

Episodic memory facilitates flexible decision-making via access to detailed events

Received: 13 March 2025

Accepted: 20 November 2025

Published online: 23 January 2026

 Check for updates

Jonathan Nicholas  & Marcelo G. Mattar

Our experiences contain countless details that may be important in the future, yet we rarely know which will matter and which will not. This uncertainty poses a difficult challenge for adaptive decision-making, as failing to preserve relevant information can prevent us from making good choices later on. One solution to this challenge is to store detailed memories of individual experiences that can be flexibly accessed whenever their details become relevant. By allowing us to store and recall specific events in vivid detail, the human episodic memory system provides exactly this capacity. Yet, whether and how this ability supports adaptive behaviour is poorly understood. Here we aimed to determine whether people use detailed episodic memories to make decisions when future task demands are uncertain. We hypothesized that the episodic memory system's ability to store events in great detail allows us to reference any of these details if they later become relevant. We tested this hypothesis using a novel decision-making task in which participants encoded individual events with multiple features and later made decisions based on these features to maximize their earnings. Across 5 experiments (total $n = 535$), we found that participants referenced episodic memories during decisions in feature-rich environments and that they did so specifically when it was unclear at encoding which features would be needed in the future. Overall, these findings reveal a fundamental adaptive function of episodic memory, showing how its rich representational capacity enables flexible decision-making under uncertainty.

Humans possess the remarkable capacity to remember individual events from the past in high fidelity using our episodic memory system^{1,2}. From merely a single exposure, we can recall complex experiences consisting of many different details—the taste of a meal, the layout of a restaurant, the conversation around a dinner table—even when it is unclear whether we will actually ever need any of this information again. Why do we maintain such detailed memories of past events?

One possible answer is that this ability is adaptive, because the aspects of an experience that will be relevant for future behaviour are rarely apparent when they are first encountered. For instance, when deciding whether to visit a restaurant, you may need to recall details that were not previously important during your prior meals there,

such as whether the menu can accommodate your vegan friend or if the ambiance is appropriate for a work meeting. By storing detailed representations of past experiences, our episodic memory system can allow us to access any of these details if they ever become relevant in the future. In this way, episodic memory may enable flexible decision-making by allowing us to repurpose our memories for novel goals.

Despite a long history of work on the adaptive role of episodic memory^{3–6}, computational research has only recently started to characterize the advantages it can provide for decision-making^{7–9}. A key theme in this work is that the episodic memory system is poised to address several shortcomings faced by other forms of memory, particularly

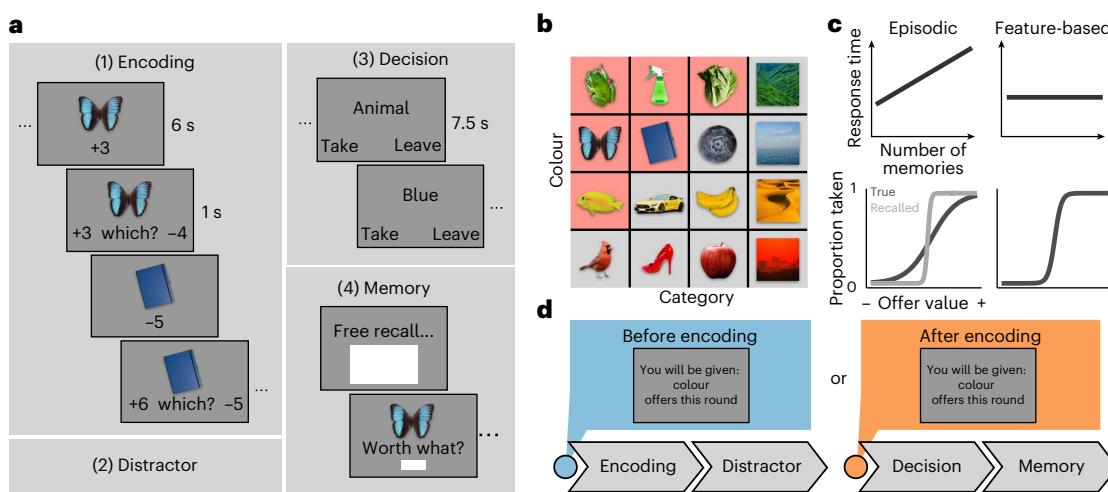


Fig. 1 | Task design. **a**, The four phases completed by participants in each round of the task used in experiments 1–4. In the first part of a round, participants completed an encoding phase that allowed them to encode individual items and an associated reward (an episode). Immediately after viewing the episode, participants completed an attention check consisting of the item alongside two options, either the associated reward that was just shown or another randomly selected reward. Participants completed six trials per round. Following encoding, participants then completed a 90-s distractor task consisting of a two-back working memory task. This was designed to prevent active rehearsal of the episodes. Following the distractor, participants completed a decision phase in which they were shown offers consisting of a single feature and were asked to either take or leave this offer. The value of each offer was the sum of all episodes described by the offer. Lastly, after the decision phase, participants completed a subsequent memory test consisting of two tasks. They were first asked to freely recall each of the items they had seen in the round. They were then shown each item and asked to provide the reward that was associated with each item. Participants completed between five and nine rounds, depending on the experiment. **b**, The full set of stimuli shown to participants in experiments 1, 3 and 4. Stimuli differed along two dimensions: colour and category. A subset of six images were sampled from these to be shown in each round (with an example

shown in red). See Supplementary Fig. 1 for the stimuli used in experiment 2. **c**, Simulated choice behaviour on the decision phase from two toy process models implementing either an episodic or feature-based decision strategy (see Methods for details). During each offer, the episodic model randomly samples individual items and their associated reward without replacement. To make choices, it then sums up the recalled rewards associated with all recalled offer-relevant items that match the offer, and then chooses to either take or leave each offer depending on whether this decision variable was positive or negative, respectively. This strategy leads to longer decisions when more memories have been retrieved (top). Its decision variable (the recalled offer value) also differs from the true offer value due to the items and associated rewards that the model actually recalls on each decision (bottom) and can be used to more strongly predict choices. By contrast, the feature-based model instead sums up the value of rewards associated with each feature at encoding time and chooses with some noise at decision time, leading to neither of these predictions. **d**, The feature uncertainty manipulation used in experiments 3 and 4. On half of the rounds, participants were informed either before or after encoding about which type of offer they would be given in the future. For experiment 3, the screen showed either colour or category offers, and for experiment 4, it showed a specific instance of either colour or category (for example, red, blue, object and so on).

incremental learning, which aids decision-making by allowing us to gradually estimate the value of different choice options over many repeated encounters^{10–16}. Importantly, incremental learning suffers from a well-known problem called the curse of dimensionality, which arises because computational and memory demands rapidly scale with the richness of what must be learned¹⁷. For instance, when learning about choice options with multiple features, an agent capable of only incremental learning would need to track and update values for each independently.

One way humans may circumvent this issue is by selectively learning about only the most relevant features of experience while ignoring information about the rest¹⁸. Using selective attention in this way can allow for efficient decisions in high-dimensional environments—when a feature is deemed relevant, its value is tracked and incrementally updated, requiring only a simple retrieval at choice time. Such a strategy is highly effective under circumstances in which feature relevance can be reliably inferred, and people probably employ it when this is the case^{19,20}. But environments like this are, in reality, rare. Furthermore, augmenting incremental learning with selective attention enforces rigidity, ultimately harming future choices that may depend on features that were initially ignored. There are also fundamental limits to the number of features that can be reasonably attended to and tracked simultaneously²¹, making this strategy increasingly impractical in the real world. These limitations are precisely the types of problems that episodic memory's detailed representations are best equipped to address.

A separate but related challenge is that incremental learning is most successful when experiences are repeated, yet actual experiences are unlikely to be encountered more than a single time. Episodic memory, by contrast, allows us to encode and retrieve individual experiences, making it naturally suited to real-world environments where experiences rarely repeat⁸. Indeed, this capacity for one-shot learning has dominated research on episodic memory's role in decision-making, where much work has demonstrated that humans can effectively guide their decisions by retrieving and leveraging individual past experiences^{22–27}. Yet, despite this progress, whether the episodic memory system's ability to encode detailed experiences provides its own advantages for decision-making remains largely unknown. Our primary goal was to address this gap.

Here we hypothesized that (1) people use episodic memory to access details from past events for decision-making in feature-rich environments, and that (2) this strategy enables flexibility when it is unclear which details will be needed in the future. To directly test these ideas, we developed a decision-making task in which participants were asked to encode individual episodes consisting of items that varied across multiple features (for example, colour and category) and an associated reward (Fig. 1a,b). After encoding these episodes, participants made value-based decisions in which they were shown offers consisting of specific features (for example, red or animal) and were then asked to either take or leave each offer. To determine the value of these offers, participants needed to sum the rewards associated with all items that had the offered feature in common (for example, all red things).

