

The Precision of Hippocampal Representations Predicts Incremental Value Learning across the Adult Lifespan

 Camilla van Geen,¹ Karolina M. Lempert,^{1,2} Michael S. Cohen,^{1,3} Kameron A. MacNear,⁴ Frances M. Reckers,¹ Laura Zaneski,¹ David A. Wolk,⁵ and Joseph W. Kable¹

¹Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, ²Gordon F. Derner School of Psychology, Adelphi University, Garden City, New York 11530, ³Department of Psychology, University of Chicago, Chicago, Illinois 60637, ⁴Department of Psychology, University of Illinois Urbana-Champaign, Champaign, Illinois 61820, and ⁵Hospital of the University of Pennsylvania, Philadelphia, Pennsylvania 19104

Correctly assigning value to different options and leveraging this information to guide choice is a cornerstone of adaptive decision-making. Reinforcement learning (RL) has provided a computational framework to study this process, and neural signals linked to RL have been identified in the striatum and medial prefrontal cortex. More recently, hippocampal contributions to this kind of value learning have been proposed, at least under some conditions. Here, we test whether the hippocampus provides a signal of the option's identity that aids in credit assignment when learning about several perceptually similar items and evaluate how this process differs across the lifespan. A sample of 251 younger and older adults (158 F; 93 M), including a subset ($n = 76$) with simultaneous fMRI, completed an RL task in which they learned the value of four houses through trial-and-error. Older adults showed decreased accuracy, accompanied by reduced neural signaling of value at choice but not feedback. Using representational similarity analysis, we found that the precision with which choice options were represented in the posterior hippocampus during choice predicted accuracy across age groups. Despite previous evidence for neural dedifferentiation with age, we found no support for a "blurring" of these stimulus representations in older adults. Rather, we observed reduced connectivity between the posterior hippocampus and the medial PFC in older adults, and this connectivity correlated with accuracy. Taken together, these findings identify a hippocampal contribution to incremental value learning, with age-related declines driven by reduced hippocampus–mPFC connectivity rather than a systematic change in the precision of hippocampal representations.

Key words: aging; decision-making; hippocampus; memory

Significance Statement

Traditionally, the hippocampus has been linked to declarative memory, while reinforcement learning (RL) has been associated with striatal systems. Here, we challenge this dichotomy by showing that even simple trial-and-error RL depends on the hippocampus: the precision with which it represents perceptually similar choice options predicts choice accuracy across the lifespan. In older adults, we find that impairments in RL performance are not related to degraded hippocampal representations but rather to reduced connectivity between the hippocampus and medial prefrontal cortex. These findings support a more integrated view of decision-making in which mnemonic and value-based systems interact even in basic learning contexts and highlight age-related disruptions in the transfer of information across brain regions.

Introduction

One major focus of decision neuroscience has been to understand the neural mechanisms underlying reinforcement learning (RL), our ability to estimate the value of choice options through

trial-and-error updating (Sutton and Barto, 1998; Niv, 2009). RL posits that the main driver of learning is the difference between the realized and expected rewards, known as reward prediction error (RPE). Neural activity that tracks RPE has been identified in the brain at the single neuron and population level and reflects dopaminergic signaling in the midbrain (Schultz et al., 1997; Schönberg et al., 2007; Niv et al., 2012; O'Doherty et al., 2015). Given the ubiquity of RL-like processes in our daily lives, understanding how they evolve throughout the lifespan is critical. Here, we systematically examine age-related differences in RL in a large sample, both neurally and computationally.

In the past, studies that measured the effects of healthy aging on RL have yielded mixed results. Several early studies reported

Received July 28, 2025; revised April 16, 2026; accepted April 23, 2026.

Author contributions: K.M.L., M.S.C., D.A.W., and J.W.K. designed research; K.M.L., M.S.C., K.A.M., F.M.R., and L.Z. performed research; C.v.G. and M.S.C. analyzed data; C.v.G. wrote the paper.

This work was funded by grant RF1AG058065 from the National Institute on Aging.

The authors declare no competing financial interests.

Correspondence should be addressed to Joseph W. Kable at kable@psych.upenn.edu.

This paper contains supplemental material available at: <https://doi.org/10.1523/JNEUROSCI.1459-25.2026>

<https://doi.org/10.1523/JNEUROSCI.1459-25.2026>

Copyright © 2026 the authors

poorer RL performance in older adults, linked to reduced striatal activity and dopaminergic decline (Volkow et al., 1996; Chowdhury et al., 2013; Samanez-Larkin et al., 2014). More recent studies, however, have found that older adults show no deficits in RL performance and intact value signaling in the brain (Lighthall et al., 2018; Daniel et al., 2020). Yet another study found that older adults do perform less well on simple RL tasks, but that this is primarily due to noisier value representations in prefrontal cortex during choice rather than reduced prediction error signaling (de Boer et al., 2017). More generally, the strength of the association between age and dopamine availability has been called into question (Karrer et al., 2017; Seaman et al., 2019). Thus, both the extent to which healthy aging leads to impairments in incremental value learning and where they reflect reduced RPE signaling remain open questions.

At the same time, recent work suggests RL also depends on mnemonic processes in the hippocampus (Biderman et al., 2020). Specifically, information from the hippocampus may serve as an input for credit assignment, allowing for a mapping between the choice option and its value (Foerde et al., 2013; Duncan et al., 2018; Ballard et al., 2019). Under this framework, distinct neural representations of choice options—maintained by mechanisms like pattern separation (Bakker et al., 2008; Yassa and Stark, 2011; Rolls, 2013)—should facilitate accurate value learning. As people age, these kinds of processes, which are responsible for the maintenance of detailed memories, tend to decline. For instance, work using the mnemonic similarity task (MST) has found that older adults are more likely than younger adults to mistake a perceptually similar lure for an item they have already seen (Stark et al., 2015, 2019). Neurally, older adults show altered processing in portions of the hippocampus that help maintain fine-grained perceptual details, potentially leading to the “blurring” of memory that occurs with age (Yassa et al., 2011). If specific information about choice options provides an input for credit assignment during RL, and if this information becomes less precise with age, this provides an alternative hypothesis for poorer RL performance. Rather than resulting from reduced RPE signaling, RL deficits in older adults may be related to altered processing in the hippocampus and its connections with value-responsive parts of the brain. To our knowledge, however, age-related differences in hippocampal contributions to RL have yet to be characterized.

Here, we set out to address these questions in a sample of older and younger adults performing an RL task in which they learned the value of four perceptually similar items. Of 251 participants, a subset of 76 performed this task while undergoing fMRI scanning, allowing us to measure the neural mechanisms underlying learning. Our specific questions were threefold. First, we asked whether older adults were impaired at RL. Second, we leveraged neural data and computational modeling to identify whether any deficits that emerged were related to disruptions in learning or in choice. Third, we asked whether performance on this task was linked to the precision with which choice options are represented in the hippocampus. Overall, we found that older adults do display subtle deficits in RL performance, linked to disrupted processing during choice but not during learning. Furthermore, across the lifespan, the precision with the choice options are represented in the posterior, but not anterior, hippocampus predicts accurate choice. Finally, connectivity between the posterior hippocampus and frontocortical circuits is reduced in older adults, highlighting a potential reduction in communication between the brain regions. Together, these findings show that accurately assigning value to perceptually

similar choice options may rely on an interplay between mnemonic processing in the hippocampus and frontocortical value signals and that aging may lead to impaired RL performance by altering this process.

Materials and Methods

Participants

The participants in this study are largely overlapping with those in Lempert et al. (2022) and van Geen et al. (2025). In previous publications, we report behavioral and fMRI data from these participants while they made decisions from a single episode. Here, we turn to the second of the two tasks they completed, in a separate experimental session, which measured their ability to learn value estimates incrementally from noisy feedback (see Reinforcement learning task section for details).


