

Memory precision and age differentially predict the use of decision-making strategies across the lifespan.

Sharon M Noh^{1★}, Umesh K Singla^{2,3}, Ilana J Bennett⁴, Aaron M Bornstein^{1,5★}

1. Department of Cognitive Sciences, The University of California, Irvine, Irvine, CA, USA

2. Princeton Neuroscience Institute, Princeton, NJ, USA

3. Department of Computer Science and Engineering, University of California, San Diego, La Jolla, CA, USA

4. Department of Psychology, The University of California, Riverside, Riverside, CA, USA

5. Center for the Neurobiology of Learning and Memory, The University of California, Irvine, Irvine, CA, USA

★ To whom correspondence should be addressed: nohsm@uci.edu, aaron.bornstein@uci.edu

Acknowledgements

The authors wish to thank Mohit Nadkarni for extensive assistance in data collection and preprocessing, Gloria Cheng, Tommie Hyunh, Angel Liu, and Vinayak Nadig for assistance in data collection, Craig EL Stark for helpful discussions, and Tyson Kerr for programming assistance. Funding was provided by a Hellman Foundation fellowship (to IJB), NIH R21AG072673 (to AMB), and F32AG072836 (to SMN).

Respective contributions

SMN, IJB, and AMB designed the experiments; SMN and IJB performed the experiments; SMN, UKS, and AMB analyzed the data; SMN, IJB and AMB wrote the manuscript. All authors reviewed the manuscript.

Data availability statement

All de-identified data and the final versions of all MATLAB model and analysis code, as well as R scripts used to generate each data figure, will be made freely available at the UCI CCNL GitHub repository (<https://github.com/uciccnl>) upon publication.

Competing Interests

The authors declare no competing interests.

Abstract

Memory function declines in normal aging, in a relatively continuous fashion following middle-age. The effect of aging on decision-making is less well-understood, with seemingly conflicting results on both the nature and direction of these age effects. One route for clarifying these mixed findings is to understand how age-related differences in memory affect decisions. Recent work has proposed *memory sampling* as a specific computational role for memory in decision-making, alongside well-studied mechanisms of reinforcement learning (RL). Here, we tested the hypothesis that age-related declines in episodic memory alter memory sampling. Participants (total

N=361; ages 18-77) performed one of two variants of a standard reward-guided decision experiment with additional trial-unique mnemonic content and a separately-administered task for assessing memory precision. When we fit participants' choices with a hybrid computational model implementing both memory-based and RL-driven valuation side-by-side, we found that memory precision tracked the contribution of memory sampling to choice. At the same time, age corresponded to decreasing influence of RL and increasing perseveration. A second experiment confirmed these results and further revealed that memory precision tracked the specificity of memories selected for sampling. Together, these findings suggest that differences in decision-making across the lifespan may be related to memory function, and that interventions which aim to improve the former may benefit from targeting the latter.

Introduction

It is widely accepted that memory, at least that for individual experiences, declines with age ^{1,2}. The results on decision-making, however, are more mixed. While some studies identify a decrease in the ability of older individuals (and/or those experiencing cognitive decline) to engage in multi-step planning ³⁻⁵ - a function which, like episodic memory, is critically linked to an intact hippocampal formation ⁶ - others report that older individuals have spared, and sometimes improved, decision-making abilities, such as in a resistance to sunk costs ⁷. Still others have identified increases in choice randomness with age ⁸, which at first blush seems incompatible with reports of increasing perseveration - repetition of the same choice irrespective of learned values ⁹⁻¹¹. The field has yet to settle on a unifying explanation for these seemingly disparate phenomena.

One route to a synthesis may be via understanding the role of memory in decisions. A number of findings have shown that memories for individual experiences can bias decision-making even in repeated selection tasks, where individuals gain extensive experience with the probabilistic outcomes of choice options ^{12,13}. The role of memory in decisions via *sampling* (i.e. the selective retrieval of memories of similar past choices during decision-making, to estimate the value of one action or another^{14,15}) has been shown to capture the influence of recent rewarding experiences on choice ¹⁶, while also serving an adaptive role in decisions when familiarity with an environment is low or uncertainty about the structure of the environment is high¹⁷. To illustrate the difference between these routes to value estimation, consider choosing between two restaurants: One, an old favorite that you have been to many times; the other, a brand-new shop of a well-known type - e.g. a slice pizza stand. For the former, you can rely on your repeated experiences to estimate the value of dining there tonight. For the latter, you must draw on memories of similar experiences, and extrapolate from there. When evaluating repeated choice options, and assuming that recent memories are more likely to be retrieved, these two approaches give similar results on average, but can diverge, especially when the content of memory contains more than just recent experience.

Building on the findings that each system contributes to behavior, work has begun to examine what factors influence the relative use of one or the other. A recent study in young adults observed that the relative uncertainty of values estimated using memory sampling and reinforcement learning indexed how strongly choices depend on each system on a given trial¹⁸. In other words, when one system had more variable estimates of option values on a given trial, it had less influence on choice. This finding concords with theoretical frameworks in which multiple learning systems contribute to behavior in proportion to the relative precision of their estimates of upcoming stimuli ¹⁹⁻²². An open question is whether this mixture of decision processes is altered across the lifespan,

perhaps as a result of relative declines in the fidelity of memory representations or neural circuits communicating between valuation systems and action selection regions^{23–26}.

Therefore, the precision of memory representations may play a critical role in its influence on choice. A link between memory precision and memory-guided decision-making may be crucial to understanding age-related differences in choice patterns. This is because a key aspect of memory that declines with aging is the ability to *pattern separate*, or represent similar mnemonic stimuli with sufficiently distinct, non-overlapping memory traces^{5,27–30}. Individuals with impaired pattern separation may encode experiences in a manner that leads to impaired retrieval, perhaps due to the relative imprecision of these representations. Pattern separation is used to store unique memories of even highly similar information without interference, and age-related reductions in pattern separation result in less precise, lower-fidelity memory in older adults.

Here, we examined whether a behavioral measure of pattern separation - mnemonic discrimination ability in the Mnemonic Similarity Task (MST;²⁷) - indexed the contribution of memory sampling to behavior in a repeated choice task where outcomes were linked with trial-unique memoranda^{16,31}. A lifespan sample of individuals (ages 18-77; total $N=361$) performed a series of choices between three options, each of which had steadily varying probabilities of paying reward. After selecting an option, they were then asked to encode a trial-unique object photograph in concert with one of six scene images. Their ultimate reward in the task depended both on their ability to identify and track the winningest choice option, and also on their ability to later successfully recall the object and scene pair that was associated with a randomly-selected subset of the choice outcomes. Critically, the experiment included a second phase during which individuals were presented with memory probes that incidentally reminded them of past choices. These probes have previously been shown to influence subsequent choice via intrusions of value representations linked to the reminded trials^{16,31} and associated contexts³¹ in younger adults. It is important to note that this effect of incidental reminder probes on choice is not adaptive - individuals would do best to attend only to recent rewards, and not the values reminded by the probes. Therefore, this aspect of the task serves as a strong test of the hypothesis that action selection incorporates value information from recent memory, independent of or alongside continual reinforcement learning-based approaches^{17,21,32}. To date, memory sampling has not been studied in older adults.

We fit each participant's series of choices with a computational model implementing both memory sampling and reinforcement learning, and measured the degree to which

the values estimated by each process influenced the individual's choice behavior. We then examined the relationship between their reliance on each process and their ability to discriminate similar objects during the MST, administered in a separate session. We predicted that, independent of age, the latter measure of mnemonic discrimination ability would index the average precision of an individual's trial-level outcome memories in the choice task, and thus their use of memory sampling. We further predicted that age would negatively influence choice performance.

Consistent with theories of uncertainty-weighted arbitration between decision systems^{19–21}, which predict that decision systems with more precise estimates will have greater influence on choices, we found that higher memory precision predicted a greater influence of memory sampling-derived values on behavior, while not affecting an individual's sensitivity to values learned by recent reinforcement. At the same time, age increased the appearance of “noisy” choices, reflected in a decreased reliance on reinforcement learning, while also increasing the influence of perseveration.

To further understand how memory precision altered the influence of memory content on decisions, we next performed a variant of the experiment in which individuals were presented with perceptually aliased contexts - three matched pairs of scenes with similar visual content, but distinct choice-reward associations. We then tested whether individual memory precision corresponded to the influence on choice of the *target* context (the scene image associated with the memory probe) or the *gist* context (the target as well as its matched pair lure context, which had distinct choice-reward values). Consistent with the idea that uncertainty in memory representations is itself sample-based, rather than inherent to the representation, we found that lower memory precision was associated with a greater influence of gist-level memories.