Importantly, this task could be solved using two different strategies, corresponding to either using episodic memory at choice time (an episodic strategy) or using incremental learning and selective attention at initial encoding (a feature-based strategy). To accomplish the former, participants could use episodic memory to compute an offer's value 'on-the-fly' during their decisions by retrieving and summing the rewards associated with each offer-relevant item. Alternatively, they could instead attend to individual features during encoding and precompute a sum for each. While the episodic strategy offers more flexibility because episodes can be repurposed according to present demands, it comes at the expense of greater computation at choice time. Likewise, the feature-based strategy may yield more efficient decisions because it removes the need to reference individual episodes during decision-making, but sacrifices flexibility. This is one instance of a broader trade-off between computational efficiency and flexibility seen across learning and decision-making systems in the brain^{8,11,28,29}.

We examined the extent to which participants engaged in these two strategies using several variants of this task across five different experiments (total $n = 535$), finding evidence that humans use episodic memory to flexibly access features of past experience during decision-making specifically when future task demands are unknown.

Results

Episodic memory allows access to details from past events for decision-making

In experiment 1, we first aimed to test whether people primarily rely on episodic memory rather than feature-based selective attention when making decisions in environments with multiple features. In this experiment, participants completed five rounds, each consisting of four phases (Fig. 1a,b). Participants first encoded six individual episodes, where each episode consisted of an item and its associated reward. This encoding phase was then followed by a brief two-back working memory task to prevent active rehearsal. Immediately following this task, participants then made value-based decisions about features of earlier encoded episodes. At the end of each round, participants were asked to first freely recall all items they had seen in that round and then report each item's associated reward.

We hypothesized that the computational demands of tracking multiple features simultaneously would make a feature-based strategy impractical, leading participants to instead use episodic memory to compute offer values at the time of choice. We tested this hypothesis in two independent samples to ensure the replicability of our findings. Before addressing our primary question, we first examined whether participants learned to make effective decisions in the task. At the group level, participants in both samples tended to take positive offers and reject negative offers (sample A: $M_{accuracy} = 63.6\% \pm 1.5\%$, $\beta_0 = 0.46$, 95% highest-density interval (HDI) 0.33–0.59; sample B: $M_{accuracy} = 60.8\% \pm 2.3\%$, $\beta_0 = 0.36$, 95% HDI 0.17 to 0.55; where M is the mean, and β_0 is the regression coefficient), despite substantial interindividual variability (Fig. 2a). Participants' choices therefore reflected their ability to compute the value of each offer by summing over individual experiences.

Our next goal was to examine whether participants maintained memory traces for the individual episodes they encoded in the first part of each round. To accomplish this, our experiment required participants to complete an additional memory phase immediately after the decision-making phase of each round (Fig. 1a). This memory phase consisted of two parts in which participants were asked to remember both elements of the episodes they had seen: they were first asked to freely recall each of the six items from a round, and then to recall the reward that was associated with each item. Participants had robust memory for the individual episodes, remembering between three and four items, on average (sample A: $M_{recall} = 63\% \pm 2.5\%$; sample B: $M_{recall} = 64.7\% \pm 2.8\%$; Fig. 2b), which was well above chance-level recall (sample A: $\beta_0 = 0.26$, 95% HDI 0.20 to 0.31; sample B: $\beta_0 = 0.27$, 95%

HDI 0.22 to 0.33). Item recalls exhibited classic properties of episodic memory, with items presented close together during encoding being more likely to be recalled consecutively (temporal contiguity effect³⁰; Supplementary Fig. 2), providing further evidence that participants engaged episodic memory during the task. Participants also accurately remembered the rewards associated with each item, showing a strong positive relationship between remembered and actual rewards (sample A: $\beta_{reward} = 0.52$, 95% HDI 0.44 to 0.59; sample B: $\beta_{reward} = 0.47$, 95% HDI 0.37 to 0.56; Supplementary Fig. 3). These results demonstrate that participants formed and retained strong memories for each episode beyond the decision-making phase and that individual memories were available for potential recall at choice time.

Next, to disambiguate between the episodic and feature-based strategies, we used participants' responses on the memory phase to analyse their choice behaviour. First, we reasoned that recalling an individual episode should take time^{30–32} and that, accordingly, the amount of time it takes to make a decision should scale with the number of episodes that are referenced (Fig. 1c). Importantly, the feature-based strategy makes no such prediction, as only a single item (a precomputed offer value) must be retrieved at choice time. To test this idea, we first used participants' free recall data to determine the total number of items that they accurately recalled on each round of the task. We then examined whether they took longer to respond to offers during rounds on which they recalled more items overall. As predicted, participants took longer to decide when they subsequently recalled more memories (sample A: $\beta_{nMemories} = 0.05$, 95% HDI 0.02 to 0.08; sample B: $\beta_{nMemories} = 0.07$, 95% HDI 0.03 to 0.11; Fig. 2c). We next conducted a complementary analysis examining whether decision response times were specifically related to the number of offer-relevant memories recalled—that is, only those memories whose features matched each offer (for example, only recalled red items for offers about red things). Interestingly, we found no consistent relationship between response times and the number of offer-relevant memories recalled (see Supplementary Table 1 for results across all experiments). Together, these results suggest that participants did not retrieve exclusively the memories needed for each decision. Rather, participants broadly retrieved their memories, presumably selecting from them only the trial-relevant information before making a choice.

Having observed response time patterns suggesting that participants accessed individual memories during their decisions, we moved to examine the actual choices they made. We reasoned that if their decisions were based on the rewards they remembered being associated with each item, as predicted by an episodic strategy, their choices should be sensitive to the summed value of these remembered rewards (Fig. 1c). To test this, we used participants' responses on the reward memory portion of the memory phase to determine the recalled value of each offer. Specifically, we summed over the reported remembered reward of each offer-relevant item that was also recalled during the free recall phase. By contrast, we expected that evidence for a feature-based strategy would manifest as choices being primarily driven by the true value of each offer, independent of what participants explicitly remembered. This prediction follows from the nature of feature-based learning: if participants precomputed feature values during encoding, these cached values would be immune to later forgetting or distortions of individual episodes. To arbitrate between these possibilities, we fitted two logistic regression models to participants' choices: one that used the true offer value to predict each choice, and another that used recalled offer value to predict each choice. We then compared the out-of-sample predictive accuracy of each model using cross-validation.

While participants' choices were sensitive to both true offer value (sample A: $\beta_{true} = 0.73$, 95% HDI 0.56 to 0.91; sample B: $\beta_{true} = 0.61$, 95% HDI 0.32 to 0.92) and recalled offer value (sample A: $\beta_{recalled} = 0.89$, 95% HDI 0.69 to 1.11; sample B: $\beta_{recalled} = 0.83$, 95% HDI 0.55 to 1.15), recalled offer value was slightly more effective at predicting held-out choices

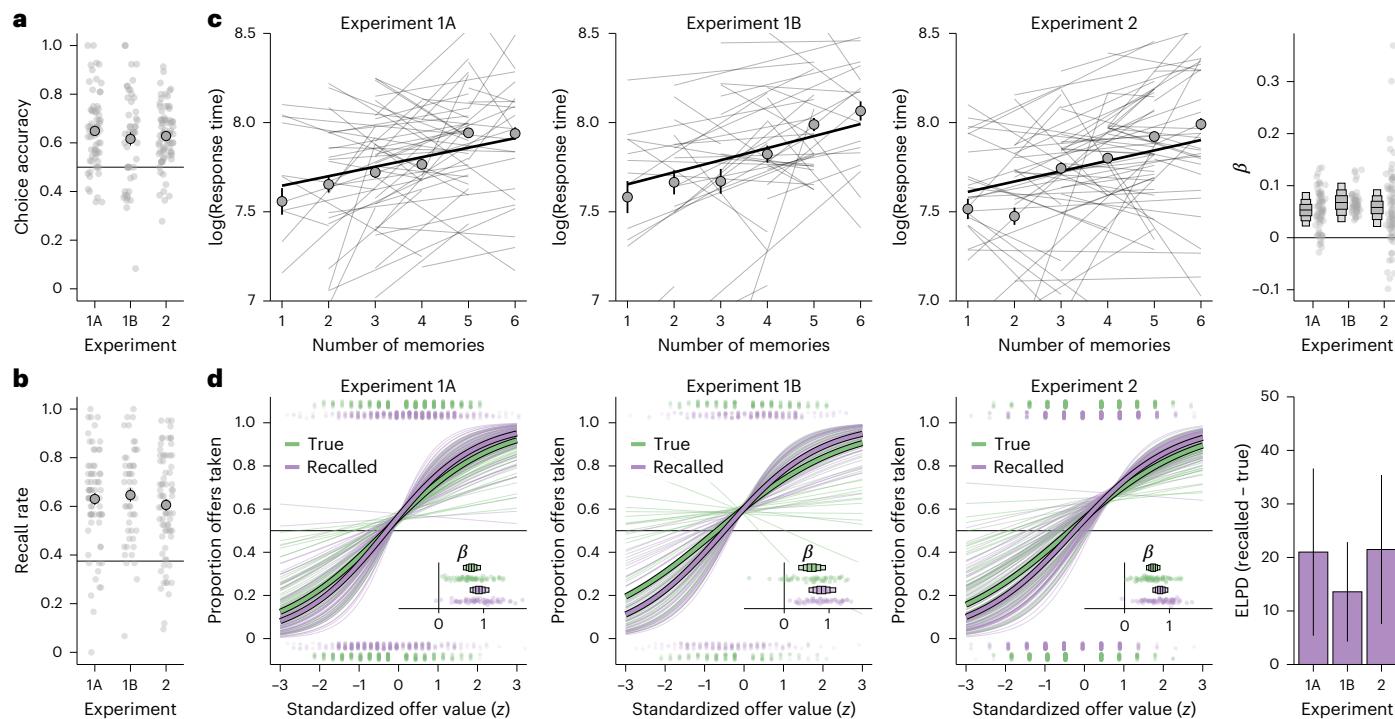


Fig. 2 | People use episodic memory to make decisions in multifeature environments.

