free control). Because the age-related differences we observed in episodic sampling were specific to the probed context, we hypothesized that sensitivity to the probed context, but not sensitivity to the probed trial, would predict planning across development. When we did not find a significant association between those two, we conducted subsequent exploratory analyses to test whether the other two distant-experience predictors could explain individual variation in planning independent of age. Sensitivity to distant experience was indexed as in the Mnemonic Similarity Task analyses, using participant-level random slopes from the main regression.

Contrary to our hypothesis, the influence of probed reward–but not probed choice or context (Tables S8,S9,S10)–predicted greater forward planning, with the effect being strongest in older participants (Figure 3B; probed trial reward x age: b = 0.81, SE = 0.38, p = .038; probed trial reward: b = 0.59, SE = 0.31, p = .066; age: b = 0.048, SE = 0.037, p = .20). Although developmental variation in episodic sampling did not explain differences in planning, these results suggest that in older participants, the retrieval of specific past rewards may support prospective planning, consistent with recent theoretical proposals linking hippocampal replay to model-based control (51, 52).

Discussion

Adults’ decisions are shaped not only by recent experience but also by memories of the distant past. Whether children and adolescents also rely on such memories remains unclear. To address this gap, we had participants aged 8 to 25 years old complete a two-day, three-option decision-making task that dissociated two potential mnemonic contributions to choice: the influence of a single retrieved episode (i.e., a trial) and that of multiple episodes linked by a shared context. On the first day, each choice trial was tagged with a unique object image. On the second day, these images served as incidental reminders of specific trials from the previous day. Participants across our age range were biased towards the option chosen on the reminded trial, indicating that episodic memory contributes to decision making from an early age. However, sensitivity to the episodic context was only evident in older participants, who preferred the option most often rewarded across the entire reminded context. Our results suggest that episodic memory supports decision making in at least two ways, each with a distinct relationship to age.

Children as young as eight years old drew on individual episodes to guide their decisions. While item memory reaches relative maturity by middle childhood (19, 54), retrieving a single episode still does not guarantee that it will be used to inform choice. Moreover, there are reasons to expect that children might have been less adept in using this information. In adults, memory-guided decision making is associated with increased hippocampal-ventromedial prefrontal cortex connectivity (55), and connectivity between these regions continues to strengthen into young adulthood (56), making our result, in some ways, surprising. One potential explanation is that, across development, different neural mechanisms may support memory-guided choice. Consistent with this notion, hippocampal-striatal coupling is heightened in adolescents relative to adults during value-based decision making (57). This raises the possibility that, in childhood and adolescence, interactions between subcortical regions coordinate mnemonic, value, and action information, supplanting later-developing contributions of the prefrontal cortex to memory-guided behavior (58).

While children's choices were influenced by individual episodes, they were not influenced by the broader context. What might explain this dissociation? Our data suggest memory precision as one key factor. Replicating prior work (29–37), we found that memory precision increased with age and fully accounted for developmental variation in context sensitivity. Adults may encode episodes with richer contextual detail, providing them with more anchors for retrieval, and increasing the likelihood that they could rapidly reinstate related episodes at decision time (43).

Differences in memory structure likely also contribute to the age-related increase in the sensitivity of choice to contextual memory. Across childhood and adolescence, improvements in item-context binding (19, 33, 34, 59–62) support the organization of memories around contextual features, including time and location (24, 63–65). Episodes that share contextual features can be integrated into coherent representations–a process known as memory integration. Children are capable of memory integration but tend to do so only at retrieval and when prompted (66–68), whereas adults integrate related episodes spontaneously at encoding (69). Adults, with an integrated structure already in place, are able to rapidly reinstate clusters of related episodes at decision time; children, lacking this structure, may only be able to retrieve the episode directly evoked by the probe. Together, age-related differences in both memory precision and structure provide plausible accounts for why sensitivity to evoked contextual information is present in older participants but not younger. Adults' high-fidelity memories preserve rich contextual detail, allowing episodes to be organized into structured representations that can be rapidly reinstated when needed.

Hippocampal replay, the sequential reactivation of previously experienced states, has been implicated in planning (51, 52), and planning improves across development (48–50). We therefore asked whether memory-guided decision making explains age-related differences in planning. We first tested whether model-based control was predicted by context sensitivity, as it was the index of memory-guided decision making we found to increase with age, but found no significant association. We then asked whether variation in planning *independent of age* was related to other indices of memory-guided decision making. Greater sensitivity to reward on the probed trial was associated with increased model-based control, but only in adults. This may reflect developmental limitations in the ability to integrate episodes. Although children can draw on a single episode to guide choice, they may fail to assemble multiple episodes into coherent, forward-looking plans (10).