Taken together, these findings support the idea that age-related declines in memory precision lead to changes in decision profiles, which can be decomposed with normatively-motivated process models. They further suggest that individuals who exhibit noisier choices with age may benefit from interventions that target the precision of their memory representations. Lastly, our results shed light on the fundamental mechanisms that guide the weighting of different decision strategies in individuals across the lifespan.

Methods

Experiments 1 and 2 were identical in procedure other than the background scene images (*contexts*) presented during the learning phase of the three-armed bandit task. As such, we have combined the method across both experiments and specified the differences across experiments where applicable.

Participants

Participants across the lifespan were recruited online on Amazon Mechanical Turk via the CloudResearch interface ($N=295$, ages 19-77, mean(sd) age = 46.7(16.5), 144 Male, 149 Female, 1 Other, 1 Unknown) and in-person ($N=156$, ages 18-45, mean(sd) age = 24.2(7.0), 56 Male, 97 Female, 1 Other). The majority of data collection occurred during the COVID-19 pandemic, which limited our ability to collect in-person data, particularly those in older age groups. For online samples, we set eligibility filters through Amazon Mechanical Turk such that only those who have completed at least 50 MTurk HiTs with over 95% approval rating were able to participate. All experimental procedures were carried out in accordance with the guidelines of the University of California, Irvine and University of California, Riverside, Institutional Review Boards (IRBs). All participants indicated their consent to participate in this experiment according to disclosures approved by the IRBs. Regardless of collection modality, all participant data underwent rigorous quality control by independent raters blind to the study hypotheses. Raters were instructed to identify individuals whose data indicated they were inattentive to the experiment or appeared to be automating responses for more than one quarter of the trials. Raters examined reaction times and response patterns for each participant and coded each participant as either 0 - do not exclude, 1 - consider excluding, or 2 - definitely exclude. Raters also provided reasonings for exclusion. Some reasons for exclusion included: high frequency of no-response trials, unrealistically low or high reaction times (which may have been indicative of connectivity or computer issues or automated respondents - e.g. bots or scripts), "button-mashing" (pressing the same button/choice for the majority of the experiment), or not sampling all 3 choices throughout the task (for instance, alternating between only 2 choice options for the majority of the task). Ratings from 2 independent raters were summed, and participants with scores greater than 2 were excluded. In other words, we excluded participants whom at least one rater identified as 'Definitely' excluded (Score 2) and the other rater scored at least 'Maybe' excluded (Score 1 or 2). Of the 451 participants originally recruited across both experiments, 83 participants were excluded on this basis.

We also excluded participants for whom the learning model did not fit better than chance ($p(\text{choose}) = \frac{1}{3}$ for all trials), as the parameters for these participants would not distinguish from a participant who made random choices. An additional 7 participants were excluded this way (all from Experiment 2), which left 361 participants (across Experiments 1 and 2) for our analyses: 226 participants were retained in Experiment 1 (112 online / 114 in-person, 88 Male / 137 Female / 1 Other, ages 18-77, mean(sd) age = 38.6(18.8)), and 135 participants were retained in Experiment 2 (117 online / 18 in-person, 66 Male / 68 Female / 1 Other, ages 18-74, mean(sd) age = 40.6(16.2)). Participants had to complete a brief tutorial and screener of the three-armed bandit task

to be eligible for the full experiment. The screener consisted of an interactive tutorial of the task and practice trials of the learning phase to ensure participants understood the instructions and were paying attention to each trial. Individuals had 2 opportunities to complete the tutorial and 10 practice trials with a minimum 70% accuracy to continue on to the full experiment. Those who were ineligible to continue on to the main task were compensated \$2 for the time spent on the tutorial. Participants received monetary compensation for their participation: \$12 for completing the experiment, and an additional bonus of up to \$8 obtained throughout the task.

Participants were recruited for two separate sessions to complete their two tasks (one of the choice tasks and the Mnemonic Similarity Task). During early phases of data collection, the order of recruitment for each task (choice task and MST) were counterbalanced in case there may be differences in retention as a function of which task was performed first. Our pilot data revealed that roughly 5 out of every 8 participants generally returned for a second session, regardless of which experiment was performed first. For participants who were recruited first through the choice experiment, participants who completed the entire experiment were invited back to complete the Mnemonic Similarity Task (MST)²⁷. For other participants, those recruited to complete the Mnemonic Similarity Task first and those who completed the MST were invited back to complete one of the three-armed bandit tasks. Ultimately, 130 participants from Experiment 1 (48 Male / 82 Female / 1 Other, ages 18-77, mean(sd) age = 37.0(20.5)) and 97 participants from Experiment 2 (45 Male / 52 Female, ages 18-73, mean(sd) age = 40.3(16.7)) completed both the MST and three-armed bandit tasks. Participants received a fixed fee of between \$3 and \$5 for their participation on the MST, depending on the experiment modality and date of data collection (during the course of data collection we adjusted our per-hour participant compensation to account for inflation).

To examine the possibility that there may be significant differences between our online and in-person participant samples influencing our main analyses, we examined the distribution of each model parameter of interest (\square_{sampler} , \square_{TD} , $\square_{\text{perservation}}$) for comparable online and in-person participants (younger adults between the ages of 18-40). For each parameter, we conducted an independent samples t-test to verify that the parameter distribution did not differ between online and in-person samples. All comparisons showed there was no significant difference in parameter distributions as a function of online or in-person data collection ($p > 0.2$ for all 3 t-tests), alleviating possible concerns that data collection source may be significantly influencing our variables of interest.

Three-Armed Bandit Task

Tutorial/Practice. Participants first completed a practice tutorial screener to determine eligibility for the main task. Participants engaged in an interactive tutorial during which

they practiced the various aspects of the task: 1) pressing keys corresponding to each card deck option of the 3-armed bandit task in a “virtual casino”, 2) experiencing the probabilistic payoff structure of each option, 3) making choices and remembering the casino room in which various item rewards (“lottery tickets”) were received within the trial’s time window, and 4) practicing the memory probe recognition trials in Phase 2. The instructions unfolded slowly in order to give participants adequate opportunities to familiarize themselves with all aspects of the task, and there was a final quiz. If participants did not successfully complete the quiz with at least 70% accuracy, they were given the opportunity to repeat the tutorial and quiz. Those who failed a second time were ineligible to continue on to the main task.

The main task followed protocols outlined in the method section of previous work that used this three-armed bandit task³¹. Additional information on the task can be found there.

Phase 1 (Learning Phase). Participants completed 180 choice trials of a 3-armed bandit task that occurred across 6 different contexts (“casino rooms”; Fig. 1A). Participants were told they are playing a card game across 6 casino rooms and that they should try to maximize reward by picking one of three card decks that most often leads to reward. The context was presented as a background image on the screen, which changed every 30 trials. In Experiment 1, the contexts presented in Phase 1 consisted of 6 unique outdoor scenes (Fig. 1B, “Experiment 1”). In Experiment 2, there were 2 contexts for each of 3 outdoor scene categories (i.e., 2 beach, 2 mountain, 2 forest). One context from each category was sampled first (e.g., beach #1, forest #1, mountain #1) before sampling from the same-context foils (e.g., forest #2, mountain #2, beach #2). There was an additional constraint that two rooms from the same category (e.g., forest #1 and forest #2) were not assigned to adjacent contexts (i.e., context 3 and 4).

Participants were told that the deck that most often leads to reward can change over time, so the optimal strategy is to periodically sample all choices to ensure that they are picking the best deck. Critically, it was emphasized that the decks “and dealer” followed participants between rooms - this feature was emphasized by the actual payout probabilities, described below. Participants had up to 2.5 seconds to make their choice by using the “1”, “2”, or “3” key on their keyboard to select one of the 3 presented card decks. After making their choice, participants were shown the chosen deck in isolation for 0.5 seconds, followed by a trial unique object image for 2 seconds, followed by a \$1 or \$0 reward for 1.5 seconds (Fig. 1A). Each trial was tagged with a trial-unique object image to later be used as an item-memory probe and remind participants of rewards received on a given trial. Participants were told that the trial-unique objects are “lottery tickets” that they will have to recognize to receive their associated rewards later. Participants learn to follow which deck most frequently yields reward through

trial-and-error, while simultaneously learning to associate the different contexts (casino rooms) with trial-unique items (lottery tickets) experienced across the 180 trials.