a, Participants' overall choice accuracy during the decision phase, shown for experiment 1 (samples A and B) and experiment 2. Chance performance is represented by the horizontal line. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average choice accuracy for individual participants. **b**, Participants' rate of accurately recalling items seen in each round during the free recall portion of the memory phase. Recall rates were defined as the proportion of items that were accurately recalled relative to all items that were shown. Chance performance is represented by the horizontal line. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average recall rates for individual participants. **c**, Participants' decision response times (log transformed) as a function of the number of memories that they accurately recalled on each round. Large lines represent group-level fits of a mixed-effects linear regression model, with fits to individuals plotted as smaller lines. Points represent group-level averages with uncertainty represented as standard error of the mean. Far right: fixed-effects and random-effects slopes (β) for regression model fits. Boxes represent the fixed-effects posterior distribution,

with horizontal central lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Points represented the random-effects slopes for each participant. **d**, Left: the proportion of offers that were taken as a function of summed true offer value and recalled offer value. Offer values are z-scored to facilitate comparison. Points display raw choices, where 0 indicates a leave response and 1 indicates a take response. Inlays show the mixed effects slopes for each predictor, where boxes represent the fixed-effects posterior distribution with horizontal central lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Individual points in the inlays represent random slopes for each participant. Far right: results of comparing the true and recalled offer value models. Model fit was first assessed using tenfold cross-validation. The ELPD was then computed. Higher ELPD values indicate a higher likelihood of accurately predicting new data. Here, the difference in ELPD between models is shown such that positive values provide more support for the recalled offer value model. Uncertainty in the comparison is computed as in ref. 72, with error bars shown as the standard error of the difference. Statistics in all figures were computed with $n = 67$ (experiment 1, sample A), $n = 50$ (experiment 1, sample B) and $n = 75$ (experiment 2).

in both samples (sample A: expected log pointwise predictive density (ELPD)_{recalled-true} 21.00 ± 15.63 , sample B: ELPD_{recalled-true} 13.58 ± 9.30 ; Fig. 2d and Supplementary Table 2). This result indicates that participants tended to rely more heavily on the information contained in individual episodes, namely the identity and value of items, to make their decisions.

One limitation of this experiment is that it does not distinguish between two fundamentally different ways that information could be represented in memory. One possibility, which we have suggested so far, is that experiences are stored as integrated episodes, where multiple features are bound together into a single conjunctive representation (for example, 'butterfly')^{33,34}. Alternatively, individual experiences could be stored as separate features (for example, 'blue' and 'animal' as independent elements)³⁵. Although both approaches could support value computation in our task, the latter becomes increasingly challenging as more features must be maintained and retrieved. To test this idea, we conducted a second experiment that was nearly identical to experiment 1, but where we doubled the number of features associated with each stimulus (see Table 3 for a summary of differences between these experiments and Supplementary Fig. 1 for all stimuli used in

this experiment). Specifically, in this experiment the stimuli varied across four binary features: texture (solid or pattern), location (land or sea), animacy (animal or object) and size (small or large). Offers then consisted of one of the feature levels (for example, solid or sea) where the value of each offer was the sum of all stimuli that had the offered feature level in common (for example, all solid things).

We reasoned that, if participants relied upon memories comprised of separate features, this increased complexity would impair performance. Conversely, if participants encoded integrated episodes, the natural binding of features should preserve performance despite the additional complexity. As predicted, performance on experiment 2 was comparable to experiment 1 and all results replicated (Fig. 2 and Supplementary Table 3), supporting the conclusion that participants relied on integrated episodic memories rather than separate feature memories to complete the task. To formally assess replicability, we pooled data across experiments 1 and 2 and tested for sample-specific deviations from the overall effects. No deviations from the pooled effects were observed across any sample (Supplementary Table 4), providing strong evidence for the reliability of these effects.

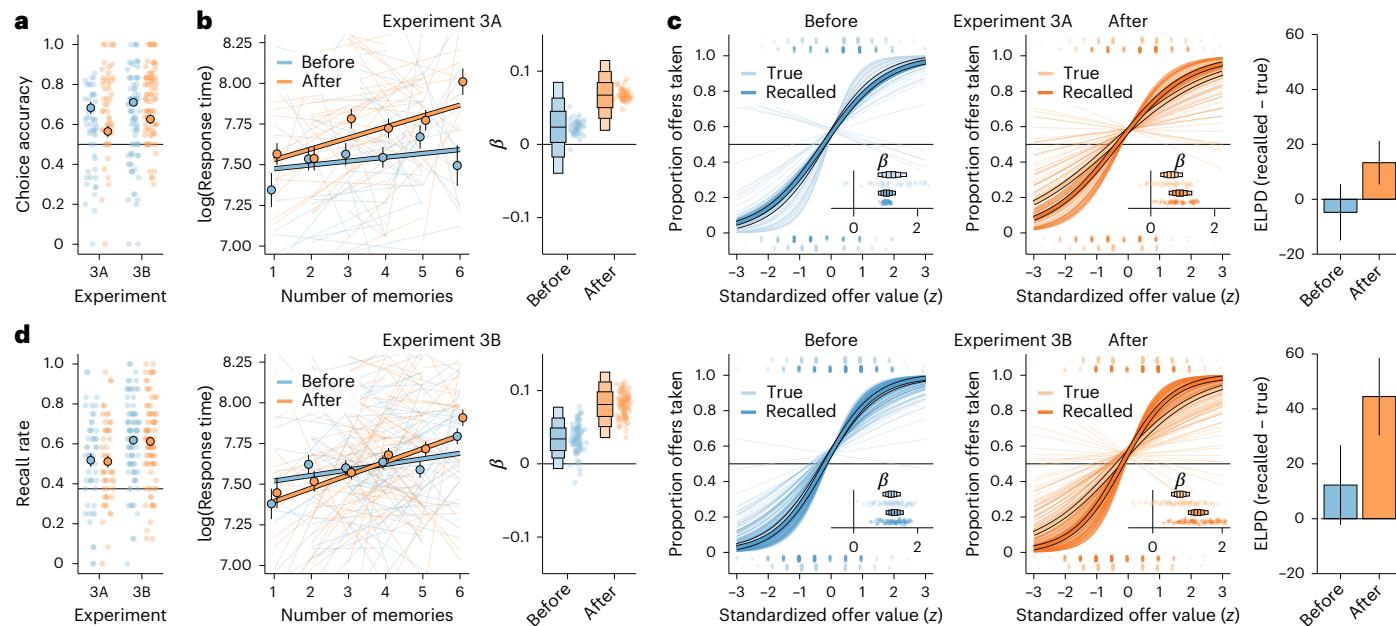


Fig. 3 | Episodic memory is primarily used for decisions when it is unclear what features are important. **a**, Participants' choice accuracy during the decision phase, shown for experiment 3 (samples A and B). Decisions made when the future relevance of features was known at encoding (the before condition, in blue) are shown separately from those made when this information was unknown at encoding (the after condition, in orange). Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average choice accuracy for individual participants. **b**, Participants' decision response times (log transformed) as a function of the number of memories they accurately recalled on each round for sample A (top) and sample B (bottom). Left: large lines represent the group-level fit of a mixed-effects regression, with fits to individuals plotted as smaller lines. Points represent group-level averages with standard error of the mean. Right: fixed-effects slopes and random-effects slopes (β) for regression model fits. Boxes represent the fixed-effects posterior distribution, with horizontal lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Points represent random-effects slopes for each participant. **c**, The relationship between choices and offer value for sample A (top) and sample B (bottom). Left: the proportion of offers taken as a function of summed true offer value and recalled offer value for the before and after conditions. Offer values are z-scored to facilitate comparison. Points display raw choices, where 0 indicates a leave response and 1 indicates a take response. Inlays show the mixed-effects slopes for each predictor, where boxes represent the fixed-effects posterior distribution, with horizontal central lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Individual points in the inlays represent random slopes for each participant. Right: model comparison showing the difference in ELPD, with error bars representing the standard error of this difference. **d**, Participants' rate of accurately recalling items seen in each round during the free recall portion of the memory phase. Chance performance is represented by the horizontal line. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average recall rates for individual participants. Statistics in all figures were computed with $n = 50$ (experiment 3, sample A) and $n = 84$ (experiment 3, sample B).