In adults, reinforcement learning and episodic memory systems are known to reciprocally interact. Reward and motivational states enhance memory and consolidation (70–76), while episodic memory is recruited to support value-based decisions (11–13, 15, 16, 77, 78). Recent research has extended these findings developmentally, characterizing how the relationships between episodic memory, reward, and choice vary with age (79). A majority of these studies have focused on how reward and choice influence memory (57, 80–85). Our findings complement this work by demonstrating that episodic memory also shapes decision making in children, adolescents, and young adults. Significantly, the form of memory's contribution differs with age, tracking known changes in the precision and structure of memory (24, 31, 32, 36, 63). A growing body of work has characterized age-related variation in model-free reinforcement learning, another strategy for option evaluation (86–93). Our findings raise the question–do these age-related differences in model-free reinforcement learning stem from differences in memory? Others have similarly asked whether the development of memory could explain known developmental changes in cognitive flexibility (94). These questions motivate future longitudinal studies testing whether improvements in episodic memory scaffold the development of decision-making strategies. Our work opens up new directions for understanding how episodic memory and decision making processes jointly develop.

Materials and Methods

Participants. We recruited 124 participants between the ages of 8 and 25 (M age = 16.71, SD = 5.06, 62 female) for an online study consisting of four sessions. The encoding phase of the memory and decision-making task was completed in Session 1, and the test phase in Session 2 the following day. Participants then completed the memory precision task in Session 3, one to two weeks later, and the forward planning task in Session 4, an additional one to two weeks after that.

Participants were recruited from the Hartley lab database which includes individuals recruited through ads on social media (e.g., Facebook and Instagram), word of mouth, local science fairs, and flyers on New York University's campus. Participants who had not previously completed an in-person study with our lab completed a brief Zoom call with a researcher. During this call, participants (and their parent or guardian, if the potential participant is under 18 years) were required to be on camera and confirm their full name and date of birth. Adult participants and parents of child and adolescent participants were additionally required to show photo identification.

Participants were recruited to ensure even coverage of the age range. Participants reported no history of psychiatric or learning disorders. Based on self- or parent-report, 43.55% of participants were White, 33.06% were Asian, 13.71% were more than one race, and 9.68% were Black. Additionally, 12.01% of the sample identified as Hispanic.

Eighteen participants were excluded for their behavior in the main task (M age= 15.20 years, SD = 4.44, range = 8.77 to 25.01, 8 females). Exclusion criteria included: failing to respond on more than 10% of trials in either the encoding or test phase (4 excluded), selecting one of the three options on fewer than 5% of trials in either phase (8 excluded), responding on average in less than 250 milliseconds (1 excluded), interacting with their browser more than 10 times in either phase (10 excluded). Data from one additional participant were not analyzed because of technical issues during data collection. This left a final sample of 106 participants (mean age= 16.96, SD= 5.09, range=8.03 to 25.97, 54 females). Our target sample size of 100 was based on power analyses using the effect sizes obtained from previous variants of the task. Bornstein & Norman (2017) performed a one-sample t-test of participants' individual estimates of the influence of context memory on choice against 0. With their obtained effect size, d=0.79, a sample size of 15 would be needed to achieve 80% power at α=.05. Noh et al. (2023) examined individual differences in memory-guided decision making, finding a moderately strong correlation (ρ_τ=0.22) between memory precision and the influence of episodic memory on choice. To achieve 80% power at α=.05, 66 participants would be necessary. Our sample size surpasses both these targets and is larger than many previous studies of age-related variation in memory and decision-making strategies (95–98).

A subset of the included participants (N=67) also completed the mnemonic similarity task, indexing memory precision, and the two-step task, indexing forward planning. For each of the four sessions, participants were compensated with a $10 Amazon gift card and had the opportunity to earn up to $2 in bonus payment based on task performance.

Memory and decision making task.

Instructions. Before beginning the task, participants received written, audio, and interactive instructions, and completed practice choice and memory trials. They also completed an instruction comprehension quiz and could not begin the task until they answered all four questions correctly.