As illustrated in Figure 1B, 251 people (137 older adults and 114 younger adults; 158 F; 93 M), participated in this study. Of these 251, 76 participants (38 younger and 38 older adults) did the task while undergoing fMRI scanning. We had to exclude four of these 76 participants due to technical issues with the behavioral task, and one person for excessive head movement (see below, fMRI methods, for details on motion exclusions). In total, this resulted in a sample of 71 fMRI participants (37 younger and 34 older adults). Among the behavioral-only group ($N=175$), we collected data from 97 of these participants in-person before the start of the Covid-19 pandemic, after which we switched to collecting data remotely. To approximate the experience of in-person data collection, a member of our research team was on a video call with each participant while they completed the task; in total, we collected data from 78 participants via video call. We found no systematic differences in choice accuracy between the in-lab versus remote groups ($\beta_{\text{Method}} = 0.006$, $p = 0.82$; $\beta_{\text{Method} \times \text{Age}} = 0.005$, $p = 0.89$) and no systematic differences between participants who completed the study inside versus outside the scanner ($\beta_{\text{Scan}} = -0.03$, $p = 0.14$; $\beta_{\text{Scan} \times \text{Age}} = 0.005$, $p = 0.88$). The demographic characteristics of each group are described in Table 1.

We also collected behavioral data from a separate sample of middle-aged adults, whose ages spanned from 35 to 60 years old (mean age = 45 ± 7.3 ; 53 F, 24 M). Performance in this sample was in between the younger and older adults' and is illustrated in Supplementary Figure S1. Because we did not collect fMRI data from middle-aged adults, we focus the results of the paper on comparisons between younger and older adults. All of the older adults in our sample were recruited from the Clinical Core of the University of Pennsylvania's Alzheimer's Disease Research Center (ADRC). As part of their involvement in the center, participants undergo annual psychometric testing as defined by the Uniform Data Set 3.0 (UDS) and medical/neurological examination. Older adults who participated in this study were all designated as cognitively normal based on this assessment and consensus conference determination. The young adult sample was recruited through the Penn MindCORE online participant recruitment system as well as through Facebook ads. All participants provided informed consent, and the procedures were approved by the University of Pennsylvania's Institutional Review Board.

Reinforcement learning task

In our reinforcement learning (RL) task, participants had to incrementally learn the value of two sets of four stimuli: four houses and four faces, which were interleaved. In previous work, we have identified social biases—such as how generous a face looks—that differentially affect choices across the age range (Lempert et al., 2022; van Geen et al., 2025). Consistent with this, we observed a significant interaction between stimulus type and trial number in predicting choice accuracy in this dataset, such that learning slopes were shallower for face than house trials ($\beta_{\text{Trial} \times \text{StimType}} = 0.0043$, $p = 0.001$). This highlights that social processes likely disrupted value learning on face trials. However, for the four faces used in this task, their perceived generosity and how much they actually gave were perfectly anticorrelated, making it impossible to cleanly separate the contribution of social perceptions and feedback on learning. Given these confounds, the focus of this paper is on age-dependent differences in RL in the nonsocial domain, where we can more clearly

A Decision-making Task



B Participants




Figure 1. **A**, Schematic of the reinforcement learning (RL) task. Participants have 4 s to choose between two of four possible houses and then receive feedback on the outcome of their choice. Each house is assigned a different value ranging from \$1 to \$5, which participants receive 60% of the time. On 40% of trials, the outcome associated with each house is randomly generated, ensuring the value must be incrementally learned. **B**, Diagram of our participant sample. All participants completed the behavioral task, but only a subset underwent fMRI scanning while doing so.

Table 1. Demographic characteristics of the two participant groups—those who participated only in the behavioral version of the task and those who did so while undergoing fMRI scanning

Age group	Age range	Age details	Sex
Behavior only	Younger adults	18–34 (mean: 22.3)	54 F, 22 M
Behavior only	Older adults	61–92 (mean: 71.3)	60 F, 39 M
Behavior + fMRI	Younger adults	21–37 (mean: 26.4)	23 F, 15 M
Behavior + fMRI	Older adults	61–86 (mean: 72.6)	21 F, 17 M

attribute performance to purely mnemonic and reinforcement learning processes.

On each trial, participants were given 4 s to make a choice between two stimuli (Fig. 1A). Once they indicated their selection with a key press, they received feedback about how much the image was worth. Participants were told that their goal was to maximize the amount of money they received from each choice. To do so, they had to estimate the value of each stimulus and choose the one that was worth more. The task included four houses and four faces in total and outcomes were generated as follows: each stimulus was assigned a “true” value of either \$1, \$2, \$4, or \$5. Then, in the RL task, choice options were worth that amount 60% of the time, while outcomes were randomly selected from the three other values on 40% of trials. We chose 60% as the probability of receiving the true outcome, as this yielded the greatest variation in performance in pilot studies. In total, participants completed 120 trials of this task, divided into three separate blocks of 40 trials each. Each of the six possible face and house pairs was presented 10 times. In between trials, participants saw a fixation cross for a randomly determined amount of time between 1 and 5 s. At the end of the task, participants received the payout from one trial selected at random, which served as an additional bonus to incentivize accurate decision-making.

Behavioral data analysis

All behavioral analyses were conducted in R (version 4.2.3).

To examine age-related effects on overall RL performance, we computed the average proportion of correct choices across groups for each learning block (Fig. 3A). We coded a choice as correct if participants chose the higher-value house, as defined by its “true” value. To test whether performance accuracy was above chance by the end of learning, we conducted a one-sample *t* test comparing the proportion of correct choices for each participant on the last block to chance (corresponding

to 0.5, since there are two possible choices). We also ran an independent-samples *t* test on participants’ accuracy scores on the last block to assess whether age group significantly modulated task performance. Finally, we used the “lme4” package in R (Bates et al., 2013) to predict trial-wise accuracy as a function of block number, age, and the interaction between age and block number. In this analysis, we also included a random intercept for each participant to account for individual differences in baseline performance.

RL model

Model description

To assess whether there were systematic age-related differences in participants’ learning rates or choice consistency, we fit a Q-learning model to the data following the hierarchical Bayesian approach described in van Geen and Gerraty (2021). Each house was assigned a separate Q-value, which was initialized at 3 (midpoint between all possible “true” values). The Q-value of the chosen house was then updated after each trial as the sum of the predicted value of the chosen option ($Q_{s,i}^t$) and a reward prediction error ($R_s^t - Q_{s,i}^t$), weighted by a learning rate. In the best fitting version of this model (see below, Model comparison), each participant was assigned two different learning rates: one to update the Q-value in response to a positive prediction error (α_s^+) and one to update the Q-value in response to a negative prediction error (α_s^-). More formally, if *s* refers to subject, *t* to trial, and *i* to stimulus identity, this update corresponds to the following:

$$Q_{s,i}^{t+1} = \begin{cases} Q_{s,i}^t + \alpha_s^+ (R_s^t - Q_{s,i}^t), & \text{if } R_s^t - Q_{s,i}^t \geq 0 \\ Q_{s,i}^t + \alpha_s^- (R_s^t - Q_{s,i}^t), & \text{if } R_s^t - Q_{s,i}^t < 0. \end{cases}$$

We then modeled the likelihood of choosing one stimulus or the other as a Bernoulli distribution governed by parameter θ . θ corresponds to a SoftMax transformation of the Q-values, weighted by a subject-specific β . Because there were only two choice options on each trial, the SoftMax simplifies to a logistic transform. Higher values of β correspond to a greater sensitivity to the difference in value between the two options, leading to a greater tendency toward choosing the higher-value image. We will refer to this parameter as choice consistency in the following sections. Formally, characterizing choice as either 0 (choose stimulus on the right) or 1 (choose stimulus on the left), we can model the choice probability as follows:

$$P(c = 1 | \alpha_s, \beta_s) = \text{Bernoulli}(\theta),$$

$$\theta = \frac{1}{1 + e^{-\beta_i(Q_{s,i}^i - Q_{s,j}^i)}}.$$

Here, i and j index the two stimuli shown on a given trial. This equation represents the likelihood of an observed choice according to the Q-learning model, given a set of parameter values.

To extract participant-level parameter estimates, we fit the model hierarchically using Hamiltonian Markov chain Monte Carlo (MCMC) in Stan (Carpenter et al., 2017) with 4 chains and 4,000 iterations (including 2,000 as burn-in). This procedure yielded posterior distributions across all 4,000 samples for each parameter for each participant. We then extracted point estimates from these distributions by calculating the mean over the model-derived posteriors. To test for age differences, we conducted independent-sample t tests on the estimates of learning rate and choice consistency across age groups. Because the model was fit hierarchically, we simultaneously estimated group-level and participant-level parameters and used empirical priors from the group-level estimates for regularization (this method is described in detail in van Geen and Gerraty, 2021). We included the data from both younger and older adults in the same model, meaning that we did not make any assumptions about participants belonging to two different groups. This approach improves model performance because it maximizes the amount of data available for model fitting while also helping to ensure the robustness of any age differences we detect.