Phase 2 (Probe Phase). During the probe phase, participants completed an additional 120 choice trials of the 3-armed bandit task. However, during this phase, no context or trial-unique item was presented in any of the trials. Instead, 60 memory probe trials were pseudorandomly interspersed amongst the choice trials. During a memory probe trial, participants were shown an object from Phase 1 to serve as a reminder of its associated reward and context from the learning phase. Participants were instructed to continue to make choices in this “unfinished” room and follow the deck most frequently associated with reward. They were also instructed that they may occasionally be shown an object image and asked to recognize if it were one of the lottery tickets they collected from Phase 1. Participants had up to 3 seconds to determine if the object was previously seen. Each correctly identified lottery ticket added a small monetary bonus (+\$0.05), and each incorrectly identified ticket resulted in a small monetary deduction (-\$0.05).

Reward Structure. In Phase 1 and 2, each card deck followed a probabilistic reward structure that changed throughout the task (Fig. 1B). Each deck was assigned to initial payout probabilities (π , the odds the deck would return a \$10 reward, rather than \$0) of 60%, 30%, or 10% randomly, without replacement. These payout probabilities then slowly drifted across trials (t) within each context according to a decaying Gaussian random walk with reflecting bounds at 5% and 95% that was centered at the target probability (θ_i) assigned to each deck (i)³¹:

$$\pi_{i,t+1} = \lambda\pi_{i,t} + (1 - \lambda)\theta_i + v \quad (1)$$

Parameters were set as follows: λ (stickiness) was 0.6, v (diffusion noise) was a zero-mean Gaussian with SD 8. To confirm to participants that the decks “carried with them” across rooms, the stickiness parameter was temporarily set to 0.95 for the first three trials of each room.

The key feature of the payoff structure was as follows: the target payout probabilities θ_i were reassigned after the first 10 trials within a given context and persisted for 30 trials before resetting again. Critically, item probes were selected only from these first ten trials, resulting in the memory probes having a reward distribution that is dissociable from the reward distribution of the context within which the item was presented. For both experiments, four payoff time series were pre-generated according to this procedure and randomly assigned to participants in the task. In Experiment 2, there was an

additional constraint that the reward distribution for one context (e.g., forest #1) was distinct from its same-category foil (e.g., forest #2), thereby allowing us to determine if evoked memory content from memory probes in Phase 2 is consistent with memory for 1) the target context, 2) the same-category lure, or 3) a gist representation of the category (target + lure).

Analysis of choice behavior.

Regression analysis.

Our initial analysis of interest examined how reminder probes - and the trials and contexts they referred to - affected choices in Phase 2. Following our previously employed procedure³¹, we constructed a three-part logistic regression model, with each part corresponding to one of the three card decks. For each deck and each trial, we modeled: for all trials, the identity of the deck chosen on the previous trial (DI : 1 for the given deck, 0 for others) and the recent rewards up to three steps back (DR : 1 if the given deck was chosen and rewarded on trial $t-k$, $k=1..3$, 0 otherwise); for trials immediately following the memory probes, the deck identity (EI : 1 if the same deck as probed, 0 otherwise), the deck-specific value received (ER : 1 if the same deck and rewarded, 0 otherwise) on the individual trial reminded by a memory probe, and the average reward of the probed deck, across the reminded context room (EC ; Eqn. 2).

$$EC^{C,i} = \frac{\# \text{ of choices of } i \text{ resulting in } \$10 - \# \text{ of choices of } i \text{ resulting in } \$0}{\# \text{ of times } i \text{ chosen in } C} \quad (2)$$

This quantity, *evoked context reward* ($EC^{C,i}$) reflects the expected value of samples of this deck i from the reminded room context C . The total regression model was thus:

$$C_{i,t} \approx \beta^{DI_i} DI_{i,t-1} + \sum_{k=1}^3 \beta^{DR_i} DR_{i,t-k} + \beta^{EI_i} EI_{i,t-1} + \beta^{ER_i} ER_{i,t-1} + \beta^{EC_i} EC_{i,t-1} \quad (3)$$

where $C_{i,t}$, the dependent variable, reflects whether deck i was chosen on trial t (1 or 0).

The resulting design matrix thus had seven columns, and 360 rows - one for each deck and choice trial combination. For the second experiment, a second analysis was run which included an eighth column representing the “gist-level” evoked context reward computed across the trials from the combined target evoked context and its paired *lure* context room. This regressor was orthogonalized against the target context regressor using Gram-Schmidt orthonormalization as implemented by the SPM8 function `spm_orth`³³. The coefficients produced by all regressions were tested against zero across the population by two-tailed t-test.