Encoding Phase. We adapted a three-armed restless bandit task from a prior adult study (15). In our version, participants acted as captains of a pirate ship (Figure 1AB). The ship traveled to a series of islands, each distinguished by a unique background image (i.e., the context). As the captain, participants chose, within a three-second response window, which of the three crew members—Red Beard, Black Beard, or White Beard—would attempt to rob the next passing cargo ship. After each choice, a trial-unique image of the cargo, always an object, was displayed for 1 second. Participants were instructed to encode this object image in conjunction with the island's background image. Participants then received feedback that indicated whether the robbery was successful. A pile of gold coins signaled that the crew member was successful and a red "X" signaled that they were not. Feedback remained on screen for 1.5 seconds. The three crew members differed in their probability of success, and these probabilities changed over time according to Equation 1 below (see Reward Structure). Participants were incentivized to track the success of crew members over time, because their bonus payment was tied to the amount of gold they collected. The encoding phase comprised 240 choice trials across six islands, with islands changing every 40 trials. Upon leaving an island, participants were offered an optional break of up to two minutes.

Reward Structure. Each crew member's probability of success drifted over time. At the beginning of each phase, crew members were assigned unique initial success probabilities of 60%, 40%, or 10%. These probabilities drifted independently according to a decaying Gaussian random walk with reflecting bounds at 5% and 95%, centered on the target probability (θ_i) assigned to each crew member i.

$$\pi_{i,t+1} = \lambda\pi_{i,t} + (1-\lambda)\theta_i + \nu \qquad [1]$$

Where λ (stickiness) was set to 0.6, ν (diffusion noise) was a zero-mean Gaussian with SD = 8.

On each island (except the first in each phase), crew members' target reward probabilities (θi) were reshuffled after the first 10 trials and remained fixed for the subsequent 40 trials until the next island transition. Critically, memory probes presented during the test phase were drawn exclusively from these first 10 trials before reshuffling. This allowed us to dissociate the reward outcome associated with a probed trial from the overall distribution of rewards within that context.

Test Phase. The following day, participants completed 180 additional choice trials of the three-armed bandit task (Figure 1C). In the test phase, background and trial-unique object images were absent — participants were told that "fog" on the island obscured this information. Interleaved pseudorandomly among the choice trials were 60 recognition memory trials. On each recognition memory trial, an object image was displayed, and participants had up to 3 seconds to indicate their confidence in having seen it the previous day (definitely saw, maybe saw, maybe did not see, definitely did not see). Correct responses earned additional bonus pay. A majority of the images (up to 50 of 60) were "old", having appeared during the encoding phase. Thus, they served as reminders of specific past trials. Probe images were drawn exclusively from the first 10 trials of each island (with the exception of the very first island). We excluded trials on which a participant did not respond because no trial-unique object would have been displayed on such trials, and hence there was no image that could serve as a memory cue. Because probe images were associated with the earliest trials within each context, they benefit from the primacy effect (99), increasing the likelihood that participants would remember them during the test phase. The probed trials bore no relation to their surrounding choice trials, and participants were given no information regarding the relationship between the recognition memory and choice trials.

Source Memory. Source memory was assessed immediately after the test phase. On each trial, participants were shown an object image that had previously served as a memory probe and had three seconds to select the island on which they originally encountered the object.

Analyses. Our primary goal was to examine how episodic memory shapes decision making across development. To address this, we fit a logistic regression model predicting participants' choices from the test phase. Fixed effects included recent experience (option chosen on the previous trial; outcomes one, two, and three trials back), distant experience from the prior day (choice and outcome from the probed trial; outcomes averaged across the probed context), continuous age, and interactions between age and all other predictors. Age was z-scored across the full sample prior to entry in the model.

We constructed the predictor matrix (Table S1) following the approach of Bornstein & Norman (2017). The following procedure was repeated once for each of the three crew members, yielding up to 540 rows (180 choice trials x 3 options). The number of rows could be less than this because we excluded from analysis any choice trials in which the memory probe, on the preceding recognition memory trial, was incorrectly judged to be "new" despite having been presented on the previous day. Confidence (i.e., "definitely" vs. "maybe") was not taken into account.

To begin, one crew member was treated as the option of interest. We included seven main effects. The first four predictors coded for the influence of recent experience. The first predictor, reflecting perseveration, was coded as 1 if the option of interest was the same as the option selected on the last trial and was coded as 0 otherwise. The following three predictors reflected the influence of rewards received one, two, and three trials back. If the option of interest was chosen and rewarded t trials ago, then the predictor was coded as 1. If chosen but not rewarded, then it was coded as -1, and if it was not selected, then 0.