Model comparison

The Q-learning model we report in the Results section contains three free parameters: an inverse temperature or choice consistency parameter that reflects how strongly Q-value guides choice and two separate learning rates, one that updates value estimates in response to positive prediction errors and another that does so in response to negative prediction errors. We confirmed that this model outperforms three alternatives:

1. A policy gradient model that maintains a preference parameter w for each of the six possible stimulus pairs, similar to Li and Daw (2011). On each trial, w is updated based on the difference between the received reward and a fixed reference point, which we set to 3 as the mean of the possible reward values (1, 2, 4, and 5). Crucially, because w is pair-specific, learning about a stimulus in one pair

context does not transfer to other pair contexts, unlike in Q-learning. Using leave-one-out cross-validation, we find that the policy gradient model substantially underperforms relative to the Q-learning model ($\Delta\text{ELPD} = -695.9$, $\text{SE} = 47.9$) and that Q-learning provides a better fit than the policy gradient model for most participants (132 out of 224).

2. A variant of the Q-learning model that includes a stickiness parameter alongside choice consistency, as in Noh et al. (2023) for example. This stickiness parameter captures people's tendency to repeat their previous button press, regardless of value, and is implemented as an additional weight parameter in the SoftMax decision function as follows:

$$P(c = 1 | \alpha_s, \beta_s, \beta_{\text{stick},s}) = \text{Bernoulli}(\theta),$$

$$\theta = \frac{1}{1 + e^{-[\beta_s(Q_{s,i}^i - Q_{s,j}^i) + \beta_{\text{stick},s}(S)]}},$$

where S is coded as +0.5 if $c = 1$ on the previous trial and -0.5 if $c = 0$ on the previous trial. Using leave-one-out cross-validation, we find that the Q-learning model with stickiness underperforms relative to the baseline Q-learning model ($\Delta\text{ELPD} = -86.5$, $\text{SE} = 30$). Additionally, we see no age-related differences in the stickiness parameter, suggesting that it is not capturing meaningful age-related performance differences ($p = 0.76$).

3. A version of the Q-learning model that only includes a single learning rate, regardless of the valence of prediction error. The version with a single learning rate produced worse out-of-sample predictive accuracy, indicating worse fit and supporting the use of asymmetric learning rates ($\Delta\text{ELPD} = -239.8$, $\text{SE} = 21$). Model comparison across the four variants is plotted in Figure 2A.

Parameter recovery

To assess the recoverability of the model's parameters, we generated 1,000 synthetic datasets, each containing 224 subjects matching our empirical sample size. Each dataset was generated by first sampling group-level parameters uniformly from ranges defined by the minimum and maximum individual-level estimates observed in the data [choice consistency (β), range: 0.54–1.76; positive learning rate (α^+), range: 0.02–0.69; negative learning rate (α^-), range: 0.005–0.57]. We also




Figure 2. Model validation. **A**, Model comparison between the baseline Q-learning model with choice consistency and two learning rates using leave-one-out cross-validation. Error bars represent ± 2 SE. The baseline model outperforms a policy gradient model, a model with an additional stickiness parameter, and a model with a single learning rate. **B**, Parameter recovery from 1,000 synthetic datasets of 224 subject each (matching our empirical sample size for the full behavioral sample). Top, Correlation between true and recovery group-level parameters; bottom, correlation between true and recovered individual-level parameters.

sampled reasonable standard deviations for individual-level variability around these group means (β SD: 0.1–0.4; α^+ SD: 0.05–0.15; α^- SD: 0.05–0.15). Next, for each of the 224 subjects within a dataset, we sampled individual-level parameters by adding Gaussian noise (with the sampled standard deviations) to the group means. Finally, we simulated each subject's trial-by-trial choices using the sampled parameters and the actual trial sequence from the experiment. This procedure allows us to span the empirically plausible parameter space while preserving the hierarchical structure. We then fit the hierarchical Bayesian model to these simulated datasets and assessed parameter recovery by correlating the estimates with the true generative parameters. We find excellent parameter recovery for all three of the model's parameters at both the group and the individual levels ($r > 0.85$; Fig. 2B).

fMRI methods

Data acquisition and preprocessing

Magnetic resonance images were acquired with a 3 T Siemens Magnetom Prisma scanner and a 64-channel head coil. A three-dimensional, high-resolution structural image was acquired using a T1-weighted, magnetization-prepared, rapid-acquisition gradient-echo pulse sequence (voxel size, $0.8 \times 0.8 \times 0.8$ mm; matrix size, 241×286 ; 241 axial slices; repetition time, 3,000 ms; echo time, 30 ms). B0 field maps were acquired using a gradient echo sequence (voxel size, $2 \times 2 \times 3$ mm; matrix size, 96×96 ; 72 slices; repetition time, 1,270 ms; echo time, 7.46 ms; flip angle, 60°). While participants completed the task, functional images were acquired using a T2*-weighted gradient echoplanar imaging pulse sequence (voxel size, $2 \times 2 \times 2$ mm; interslice gap, 0.15 mm; matrix size, 98×98 ; 72 oblique axial slices; repetition time, 1,500 ms; echo time, 30 ms). Slices were angled $+30^\circ$ with respect to the anterior commissure–posterior commissure line to reduce signal dropout in the OFC (Weiskopf et al., 2006).

We preprocessed the data using *fMRIPrep* 20.2.2 (Esteban et al., 2019), which is based on *Nipype* 1.6.1 (Gorgolewski et al., 2011). All BOLD runs were motion corrected, slice time corrected, b0-map unwrapped, registered, and resampled to a Montreal Neurological Institute (MNI) 2 mm template. After preprocessing, we performed spatial smoothing in FSL using a Gaussian kernel with a full-width at half-maximum (FWHM) of 5 mm.

fMRI analyses

Univariate. We conducted all univariate fMRI analyses in FSL, using generalized linear models (GLMs) to measure changes in neural activity in response to our conditions of interest. For each model, we also included confound regressors to account for motion-related artifacts. These consisted of six realignment parameters (three translational and three rotational), as well as the temporal derivative, quadratic terms, and the quadratic of the temporal derivative of each. Additionally, we performed spike regression by creating as many unit impulse functions as there were volumes with framewise displacement >0.9 mm for each participant. Each of these regressors had a value of 1 at the timepoint where framewise displacement was larger than 0.9 mm and a value of 0 elsewhere (Ciric et al., 2017). The number of spike regressors varied substantially across participants, ranging from 0 to 74 with a median of 2. Given this wide range, we additionally excluded runs altogether if $>10\%$ of TRs had a framewise displacement of >0.9 mm (equivalent to any run with >27 spike regressors), and we excluded participants altogether if the majority of their runs (2 or more) fit this criterion. In practice, this resulted in the full exclusion of one participant and the exclusion of one run for an additional three people. Group-level maps were computed using a FLAME1 + FLAME2 mixed effects model, thresholded at $p = 0.001$, and corrected for multiple comparisons using parametric cluster-based correction.

We used a recently developed python-based pipeline (<https://github.com/aliceux/fmri-pipeline>) to help automate the analyses in FSL. As described in the Results section, we modeled neural activity separately for the 4 s choice phase and the 2 s feedback phase. For the results reported in Figure 3, we included the following boxcar regressors of interest in our GLM: (1) choice phase, (2) parametric modulator during the choice phase for the Q-value of the chosen option, (3) feedback phase,

and (4) parametric modulator during the feedback phase for reward prediction error. To make sure that we modeled neural activation comprehensively, we also included the same set of regressors for the trials on which participants chose between two faces, although these trials are not the focus of this paper. To verify that age-related differences in reaction time were not driving age differences in univariate results, we also ran a version of this analysis that included reaction time on each trial as a parametric regressor.

We then assessed whether there were any effects of age on the strength of neural signaling by combining a priori regions of interest (namely, the striatum and mPFC) with the results from the univariate analyses. In the case of prediction error signaling, since neural activity tracked with prediction error in a large cluster that extended beyond the striatum, we created a functional ROI by intersecting the univariate activation with value-responsive regions of the striatum identified in Bartra et al. (2013) (their Fig. 9: 5-way conjunction analysis designed to detect brain regions that track subjective value). In the case of Q-value signaling, we followed a similar approach and generated the fROI by intersecting the univariate activation with regions of the mPFC identified in Bartra et al. (2013), their Figure 9. Within each fROI, we then used *fslmeans* to extract subject-specific betas that reflect the magnitude of neural activation in response to the variable of interest. We then performed an independent-samples *t* test on the extracted coefficients in order to test for age differences. Given the difference in size between our behavioral and fMRI samples, we further evaluated age differences following the bootstrapping procedure described below.