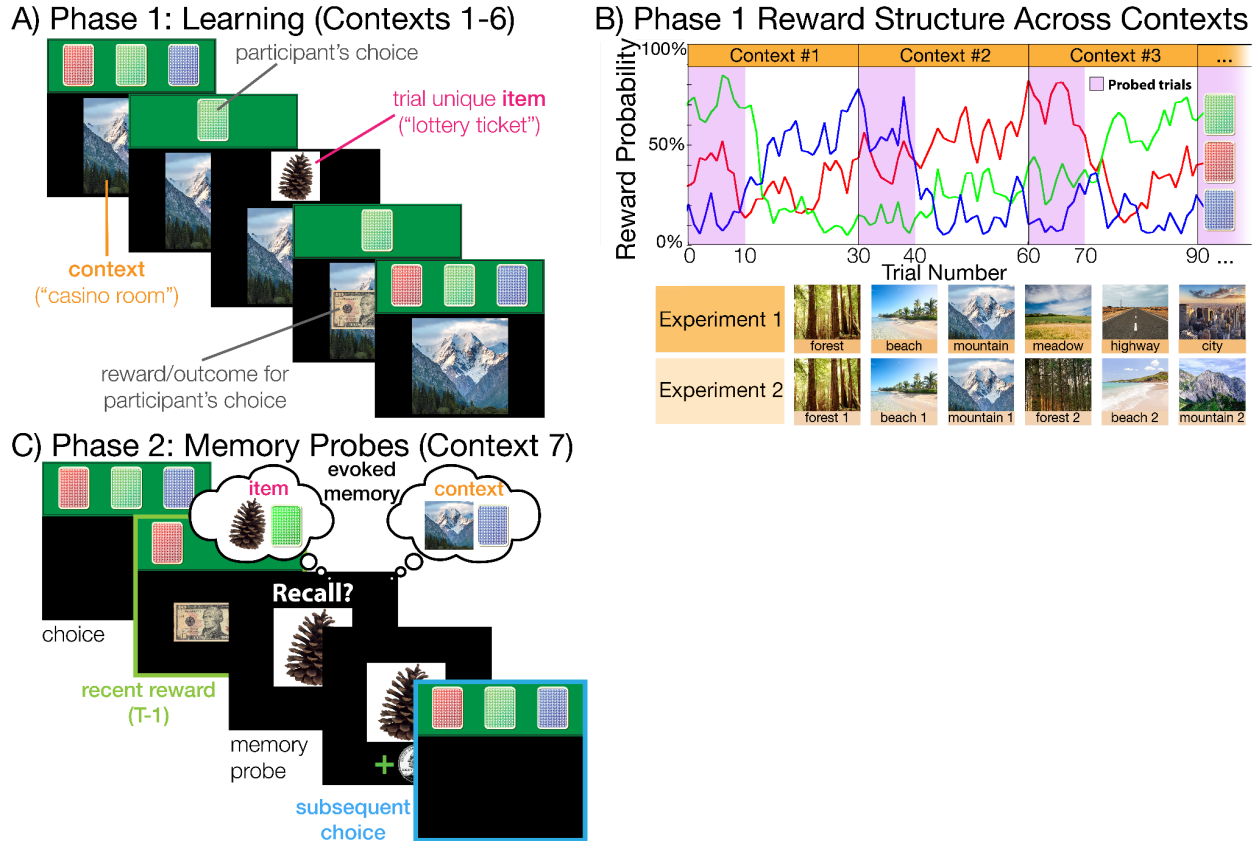


Figure 1: Three-armed restless bandit task. Across two experiments, participants performed a three-choice decision-making task with uncertain reward contingencies and trial-unique memoranda. Task was adapted from Bornstein & Norman (2017)³¹. **(a) Learning phase.** Participants first performed 180 choice trials, split into six consecutive “casino rooms.” Each “room” lasted for 30 trials and consisted of three differently-colored card decks on a green background at the top of the screen, and a scene image at the bottom of the screen. The scene remained on screen for all 30 trials of each room. After selecting a card deck, the top card was turned over to reveal a trial-unique item photograph. Then, a reward amount - either a picture of a US \$10 bill, or a phase-scrambled version of the same, the latter indicating a \$0 reward for that trial. Participants were told to treat each trial as equally important, and that their total payout would depend on the choice made on one randomly-selected trial. They were also told that the decks and “dealer” would follow them through each room of the casino. **(b) Reward structure.** Each card deck had a different, steadily-varying, probability of paying out a reward. Decks were initialized to 60%, 30%, and 10% randomly for each participant - these payouts then varied across trials according to a gaussian random walk tending towards one of those three probabilities as its center. After trial 10 in each room, the centers were rotated such that the new highest-paying deck was different from the previously highest-paying deck. The payout rankings continued uninterrupted into the next room, reinforced by the payout of the highest deck being set to 100% for the first three trials of that room. Though payout order did not change between rooms, context images did. In **Experiment 1**, each room had a highly distinct image as a background. In **Experiment 2**, two rooms were *perceptually aliased*, with background images that were highly visually similar (i.e. two forests, two beaches, two mountains). Critically, the payouts for the aliased rooms were opposed to each other - the highest paying deck in the first room of each kind was different from the highest paying deck in the second. **(c) Memory probe phase.** Participants were told that the remaining 120 trials would be performed in an “unfinished” room of the casino, which had no background image. The payouts continued to drift slowly and rotate every 30 trials, however there were

no further room boundaries. Here, choices no longer resulted in an item image - only a reward outcome. Interspersed among the 120 choice trials were 60 *memory probe trials*, 50 of which presented an image previously observed, and 10 of which presented a novel image. The key measure of interest was the degree to which participant choices on trials following the memory probe were biased towards choices made in the room brought to mind by the probed image.

Computational Model.

The *hybrid model* consisted of a participant-specific mixture of two value-estimation processes, each drawing on previous experience in different ways^{16,31}. The first process, *reinforcement learning*, is an incremental, error-driven updating process which learns an expected value, $V_{t,RL}(x)$ for each card deck, x , at each timepoint t (Eqn. 4;³⁴). This process had one free parameter, the learning rate (α_{RL}), which was fit to each participant individually, and captured the degree to which the expected value was updated by the difference between the reward received for the choice made on that trial ($R_t(x)$) and the value to be expected on the basis of previous experience ($V_{RL,t-1}(x)$):

$$V_{RL,t}(x) = V_{RL,t-1}(x) + \alpha_{RL} [R_t(x) - V_{RL,t-1}(x)] \quad (4)$$

Choice values were initialized to zero, and were not updated for options not chosen on a trial.

The second process, *memory sampling*, estimated values on the basis of a single “sample” drawn from past experiences with each choice option. The probability of a particular past experience serving as the value was proportional to how recently in the past it was selected, with recency weighted according to the decay parameter (α_{sample}):

$$P(V_{sample}(x) == R_i) = \alpha_{sample} (1 - \alpha_{sample})^{t-i} \quad (5)$$

For the model-fitting procedure, the value was computed as the expectation-weighted average value across all combinations of possible samples for each choice option¹⁶.

These values were transformed into action probabilities $P_t(x)$ via a *softmax* action-selection function, with separate *inverse temperature* parameters (β_{sample} , β_{RL}) reflecting how sensitive action selection is to the value differences estimated by that process, for that participant. A third parameter (β_p) reflected the influence of perseveration, or “stickiness” of recent choices, independent of the estimated action values.

$$P_t(\text{choose } x) = \frac{\exp[\beta_p I(x_t = x_{t-1}) + \beta_{\text{sample}} V_{\text{sample},t}(x) + \beta_{RL} V_{RL,t}(x)]}{\exp[\sum_{x=1}^3 \beta_p I(x_t = x_{t-1}) + \beta_{\text{sample}} V_{\text{sample},t}(x_t) + \beta_{RL} V_{RL,t}(x_t)]} \quad (6)$$

The resulting temperature parameters were treated as the variables of interest for our analyses below.

Models were fit using maximum a-posteriori estimation via unconstrained optimization (MATLAB function `fminunc`). Parameters were input to the optimizer as unbounded real numbers, and then logistic-transformed within the likelihood function to appropriate bounds: $\alpha_{\text{sample}}, \alpha_{RL} \sim [0, 1]$; $\beta_{\text{sample}}, \beta_{RL} \sim [0, 20]$; $\beta_p \sim [-3, 3]$. Randomly-sampled starting points were selected for each run of the optimizer until the minimum observed value did not change for five consecutive runs.

The following weakly informative priors, truncated within the ranges specified above, were used in the final likelihood calculation: $\alpha_{\text{sample}}, \alpha_{RL} \sim \text{Beta}(1, 1, 1, 1)$; $\beta_{\text{sample}}, \beta_{RL} \sim \text{Normal}(0, 10)T[0, 20]$; $\beta_p \sim \text{Normal}(0, 10)T[-3, 3]$.

Mnemonic Similarity Task

Participants were invited to return to complete the Mnemonic Similarity Task²⁷ in a separate session (Fig. 2). The Mnemonic Similarity Task was used as an individual difference measure of memory precision for each participant, which we reasoned would play a role in the degree to which participants would be influenced by evoked memories following probe trials in the second phase of the choice task. In the study phase, participants view 128 object images in an incidental encoding task in which participants are tasked with judging whether each image is an indoor or outdoor object. In a surprise recognition test, participants are shown 192 object images, with one-third consisting of exact repetitions of the same object that was shown during study (Fig. 2B, “repetition”), one-third consisting of perceptually similar objects (Fig. 2B, “lure”), and one-third consisting of objects that were never shown during the study phase (Fig. 2B, “foil”). Participants are asked to judge whether each object is old, similar, or new relative to what was shown during the study phase. A lure discrimination index (LDI) was computed for each participant using the following formula: $\text{LDI} = p(\text{“similar”}|\text{Lure}) - p(\text{“similar”}|\text{Foil})$. The LDI is widely used to measure individual and age-related differences in mnemonic discrimination, with higher scores being associated with better memory precision^{27,28,35–38}.

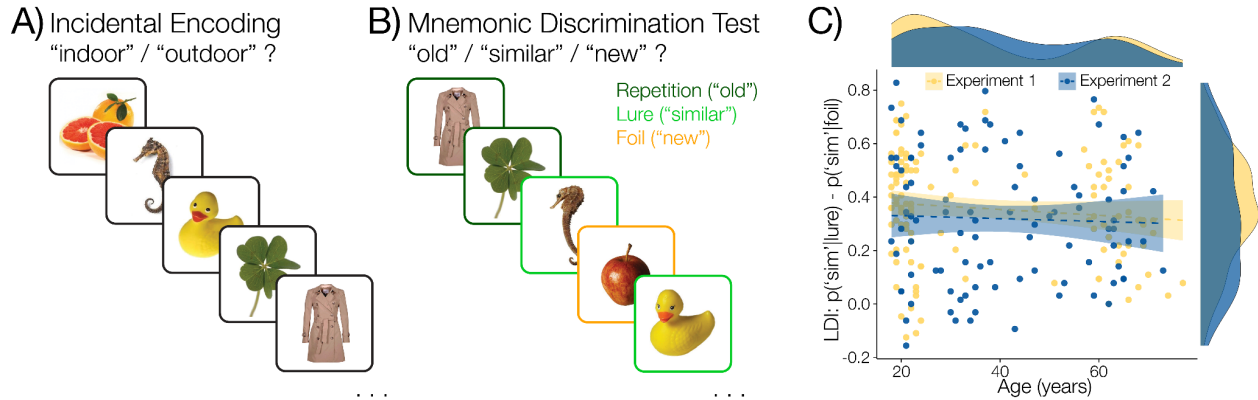


Figure 2. Mnemonic Similarity Task. (A) Participants view a sequence of objects during an incidental encoding phase in which participants classify each object as an indoor or outdoor object. (B) In a surprise discrimination test, participants view a series of objects and are asked to determine whether each object is “old,” “new,” or “similar” relative to the items that were shown during the encoding phase. (C) Correlation between age and lure discrimination scores (LDI) for each experiment. Shaded bands indicate 95% confidence intervals around the trendline.

Parameter correlations

The critical test of our hypotheses involved correlations between parameters extracted from computational model fits, described above, and external measures of interest (age, LDI). Because the hypothesized relationship between model parameters and these measures is ordinal, and because the variables exhibit “ties”, we report the rank-based Kendall’s τ_b correlation statistic³⁹. Comparisons between correlations (i.e. across experiment populations) were assessed by z-test after first converting to the corresponding linear R value using the mapping $R = \sin(0.5 * \pi * \tau)$ ⁴⁰.