The next three predictors captured the influence of experiences from the prior day. The first (choice on probed trial) was coded as 1 if the option of interest matched the option chosen on the probed trial, and 0 otherwise. The second (reward on probed trial) was coded as 1 if the option was chosen and rewarded, -1 if chosen and unrewarded, and 0 if not chosen. The third (reward across probed context) was calculated as the number of times the option was chosen and rewarded minus the number of times it was chosen and unrewarded, divided by the total number of choices of that option within the context; if the option was never chosen in the context, this value was set to 0. When a choice trial did not immediately follow a recognition probe trial, or when the probe was "new," all memory predictors were set to 0. Confidence judgments (definitely vs. maybe) were not included in the main analyses. Results from an analysis including confidence can be found in the Supplementary Materials (S6). Including confidence does not alter any of the significant effects we report here. In total, the model included eight predictors, seven interaction terms, and an intercept.

The regression was conducted in R using the lme4 package (Bates et al, 2015). It included participant-level random intercepts and random slopes for the effects of reward one, two, and three trials back, as well as for the effects of choice on the probed trial, reward on the probed trial, and reward history across the probed context. To minimize Type I error, we initially specified the maximal model (Barr et al. 2013). If the model failed to converge, we iteratively simplified the model by removing random slopes until convergence was achieved.

Mnemonic Similarity Task. The task was split into an incidental encoding and test phase and was performed in a separate session from the memory and decision-making task (Figure S4A). During the encoding phase, participants performed a cover task in which they were shown an object and while it remained on screen, were asked to identify it as belonging indoors by pressing the 'v' or outdoors by pressing the 'n' on their keyboard. 64 objects were presented sequentially in a randomized order for 2 seconds each followed by a 500 ms inter-stimulus interval.

During the test phase, participants were informed that they would be shown some objects they had previously seen and some objects that were novel. When an object was displayed, they had 3 seconds to decide if the object was identical to one from the encoding phase ('old'), similar to one they had seen before ('similar'), or a novel object dissimilar to any objects encountered during the encoding phase ('new'). 32 old objects, 32 similar lures, and 32 new foils were presented in a randomized order and each choice was followed by a 500 ms inter-stimulus interval. Participants responded using their keyboard, pressing 'v' for old, 'b' for similar, and 'n' for new. We selected images from the Stark lab database (https://github.com/celstark/MST/tree/master/Set%206) based on our subjective assessment of children's familiarity with the objects. Lure images belonged to five bins of varying similarity as rated by a separate group of adult participants.

Analysis. We used the lure discrimination index (LDI) as a measure of an individual's pattern separation abilities or memory precision. It was computed as the difference between the rate of "similar" responses given for lure items and "similar" responses given to foils. We also computed a general recognition memory score (REC). This was taken as the difference between the rate of "old" responses given for repeat items and "old" responses given to foils.

The Two Step Task. In a separate session, participants completed a version of the original task (53) adapted for developmental populations (48), specifically for online data collection.

Participants were tasked with collecting as much space treasure as possible from aliens residing on different planets (Figure S5B). Each trial consisted of two decision stages. At the first stage, participants chose between two spaceships that would probabilistically transition to a red or purple planet. Each spaceship had a planet it would more commonly transition to (70% versus 30% of the time). Participants were informed that some transitions between spaceships and planets would be more frequent, but they were not informed of the exact transition structure, requiring them to learn this through experience. Once on a planet, they would then select one of two aliens to ask for treasure (the second-stage decision). Aliens provided treasure with some probability that slowly changed over time. Participants completed 200 of these trials.

Analysis. Using a previously described analytical approach (Nussenbaum et al. 2020), we examined participants' use of model-free and model-based strategies by running a mixed effects logistic regression predicting choice from the transition type (common vs. rare) and reward received on the previous trial. If using a model-free strategy, a participant should be insensitive to the previous trial's transition type. On the next trial, they should be more likely to repeat their previous first-stage choice if it led to reward and switch to the other option if it did not. In contrast, a participant using a model-based strategy would be sensitive to transition type. They should be more likely to repeat their first-stage choice under two scenarios — if it was common transition trial and they received a reward or it was a rare transition and they did not receive reward. If the opposite of either scenario occurred, then they should be more likely to switch their choice. Thus, we used the random slopes of the reward x transition type interaction effect as a measure of individuals' model-based strategy use.

Data, Materials, and Software Availability. All data and code are available at https://github.com/noraharhen/HarhenBornsteinHartley2025_ContextBandits.