Bootstrapping. Our goal in this paper was to integrate the behavioral and fMRI findings; however, a challenge in doing so is that the behavioral sample ($N = 251$) is much larger than the fMRI sample ($N = 72$). This difference in sample size makes it difficult to interpret age differences in the fMRI sample when they are not large enough to reach statistical significance. For example, the age differences we observed in behavior were only robust when we considered the entire sample, and not when we restrict this analysis to the fMRI group (see Supplementary Fig. S2 for fMRI-only behavioral results). Is this null result likely due to a lack of statistical power, or should we interpret it as evidence against age differences in the fMRI sample? To address this question, we devised a statistical test based on a bootstrapping procedure to assess whether age differences might have been reliably detectable had the fMRI sample been of the same size as the behavioral one.

To do this, we generated 10,000 hypothetical datasets by sampling data points (with replacement) until the simulated fMRI sample size was equivalent to our behavioral sample (137 older adults and 114 younger adults). For each generated dataset, we performed an independent-sample *t* test comparing the two age groups. This procedure generated a distribution of 10,000 *t* statistics, approximating the strength of the differences in a would-be sample of the same size as our behavioral sample. We then conducted a permutation test using the same sampling procedure but shuffling the age labels across subjects. We compared the null distribution to the bootstrapped distribution and computed a *p* value as the proportion of permuted *t* statistics that were equal to or exceeded the absolute value of the mean *t* statistic of the bootstrapped distribution.

This test estimates the likelihood that the age differences we observed in the fMRI sample would have reached statistical significance if we were to have an fMRI sample the size of our behavioral sample. We applied this procedure to both behavioral outcomes (performance accuracy and choice consistency) and neural activation, allowing us to assess whether any null results in the fMRI-only sample were attributable to reduced power rather than qualitatively different effects.

Representational similarity analysis. The process by which we identified stimulus-specific patterns of neural activity for each participant is illustrated in Figure 4A. First, we used *nibeta-series* to extract beta series from the functional MRI data, using the same set of confound regressors as in the univariate analyses. This procedure isolated patterns of neural activity during every trial of every run. Then, because the task design is such that two stimuli are always on the screen at the same time for each trial, we used a regression-based approach to estimate the neural response to each of the four possible houses independently (Fig. 4A, step 2). We created a stimulus design matrix for each scan run that was

composed of as many rows as there were trials (40) and as many columns as there were stimuli (four). We populated the stimulus design matrix with ones and zeros, with ones indicating the presence of that stimulus on that particular trial. Using ordinary least squares (OLS), we then regressed the neural activity onto the stimulus design matrix, allowing us to derive an activation coefficient for each voxel in response to each stimulus. Importantly, this only extracts stimulus-specific patterns of neural activity at the run level, rather than on each trial. Next, we correlated the run-level coefficients for each stimulus across runs to obtain a same-stimulus correlation (r_{same} ; e.g., house 1 in run 1 with house 1 in run 2). We then computed r_{diff} as the average correlation between each stimulus's representation and those of the other stimuli across runs (e.g., the average of house 1 in run 1 with houses 2, 3, and 4 in run 2). Because all four stimuli were present in every run, the same stimulus pairs contributed to r_{diff} consistently across runs. The difference between these two values ($r_{\text{same}} - r_{\text{diff}}$) is our main differentiation metric and captures the extent to which each person's brain represents different stimuli differently, relative to a within-participant baseline.

In this analysis, we were specifically interested in patterns of neural activity in the hippocampus, and its anterior and posterior portions independently. To localize neural activity to these regions of interest (ROIs), we multiplied the beta series with masks that contained zeros in all but the relevant brain areas. We used anterior and posterior hippocampal ROIs that were defined from a probabilistic atlas of the medial temporal lobe (Hindy and Turk-Browne, 2016).

Finally, because the similarity metric described above crucially relies on computing correlations in activity patterns across runs, we did not apply the 10% motion threshold that led to the exclusion of entire runs in the univariate analyses—rather, we included all runs for all participants, excluding the one participant with excessive head motion on every single run. However, we verified that all effects hold if we additionally exclude the three participants who each had excessive head motion on one of the runs.

Functional connectivity. To measure age effects on functional connectivity, we conducted a seed-based functional connectivity analysis in FSL. Based on our multivariate findings, we were specifically interested in understanding whether the posterior hippocampus was differentially connected to the rest of the brain in younger versus older adults. Thus, the GLM in this analysis consisted of a single regressor of interest, the timeseries from the posterior hippocampus, as well as the same regressors of no interest that we used in other analyses to account for motion. We then tested for whole-brain age differences using a FLAME1 + FLAME2 mixed effect model, thresholded at $z = 3.1$, and corrected for multiple comparisons using parametric cluster-based correction.

Given our univariate results, we were specifically interested in the magnitude of this effect in the mPFC. To quantify this, we used *fslmeans* to extract subject-specific betas in an ROI we defined anatomically using the Harvard-Oxford Atlas and tested to what extent age determines this brain region's connectivity with the posterior hippocampus. Given the high variability in the individual estimates, we applied a consistent outlier exclusion criterion across all connectivity analyses (including the controls described below): a modified Z-score exceeding 3.5 and significance on Grubbs' test with $p < 0.001$. This criterion led to the exclusion of three participants overall. We then ran an independent-samples t test on the extracted coefficients and compared the strength of connectivity between the posterior hippocampus and relevant portions of the mPFC in younger versus older adults.

To assess the specificity of this result, we took two approaches. First, we repeated the ROI connectivity analysis in two control regions: the motor cortex (precentral gyrus) and visual cortex (calcarine sulcus), as defined by the AAL atlas. For each region, we extracted subject-specific connectivity estimates with the posterior hippocampus using *fslmeans*, tested for age differences using an independent-samples t test, and examined the relationship between connectivity and behavioral outcomes (choice accuracy and model-derived choice consistency). Second, we conducted an exploratory whole-brain analysis to identify regions where posterior hippocampal connectivity positively predicted choice performance across participants. Individual connectivity maps were entered into a general linear model in FSL's *randomise*, with choice accuracy

as the regressor of interest. We ran the analysis with 5,000 permutations using threshold-free cluster enhancement (TFCE). As we did not find any significant effects at corrected thresholds, we report the results of this analysis at a lenient threshold of $p < 0.2$ to illustrate the location of the strongest effects.

Results

Behavior

First, we verified that participants in our full behavioral sample ($n = 251$) learn to choose the higher-value image over time. To do so, we computed the proportion of trials in each block on which participants chose the higher-value option. For both younger and older adults, choice accuracy increases significantly as learning progresses and is, on average, well above chance by the end of the task ($t_{(244)} = 17.1, p > 0.0001$; Fig. 3A). However, despite similar learning trajectories ($\beta_{\text{Block} \times \text{Age}} = 0.008; p = 0.25$), older adults make fewer correct choices by the end of learning, as illustrated by lower choice accuracy in block 3 ($t_{(244)} = -3.7, p = 0.0002$; Fig. 3A). Older adults also responded significantly more slowly than younger adults (linear mixed-effects model controlling for block and testing context: $\beta = -332$ ms, $p < 0.001$).

To better understand the mechanisms underlying these age differences, we fit a Q-learning model to the data that characterizes task performance based on a choice consistency parameter (β or inverse temperature) and two separate learning rates: one that updates value estimates after positive reward prediction errors and another that updates value estimates after negative reward prediction errors (see Materials and Methods for more details, including model comparison and parameter recovery). From this model, we find that older adults' reduced accuracy is captured by significantly lower choice consistency ($t_{(222)} = -3.34, p = 0.0009$; Fig. 3B) rather than age differences in learning rates ($\alpha: t_{(222)} = 1.58, p = 0.12$; $\alpha^-: t_{(222)} = 0.21, p = 0.82$; Fig. 3B).