between choices and offer value for sample A (top) and sample B (bottom). Left: the proportion of offers taken as a function of summed true offer value and recalled offer value for the before and after conditions. Offer values are z-scored to facilitate comparison. Points display raw choices, where 0 indicates a leave response and 1 indicates a take response. Inlays show the mixed-effects slopes for each predictor, where boxes represent the fixed-effects posterior distribution, with horizontal central lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Individual points in the inlays represent random slopes for each participant. Right: model comparison showing the difference in ELPD, with error bars representing the standard error of this difference. **d**, Participants' rate of accurately recalling items seen in each round during the free recall portion of the memory phase. Chance performance is represented by the horizontal line. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average recall rates for individual participants. Statistics in all figures were computed with $n = 50$ (experiment 3, sample A) and $n = 84$ (experiment 3, sample B).

Together, these results demonstrate that, when faced with choices that could be based on many different features, participants primarily used an episodic strategy that involved retrieving and computing over individual memories at choice time. This finding suggests that, despite the increased computational demands during decision-making, people prefer to maintain detailed episodic memories that can be flexibly accessed rather than attempting to track and update values for multiple features simultaneously. This preference may reflect the difficulty of implementing a feature-based strategy when faced with the curse of dimensionality, as suggested by previous theoretical work⁸.

Episodic memory enables flexible decision-making when it is unclear which details are important

Our findings so far demonstrate that people rely on episodic memory to make decisions in multifeature environments, in part because episodic memories provide a natural solution to the curse of dimensionality. However, this observation alone does not fully explain why episodic memory might be specifically advantageous for decision-making. To address this, we next hypothesized that episodic memory's key benefit lies in its ability to enable flexible decisions when future task demands are uncertain—a common situation in the real world. This hypothesis predicts that people should shift away from using episodic memory when they can anticipate which features will be relevant for their upcoming decisions, as this foreknowledge would make a feature-based strategy more viable. To test this prediction directly, we next conducted

an experiment where we manipulated whether participants knew in advance which features would be relevant for their decisions.

In experiment 3, we contrasted two conditions: one where participants learned which features would be relevant only at choice time (after encoding), and another where participants knew before encoding which features would be needed for their upcoming decisions (Fig. 1c). Specifically, in this new before condition, participants were told before encoding that they would later receive offers about colour or category, but not both. This manipulation created conditions where the feature-based strategy is more feasible, as participants could safely ignore irrelevant features (for example, category when only colour offers would be made). This advantage was absent in the after condition, where feature relevance remained uncertain during encoding. We predicted that people would primarily rely on an episodic strategy in the after condition, but would shift towards a feature-based strategy when feature relevance was known in advance. We again tested this prediction across two independent samples to ensure the replicability of our findings.

Participants responded accurately, primarily taking positive and leaving negative offers both during rounds in which feature relevance was communicated before encoding (sample A: $M = 68.5\% \pm 2.9\%$, $\beta_0 = 0.86$, 95% HDI 0.62 to 1.33; sample B: $M = 71.7\% \pm 1.9\%$, $\beta_0 = 1.00$, 95% HDI 0.81 to 1.21) and after encoding (sample A: $M = 59.7\% \pm 2.9\%$, $\beta_0 = 0.43$, 95% HDI 0.20 to 0.66; sample B: $M = 63.3\% \pm 2.2\%$, $\beta_0 = 0.62$, 95% HDI 0.41 to 0.83; Fig. 3a). Notably, participants chose more accurately in the before condition (by an average of 8.8% and 8.4% in each

sample, respectively), suggesting that there were clear benefits to performance when feature relevance was known during encoding. Importantly, however, performance in the after condition was on par with our prior experiments.

We next aimed to test our primary hypothesis. As predicted, in the after condition, participants took longer to make decisions when they recalled more memories (sample A: $\beta_{n\text{Memories}} = 0.08$, 95% HDI 0.04 to 0.12; sample B: $\beta_{n\text{Memories}} = 0.06$, 95% HDI 0.02 to 0.11; Fig. 3b). Yet, we found little evidence for this relationship in the before condition, when participants were aware of which features they would need for their future decisions (sample A: $\beta_{n\text{Memories}} = 0.02$, 95% HDI -0.04 to 0.07; Sample B: $\beta_{n\text{Memories}} = 0.03$, 95% HDI -0.01 to 0.07). We then further examined the extent to which true and recalled offer value predicted participants' choices in both conditions. We expected recalled offer value to be a better predictor of choices in the after condition, but not in the before condition, if participants retrieved episodes during decisions made in the former but not the latter. In line with this prediction, recalled offer value better predicted held-out choices in the after condition but not in the before condition in both samples (Fig. 3c and Table 1; see Supplementary Table 2 for the performance of all models). In models pooling across samples, we further assessed the interaction between these conditions, finding evidence for a greater effect of recalled memories on decision times in the after condition relative to the before condition ($\beta_{n\text{Memories} \times \text{Condition}} = 0.07$, 95% HDI 0.04 to 0.10; Supplementary Table 5). Consistent with this finding, recalled offer value also provided comparably better predictions of held-out choices ($\beta_{\text{ELPD}_{\text{recalled-true}} \times \text{Condition}} = 0.06$, 95% HDI 0.03 to 0.10). These results indicate that participants primarily referenced individual episodes during decision-making when it was unclear at encoding which features would be needed for future decisions.

To better understand the source of these differences in strategy recruitment, we next examined participants' performance on the memory phase. Specifically, while participants shifted away from using episodic memory when feature relevance was known before decision-making, it is unclear whether this effect emerged from changes in how experiences were initially encoded or from whether they were later accessed during decision-making. We reasoned that, if participants strategically modified their encoding based on known feature relevance, they should show impaired subsequent memory for individual episodes in the before condition, as they would be focused primarily on computing and maintaining feature-level values rather than encoding complete episodic memories. However, another possibility is that participants instead continued to encode episodic memories alongside precomputing feature values, consistent with evidence that these distinct systems often operate in parallel during learning^{22,25,36-38}. In this case, we expected to see comparable memory performance between conditions, with task demands influencing whether these memories were later accessed at choice time rather than whether they were initially stored.

Across both samples, we found that participants maintained strong memories of the episodes encountered in each round regardless of condition. Participants accurately recalled individual items with no differences in recall rates between conditions (sample A: $\beta_{\text{after-before}} = -0.01$, 95% HDI -0.10 to 0.08; sample B: $\beta_{\text{after-before}} = -0.01$, 95% HDI -0.07 to 0.06; Fig. 3d and Table 2), and they further showed equally similar memory for the rewards associated with each item (sample A: $\beta_{\text{after-before}} = -0.06$, 95% HDI -0.18 to 0.06; sample B: $\beta_{\text{after-before}} = -0.05$, 95% HDI -0.15 to 0.05; Table 2 and Supplementary Fig. 3). These findings demonstrate that participants encoded complete episodic memories regardless of whether they knew which features would be relevant for future decisions.

Together, these results provide evidence that people selectively use episodic memory for decision-making when feature relevance is unknown during encoding. When participants were aware of which features would be important for decision-making before encoding, our results suggested that they no longer retrieved individual memories

Table 1 | Experiment 3 model comparison results showing fixed-effects estimates (β) with 95% HDIs and difference in ELPD between true and recalled offer value models

Condition	Sample	Model	β (95% HDI)	ELPD _{recalled-true} ± standard error
Before	A	True	1.17 (0.76–1.65)	-4.74 ± 10.22
		Recalled	1.02 (0.77–1.31)	
	B	True	1.17 (0.92–1.46)	12.26 ± 14.50
		Recalled	1.27 (1.02–1.55)	
After	A	True	0.63 (0.31–0.99)	13.32 ± 7.96
		Recalled	0.91 (0.58–1.28)	
	B	True	0.79 (0.54–1.06)	44.50 ± 14.13
		Recalled	1.29 (1.03–1.59)	

at choice time. This shift towards the feature-based strategy is sensible—it reduces computational demands at decision time by allowing direct access to precomputed feature values rather than requiring the retrieval and integration of multiple episodic memories. This approach led to improved performance on expected decisions because the episodic strategy can introduce noise (for example, through errors in memory retrieval), especially under time constraints when only a subset of memories might be accessible. Yet, we also found that participants maintained detailed episodic memories in both conditions, suggesting that the differences we observed during decision-making emerged from how information was accessed at choice time rather than how it was initially encoded. We next aimed to determine whether this parallel maintenance of episodic memories provided its own adaptive benefits for decision-making.