To determine the specificity of LDI~parameter relationship, correlations were also examined after regressing out the linear relationship with age. This computation was performed by first estimating the residuals of the regression model $X \sim \beta * \text{age}$, and performing the corresponding correlation against the resulting values.

Results

Experiment 1.

The influence of memory intrusions on choice is observable across the lifespan. Following the preceding study which examined behavior in this task in young adults (Fig. 1;³¹), our measure of interest was performance on choice trials following recognition memory probes. Previous research using this decision-making task has shown that memory probes result in intrusive recollection of past choices such that evoked item memory content and evoked context memory content have separable contributions to subsequent choice behaviors, independently of recently received rewards³¹. Our initial question of interest was whether this general pattern of findings

was replicable in our lifespan sample of participants. Following previous work, we ran a multiple regression to model the contributions of each of our variables of interest on choice behavior. As previously observed in young adults^{16,31}, there was a significant influence of both the reminded trial and reminded context on choices following the memory probe, while controlling for the influence of recent rewards (Fig. 3A; *item*: $t(226) = 2.09$, $p = 0.038$; *ctx*: $t(226) = 2.32$, $p = 0.021$).

Previous work has documented a decline in memory specificity with age⁴¹. Such an effect could dampen the intrusive influence of retrieved values on decisions. On the other hand, other work has shown that memory's effect on value-guided choice reflects an elaborated retrieval process, and so the presence of memory-biased choice would itself necessitate the retrieval of specific information⁴². Therefore, we examined whether memory effects were modulated by either chronological age or cognitive decline, the latter indexed by the Lure Discrimination Index (LDI) computed on performance during the separately-administered Mnemonic Similarity Task (Fig. 2). The LDI, which measures the participants' selective accuracy in identifying similar lures, is widely used as a behavioral measure of the specificity of encoded memories, and is known to decrease with age and more specifically with the onset of age-related cognitive decline²⁸. Intriguingly, and consistent with the idea that value itself is the specific memory content of interest in this task⁴², we found that neither age nor LDI modulated the effect of item or context-based intrusions ($age * item$: $r_t(225) = -0.039$, $p = 0.391$; $LDI * item$: $r_t(128) = -0.925$, $p = 0.355$; $age * ctx$: $r_t(225) = 0.050$, $p = 0.271$; $LDI * ctx$: $r_t(128) = -0.583$, $p = 0.560$).

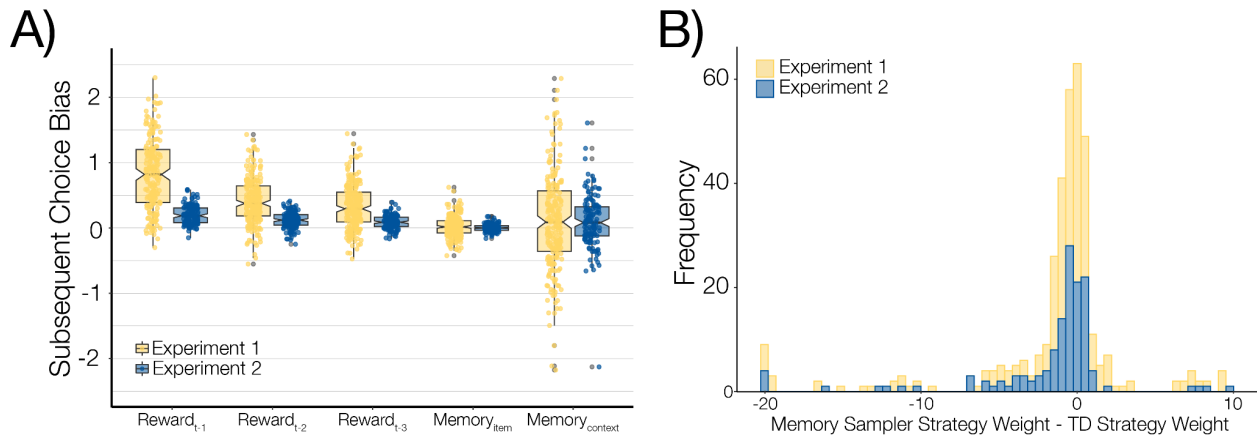


Figure 3. Influence of recent reinforcement and memory on choices. (A) Recent reinforcement and reminded memories both influence choices. Boxplots of logistic regression beta estimates (log relative choice odds) quantifying the contributions of reinforcement learning (recent reward history) and memory content (specific item memory vs. context memory) on subsequent choices following memory probes in Phase 2. In both Experiments, we observed a significant and consistent influence of both recent rewards and memory probes on subsequent choices. (Experiment 1: Yellow; Experiment 2: Blue.) **(B) Subjects**

exhibit a mixture of Memory sampling and Temporal-Difference Reinforcement Learning strategies. Across participants, the relative use of temporal-difference reinforcement learning and memory sampling varied. The bulk of subjects exhibited a near-equal mixture of the two, but overall recent reinforcement was a greater influence on choices (Experiment 1, yellow: $t(225) = -3.43, p < .001$; Experiment 2, blue: $t(134) = -3.75, p < .001$).

Memory sampling and RL are separately modulated by memory precision and age. Having observed that memory intrusions were consistent across the lifespan and the range of memory precision scores, we next examined factors that guided the use of behavioral strategies to perform the choice task. To answer this question, we fit to choice behavior a hybrid computational model implementing two distinct approaches to value-based decision-making. The key measure of interest in this model is the degree to which an individual's choices reflect values estimated by *memory sampling*^{15,16} and also values estimated using standard temporal-difference reinforcement learning⁴³. Models were fit to choices in the first six rooms of the “casino”, absent the influence of memory probes. The result of this model fitting is that the reliance on each strategy is reflected in the inverse softmax temperature parameters that best describe each participant's behavior as a function of the values estimated by each process. Overall, we found that most subjects exhibited a near-equal mixture of TD and memory sampling, but that TD was a greater contributor to behavior across the population ($t(225) = -3.43, p < .001$; Fig. 3B). Consistent with the hypothesis that the use of memory-based decision strategies is guided by the relative uncertainty of memory representations^{12,19–21}, we found that, at a subject-level, increased memory precision (as measured by LDI) predicted the use of a memory sampling strategy ($r_{\tau}(128) = .22, p < .001$); this effect remained even when regressing out age ($r_{\tau}(128) = .13, p = .03$). Notably, LDI did not predict the use of the reinforcement learning strategy ($LDI * \beta_{TD}: r_{\tau}(128) = .01, p = .9$), but, consistent with previous work suggesting increased choice noise and perseveration with age⁴⁴, we found that age predicted less reliance on reinforcement learning ($r_{\tau}(128) = -.11, p = .019$; Fig. 4A), as well as a greater influence of perseveration ($r_{\tau}(128) = .11, p = .017$; Fig. 4B).

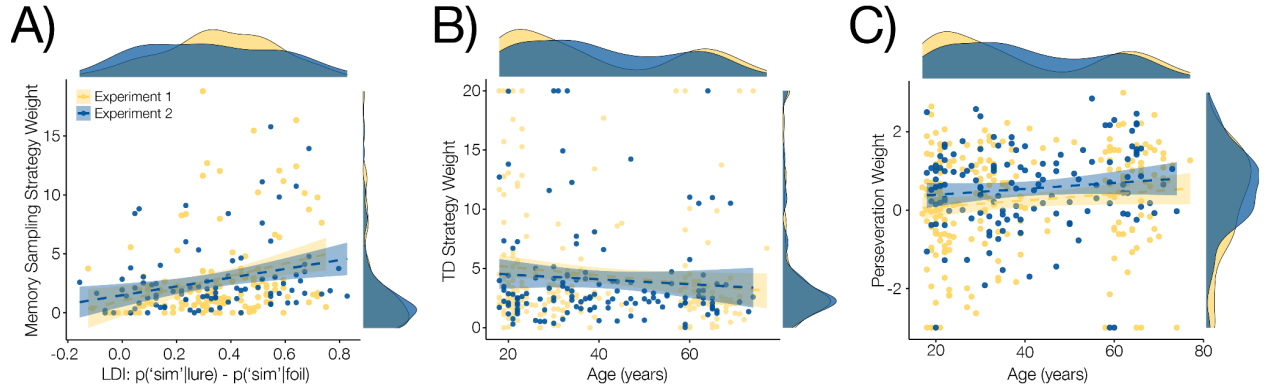


Figure 4. Memory precision and age differentially modulate the influence of experience on decisions. Both Lure Discrimination Index and Age separately tracked individuals' use of distinct experience-derived choice influences. **(A) Memory sampling increases with memory precision.** Consistent with theoretical and empirical findings that uncertainty is a critical factor in the weight of decision strategies, increased memory precision was associated with an increase in the use of the memory sampling strategy in choice: $r_{\tau}(225) = .228, p < .001$. **(B) Choices are less sensitive to reinforcement learning with age.** Consistent with previous findings that choice noise increases with age, participants exhibited less influence of Reinforcement Learning-derived values with age. $r_{\tau}(359) = -.087, p = .015$. **(C) Perseveration increases across the lifespan.** Consistent with previous findings, participants exhibited a greater tendency towards perseverative responding with age. $r_{\tau}(359) = .101, p = .0048$. All correlations plotted for each experiment individually (Experiment 1: Yellow, Experiment 2: Blue. Shaded bands represent 95%CI around the trendline), and statistics reported for the combined sample. Correlations were not different between experiment samples ($LDI * \beta_{sampler} : z = -0.21, p = .417$; $Age * \beta_{TD} : z = -0.704, p = .241$; $Age * \beta_{persev} : z = 0.446, p = .328$).