For the subset of participants that completed the task in the fMRI scanner, we find that the same numerical patterns emerge (Supplemental Fig. S2A). Due to the reduced sample size, however, the effects of age on performance accuracy and choice consistency are not significant in the fMRI-only group ($t_{(70)} = 1.64; p = 0.104$ and $t_{(70)} = 1.60; p = 0.11$, respectively). To assess whether the absence of significant differences is likely due to reduced power, we used bootstrapping to estimate whether the behavioral effects observed in the fMRI sample would reach significance if the sample were matched in size to the behavioral cohort. These analyses indicate that both effects would reach significance with a matched sample size (accuracy: $p = 0.0014$; choice consistency: $p = 0.0034$; Supplemental Fig. S2B).

fMRI: univariate

Our behavioral findings suggest that age differences in RL performance are likely related to a noisier decision process rather than impaired value updating. To further assess this claim, we turned to the neural data and tested whether there are significant age differences in neural activity during feedback or during choice. Focusing first on the 2 s of feedback, we found that model-derived reward prediction error (RPE) tracks with neural activity in a variety of brain regions including the striatum and mPFC (Fig. 4A). In a functional ROI (Fig. 4A, colored yellow) identified based on the intersection between the univariate activation and the value-responsive regions of the striatum identified in Bartra et al. (2013), age differences in RPE signaling are not significant ($t_{(69)} = 0.82, p = 0.41$). Bootstrapping the fMRI data to estimate the effect in a sample the size of our behavioral sample (see Materials and Methods), we found that this null result is unlikely




Figure 3. *A*, Despite similar learning trajectories, older adults make fewer correct choices at the end of learning compared with younger adults. *B*, This subtle age-related impairment can be captured by decreased choice consistency in a standard Q -learning model. In this same model, there are no age differences in learning rates from both positive and negative reward prediction errors (α^+ and α^- , respectively).




Figure 4. *A*, Neural activity in value-responsive brain regions including the ventral striatum tracks with reward prediction error (RPE) across age groups. There are no significant age differences in the magnitude of this activation in a fROI defined as the intersection between the univariate activation and outcome-sensitive regions of the striatum identified in a large meta-analysis (Bartra et al., 2013). This is true in the actual fMRI sample and a bootstrapped group of participants whose sample size matches that of the behavioral data (see Materials and Methods for more details on this approach). *B*, Neural activity in the mPFC tracks with the Q -value of the chosen option across the full sample of participants. However, in a fROI of the mPFC, our bootstrapping analysis suggests that activation may be stronger in the younger than the older adult group. In both *A* and *B*, the fROIs from which we extracted participant-specific coefficients are overlaid in yellow and red, respectively.

to be due to the smaller sample size in the fMRI data (permutation test: $p = 0.10$). In fact, if anything, the trend is toward stronger RPE signaling in older adults. This set of findings parallels the absence of age differences in learning rate estimates derived from the Q -learning model. Both the neural and behavioral data suggest that deficits in RPE signaling during feedback are not at the root of older adults' lower performance.

Instead, older adults may represent Q -value less strongly, leading to more stochastic decisions. Turning to the choice phase, we found that consistent with previous work (Knutson et al., 2001; Niv, 2009), neural activity in the medial prefrontal cortex (mPFC) parametrically tracks the Q -value of the chosen option (Fig. 4*B*). To assess age differences in this effect, we again computed a fROI (Fig. 4*B*, red) by taking the intersection between

the univariate activation and the parts of mPFC that were identified as value sensitive in Bartra et al. (2013). Focusing on this fROI, we found that Q-value signaling is marginally stronger in the younger than the older adult group ($t_{(69)} = -1.85$, $p = 0.055$; Fig. 4B). When we used bootstrapping to estimate this effect in a sample the size of our behavioral sample, we found that this effect would likely reach statistical significance in a larger sample (permutation test: $p < 0.0001$). Overall, this set of fMRI findings is consistent with the notion that dampened value representations during choice may underlie noisier performance in older adults.

fMRI: representational similarity analysis

So far, we have identified a possible neural correlate of older adults' lower choice consistency: a tendency toward reduced Q-value signaling in the mPFC. It remains to be seen, however, what contributes to the strength of this Q-value representation and whether age may affect this process. One possibility is that accurate value signaling relies on the retrieval of information about choice options from the hippocampus. If this is true, we might expect the precision with which each stimulus is represented in the hippocampus to predict value learning and

consequent choice accuracy. Conversely, if these representations become blurred, participants may have more difficulty with credit assignment, leading to more stochastic choice. Given that this hypothesis relies on the maintenance of detailed and specific memories, we focused our analyses on the posterior portion of the hippocampus, which has been shown to be preferentially implicated in memory precision and pattern separation (Yassa and Stark, 2011; Yassa et al., 2011; Robin and Moscovitch, 2017; Wanja et al., 2021).

We next asked whether stimulus differentiation in the posterior hippocampus is important for successful credit assignment and consequently accurate choice. In line with this hypothesis, we found a positive correlation between stimulus differentiation in the posterior hippocampus and choice performance across participants ($\rho = 0.30$, $p = 0.012$; Fig. 5C). We also found a positive relationship between stimulus differentiation and each participants' choice consistency parameter from the Q-learning model ($\rho = 0.29$, $p = 0.016$; Fig. 5C). We verified that goodness of model fit was not a confound in this analysis by adding participant-specific log-likelihood as a covariate in a linear regression predicting choice consistency from neural




Figure 5. **A**, Overview of the method that allows us to identify stimulus-specific patterns of neural activity. First, we extract patterns of neural activity for each trial for every run. Then, in order to identify stimulus-specific activation, we regress neural activity for each run onto a stimulus design matrix using OLS. Finally, we correlate the beta coefficients for each stimulus across runs and compare the strength of same-stimulus versus different-stimulus correlations. **B**, Distribution of average neural differentiation across our sample. In the posterior hippocampus, patterns of neural activity at the group level are more similar when the same stimulus is repeated across runs than when different stimuli are repeated (mean differentiation > 0). **C**, Left, Stimulus differentiation in the posterior hippocampus predicts choice accuracy across participants. Right, Stimulus differentiation in the posterior hippocampus also predicts choice consistency, as captured by the β parameter in the Q-learning model.

differentiation. We found that the relationship between neural differentiation and inverse temperature remained positive and significant ($\beta = 0.13$; $p = 0.027$) while the relationship between neural differentiation and log-likelihood was not ($\beta = 0.23$; $p = 0.11$). Furthermore, these relationships did not hold in the anterior portion of the hippocampus, which we would expect to be less important for maintaining stimulus-specific representations ($\rho = 0.08$, $p = 0.50$, and $\rho = 0.06$, $p = 0.57$ respectively). These findings suggest that stimulus representations in the posterior hippocampus could provide an input for value learning and facilitate Q-value computations.

Given the results above, we might expect that older adults represent different stimuli less distinctly, leading to more imprecise valuation and lower performance. To test this hypothesis, we asked whether there were any age differences in stimulus differentiation in the posterior hippocampus or in its relationship to choice. Interestingly, the relationship between stimulus differentiation and choice performance holds across the age range, with no significant interaction between age and the strength of the effect ($\beta_{\text{Diff} \times \text{Age}} = 0.25$, $p = 0.43$). Furthermore, there was no main effect of age on relative stimulus differentiation ($t_{(68)} = -1.19$, $p = 0.24$). When we used bootstrapping to estimate age differences in stimulus differentiation in a sample the size of our behavioral sample, we found that a larger sample would possibly reveal stronger, rather than weaker, neural differentiation in the posterior hippocampus of older adults (permutation test: $p = 0.025$). Thus, the extent to which stimuli are distinctly represented in the posterior hippocampus does not seem to decrease with age.

fMRI: functional connectivity

Although stimulus differentiation appears intact, we see possible dampening in the extent to which the mPFC activity tracks Q-value in older adults. Mechanistically, these findings suggest that age may be linked to a disruption in how stimulus information from the posterior hippocampus impacts an evaluative process performed in the prefrontal cortex.

As a test of this hypothesis, we conducted a connectivity analysis looking at where neural activity is more functionally

coupled with the posterior hippocampus in younger than older adults. Given our hypothesis, we were particularly interested in measuring connectivity between the posterior hippocampus and the mPFC, as defined by an anatomical mask. Figure 6A shows that, as predicted, there is greater connectivity between the posterior hippocampus and the mPFC in the younger than older adult group ($t_{(70)} = -2.29$, $p = 0.026$). One young adult participant was a clear outlier in this analysis, identified by both a modified Z-score exceeding 3.5 and significance on Grubbs' test with $p < 0.001$. To ensure that this outlier was not driving the significant age difference, we verified that the effect was robust to excluding them ($t_{(69)} = -2.07$, $p = 0.043$). Finally, if the transfer of information between the posterior hippocampus and mPFC supports accurate value learning, we might expect the connectivity between the two brain regions to be correlated with choice accuracy and consistency. Supporting this hypothesis, we found that individual differences in functional connectivity between the two brain regions predict more accurate and consistent choice ($\rho = 0.25$, $p = 0.033$, and $\rho = 0.35$, $p = 0.002$ for accuracy and choice consistency respectively; Fig. 6B).