Episodic memory maintains access to details if they become unexpectedly relevant

The results of experiment 3 suggest that, while participants appeared to rely on a feature-based strategy when feature relevance was known in advance, they still maintained detailed episodic memories. One prediction that follows from this observation is that participants should still be able to use their episodic memories to maintain access to features they had initially deemed irrelevant in order to inform their decisions. By contrast, if participants abandoned episodic encoding entirely, they should struggle to access information about these irrelevant features.

We designed a fourth experiment to test this idea. In experiment 4, after participants completed eight rounds in which the manipulation introduced in experiment 3 was implemented, they were asked to complete an additional final round in which all possible offers were presented. Importantly, participants were shown these unexpected offers regardless of whether they had been told before the encoding phase of this final round that they would see only a subset of offers later on. This allowed us to examine performance when the final round was completed under the before condition for both expected and unexpected offers. By definition, all offers shown during the after condition were unexpected. We additionally modified the task to create conditions that would encourage greater reliance on the feature-based strategy. To accomplish this, we simplified the decision phase by presenting only a single offer type per round (for example, 'red') before the final round. This modification meant that participants in the before condition needed to track only one specific feature value rather than

Table 2 | Memory performance across experiments 3 and 4

Experiment	Sample	Condition	Recall rate ± standard error	β_0 (95% HDI)	β_{reward} (95% HDI)
3	A	Before	51.8% ± 3.2%	0.14 (0.08–0.21)	0.55 (0.48–0.62)
		After	51.1% ± 2.9%	0.14 (0.08–0.19)	0.49 (0.39–0.57)
	B	Before	61.8% ± 2.2%	0.24 (0.20–0.29)	0.53 (0.46–0.60)
		After	61.2% ± 2.3%	0.24 (0.19–0.28)	0.48 (0.41–0.55)
4	–	Before	62.5% ± 1.9%	0.25 (0.21–0.29)	0.54 (0.49–0.59)
		After	61.9% ± 1.9%	0.23 (0.19–0.27)	0.48 (0.42–0.54)

Recall rate shows the mean percentage of items recalled (± standard error). β_0 represents the intercept of a model predicting recall performance relative to chance. β_{reward} represents the slope of a model predicting the relationship between remembered and actual rewards.

multiple instances of the same feature (for example, just ‘red’ instead of all colours), making the feature-based strategy exceptionally efficient. We predicted that this change would lead participants to rely more heavily on the feature-based strategy in the before condition relative to experiment 3.

We further predicted that participants should show a specific pattern of choice behaviour if they maintained access to their episodic memories despite this increased commitment to the feature-based strategy. First, in response to unexpected final round offers completed under the before condition, we predicted that participants’ accuracy should drop relative to their performance on previous rounds. This is because they would no longer be able to rely on the feature-based strategy, which provided clear benefits to performance (improving accuracy by 8–9% in experiment 3 when feature relevance was known in advance). Second, we predicted that, to make these unexpected decisions, participants would instead need to fall back on using their episodic memories, leading to accuracy levels similar to their performance in the after condition.

We first examined overall performance on rounds before the final round, finding that, while participants responded accurately in both conditions (before: $M = 78.4\% \pm 2.0\%$, $\beta_0 = 1.57$, 95% HDI 1.29 to 1.90; after: $M = 55.2\% \pm 2.2\%$, $\beta_0 = 0.52$, 95% HDI 0.32 to 0.74; Fig. 4a), they performed substantially better in the before condition relative to both samples in experiment three. Participants further showed equivalent memory performance across both conditions, with no differences in recall rates ($\beta_{\text{after-before}} = -0.02$, 95% HDI -0.07 to 0.04; Fig. 4b and Table 2) or in memory for associated rewards ($\beta_{\text{after-before}} = -0.06$, 95% HDI -0.14 to 0.02; Table 2 and Supplementary Fig. 3). These results suggest that, although participants were more capable at deciding in the before condition when they had to compute only the value of a single feature, they still separately maintained episodic memories.

We next asked whether this increase in performance in the before condition was due to participants’ greater reliance on the feature-based strategy. In line with this interpretation, we found that during the before condition there was virtually no evidence of a positive relationship between the number of memories recalled and decision response times ($\beta_{\text{nMemories}} = -0.02$, 95% HDI -0.07 to 0.02; Fig. 4c), and participants’ choices were better predicted by the true offer value ($\beta_{\text{true}} = 2.74$, 95% HDI 2.03 to 3.65) than the recalled offer value ($\beta_{\text{recalled}} = 2.04$, 95% HDI 1.49 to 2.75, $\text{ELPD}_{\text{recalled-true}} = -26.54 \pm 12.69$; Fig. 4d and Supplementary Table 2). Formal comparisons confirmed that, relative to experiment 3, in experiment 4 there was a weaker relationship between the number of recalled memories and response times in the before condition ($\beta_{\Delta} = -0.05$, 95% HDI -0.11 to -0.006), and the true

offer value also exerted a stronger effect on choice ($\beta_{\Delta} = 0.68$, 95% HDI 0.38 to 0.99). By contrast, in the after condition, participants showed longer decision response times when more memories were recalled ($\beta_{\text{nMemories}} = 0.09$, 95% HDI 0.05 to 0.12). While participants’ choices in this condition were numerically more sensitive to the recalled offer value ($\beta_{\text{recalled}} = 0.90$, 95% HDI 0.59 to 1.27) than the true offer value ($\beta_{\text{true}} = 0.79$, 95% HDI 0.52 to 1.10), both models were equally capable of predicting held-out choices ($\text{ELPD}_{\text{recalled-true}} = 1.33 \pm 6.97$; Fig. 4d). This equivalent model performance probably reflects the limited data available for cross-validation, as participants made substantially fewer decisions per condition in this experiment compared with experiment 3.

Finally, we turned to test the primary question of this experiment: whether in the final round participants maintained access to information about the features they were told would be irrelevant for future decisions. As predicted, participants who completed their final round in the before condition showed impaired performance on unexpected offers compared with their performance on prior rounds in this condition ($M = 57.3\% \pm 3.7\%$, $\beta_0 = 0.14$, 95% HDI 0.04 to 0.24; Fig. 4a), but maintained high accuracy on expected offers ($M = 78.4\% \pm 5.8\%$, $\beta_0 = -0.07$, 95% HDI -0.20 to 0.06). Critically, participants’ performance on unexpected offers matched that of the prior rounds in the after condition ($\beta_0 = -0.02$, 95% HDI -0.11 to 0.07), suggesting they could successfully fall back on episodic memories when the feature-based strategy was insufficient. Surprisingly, we also found that participants who completed their final round in the after condition exceeded their prior performance in this condition ($M = 66.7\% \pm 3.0$, $\beta_0 = -0.13$, 95% HDI -0.21 to -0.05). This improvement may stem from the increased variety of offer types shown in the final round, which provided participants with more opportunities to make decisions about items they successfully remembered, whereas previous rounds with fewer offer types may have tested only items for which their memory was weaker.

The results of experiment 4 suggest that, even under conditions that strongly encouraged reliance on a feature-based strategy, participants maintained detailed episodic memories that they could access when needed. While participants showed clear evidence of using feature-based computations in the before condition, they were still able to rely upon their episodic memories when faced with unexpected offers, performing comparably to the after condition. This pattern of results provides evidence that episodic memory may serve as a ‘backup’ to aid decisions when initially irrelevant features become unexpectedly relevant.

Episodic memory recall becomes more targeted under realistic decision demands

An intriguing aspect of our results so far is that, when participants used episodic memory for decision-making, their decision times scaled with the total number of memories they recalled in each round rather than just those relevant to each offer, suggesting that retrieval during decision-making was not preferentially biased towards relevant memories. Such a broad and unfocused memory search seems poorly adapted to real-world contexts where the space of possible memories is vast, raising the possibility that our relatively small memory sets may not have provided sufficient pressure for participants to develop efficient filtering strategies. In addition, our task’s simplified feature space, while allowing us to characterize episodic memory use in a controlled setting, remains far simpler than real-world experience, which can consist of a vast and nearly unlimited number of features. Although participants’ choice behaviour was largely unaltered by the addition of new features in experiment 2, this setting was still far from naturalistic, and it remains possible that further increases to feature dimensionality could reveal more subtle differences. It is also important to note that structuring our experiments as a series of rounds—where participants knew they would be tested on their memory after each decision phase—may have encouraged them to maintain episodic memories to improve memory test performance rather than for their utility for decision-making.