Experiment 2.

Age-related memory decline predicts the intrusion of gist, rather than specific-context memories on choice.

Experiment 1 demonstrated that the effect of memory intrusions on choice was not a function of age or memory specificity, but that both factors contributed to the relative influence of decision strategies on choice. We thus conducted a follow-up experiment examining whether a more fine-grained distinction, about the *content* of intrusive memories, might be revealed by a direct test of memory specificity. Specifically, motivated by previous research suggesting that age-related memory decline is associated with a greater reliance on 'gist' representations, rather than specific memories^{45–48}, we examined whether the content of probe-triggered memories was more 'gist'-like in individuals with lower memory precision. In our previous study using this task³¹, we found that neuroimaging measures of specific-scene reinstatement on each trial—which indicated the degree to which an individual reinstated alternative, rather than target, contexts on that trial—modulated the value that guided subsequent choice. This finding is consistent with the idea that samples are drawn from the context

reinstated at the time of each decision. Therefore, we reasoned that making the distinction between scene contexts more difficult might lead individuals with lower memory precision to sample, at least partly, from the alternate, 'lure' room.

In Experiment 2, individuals again performed a three-armed restless bandit task across six distinct room "contexts." However, in this experiment the contexts were split into perceptually aliased pairs - two scene images each of mountains, beaches, or forests (Fig. 1B). Critically, we controlled the payoff probabilities such that the best-performing card deck in the first room of a given type was not the best-performing in the second room. This allowed us to distinguish the effect of specific-context intrusions (e.g. 'this particular beach') versus intrusions based on a more 'gist'-like representation (e.g. 'beaches'). We first analyzed whether, consistent with Experiment 1 and previous observations, behavior was significantly modulated by reinstated target context intrusions, across the lifespan. Indeed, across the population, participants showed an effect of context-memory intrusions ($t(133) = 2.81, p = 0.006$). We next examined whether age or LDI correlated with the influence of the target context relative to its same-category lure (the other image of the same type ('beaches')). As expected, including the (orthogonalized, see *Methods*) lure context reward in the regression competed with variance for the target context, reducing the effect of the former (Fig. 5A), suggesting that, across the population, participants relied on a mixture of both target and lure context representations (i.e., a 'gist' category level representation of the context). We then examined the relative reliance on one (target context) versus the other (lure context), at a participant level, as a function of age and LDI. Consistent with the idea that age-related cognitive decline in memory precision alters decision-relevant representations, we found that the influence of the target context on choice was correlated with LDI score ($LDI * Target - Lure: r_{\tau}(95) = 0.148, p = 0.0329$), but not age ($Age * Target - Lure: r_{\tau}(95) = 0.070, p = 0.317$; difference between correlations: $z = 0.854, p = 0.197$); the relationship between LDI and target context influence held even after controlling for age ($r_{\tau}(95) = .144, p = .038$).

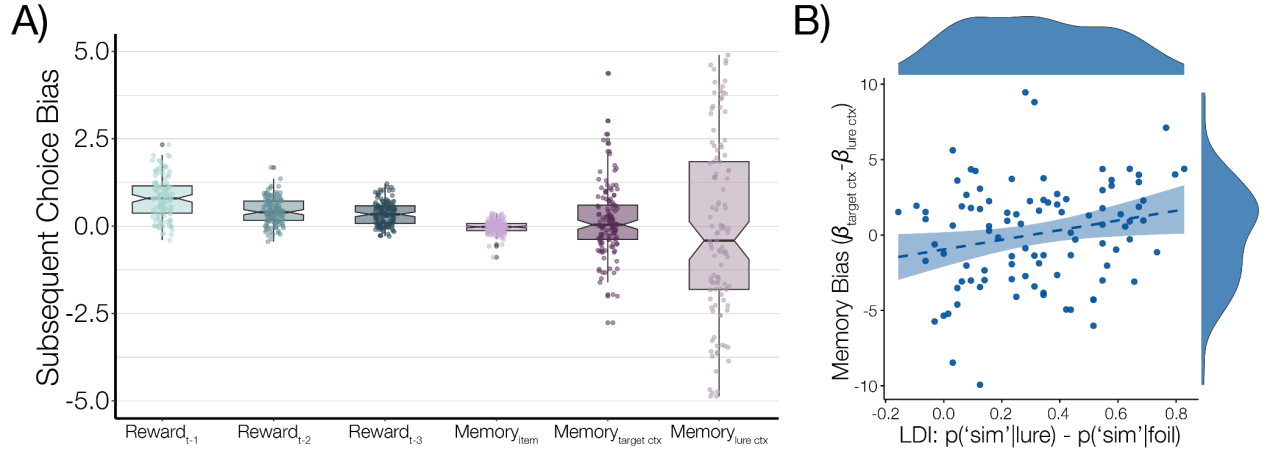


Figure 5. Experiment 2: Two kinds of reminded context influence choices. (A) Specific and gist context memory compete for influence on choice. Boxplots of logistic regression beta estimates measuring the contributions of reinforcement learning (recent reward history) and memory content (specific item memory, evoked target context memory, and evoked lure context memory) on subsequent choices following memory probes in Phase 2. Target and Lure contexts compete for variance in this analysis, reducing the effect of context observed in Fig. 3 (Blue). **(B) Memory precision tracks the influence of Specific and Lure context memory.** Across subjects, individual estimates of memory precision (LDI) predict a greater influence of memory intrusions for the target reminded context relative to its same-category lure ($r_{\tau}(95) = 0.148$, $p = 0.0329$).

Strategy use. Again, participants placed greater weight on the temporal-difference learning strategy than the memory sampling strategy ($t(134) = -3.75$, $p < .001$; Fig. 3B), and memory precision, but not age, predicted the use of the memory sampling strategy

($LDI * \beta_{sampler} : r_{\tau}(95) = .238$, $p < .001$; $age * \beta_{sampler} : r_{\tau}(133) = -.027$, $p = .641$; $z = 3.149$, $p = .001$). In this smaller sample, neither the correlation between age and TD strategy ($Age * \beta_{TD} : r_{\tau}(133) = -.0582$, $p = .321$; Fig. 4B, blue) nor between age and perseveration ($Age * \beta_{persever} : r_{\tau}(133) = .0766$, $p = .192$; Fig. 4C, blue) was statistically significant. Further, the differences between these correlations and those observed in Experiment 1 were not themselves significant ($Age * \beta_{TD} : z = -0.704$, $p = .241$; $Age * \beta_{persever} : z = 0.446$, $p = .328$)⁴⁹. Combining across experiments, both effects of age on strategy use were statistically significant ($Age * \beta_{TD} : r_{\tau}(359) = -.087$, $p = .015$; $Age * \beta_{persever} : r_{\tau}(359) = .101$, $p = .0048$; Fig. 4B,C).

Discussion

Individuals across the lifespan are often called on to make value-based decisions under uncertainty that have important consequences for themselves and others. However,

despite extensive characterization of age-related differences in decision-making capacities, little is known about the mechanisms by which these arise. Here, we examined the hypothesis that age-related differences in decisions may be linked to age-related differences in memory, which have been studied for longer and are better understood at a mechanistic level.