ACKNOWLEDGMENTS. We thank the Hartley and Bornstein Labs for their feedback on the task design and analysis, and Susan Benear, Rheza Budiono, Alana Jaskir, Kate Nussenbaum, and Wangjing Yu for their feedback on the manuscript. This work was supported by the National Institute of Mental Health (R01 MH126183 to C.A.H. and F31 MH134620 to N.C.H.), the Department of Defense (National Defense Science and Engineering Graduate Fellowship to N.C.H.), the Brain and Behavior Foundation (NARSAD Young Investigator Award to A.M.B.), the NYU Developing Brain Innovation Fund (to C.A.H.) and the NYU Vulnerable Brain Project (to C.A.H.).

1. D Hassabis, D Kumaran, C Summerfield, M Botvinick, Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017).
2. M Botvinick, et al., Reinforcement learning, fast and slow. *Trends Cogn. Sci.* **23**, 408–422 (2019).
3. D Kumaran, D Hassabis, JL McClelland, What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends Cogn. Sci.* **20**, 512–534 (2016).
4. M Lengyel, P Dayan, Hippocampal contributions to control: The third way. *Adv. Neural Inf. Process. Syst.* **20** (2007).
5. C Blundell, et al., Model-free episodic control. *arXiv [stat.ML]* (2016).
6. A Pritzel, et al., Neural episodic control. *ICML* **abs/1703.01988**, 2827–2836 (2017).
7. Z Lin, T Zhao, G Yang, L Zhang, Episodic memory deep Q-networks. *arXiv [cs.LG]* (2018).
8. SJ Gershman, ND Daw, Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annu. Rev. Psychol.* **68**, 101–128 (2017).
9. N Biderman, A Bakkour, D Shohamy, What are memories for? the hippocampus bridges past experience with future decisions. *Trends Cogn. Sci.* **24**, 542–556 (2020).
10. DL Schacter, DR Addis, RL Buckner, Remembering the past to imagine the future: the prospective brain. *Nat. Rev. Neurosci.* **8**, 657–661 (2007).
11. GE Wimmer, D Shohamy, Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science* **338**, 270–273 (2012).
12. KD Duncan, D Shohamy, Memory states influence value-based decisions. *J. Exp. Psychol. Gen.* **145**, 1420–1426 (2016).
13. R Botvinik-Nezer, A Bakkour, T Salomon, D Shohamy, T Schonberg, Memory for individual items is related to nonreinforced preference change. *Learn. Mem.* **28**, 348–360 (2021).
14. J Nicholas, ND Daw, D Shohamy, Uncertainty alters the balance between incremental learning and episodic memory. *Elife* **11** (2022).
15. AM Bornstein, KA Norman, Reinstated episodic context guides sampling-based decisions for reward. *Nat. Neurosci.* **20**, 997–1003 (2017).
16. AM Bornstein, MW Khaw, D Shohamy, ND Daw, Reminders of past choices bias decisions for reward in humans. *Nat. Commun.* **8**, 15958 (2017).
17. K Nussenbaum, CA Hartley, Reinforcement learning across development: What insights can we draw from a decade of research? *Dev. Cogn. Neurosci.* **40**, 100733 (2019).
18. J Sluzenski, NS Newcombe, SL Kovacs, Binding, relational memory, and recall of naturalistic events: a developmental perspective. *J. Exp. Psychol. Learn. Mem. Cogn.* **32**, 89–100 (2006).
19. T Riggins, Longitudinal investigation of source memory reveals different developmental trajectories for item memory and binding. *Dev. Psychol.* **50**, 449–459 (2014).
20. PJ Bauer, et al., It's all about location, location, location: children's memory for the "where" of personally experienced events. *J. Exp. Child Psychol.* **113**, 510–522 (2012).
21. TC Lorsbach, JF Reimer, Feature binding in children and young adults. *J. Genet. Psychol.* **166**, 313–327 (2005).
22. DM Demaster, S Ghetti, Developmental differences in hippocampal and cortical contributions to episodic retrieval. *Cortex* **49**, 1482–1493 (2013).
23. JK Lee, C Wendelken, SA Bunge, S Ghetti, A time and place for everything: Developmental differences in the building blocks of episodic memory. *Child Dev.* **87**, 194–210 (2016).
24. L Picard, I Reffuveille, F Eustache, P Piolino, Development of autonoetic autobiographical memory in school-age children: genuine age effect or development of basic cognitive abilities? *Conscious. Cogn.* **18**, 864–876 (2009).
25. ME Lloyd, AO Doydum, NS Newcombe, Memory binding in early childhood: evidence for a retrieval deficit. *Child Dev.* **80**, 1321–1328 (2009).
26. C Coughlin, et al., Asynchronous development of memory integration and differentiation influences temporal memory organization. *Dev. Sci.* **27**, e13437 (2024).
27. IR Olson, NS Newcombe, Binding together the elements of episodes: Relational memory and the developmental trajectory of the hippocampus. *The Wiley handbook on development children's memory* (2014).
28. S Ghetti, Development of item–space and item–time binding. *Curr. Opin. Behav. Sci.* **17**, 211–216 (2017).