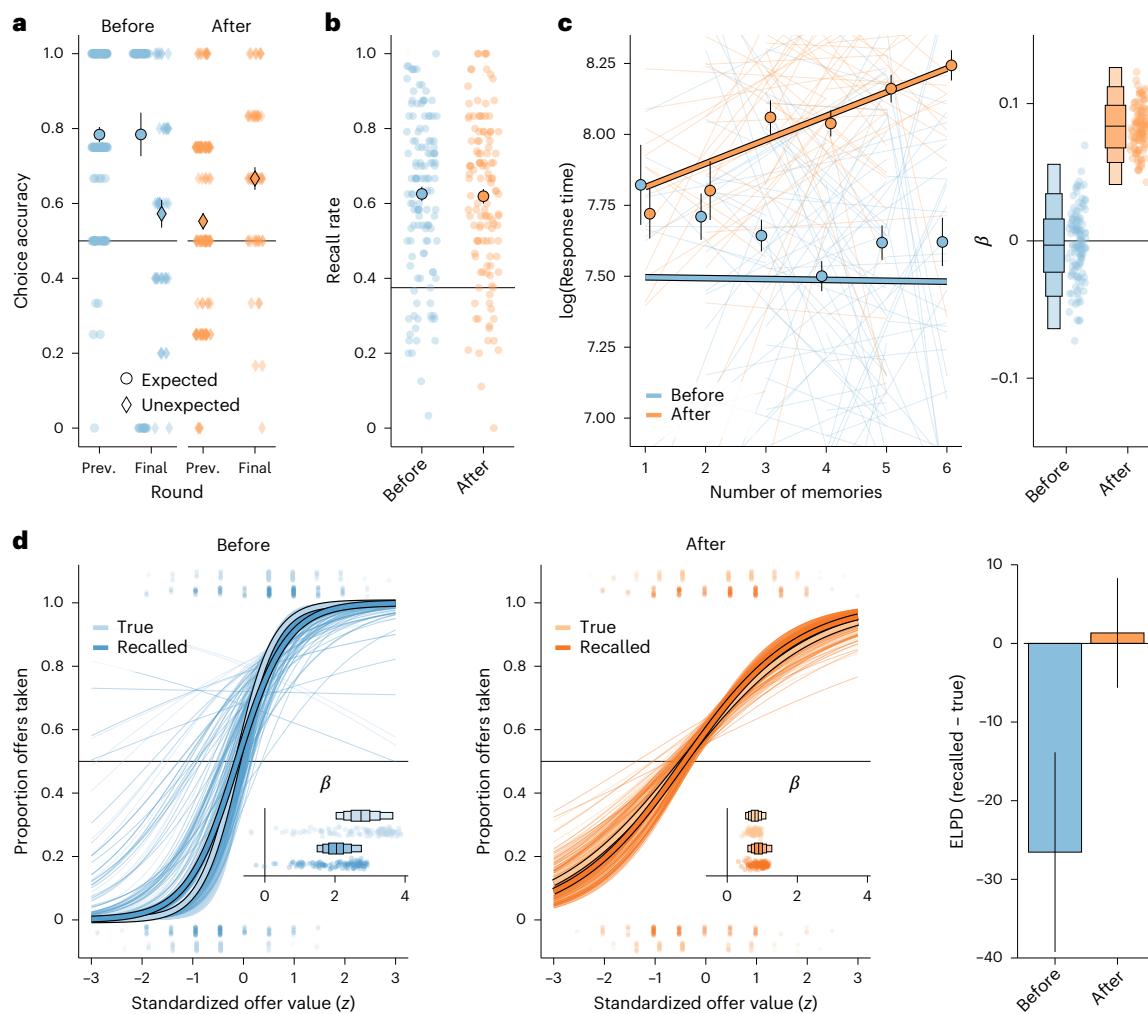


Fig. 4 | Episodic memory is used to make choices about unexpected offers. **a**, Experiment 4 choice accuracy, separated by all rounds before the final round (labelled as previous (Prev.) rounds) and the final round in which participants were asked to make decisions about both offers they expected (in the before condition) as well as those that were unexpected (in both conditions with $n = 65$ in the before condition and $n = 60$ in the after condition). Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average choice accuracy for individual participants. By design, only a single offer was expected in the before condition. **b**, Participants' rate of accurately recalling items seen in each round of experiment 4 during the free recall portion of the memory phase. Chance performance is represented by the horizontal line. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average recall rates for individual participants. **c**, Experiment 4 participants' decision response times (log transformed) as a function of the number of memories they accurately recalled on each round. Left: large lines represent the group-level fit of a mixed-

effects regression, with fits to individuals plotted as smaller lines. Points represent group-level averages with standard error of the mean. Right: fixed-effects slopes and random-effects slopes for regression model fits. Boxes represent the fixed-effects posterior distribution, with horizontal lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Points represent random-effects slopes for each participant. **d**, Left: the proportion of offers taken as a function of summed true offer value and recalled offer value for the before and after conditions. Offer values are z-scored to facilitate comparison. Points display raw choices, where 0 indicates a leave response and 1 indicates a take response. Insets show the mixed-effects slopes for each predictor, where boxes represent the fixed-effects posterior distribution with horizontal central lines representing the mean and boxes representing 50%, 80% and 95% HDIs. Individual points in the inlays represent random slopes for each participant. Right: model comparison showing the difference in ELPD, with error bars representing the standard error of this difference. Unless otherwise noted, statistics in all figures were computed with $n = 125$.

We designed a fifth experiment to investigate (1) whether our findings would translate to a more naturalistic setting in which these constraints were relaxed, and (2) whether this would encourage participants to develop more selective retrieval strategies that improve decision-making. To address these questions, we used a larger memory set, naturalistic stimuli and a continuous format with a surprise memory test. This task followed the structure of a single extended round in our first experiment, where participants first encoded 21 episodes that each consisted of a naturalistic image associated with a binary (+1 or -1) reward (Fig. 5a). Following a distractor phase identical to that used in our previous experiments, participants were then given 30 offers about various features of the episodes (for example vehicles,

birds and things that fly; Supplementary Fig. 4). Finally, participants completed a surprise memory test probing both free recall of the items and memory for the reward associated with each.

We predicted that participants would demonstrate the signatures of episodic memory retrieval during decision-making that we established in our prior experiments. Given the larger memory pool, we further expected that the number of offer-relevant memories recalled by participants would emerge as a stronger predictor of response times than in previous experiments, reflecting more targeted retrieval processes.

Participants performed well on this task, achieving $70.6\% \pm 1.0\%$ accuracy, which was well above chance ($\beta_0 = 1.02$, 95% HDI 0.86 to 1.18;