To assess the specificity of this effect, we took two approaches. First, we examined connectivity between posterior hippocampus and two control regions: the visual and motor cortices. We found no significant age difference in posterior hippocampal connectivity with the motor cortex ($t_{(69)} = -1.46$, $p = 0.15$), and connectivity with the motor cortex did not predict choice accuracy or consistency ($\rho = 0.06$, $p = 0.61$, and $\rho = 0.20$, $p = 0.093$, respectively). Although there was a significant age difference in posterior hippocampal connectivity with the visual cortex ($t_{(69)} = -3.5$, $p = 0.0005$), this connectivity did not predict accuracy or consistency ($\rho = 0.11$, $p = 0.35$, and $\rho = 0.16$, $p = 0.20$, respectively). Second, we ran a whole-brain analysis testing where brain connectivity with the posterior hippocampus predicts higher accuracy. No clusters survived a corrected whole brain threshold, and one of the largest clusters at a more lenient threshold ($p < 0.2$, TFCE corrected) was found within the mPFC (Supplementary Fig. S3). Taken together, these results highlight a potential disruption in the transfer of information between posterior hippocampus




Figure 6. *A*, Age differences in functional connectivity between posterior hippocampus and mPFC across the experimental task. There is greater functional connectivity between the two relevant brain regions in younger than older adults. *B*, Across individuals, connectivity between posterior hippocampus and mPFC predicts performance accuracy and choice consistency (the latter is derived from the computational model).

and mPFC in older adults, with downstream consequences for choice accuracy and consistency.

Discussion

Here, we aimed to address several questions: (1) Do older adults perform less well than younger adults on a reinforcement learning task? (2) Can the origins of any performance differences be traced back to altered processing during feedback or during choice? (3) Could hippocampal contributions to RL help explain individual differences in task performance, across individuals and between age groups? Our data suggest that older adults perform subtly but significantly less well in an RL task and that these deficits emerge due to reduced value signaling at choice rather than reduced prediction error signaling at feedback. Furthermore, precise representations of the choice options in the posterior hippocampus positively predict performance across the age range. The stimulus-specificity of these representations is not, contrary to what we hypothesized, lower in older adults. However, functional connectivity between the hippocampus and mPFC is, suggesting that the transfer of information rather than the information itself is disrupted with age.

Our finding that hippocampal representations of choice options predict performance expands traditional views of RL, which view it primarily as a form of stimulus-response learning that relies on the basal ganglia (Squire, 1992; Schönberg et al., 2007; O'Doherty et al., 2015; Gerraty et al., 2018). While we do see strong RPE signaling in the striatum, we also see evidence that stimulus information stored in the hippocampus may play a part in credit assignment. This perspective is in line with findings that highlight the hippocampus's role in incremental value learning, both in complex environments (Simon and Daw, 2011; Doll et al., 2012; Wimmer and Shohamy, 2012; Russek et al., 2017; Vikbladh et al., 2019; Liu et al., 2021) and also during simpler stimulus-outcome learning (Foerde et al., 2013; Davidow et al., 2016; Duncan et al., 2018; Ballard et al., 2019). In these tasks, value computation often relies on the explicit recruitment of mnemonic processes—for instance, to remember the feedback that was associated with a choice after a delay of several seconds. Similarly, in our task, participants choose between two stimuli drawn from a larger set of four options, meaning that only a subset of the choice space is visible on any given trial and the same stimulus may not reappear for up to 10 trials. Work in macaques has also found that hippocampal activity tracks the value of unavailable alternatives during decision-making (Fouragnan et al., 2019). Thus, hippocampal contributions to “model-free” RL may emerge when credit assignment relies on forming holistic representations of choice options that must be remembered across trials [but see Palombo et al. (2019) for a different perspective in amnesiacs].

That the hippocampus would be involved in both configural learning and in maintaining distinct, stimulus-specific representations as in our task points to its dual function. On the one hand, the hippocampus allows for associations and patterns to be remembered, combining features across experienced events (Brasted et al., 2003; Schapiro et al., 2012; Gilboa and Marlatte, 2017; Schlichting et al., 2017). On the other hand, the hippocampus also encodes episodic memories and maintains separate records of even very similar experiences (Scoville and Milner, 1957; Tulving and Markowitsch, 1998; Chanales et al., 2017; Korkki et al., 2021). To unite these opposing functions—pattern completion and pattern separation—researchers have identified a granularity gradient along the hippocampal long axis (Strange et al., 2014; Robin and Moscovitch, 2017). The anterior

hippocampus is thought to support more abstract, generalized processing, whereas the posterior hippocampus encodes detailed, item-specific representations. This distinction may reflect underlying subfield organization, as the dentate gyrus (DG) and CA3 are more prominent in the posterior hippocampus and constitute central nodes of a pathway that enables pattern separation (Leutgeb et al., 2007; Bakker et al., 2008; Yassa and Stark, 2011; Rolls, 2013; Schapiro et al., 2017; Wanji et al., 2021; Sučević and Schapiro, 2022). Consistent with this framework, we find that stimulus differentiation in posterior, but not anterior, hippocampus predicts accurate choice.

Although our findings highlight hippocampal contributions, they do not rule out a role for extra-hippocampal systems in RL. However, we find no evidence for an attenuation in striatal signals of RPE in older adults, consistent with the most recent research on the topic (de Boer et al., 2017; Lighthall et al., 2018; Daniel et al., 2020; Seaman et al., 2023). If anything, we and others (Daniel et al., 2020) see numerically stronger prediction error signals with age. Given that feedback-related activity in the striatum is stronger in high-effort tasks (Dobryakova et al., 2017), it could be that older adults perceive the task as more difficult. In any case, numerically stronger RPE signaling follows from assuming older adults form noisier value estimates and thus experience more prediction errors. Further corroborating this idea is the fact that in our sample, we see a negative relationship between RPE signaling in the striatum during feedback and Q-value representation at choice ($r = -0.29$, $p = 0.021$).

That said, several older papers do find reduced or incomplete prediction error signaling when older adults perform a range of probabilistic value learning tasks (Schott et al., 2007; Chowdhury et al., 2013; Eppinger et al., 2013; Samanez-Larkin et al., 2014). However, these studies either had small sample sizes ($N < 14$ in each age group) or did not include a younger adult comparison group. Furthermore, none report systematic age differences in striatal RPE signaling from whole-brain or ROI analyses that use model-derived prediction error as a regressor. Conversely, the more recent studies finding no age differences in RPE signaling (de Boer et al., 2017; Lighthall et al., 2018; Daniel et al., 2020) all have sample sizes of at least 25 participants in each group and identify neural signals of RPE using the same approach we do: deriving its magnitude from a model and using that as a regressor in ROI analyses. Thus, we believe that the preponderance of evidence favors the conclusion that there are no reliable age differences in prediction error signaling.

In line with the absence of neural differences in RPE, we also find that learning rates are not significantly reduced in older adults. The existing literature has yielded mixed results on this question. Radulescu et al. (2016), for instance, find that although older adults exhibit faster memory decay, there are no age differences in computationally derived learning rate. Cutler et al. (2021) measure to what extent people's RL parameters vary as a function of whether they are receiving rewards for themselves or others. The authors find that while young adults show higher learning rates when they are learning for themselves, older adults' learning rates do not vary by social context. As a result, younger adults have higher learning rates than older adults in the “self” (but not “other”) condition. It may be that motivational factors related to the social nature of the task are shaping these results, rather than a fundamental difference in how younger and older adults update their reward expectations. Still other studies have reported the opposite pattern, with older adults exhibiting higher learning rates than their younger counterparts (Eppinger et al., 2013; Koch et al., 2024). Taken together, these mixed findings

suggest that age-related differences in learning rates are not consistent across studies and may depend strongly on task demands and motivational context, as is also proposed in Daniel et al. (2020) and Lin et al. (2025).