We specifically focused on the role of *memory sampling*, a process of retrieving memories of previous choice instances during decision deliberation that has been shown to affect neural signals and behavioral outcomes^{13,15,16,31,50}, even during tasks for which it is maladaptive to rely on episodic memories^{16,31,51,52}. We used a previously-validated task that, in younger adults, identified separable influences on choice of values associated with *item* and *context* memory, the latter corresponding with neural markers of context reinstatement that linked it to specific retrieved experiences - even if those experiences were incorrectly attributed³¹. In Experiment 1, we observed that a lifespan sample of participants replicated previous findings obtained in young adults. Further, we extended the previous findings by showing that *memory precision* - as measured by the separately-administered Mnemonic Similarity Task²⁷ - modulated the influence of memory sampling, with greater precision leading to a greater effect of memory sampling on choice. In parallel, and in keeping with previous findings in similar repeated decision-making tasks, chronological age was associated with increasing noisiness of choices relative to values estimated using standard reinforcement learning, and a concurrent increase in perseverative responding. In Experiment 2, we delved further into the relationship between memory precision and choice, by identifying a role for memory precision in selecting which memories are sampled. Specifically, we designed a variant of the previous task in which sampled context memories could be identified as specific or 'gist'-level (e.g. 'beaches' as opposed to 'that one particular beach'), with each having distinct, opposing effects on choice. We found that lower memory precision was associated with a greater reliance on gist-based memory during memory sampling. This result concords with an extensive literature on older adults' greater reliance on gist, rather than specific, memories⁴⁷, and connects it to a recently developed computational model of action selection to explain why some memories are sampled during decision deliberation, rather than others²².

These findings suggest that age-related differences in decision-making result at least in part from interactions between systems for memory-guided and reinforcement learning-based action selection. More specifically, they link age-related differences in specific memory functions, in particular pattern separation, to age-related alterations in value-based decision-making. The finding that older adults, broadly, may be less sensitive to values estimated using reinforcement learning suggests that this decline in sensitivity could be at least partly mitigated by a greater reliance on recency-weighted memory sampling, if memory precision is spared. Examining the interactive influence of these systems on choice across the lifespan is a promising topic for future research. For

instance, these results suggest that interventions that aim to improve older adults' sensitivity to value may profitably focus on improving access to, or encoding of, specific choice-related memories that may be of critical importance to future decisions.

Our results also bear on previous findings regarding the nature of mnemonic discrimination deficits in aging. There is some evidence that age effects on discrimination are amplified for objects compared to scenes^{53,54}. On the one hand, this suggests that our scene-based context manipulation in Experiment 2 was appropriate for participants across the lifespan and may explain why age did not modulate the effect of context-based intrusions on choices. On the other hand, however, our use of a separate object-based discrimination task to assess memory precision may not capture the same pattern separation processes taxed by the context-specific value-based decision task. Future examinations of the influence of age and memory precision (LDI) may benefit from using scene, not just object, mnemonic discrimination tasks.

More broadly, the finding that memory precision guides decision strategy has implications for the study of the computational nature of goal-directed decision-making more broadly⁵⁵. For one, that memory sampling and RL appear to be independently modulated by memory precision and age supports suggestions that these systems may indeed be distinct approaches to action selection^{17,32}, with the former supporting both model-free¹⁶ as well as model-based⁶ choice. Second, the finding that memory precision indexes both the specificity of contexts for sampling and also the overall reliance on memory suggests that uncertainty is computed dynamically and adjusted for in response to the available momentary evidence²⁰⁻²², rather than cached as a controller-specific quantity dictating patterns of choice across trials⁵⁶. At a neural level, future studies could investigate whether the seemingly independent differences in RL-based and memory-based choice processes are related to differences in fronto-striatal and medial temporal systems, respectively, extending prior work that typically examines these brain-behavior relationships separately for each neural system⁵⁷.

This study is not without its limitations: most of the limitations come from potential issues with online data collection. Because this project began during a global pandemic, a large majority of our data was collected online. As a result, our data was collected in a largely unsupervised fashion, and thus there is a risk that the data quality may not be as good as data collected through in-person means. We have taken several steps to remedy this such as: 1) setting eligibility filters to ensure that we had reliable online participants who have good experience participating in online studies and surveys, 2) comparing model parameter distributions between in-person and online samples for comparable groups, 3) manually inspecting data from all participants for quality using independent measures taken from raters who were blind to the hypotheses of the experiment, and 4) setting up detailed instructions and checks throughout the experiment to promote participant engagement (details provided in the Method section).

Despite our best efforts, it is likely there are limitations that we have not considered or perfectly controlled for in our experiment. For example, we were unable to collect in-person data from older adults during the project period, largely due to increased health risks associated with bringing older adults to campus for in-person testing during the global pandemic. Thus, our older adult data exclusively came from online sources. In general, online older adult samples tend to perform better on several cognitive measures (e.g., verbal fluency, subjective memory, subjective health, etc.)⁵⁸, which may result in weaker age-effects (specifically less evidence of age-related cognitive decline) and reduced potential for generalizability of our findings. However, in our study, we are still able to see significant age-related effects in the expected patterns (i.e., evidence of age-related deficits in performance) across our variables and analyses of interest, such as age-related increases in perseveration, and age-related decline in sensitivity to reinforcement learning. We expect that these age-related effects would only get stronger in a more diverse sample. Furthermore, by incorporating an independent and more standardized measure of cognitive functioning such as the MST, we are able to capture unique individual differences in cognitive functioning that may also be influencing performance on our decision-making task. With that said, future studies using this task could aim to collect a more representative and heterogeneous older adult population, with age-related cognitive decline characterized using standard neuropsychological and physiological measures, to strengthen the generalizability of our findings.

In sum, the finding that memory and reinforcement learning each exhibit distinct patterns with age and age-related cognitive decline provides new insights into the computational basis of age-related differences in decision-making, and suggests several new avenues for further research that may yield interventions of importance for mitigating the harmful effects of cognitive aging on individuals and their communities.

References

1. Nyberg, L., Bäckman, L., Erngrund, K., Olofsson, U. & Nilsson, L. G. Age differences in episodic memory, semantic memory, and priming: relationships to demographic, intellectual, and biological factors. *J. Gerontol. B Psychol. Sci. Soc. Sci.* **51**, P234–40 (1996).
2. Craik, F. I. M. Memory Changes in Normal Aging. *Curr. Dir. Psychol. Sci.* **3**, 155–158 (1994).
3. Eppinger, B., Walter, M., Heekeren, H. R. & Li, S.-C. Of goals and habits:

- age-related and individual differences in goal-directed decision-making. *Front. Neurosci.* **7**, 253 (2013).
4. Lempert, K. M. *et al.* Aging is associated with maladaptive episodic memory-guided social decision-making. *Proc. Natl. Acad. Sci. U. S. A.* **119**, e2208681119 (2022).
 5. Noh, S. M., Cooper, K. W., Stark, C. E. L. & Bornstein, A. M. Multi-step inference can be improved across the lifespan with individualized memory interventions. *Strategies* **26**, 28.
 6. Vikbladh, O. M. *et al.* Hippocampal Contributions to Model-Based Planning and Spatial Memory. *Neuron* **102**, 683–693.e4 (2019).
 7. Bruine de Bruin, W., Parker, A. M. & Fischhoff, B. Individual differences in adult decision-making competence. *J. Pers. Soc. Psychol.* **92**, 938–956 (2007).
 8. Sojitra, R. B., Lerner, I., Petok, J. R. & Gluck, M. A. Age affects reinforcement learning through dopamine-based learning imbalance and high decision noise—not through Parkinsonian mechanisms. *Neurobiol. Aging* **68**, 102–113 (2018).
 9. Head, D., Kennedy, K. M., Rodrigue, K. M. & Raz, N. Age differences in perseveration: cognitive and neuroanatomical mediators of performance on the Wisconsin Card Sorting Test. *Neuropsychologia* **47**, 1200–1203 (2009).
 10. Lemaire, P. & Brun, F. Adults' age-related differences in strategy perseveration are modulated by response-stimulus intervals and problem features. *Q. J. Exp. Psychol.* **67**, 1863–1870 (2014).
 11. Bornstein, A. & Banavar, N. V. Multi-plasticities: Distinguishing context-specific habits from complex perseverations.
<https://psyarxiv.com/t7vsc/download?format=pdf>.