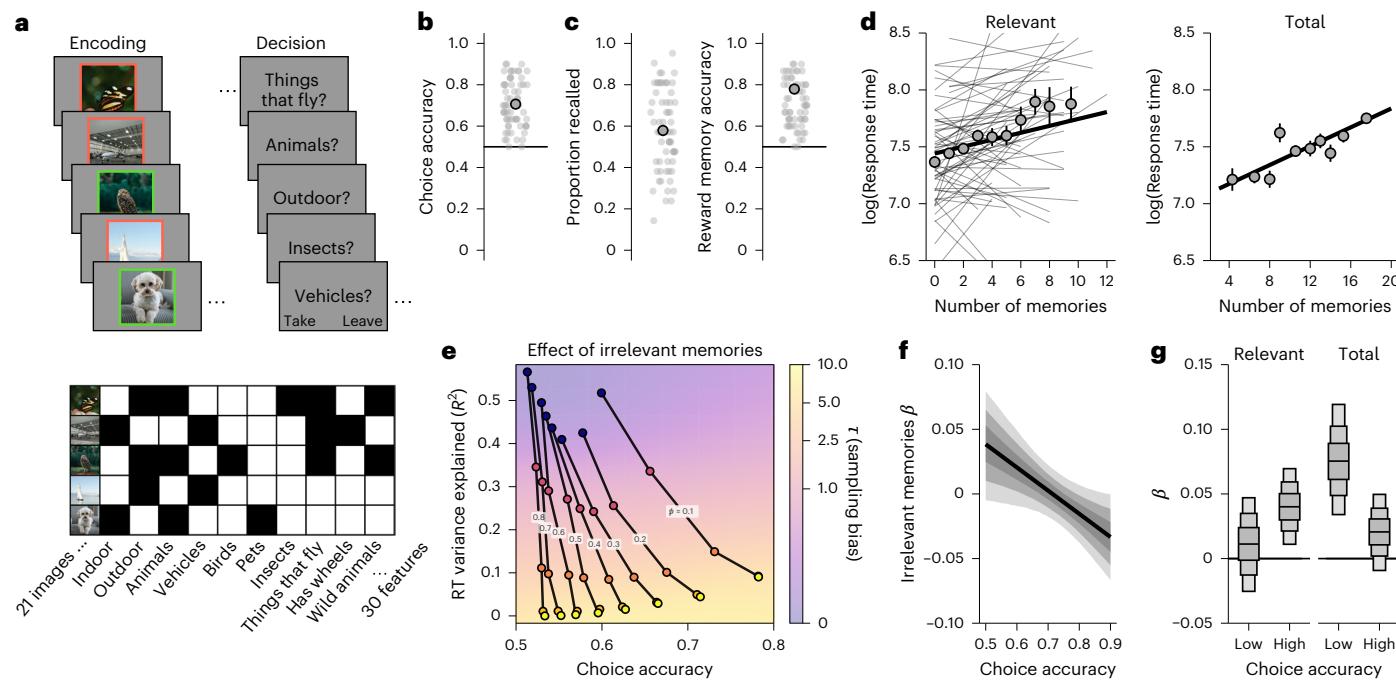


Fig. 5 | Realistic decision demands encourage participants to develop more efficient retrieval strategies. **a**, Design of experiment 5. Top: participants completed a single, long round of each of the four phases used in prior experiments. During encoding, participants were shown 21 naturalistic images with a binary reward (+1 or -1) represented by a green or red border. After the distractor phase, they were asked to make decisions about 30 offers about these episodes, which were then followed by a two-part memory test for free recall and the reward associated with each item. Bottom: example features used for the offers in experiment 5 with the images that matched each feature denoted in black. Each offer had between 1 and 12 matching images. See Supplementary Fig. 4 for all images and offers provided to participants. **b**, Experiment 5 choice accuracy. Large points represent group-level averages, with error bars representing the standard error of the mean. Small points represent average choice accuracy for individual participants. **c**, Recall performance in experiment 5. Left: participants' proportion of the 21 possible items that were accurately recalled during the free recall phase. Right: participants' accuracy on the two-alternative forced choice reward recall test. Large points represent group-level averages with error bars representing standard error of the mean. Small points represent average recall rates for individual participants. **d**, Experiment 5 participants' decision response times (log transformed) as a function of the number of offer-relevant memories (left) and overall total memories (right) that they recalled. Large lines represented the group-level fit of regression models, with fits to individual participants plotted as smaller lines. Points represent

group-level averages across ten bins with equal trials in each and standard error of the mean. Note that, because only a single memory test was completed by each participant, the total number of memories was a between-participants measure in this experiment. **e**, Results from the behaviour of simulated agents completing experiment 5. On each offer, agents sample individual items and their associated reward without replacement, controlled by two parameters. The sampling bias (τ) determines the randomness of this sampling such that larger values bias sampling towards offer-relevant items, and a bias of zero is fully random. The probability of stopping (ϕ) sets a constant probability that recall will halt after each sample and controls the overall number of memories that are sampled within the task time limit. Agents with lower τ and higher ϕ are less accurate, and the number of offer-irrelevant memories during each decision explains a larger proportion of variance in response times when recall is more random. See Methods for more details. **f**, The interaction effect of the number of offer-irrelevant memories and choice accuracy on decision response times. Participants who are less accurate overall tend to have response times that are more strongly related to their number of irrelevant memories. The dark line represents the marginal posterior mean of this effect, with 50%, 80% and 95% HDIs. **g**, Fixed-effects slopes for regression models predicting the response times of median-split low ($n = 30$) and high ($n = 33$) accuracy groups from their number of offer-relevant (left) and overall total (right) memories. The horizontal line within each box represents the posterior mean and boxes represent 50%, 80% and 95% HDIs. Unless otherwise noted, statistics in all figures were computed with $n = 63$.

Fig. 5b). On average, participants recalled 12 ± 0.5 items and demonstrated excellent memory for the rewards associated with each item ($M_{\text{accuracy}} = 77.9\% \pm 1.4\%$; $\beta_0 = 0.78$, 95% HDI 0.74 to 0.82; Fig. 5c). These results show that participants successfully adapted to the increased complexity of this task, maintaining both effective decision-making and memory recall.

Next, we examined participants' response times and choice behaviour to test our primary hypotheses. As predicted, participants took longer to make decisions when they recalled more memories that were relevant to each offer ($\beta_{\text{nMemories}} = 0.03$, 95% HDI 0.01 to 0.05; Fig. 5d), suggesting that their retrievals were guided by a more selective search through memory. Turning to choice behaviour, we found that participants' decisions were sensitive to both true offer value ($\beta_{\text{true}} = 1.03$, 95% HDI 0.84 to 1.24) and recalled offer value ($\beta_{\text{recalled}} = 1.01$, 95% HDI 0.81 to 1.21), consistent with our prior experiments. However, recalled offer value provided a negligible out-of-sample predictive advantage over true offer value (ELPD_{recalled-true} 4.25 ± 2.66 ; Supplementary Fig. 4).

probably because participants' highly accurate memory for rewards left little room for distortions in value to influence decision-making.

Interestingly, we also found that participants took longer to make choices when they recalled more memories overall ($\beta_{\text{nMemories}} = 0.04$, 95% HDI 0.02 to 0.06; Fig. 5d), similar to our previous experiments. The persistence of this effect reflects, at least in part, the fact that participants who recalled more offer-relevant memories also tended to recall more memories in general ($r = 0.41$; $\beta_{\text{nMemories}} = 0.74$, 95% HDI 0.65 to 0.82). However, a related possibility is that individual differences in memory search strategies may drive these relationships, with some participants engaging in broad, less selective retrieval that is more influenced by offer-irrelevant memories, while others use a more targeted search process focused primarily on retrieving offer-relevant memories.

Our next aim was to investigate this idea. We predicted that task performance should be related to the selectivity of retrieval during choice. We first explored this prediction computationally by simulating agents that varied in how selectively they sampled memories

during decision-making. Agents that sampled more broadly and terminated search earlier showed both poorer task performance and response times that were more strongly driven by memories irrelevant to the current offer (Fig. 5e). Consistent with this result, we found that lower-performing participants' response times were more strongly related to the number of offer-irrelevant memories they recalled ($\beta_{n\text{Memories} \times \text{Accuracy}} = -0.02$, 95% HDI -0.04 to -0.002 ; Fig. 5f). This finding suggests that participants who performed worse on the task engaged in less selective memory retrieval, as their decision times were disproportionately influenced by memories that were not useful for the current choice. To examine these individual differences more closely, we further split participants into low- and high-performing groups and directly compared whether total memories or offer-relevant memories better predicted response times in each group. Low-performing participants took longer to make decisions when they had more memories overall ($\beta_{n\text{Memories}} = 0.08$, 95% HDI 0.03 to 0.12), but their response times were not reliably related to the quantity of their offer-relevant memories ($\beta_{n\text{Memories}} = 0.01$, 95% HDI -0.03 to 0.05). As expected, high-performing participants showed the opposite pattern: they took longer to respond when they had more offer-relevant memories ($\beta_{n\text{Memories}} = 0.04$, 95% HDI 0.01 to 0.07) but their response times showed no reliable relationship with their overall pool of memories ($\beta_{n\text{Memories}} = 0.02$, 95% HDI -0.01 to 0.05). Indeed, offer-relevant memories better predicted response times out-of-sample in high-performers ($\text{ELPD}_{\text{relevant-total}} = 7.55 \pm 3.34$), confirming that these participants likely engaged in more targeted memory search during decision-making, but both memory measures were equally predictive for low-performers ($\text{ELPD}_{\text{relevant-total}} = -0.08 \pm 2.7$). These findings confirm that more selective memory retrieval is associated with better decision performance, supporting the adaptive value of targeted memory search during decision-making.

Together, the results of experiment 5 demonstrate that, when memory demands more closely approximate real-world complexity—through a larger memory pool, naturalistic stimuli and a surprise memory test—participants continued to rely on episodic memory for flexible decision-making. Yet, they differed in how they accessed their memories, and the ability to engage in targeted memory retrieval rather than broadly searching through all available memories distinguished better from worse decision-makers. These results provide evidence that, as memory demands approach real-world complexity, efficient memory search becomes a crucial determinant of decision quality.

Controlling for individual differences in effort

An alternative explanation for the relationship between memory retrieval and response times that we have uncovered is that individual differences in task engagement or effort may act as a confounding third variable. Under this account, more careful responding could lead to improved performance on both the decision and memory phases of our tasks, creating a spurious association between these measures rather than reflecting genuine episodic memory use during decision-making.

We conducted a series of complementary analyses across each of our experiments to rule out this possibility. First, we predicted that, if individual differences in effort were driving these effects, we should observe a classic speed–accuracy trade-off during decision-making. Contrary to this prediction, we found little evidence for a consistent relationship between response speed and accuracy across experiments (Supplementary Table 6). We next used performance during the distractor phase as a proxy for effort. We reasoned that, if task engagement were the common cause underlying both slower response times and better memory performance, participants with superior distractor task performance should exhibit both patterns. To test for this possibility, we included distractor performance as a covariate in our analyses. This did not meaningfully alter the strength of the observed relationship between recalled memories and response times in any experiment (Supplementary Table 7). Lastly, we decomposed each participant's number of recalled memories per round into within- and

between-participant components to isolate effects occurring within individuals from those driven by individual differences (note that this was not possible in experiment 5, which did not use a round-based structure). Within-participant effects remained largely consistent with our original analyses, with minor differences in the strength of effects in experiments 1B ($\beta_{n\text{Memories}} = 0.05$, 95% HDI -0.002 to 0.10 , 90% HDI 0.005 to 0.09) and 4 (after condition: $\beta_{n\text{Memories}} = 0.05$, 95% HDI -0.005 to 0.09 , 90% HDI 0.003 to 0.09 ; Supplementary Table 8). Furthermore, when data were pooled across comparable experiments to assess replicability, the pooled mean within-participant effect remained robust, with no substantial deviations across experiments (Supplementary Table 9).