Our finding that neural differentiation in the posterior hippocampus does not differ with age contributes to a similarly lively debate. On the one hand, the phenomenon of age-related neural dedifferentiation (Koen and Rugg, 2019) posits that as people get older, the amount of noise present in their neural activity increases, leading neurons to become less specialized and patterns of activation to become less distinct. Much of the evidence for this phenomenon comes from category-selective regions of the cerebral cortex, highlighting reduced functional specificity across cortical areas. For instance, neural activity in the fusiform face area of older adults responds less selectively to faces relative to other categories (Goh et al., 2010). Other studies quantify neural differentiation using multivariate pattern similarity within a given cortical region rather than across regions and similarly find that is reduced with age (St-Laurent et al., 2014; Srokova et al., 2024).

Whether we should expect similar decreases in the selectivity of neural response patterns in the hippocampus remains less clear. On the one hand, there is behavioral evidence for a reliable decrease in memory precision with age (Stark et al., 2015, 2019; Greene and Naveh-Benjamin, 2023) and neural evidence for DG/CA3 hyperactivity during pattern separation (Yassa et al., 2010). These findings have led to speculation about whether hippocampal activity patterns become less stimulus selective as people get older, underlying the reported “blurrier” memory (Leal and Yassa, 2018). However, as far as we know, evidence in favor of this claim based on multivariate analyses of patterns of neural activity is scarce. In general, we may not observe a decrease in hippocampal differentiation for at least three reasons: first, it may be that age-related dedifferentiation primarily reflects reduced functional selectivity across cortical regions rather than reduced precision within a region. Alternatively, the phenomenon may primarily occur in the cortex and not extend to the hippocampus. Finally, our current task may simply lack sufficient sensitivity to detect subtle changes in hippocampal differentiation, either because there are only four stimuli, or because they are not sufficiently similar.

Although hippocampal differentiation may not change with age, we observe disrupted communication from the hippocampus to mPFC. Prior work similarly reports reduced posterior hippocampus–mPFC coupling in older adults (Grady et al., 2003; Panitz et al., 2021) which, in our sample, may underlie the attenuation of value signaling in the mPFC. This mechanism would be in line with a view of the hippocampus’ contribution to RL as providing an input for credit assignment, as has been previously proposed (Ballard et al., 2019). Our findings, however, also raise an important follow-up question: if the hippocampus contributes to RL only insofar as it allows for the accurate transfer of value information to prefrontal regions, how come the precision with which it represents choice options predicts performance accuracy in older adults, for whom this process is disrupted? The simplest explanation is that despite the group-level decrease in connectivity, individual differences in the transfer of information relate to hippocampal precision. However, other alternatives could also explain these findings. For instance, it could be that in older adults, an alternate neural pathway, implicating the hippocampus and brain regions other than the mPFC, supports the link between stimulus representation and choice accuracy. Yet another possibility is that the hippocampus itself encodes not

only the identity of choice options but also information about their respective values (Knudsen and Wallis, 2021). Further adjudicating between these alternatives is an important question for future research.

References

- Bakker A, Kirwan CB, Miller M, Stark CEL (2008) Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science* 319:1640–1642.
- Ballard IC, Wagner AD, McClure SM (2019) Hippocampal pattern separation supports reinforcement learning. *Nat Commun* 10:1–12.
- Bartra O, McGuire JT, Kable JW (2013) The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage* 76:412–427.
- Bates D, Maechler M, Bolker B (2013) lme4: Linear mixed-effects models using Eigen and R syntax. <http://cran.r-project.org/package=lme4>
- Bideman N, Bakkour A, Shohamy D (2020) What are memories for? The hippocampus bridges past experience with future decisions. *Trends Cogn Sci (Regul Ed)* 24:542–556.
- Brasted PJ, Bussey TJ, Murray EA, Wise SP (2003) Role of the hippocampal system in associative learning beyond the spatial domain. *Brain* 126:1202–1223.
- Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker MA, Guo J, Li P, Riddell A (2017) Stan: a probabilistic programming language. *J Stat Softw* 76:1–32.
- Chanales AJH, Oza A, Favila SE, Kuhl BA (2017) Overlap among spatial memories triggers repulsion of hippocampal representations. *Curr Biol* 27:2307–2317.
- Chowdhury R, Guitart-Masip M, Lambert C, Dayan P, Huys Q, Düzel E, Dolan RJ (2013) Dopamine restores reward prediction errors in old age. *Nat Neurosci* 16:648–653.
- Ciric R, Wolf DH, Power JD, Roalf DR, Baum GL, Ruparel K, Shinohara RT, Elliott MA, Eickhoff SB, Davatzikos C (2017) Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *Neuroimage* 154:174–187.
- Cutler J, Wittmann MK, Abdurrahman A, Hargitai LD, Drew D, Husain M, Lockwood PL (2021) Ageing is associated with disrupted reinforcement learning whilst learning to help others is preserved. *Nat Commun* 12:4440.
- Daniel R, Radulescu A, Niv Y (2020) Intact reinforcement learning but impaired attentional control during multidimensional probabilistic learning in older adults. *J Neurosci* 40:1084–1096.
- Davidow JY, Foerke K, Galván A, Shohamy D (2016) An upside to reward sensitivity: the hippocampus supports enhanced reinforcement learning in adolescence. *Neuron* 92:93–99.
- de Boer L, Axelsson J, Riklund K, Nyberg L, Dayan P, Bäckman L, Guitart-Masip M (2017) Attenuation of dopamine-modulated prefrontal value signals underlies probabilistic reward learning deficits in old age. *eLife* 6:e26424.
- Dobryakova E, Jessup RK, Tricomi E (2017) Modulation of ventral striatal activity by cognitive effort. *Neuroimage* 147:330–338.
- Doll BB, Simon DA, Daw ND (2012) The ubiquity of model-based reinforcement learning. *Curr Opin Neurobiol* 22:1075–1081.
- Duncan K, Doll BB, Daw ND, Shohamy D (2018) More than the sum of its parts: a role for the hippocampus in configural reinforcement learning. *Neuron* 98:645–657.e6.
- Eppinger B, Walter M, Heekeren HR, Li S-C (2013) Of goals and habits: age-related and individual differences in goal-directed decision-making. *Front Neurosci* 7:253.
- Esteban O, et al. (2019) fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat Methods* 16:111–116.
- Foerke K, Race E, Verfaellie M, Shohamy D (2013) A role for the medial temporal lobe in feedback-driven learning: evidence from amnesia. *J Neurosci* 33:5698–5704.
- Fouragnan EF, Chau BKH, Folloni D, Kolling N, Verhagen L, Klein-Flügge M, Tankelevitch L, Papageorgiou GK, Aubry J-F, Sallet J (2019) The macaque anterior cingulate cortex translates counterfactual choice value into actual behavioral change. *Nat Neurosci* 22:797–808.
- Gerraty RT, Davidow JY, Foerke K, Galván A, Bassett DS, Shohamy D (2018) Dynamic flexibility in striatal-cortical circuits supports reinforcement learning. *J Neurosci* 38:2442–2453.
- Gilboa A, Marlatte H (2017) Neurobiology of schemas and schema-mediated memory. *Trends Cogn Sci (Regul Ed)* 21:618–631.