29. KL Canada, CT Ngo, NS Newcombe, F Geng, T Riggins, It's all in the details: Relations between young children's developing pattern separation abilities and hippocampal subfield volumes. *Cereb. Cortex* **29**, 3427–3433 (2019).
30. S Ghetti, DM DeMaster, AP Yonelinas, SA Bunge, Developmental differences in medial temporal lobe function during memory encoding. *J. Neurosci.* **30**, 9548–9556 (2010).
31. FR Lambert, P Lavenex, PB Lavenex, Improvement of allocentric spatial memory resolution in children from 2 to 4 years of age. *Int. J. Behav. Dev.* **39**, 318–331 (2015).
32. F Ribordy Lambert, P Lavenex, P Banta Lavenex, The "when" and the "where" of single-trial allocentric spatial memory performance in young children: Insights into the development of episodic memory. *Dev. Psychobiol.* **59**, 185–196 (2017).
33. CT Ngo, NS Newcombe, IR Olson, The ontogeny of relational memory and pattern separation. *Dev. Sci.* **21** (2018).
34. CT Ngo, Y Lin, NS Newcombe, IR Olson, Building up and wearing down episodic memory: Mnemonic discrimination and relational binding. *J. Exp. Psychol. Gen.* **148**, 1463–1479 (2019).
35. M Peng, et al., Extended trajectory of spatial memory errors in typical and atypical development: The role of binding and precision. *Hippocampus* **33**, 1171–1188 (2023).
36. L Rollins, EB Cloude, Development of mnemonic discrimination during childhood. *Learn. Mem.* **25**, 294–297 (2018).
37. S Vijayarajah, ML Schlichting, Neural maturity of encoding states supports gains to memory precision in childhood. *Child Dev.* (2025).
38. S Wang, SF Feng, AM Bornstein, Mixing memory and desire: How memory reactivation supports deliberative decision-making. *Wiley Interdiscip. Rev. Cogn. Sci.* p. e1581 (2021).
39. YS Shin, S DuBrow, Structuring memory through inference-based event segmentation. *Top. Cogn. Sci.* **13**, 106–127 (2021).
40. EA Crone, RE Dahl, Understanding adolescence as a period of social-affective engagement and goal flexibility. *Nat. Rev. Neurosci.* **13**, 636–650 (2012).
41. S Urošević, P Collins, R Muetzel, K Lim, M Luciana, Longitudinal changes in behavioral approach system sensitivity and brain structures involved in reward processing during adolescence. *Dev. Psychol.* **48**, 1488–1500 (2012).
42. B Luna, S Marek, B Larsen, B Tervo-Clemmens, R Chahal, An integrative model of the maturation of cognitive control. *Annu. Rev. Neurosci.* **38**, 151–170 (2015).
43. AD Ekstrom, AP Yonelinas, Precision, binding, and the hippocampus: Precisely what are we talking about? *Neuropsychologia* **138**, 107341 (2020).
44. CB Kirwan, CEL Stark, Overcoming interference: an fMRI investigation of pattern separation in the medial temporal lobe. *Learn. Mem.* **14**, 625–633 (2007).
45. RF Stevenson, ZM Reagh, AP Chun, EA Murray, MA Yassa, Pattern separation and source memory engage distinct hippocampal and neocortical regions during retrieval. *J. Neurosci.* **40**, 843–851 (2020).
46. CT Ngo, S Michelmann, IR Olson, NS Newcombe, Pattern separation and pattern completion: Behaviorally separable processes? *Mem. Cogn.* **49**, 193–205 (2021).
47. SS Cohen, CT Ngo, IR Olson, NS Newcombe, Pattern separation and pattern completion in early childhood. *Proc. Natl. Acad. Sci. U. S. A.* **122**, e2416985122 (2025).
48. JH Decker, AR Otto, ND Daw, CA Hartley, From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychol. Sci.* **27**, 848–858 (2016).
49. D Albert, L Steinberg, Age differences in strategic planning as indexed by the tower of london: Strategic planning. *Child Dev.* **82**, 1501–1517 (2011).
50. M Luciana, PF Collins, EA Olson, AM Schissel, Tower of london performance in healthy adolescents: the development of planning skills and associations with self-reported inattention and impulsivity. *Dev. Neuropsychol.* **34**, 461–475 (2009).
51. G Buzsáki, Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073–1188 (2015).
52. DJ Foster, Replay comes of age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).
53. ND Daw, Y Niv, P Dayan, Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711 (2005).
54. V Sprondel, KH Kipp, A Mecklinger, Developmental changes in item and source memory: evidence from an ERP recognition memory study with children, adolescents, and adults: ERP evidence of item and source memory development. *Child Dev.* **82**, 1638–1953 (2011).
55. RA Weilbächer, S Gluth, The interplay of hippocampus and ventromedial prefrontal cortex in memory-based decision making. *Brain Sci.* **7**, 4 (2016).
56. FJ Calabro, VP Murty, M Jalbrzikowski, B Tervo-Clemmens, B Luna, Development of hippocampal-prefrontal cortex interactions through adolescence. *Cereb. Cortex* **30**, 1548–1558 (2020).