12. Biderman, N. & Shohamy, D. Memory and decision making interact to shape the value of unchosen options. *Nat. Commun.* **12**, 4648 (2021).
13. Duncan, K. D. & Shohamy, D. Memory states influence value-based decisions. *J. Exp. Psychol. Gen.* **145**, 1420–1426 (2016).
14. Wang, S. & Bornstein, A. M. Mixing memory and desire. *Wiley Interdiscip. Rev. Cogn. Sci.* (08/2019).
15. Shadlen, M. N. & Shohamy, D. Decision Making and Sequential Sampling from Memory. *Neuron* **90**, 927–939 (2016).
16. Bornstein, A. M., Khaw, M. W., Shohamy, D. & Daw, N. D. Reminders of past choices bias decisions for reward in humans. *Nat. Commun.* **8**, 15958 (2017).
17. Lengyel, M. & Dayan, P. Hippocampal Contributions to Control: The Third Way. in *Advances in Neural Information Processing Systems 20* (eds. Platt, J. C., Koller, D., Singer, Y. & Roweis, S. T.) 889–896 (Curran Associates, Inc., 2008).
18. Nicholas, J., Daw, N. D. & Shohamy, D. Uncertainty alters the balance between incremental learning and episodic memory. *Elife* **11**, (2022).
19. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711 (2005).
20. Keramati, M., Dezfouli, A. & Piray, P. Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Comput. Biol.* **7**, e1002055 (2011).
21. Wang, S., Feng, S. F. & Bornstein, A. M. Mixing memory and desire: How memory reactivation supports deliberative decision-making. *Wiley Interdiscip. Rev. Cogn. Sci.* **13**, e1581 (2022).

22. Khoudary, A., Peters, M. A. K. & Bornstein, A. M. Precision-weighted evidence integration predicts time-varying influence of memory on perceptual decisions. in *Cognitive Computational Neuroscience* (2022).
23. Samanez-Larkin, G. R., Levens, S. M., Perry, L. M., Dougherty, R. F. & Knutson, B. Frontostriatal white matter integrity mediates adult age differences in probabilistic reward learning. *J. Neurosci.* **32**, 5333–5337 (2012).
24. Ness, H. T. *et al.* Reduced Hippocampal-Striatal Interactions during Formation of Durable Episodic Memories in Aging. *Cereb. Cortex* **32**, 2358–2372 (2022).
25. Zheng, L. *et al.* Reduced Fidelity of Neural Representation Underlies Episodic Memory Decline in Normal Aging. *Cereb. Cortex* **28**, 2283–2296 (2018).
26. Koen, J. D., Hauck, N. & Rugg, M. D. The Relationship between Age, Neural Differentiation, and Memory Performance. *J. Neurosci.* **39**, 149–162 (2019).
27. Stark, S. M., Kirwan, C. B. & Stark, C. E. L. Mnemonic Similarity Task: A Tool for Assessing Hippocampal Integrity. *Trends Cogn. Sci.* **23**, 938–951 (2019).
28. Stark, S. M., Yassa, M. A., Lacy, J. W. & Stark, C. E. L. A task to assess behavioral pattern separation (BPS) in humans: Data from healthy aging and mild cognitive impairment. *Neuropsychologia* **51**, 2442–2449 (2013).
29. Yassa, M. A. *et al.* High-resolution structural and functional MRI of hippocampal CA3 and dentate gyrus in patients with amnesic Mild Cognitive Impairment. *Neuroimage* **51**, 1242–1252 (2010).
30. Ballard, I. C., Wagner, A. D. & McClure, S. M. Hippocampal pattern separation supports reinforcement learning. *Nat. Commun.* **10**, 1073 (2019).
31. Bornstein, A. M. & Norman, K. A. Reinstated episodic context guides

- sampling-based decisions for reward. *Nat. Neurosci.* **20**, 997–1003 (2017).
32. Gershman, S. J. & Daw, N. D. Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework. *Annu. Rev. Psychol.* **68**, 101–128 (2017).
 33. Ashburner, J. *et al.* SPM8 manual. *Functional Imaging Laboratory, Institute of Neurology* (2012).
 34. Sutton, R. S. & Barto, A. G. *Reinforcement Learning, second edition: An Introduction*. (MIT Press, 2018).
 35. Kolarik, B. S., Stark, S. M. & Stark, C. E. L. Enriching Hippocampal Memory Function in Older Adults Through Real-World Exploration. *Front. Aging Neurosci.* **12**, 158 (2020).
 36. Lacy, J. W., Yassa, M. A., Stark, S. M., Muftuler, L. T. & Stark, C. E. L. Distinct pattern separation related transfer functions in human CA3/dentate and CA1 revealed using high-resolution fMRI and variable mnemonic similarity. *Learn. Mem.* **18**, 15–18 (2011).
 37. Yassa, M. A. & Stark, C. E. L. Pattern separation in the hippocampus. *Trends Neurosci.* **34**, 515–525 (2011).
 38. Bennett, I. J., Huffman, D. J. & Stark, C. E. L. Limbic Tract Integrity Contributes to Pattern Separation Performance Across the Lifespan. *Cereb. Cortex* **25**, 2988–2999 (2015).
 39. Daniel, W. W. Applied Nonparametric Statistics. *Technometrics* **33**, 364 (1991).
 40. Walker, D. A. JMASM9: Converting Kendall's tau for correlational or meta-analytic analyses. *J. Mod. Appl. Stat. Methods* **2**, 525–530 (2003).

41. Bowman, C. R., Chamberlain, J. D. & Dennis, N. A. Sensory Representations Supporting Memory Specificity: Age Effects on Behavioral and Neural Discriminability. *J. Neurosci.* **39**, 2265–2275 (2019).
42. DuBrow, S., Eberts, E. A. & Murty, V. P. A common mechanism underlying choice's influence on preference and memory. *Psychon. Bull. Rev.* **26**, 1958–1966 (2019).
43. O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H. & Dolan, R. J. Temporal difference models and reward-related learning in the human brain. *Neuron* **38**, 329–337 (2003).
44. Mata, R., Josef, A. K., Samanez-Larkin, G. R. & Hertwig, R. Age differences in risky choice: a meta-analysis. *Ann. N. Y. Acad. Sci.* **1235**, 18–29 (2011).
45. Abadie, M., Gavard, E. & Guillaume, F. Verbatim and gist memory in aging. *Psychol. Aging* **36**, 891–901 (2021).
46. Nolte, J., Löckenhoff, C. E. & Reyna, V. F. The influence of verbatim versus gist formatting on younger and older adults' information acquisition and decision-making. *Psychol. Aging* **37**, 197–209 (2022).
47. Grilli, M. D. & Sheldon, S. Autobiographical event memory and aging: older adults get the gist. *Trends Cogn. Sci.* **26**, 1079–1089 (2022).
48. Flores, C. C., Hargis, M. B., McGillivray, S., Friedman, M. C. & Castel, A. D. Gist-based memory for prices and 'better buys' in younger and older adults. *Memory* **25**, 565–573 (2017).
49. Gelman, A. & Stern, H. The Difference Between 'Significant' and 'Not Significant' is not Itself Statistically Significant. *Am. Stat.* **60**, 328–331 (2006).
50. Bakkour, A. *et al.* The hippocampus supports deliberation during value-based

- decisions. *Elife* **8**, (2019).
51. Bornstein, A. M. *et al.* Associative memory retrieval modulates upcoming perceptual decisions. *Cognitive, Affective, and Behavioral Neuroscience* (2023) doi:10.3758/s13415-023-01092-6.
 52. Hoskin, A. N., Bornstein, A. M., Norman, K. A. & Cohen, J. D. Refresh my memory: Episodic memory reinstatements intrude on working memory maintenance. *Cogn. Affect. Behav. Neurosci.* **19**, 338–354 (2019).
 53. Reagh, Z. M. *et al.* Greater loss of object than spatial mnemonic discrimination in aged adults. *Hippocampus* **26**, 417–422 (2016).
 54. Güsten, J., Ziegler, G., Düzel, E. & Berron, D. Age impairs mnemonic discrimination of objects more than scenes: A web-based, large-scale approach across the lifespan. *Cortex* **137**, 138–148 (2021).
 55. Morris, R. W., Bornstein, A. & Shenhav, A. *Goal-Directed Decision Making: Computations and Neural Circuits*. (Academic Press, 2018).
 56. Lee, S. W., Shimojo, S. & O'Doherty, J. P. Neural computations underlying arbitration between model-based and model-free learning. *Neuron* **81**, 687–699 (2014).
 57. Lighthall, N. R. Neural mechanisms of decision-making in aging. *Wiley Interdiscip. Rev. Cogn. Sci.* **11**, e1519 (2020).
 58. Ogletree, A. M. & Katz, B. How Do Older Adults Recruited Using MTurk Differ From Those in a National Probability Sample? *Int. J. Aging Hum. Dev.* **93**, 700–721 (2021).