- Goh JO, Suzuki A, Park DC (2010) Reduced neural selectivity increases fMRI adaptation with age during face discrimination. *Neuroimage* 51:336–344.
- Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS (2011) Nipy: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front Neuroinform* 5:13.
- Grady CL, McIntosh AR, Craik FIM (2003) Age-related differences in the functional connectivity of the hippocampus during memory encoding. *Hippocampus* 13:572–586.
- Greene NR, Naveh-Benjamin M (2023) Adult age-related changes in the specificity of episodic memory representations: a review and theoretical framework. *Psychol Aging* 38:67–86.
- Hindy NC, Turk-Browne NB (2016) Action-based learning of multistate objects in the medial temporal lobe. *Cereb Cortex* 26:1853–1865.
- Karrer TM, Josef AK, Mata R, Morris ED, Samanez-Larkin GR (2017) Reduced dopamine receptors and transporters but not synthesis capacity in normal aging adults: a meta-analysis. *Neurobiol Aging* 57:36–46.
- Knudsen EB, Wallis JD (2021) Hippocampal neurons construct a map of an abstract value space. *Cell* 184:4640–4650.
- Knutson B, Fong GW, Adams CM, Varner JL, Hommer D (2001) Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport* 12:3683–3687.
- Koch C, Zika O, Bruckner R, Schuck NW (2024) Influence of surprise on reinforcement learning in younger and older adults. *PLoS Comput Biol* 20:e1012331.
- Koen JD, Rugg MD (2019) Neural dedifferentiation in the aging brain. *Trends Cogn Sci (Regul Ed)* 23:547–559.
- Korkki SM, Richter FR, Simons JS (2021) Hippocampal–cortical encoding activity predicts the precision of episodic memory. *J Cogn Neurosci* 33:2328–2341.
- Leal SL, Yassa MA (2018) Integrating new findings and examining clinical applications of pattern separation. *Nat Neurosci* 21:163–173.
- Lempert KM, Cohen MS, MacNear KA, Reckers FM, Zaneski L, Wolk DA, Kable JW (2022) Aging is associated with maladaptive episodic memory-guided social decision-making. *Proc Natl Acad Sci U S A* 119:e2208681119.
- Leutgeb JK, Leutgeb S, Moser M-B, Moser EI (2007) Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315:961–966.
- Li J, Daw ND (2011) Signals in human striatum are appropriate for policy update rather than value prediction. *J Neurosci* 31:5504–5511.
- Lighthall NR, Pearson JM, Huettel SA, Cabeza R (2018) Feedback-based learning in aging: contributions and trajectories of change in striatal and hippocampal systems. *J Neurosci* 38:8453–8462.
- Lin W-H, Pilz KS, Herzog MH, Kunchulia M (2025) Intact reinforcement learning in healthy ageing. *Exp Brain Res* 243:185.
- Liu Y, Mattar MG, Behrens TEJ, Daw ND, Dolan RJ (2021) Experience replay is associated with efficient nonlocal learning. *Science* 372:eabf1357.
- Niv Y (2009) Reinforcement learning in the brain. *J Math Psychol* 53:139–154.
- Niv Y, Edlund JA, Dayan P, O'Doherty JP (2012) Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *J Neurosci* 32:551–562.
- Noh SM, Singla UK, Bennett IJ, Bornstein AM (2023) Memory precision and age differentially predict the use of decision-making strategies across the lifespan. *Sci Rep* 13:17014.
- O'Doherty JP, Lee SW, McNamee D (2015) The structure of reinforcement-learning mechanisms in the human brain. *Curr Opin Behav Sci* 1:94–100.
- Palombo DJ, Hayes SM, Reid AG, Verfaellie M (2019) Hippocampal contributions to value-based learning: converging evidence from fMRI and amnesia. *Cogn Affect Behav Neurosci* 19:523–536.
- Panitz DY, Berkovich-Ohana A, Mendelsohn A (2021) Age-related functional connectivity along the hippocampal longitudinal axis. *Hippocampus* 31:1115–1127.
- Radulescu A, Daniel R, Niv Y (2016) The effects of aging on the interaction between reinforcement learning and attention. *Psychol Aging* 31:747–757.
- Robin J, Moscovitch M (2017) Details, gist and schema: hippocampal–neocortical interactions underlying recent and remote episodic and spatial memory. *Curr Opin Behav Sci* 17:114–123.
- Rolls E (2013) The mechanisms for pattern completion and pattern separation in the hippocampus. *Front Syst Neurosci* 7:74.
- Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND (2017) Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput Biol* 13:e1005768.
- Samanez-Larkin GR, Worthy DA, Mata R, McClure SM, Knutson B (2014) Adult age differences in frontostriatal representation of prediction error but not reward outcome. *Cogn Affect Behav Neurosci* 14:672–682.
- Schapiro AC, Kustner LV, Turk-Browne NB (2012) Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Curr Biol* 22:1622–1627.
- Schapiro AC, Turk-Browne NB, Botvinick MM, Norman KA (2017) Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philos Trans R Soc Lond B Biol Sci* 372:20160049.
- Schlichting ML, Guarino KF, Schapiro AC, Turk-Browne NB, Preston AR (2017) Hippocampal structure predicts statistical learning and associative inference abilities during development. *J Cogn Neurosci* 29:37–51.
- Schönberg T, Daw ND, Joel D, O'Doherty JP (2007) Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *J Neurosci* 27:12860–12867.
- Schott BH, Niehaus L, Wittmann BC, Schutze H, Seidenbecher CI, Heinze HJ, Duzel E (2007) Ageing and early-stage Parkinson's disease affect separable neural mechanisms of mesolimbic reward processing. *Brain* 130:2412–2424.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599.
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatr* 20:11.
- Seaman KL, Smith CT, Juarez EJ, Dang LC, Castellon JJ, Burgess LL, San Juan MD, Kundzic PM, Cowan RL, Zald DH (2019) Differential regional decline in dopamine receptor availability across adulthood: linear and nonlinear effects of age. *Hum Brain Mapp* 40:3125–3138.
- Seaman KL, Christensen AP, Senn KD, Cooper JA, Cassidy BS (2023) Age-related differences in the social associative learning of trust information. *Neurobiol Aging* 125:32–40.
- Simon DA, Daw ND (2011) Neural correlates of forward planning in a spatial decision task in humans. *J Neurosci* 31:5526–5539.
- Squire LR (1992) Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. *J Cogn Neurosci* 4:232–243.
- Srokova S, Aktas ANZ, Koen JD, Rugg MD (2024) Dissociative effects of age on neural differentiation at the category and item levels. *J Neurosci* 44:e0959232023.
- St-Laurent M, Abdi H, Bondad A, Buchsbaum BR (2014) Memory reactivation in healthy aging: evidence of stimulus-specific dedifferentiation. *J Neurosci* 34:4175–4186.
- Stark SM, Stevenson R, Wu C, Rutledge S, Stark CEL (2015) Stability of age-related deficits in the mnemonic similarity task across task variations. *Behav Neurosci* 129:257.
- Stark SM, Kirwan CB, Stark CEL (2019) Mnemonic similarity task: a tool for assessing hippocampal integrity. *Trends Cogn Sci (Regul Ed)* 23:938–951.
- Strange BA, Witter MP, Lein ES, Moser EI (2014) Functional organization of the hippocampal longitudinal axis. *Nat Rev Neurosci* 15:655–669.
- Sučević J, Schapiro AC (2022) A neural network model of hippocampal contributions to category learning. *BioRxiv*, 2022.01.12.476051.
- Sutton RS, Barto AG (1998) *Introduction to reinforcement learning*. Vol. 2. Cambridge, MA: MIT Press.
- Tulving E, Markowitsch HJ (1998) Episodic and declarative memory: role of the hippocampus. *Hippocampus* 8:198–204.
- van Geen C, Cohen MS, Lempert KM, MacNear KA, Reckers FM, Zaneski L, Wolk DA, Kable JW (2025) Age-related differences in trust decisions: when memory fails and appearances prevail. *Soc Cogn Affect Neurosci* 20:nsaf032.
- van Geen C, Gerraty RT (2021) Hierarchical Bayesian models of reinforcement learning: introduction and comparison to alternative methods. *J Math Psychol* 105:102602.
- Vikbladh OM, Meager MR, King J, Blackmon K, Devinsky O, Shohamy D, Burgess N, Daw ND (2019) Hippocampal contributions to model-based planning and spatial memory. *Neuron* 102:683–693.
- Volkow ND, Ding Y-S, Fowler JS, Wang G-J, Logan J, Gatley SJ, Hitzemann R, Smith G, Fields SD, Gur R (1996) Dopamine transporters decrease with age. *J Nucl Med* 37:554–559.
- Wanjia G, Favila SE, Kim G, Molitor RJ, Kuhl BA (2021) Abrupt hippocampal remapping signals resolution of memory interference. *Nat Commun* 12:1–11.
- Weiskopf N, Hutton C, Josephs O, Deichmann R (2006) Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage* 33:493–504.

- Wimmer GE, Shohamy D (2012) Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science* 338:270–273.
- Yassa MA, Stark SM, Bakker A, Albert MS, Gallagher M, Stark CEL (2010) High-resolution structural and functional MRI of hippocampal CA3 and dentate gyrus in patients with amnesic mild cognitive impairment. *Neuroimage* 51:1242–1252.
- Yassa MA, Lacy JW, Stark SM, Albert MS, Gallagher M, Stark CEL (2011) Pattern separation deficits associated with increased hippocampal CA3 and dentate gyrus activity in nondemented older adults. *Hippocampus* 21:968–979.
- Yassa MA, Stark CEL (2011) Pattern separation in the hippocampus. *Trends Neurosci* 34:515–525.