57. JY Davidow, K Foerde, A Galván, D Shohamy, An upside to reward sensitivity: The hippocampus supports enhanced reinforcement learning in adolescence. *Neuron* **92**, 93–99 (2016).
58. JY Davidow, C Insel, LH Somerville, Adolescent development of value-guided goal pursuit. *Trends Cogn. Sci.* **22**, 725–736 (2018).
59. SL Benear, CT Ngo, IR Olson, NS Newcombe, Understanding relational binding in early childhood: Interacting effects of overlap and delay. *J. Exp. Child Psychol.* **208**, 105152 (2021).
60. H Yim, SJ Dennis, VM Sloutsky, The development of episodic memory: items, contexts, and relations: Items, contexts, and relations. *Psychol. Sci.* **24**, 2163–2172 (2013).
61. S Ghetti, L Angelini, The development of recollection and familiarity in childhood and adolescence: evidence from the dual-process signal detection model. *Child Dev.* **79**, 339–358 (2008).
62. S Ghetti, C Mirandola, L Angelini, C Cornoldi, E Ciaramelli, Development of subjective recollection: understanding of and introspection on memory states: Development of subjective recollection. *Child Dev.* **82**, 1954–1969 (2011).
63. P Piolino, et al., Do school-age children remember or know the personal past? *Conscious. Cogn.* **16**, 84–101 (2007).
64. KA Willoughby, M Desrocher, B Levine, JF Rovet, Episodic and semantic autobiographical memory and everyday memory during late childhood and early adolescence. *Front. Psychol.* **3**, 53 (2012).
65. C Coughlin, KE Lyons, S Ghetti, Remembering the past to envision the future in middle childhood: Developmental linkages between prospection and episodic memory. *Cogn. Dev.* **30**, 96–110 (2014).
66. PJ Bauer, P San Souci, Going beyond the facts: young children extend knowledge by integrating episodes. *J. Exp. Child Psychol.* **107**, 452–465 (2010).
67. PJ Bauer, JE King, M Larkina, NL Varga, EA White, Characters and clues: Factors affecting children's extension of knowledge through integration of separate episodes. *J. Exp. Child Psychol.* **111**, 681–694 (2012).
68. PJ Bauer, NL Varga, JE King, AM Nolen, EA White, Semantic elaboration through integration: Hints both facilitate and inform the process. *J. Cogn. Dev.* **16**, 351–369 (2015).
69. YL Shing, et al., Integrating across memory episodes: Developmental trends. *PLoS One* **14**, e0215848 (2019).
70. D Shohamy, RA Adcock, Dopamine and adaptive memory. *Trends Cogn. Sci.* **14**, 464–472 (2010).
71. VP Murty, S DuBrow, L Davachi, The simple act of choosing influences declarative memory. *J. Neurosci.* **35**, 6255–6264 (2015).
72. VP Murty, KC Dickerson, Motivational influences on memory in *Recent Developments in Neuroscience Research on Human Motivation*, Advances in Motivation and Achievement. (Emerald Group Publishing Limited) Vol. 19, pp. 203–227 (2016).
73. VP Murty, S DuBrow, L Davachi, Decision-making increases episodic memory via postencoding consolidation. *J. Cogn. Neurosci.* **31**, 1308–1317 (2019).
74. NA Ruiz, S DuBrow, VP Murty, Agency as a bridge to form associative memories. *J. Exp. Psychol. Gen.* (2023).
75. AH Sinclair, YC Wang, RA Adcock, Instructed motivational states bias reinforcement learning and memory formation. *Proc. Natl. Acad. Sci. U. S. A.* **120**, e2304881120 (2023).
76. SE Qasim, A Deswal, I Saez, X Gu, Positive affect modulates memory by regulating the influence of reward prediction errors. *Commun. Psychol.* **2**, 52 (2024).
77. A Bakkour, et al., The hippocampus supports deliberation during value-based decisions. *Elife* **8** (2019).
78. VP Murty, O FeldmanHall, LE Hunter, EA Phelps, L Davachi, Episodic memories predict adaptive value-based decision-making. *J. Exp. Psychol. Gen.* **145**, 548–558 (2016).
79. CA Hartley, K Nussenbaum, AO Cohen, Interactive development of adaptive learning and memory. *Annu. Rev. Dev. Psychol.* **3**, 59–85 (2021).
80. K Nussenbaum, CA Hartley, Reinforcement learning increasingly relates to memory specificity from childhood to adulthood. *Nat. Commun.* **16**, 4074 (2025).
81. PL Katzman, CA Hartley, The value of choice facilitates subsequent memory across development. *Cognition* **199**, 104239 (2020).
82. AO Cohen, et al., Reward enhances memory via age-varying online and offline neural mechanisms across development. *J. Neurosci.* **42**, 6424–6434 (2022).
83. CT Ngo, NS Newcombe, IR Olson, Gain-loss framing enhances mnemonic discrimination in preschoolers. *Child Dev.* **90**, 1569–1578 (2019).
84. Z Ding, et al., The effect of choice on memory across development. *J. Exp. Child Psychol.* **246**, 105982 (2024).
85. Y Fandakova, MJ Gruber, States of curiosity and interest enhance memory differently in adolescents and in children. *Dev. Sci.* **24**, e13005 (2021).
86. SL Master, et al., Distentangling the systems contributing to changes in learning during adolescence. *Dev. Cogn. Neurosci.* **41**, 100732 (2020).
87. W van den Bos, MX Cohen, T Kahnt, EA Crone, Striatum-medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cereb. Cortex* **22**, 1247–1255 (2012).
88. S Palminteri, EJ Kilford, G Coricelli, SJ Blakemore, The computational development of reinforcement learning during adolescence. *PLoS Comput. Biol.* **12**, e1004953 (2016).
89. B Eppinger, B Mock, J Kray, Developmental differences in learning and error processing: evidence from ERPs. *Psychophysiology* **46**, 1043–1053 (2009).
90. A Christakou, et al., Neural and psychological maturation of decision-making in adolescence and young adulthood. *J. Cogn. Neurosci.* **25**, 1807–1823 (2013).
91. TU Hauser, R Iannaccone, S Walitza, D Brandeis, S Brem, Cognitive flexibility in adolescence: neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *Neuroimage* **104**, 347–354 (2015).
92. AH Javadi, DHK Schmidt, MN Smolka, Adolescents adapt more slowly than adults to varying reward contingencies. *J. Cogn. Neurosci.* **26**, 2670–2681 (2014).
93. MK Eckstein, SL Master, RE Dahl, L Wilbrecht, AGE Collins, Reinforcement learning and bayesian inference provide complementary models for the unique advantage of adolescents in stochastic reversal. *Dev. Cogn. Neurosci.* **55**, 101106 (2022).
94. Y Fandakova, G Akin, Memory contributions to the development of cognitive flexibility in childhood. *Eur. Psychol.* **31**, 63–74 (2026).
95. K Nussenbaum, E Prentis, CA Hartley, Memory's reflection of learned information value increases across development. *J. Exp. Psychol. Gen.* **149**, 1919–1934 (2020).
96. AO Cohen, K Nussenbaum, HM Dorfman, SJ Gershman, CA Hartley, The rational use of causal inference to guide reinforcement learning strengthens with age. *NPJ Sci Learn.* **5**, 16 (2020).
97. GM Rosenbaum, HL Grassie, CA Hartley, Valence biases in reinforcement learning shift across adolescence and modulate subsequent memory. *Elife* **11** (2022).
98. HA Raab, C Foord, R Ligneul, CA Hartley, Developmental shifts in computations used to detect environmental controllability. *PLoS Comput. Biol.* **18**, e1010120 (2022).
99. H Ebbinghaus, Memory: A contribution to experimental psychology in *The history of psychology: Fundamental questions (pp*, ed. MP Munger. (Oxford University Press, xi, New York, NY, US) Vol. 514, pp. 168–177 (2003).