Together, these analyses provide strong converging evidence against an effort-based explanation for our findings, supporting the conclusion that participants used episodic memory during their decisions rather than exhibiting patterns driven by individual differences in task engagement.

Discussion

Our findings indicate that people flexibly use episodic memory to guide their choices in multifeature environments, particularly when future task demands are uncertain. When faced with multiple decision-relevant features (experiment 1), participants relied primarily on episodic memories to compute offer values during decision-making, as indicated by both their response times and choice patterns. This strategy persisted even as feature complexity increased, suggesting that participants stored experiences as integrated episodes rather than separate feature representations (experiment 2). When given advance knowledge of feature relevance (experiment 3), participants shifted towards a more computationally efficient feature-based strategy that involved precomputing values during encoding. Yet, despite this shift in strategy, participants continued to encode episodic memories, a parallel operation that proved useful when knowledge about previously irrelevant information was needed for decision-making (experiment 4). Finally, in an environment designed to induce pressure for participants to more efficiently recall their memories during decision-making, those who retrieved more selectively made better decisions (experiment 5). Overall, these results demonstrate that episodic memory serves as an adaptive solution to decision-making under uncertainty in complex environments, allowing us to flexibly repurpose our memories according to the demands of the present.

This work connects at least two established but largely separate literatures on memory and choice. First, a number of studies focused on decision-making have explored the ways in which individual experiences may be recalled for choice^{27,39–42}. This research, typically called 'decision by sampling', proposes that decision variables may be constructed by drawing samples from memory and explains a number of ways in which peoples' choice behaviour differs when information is learned from experience rather than instructed descriptions⁴³. Much of this work has emphasized episodic memory's value as a store for single experiences, which is a useful property when data are sparse and summary statistics cannot be reliably constructed, such as at the beginning of learning or following changes in the environment^{7,25}. Our findings address a complementary computational advantage: episodic memory's ability to store experiences in high fidelity across multiple features simultaneously. Second, other work has proposed that episodic memory plays a critical role in our ability to infer new information about the world by allowing the formation of new links between past experiences^{44–46}. An important but underappreciated part of this role is episodic memory's ability to store multiple features of experience, because each feature provides another opportunity to relate past events with one another. Here, we connect these ideas by proposing that features of episodes may enable the on-the-fly formation of new decision variables when they are required for a choice.

Our findings add to a substantial body of research which has found that the brain contains multiple memory systems that can operate

independently yet interact continuously^{11,26,37,38,47,48}, with task demands determining which is ultimately recruited for decision-making^{25,49,50}. Our results suggest that this parallel operation is adaptive—while incremental learning can efficiently track relevant statistics when future demands are known, episodic memory may serve as a crucial ‘backup’ system by preserving detailed records of experiences. This is indicated by participants’ ability to successfully encode detailed episodes in experiments 3 and 4 even while engaging in the feature-based strategy, and by their ability to use these memories to support unexpected decisions when incremental estimates were poorly calibrated. These findings align with recent computational work⁹ suggesting that episodic memory may complement incrementally learned summaries by enabling the computation of new task-relevant statistics when unexpected events occur. Further research has demonstrated that episodic encoding is enhanced for schema-incongruent or surprising events^{51,52}, which suggests that creating detailed memories is most critical when current predictions fail. However, our results indicate that some level of detailed episodic encoding may occur for most experiences, providing a baseline level of flexibility for future behaviour that may be further enhanced by surprise or prediction error signals^{53,54}.

Separately, our experiments uncovered systematic differences in how memory search strategies adapt to task demands and distinguish effective from ineffective decision-makers. In experiments 1–4, we found that participants’ decision times scaled with their total number of memories in each round rather than just those relevant to each offer, suggesting relatively broad retrieval from memory. However, in experiment 5, we found that better-performing participants were characterized by their ability to target retrieval towards decision-relevant memories. These findings suggest that, while some intrusion of irrelevant memories during retrieval is inevitable, as assumed by most computational models of free recall³⁰, people also use more efficient search strategies to improve decision-making. Future work examining how people develop selective retrieval strategies will be important to understand what factors promote more efficient memory search during decision-making.

A related next step lies in characterizing how the various organizational principles of episodic memory guide choice. Classic computational models of memory retrieval have formalized how memories are naturally organized and accessed according to their temporal proximity and semantic relationships^{55–57}. This work suggests that, when we retrieve a memory, it automatically triggers the recall of other memories that, for instance, occurred nearby in time or share meaningful features. Indeed, participants in our experiments showed a clear tendency to consecutively recall items that were presented close together during encoding (Supplementary Fig. 3). Such organizational principles are probably central to memory-guided decision-making, as indicated by findings demonstrating that choices are systematically influenced by temporally proximate experiences²², by the context present during both encoding²³ and decision-making^{23,58}, and by semantic similarity between experiences^{27,59}. While our results show that people can effectively use episodic memories for flexible decision-making, they do not reveal the specific mechanisms by which memories are accessed. Our behavioural measures were designed to be predicted by basic properties of episodic memory recall and so remain agnostic towards any particular model or sampling algorithm. Indeed, these findings could theoretically arise from many different memory sampling strategies—from random sampling to more structured retrieval guided by these organizational principles. While we developed our toy process model to demonstrate that basic probabilistic sampling of memories can reproduce observed behaviour, it is likely that it does not capture the sophisticated retrieval mechanisms that probably govern real-world memory-guided decisions. Applying formal memory models to value-based choice could provide a more precise account of how memories are actually accessed during decision-making, while potentially revealing new principles about how memory organization

shapes adaptive behaviour. Doing so will require not only experimental paradigms that more closely approximate real-world decision-making, where people draw upon vast numbers of feature-rich experiences across extended timescales, but also methodology to measure direct memory access during decision-making itself.

Related to this point, a critical feature of our experimental design was that we collected memory measures immediately following participants’ choices rather than during decision-making. We did not ask participants to directly recall items during their decisions because our aim was to assess the strategy participants relied upon without instructing them to use any strategy in particular, and we reasoned that this approach may bias them towards using their episodic memories for choice. One way to circumvent this limitation would be to record neural activity during the decision phase. For example, past approaches using magnetoencephalography have successfully decoded both the recall of individual episodes during standard memory tasks^{60,61} and the rapid replay of sequences of memories during decision-making^{62–64}. One possible future direction would be to similarly attempt to decode memory access during the decision phase of our task design. In addition to providing direct evidence for the recall of individual memories during choice, taking such an approach may provide multiple insights into the ways in which value is computed from memories, for instance by testing hypotheses about the number, temporal order and semantic relationships between recalled memories.

There are also several other limitations of our experimental approach. First, we assumed that episodes consist of perfect recollections of experienced details, when in reality these details are often substantially compressed and abstracted^{65–68}. This simplification may not capture how episodic memory actually operates in naturalistic settings, where imperfect and reconstructed details may alter its relative advantages over incremental learning. Separately, our design lacked explicit behavioural measures during encoding that may have revealed more subtle differences in strategy use, particularly because the feature-based strategy should impose greater demands during encoding while the episodic strategy should be more demanding during retrieval. Future work could address this limitation by manipulating cognitive load at each of these timepoints, or by using neuroimaging to predict subsequent recall from neural activity at encoding time. Finally, while overall performance was not altered by increasing feature complexity in experiments 2 and 5, suggesting that participants relied on integrated episodes rather than separate feature stores throughout our tasks, it remains possible that people differ in the extent to which they use either of these representations during choice. This possibility could be addressed by future experiments that manipulate feature complexity within rather than across individuals.

In conclusion, our results demonstrate that episodic memory plays an important role in enabling flexible decision-making when future task demands are unknown. By maintaining detailed representations of individual experiences, episodic memory allows us to access details from our past if they become relevant for present decisions. This flexibility comes at the cost of increased computational demands during choice, leading people to adopt more efficient strategies based on precomputing decision variables when possible. Recent work on the timing of memory-based decisions supports this view, showing that people proactively compute value from memory whenever circumstances allow them to anticipate future choice requirements⁶⁹. Together, these findings suggest that one reason why we maintain detailed memories of the past is to help us flexibly adapt to an uncertain future.

Methods

Experimental procedure

Unless otherwise noted, all procedures were identical across experiments, and differences between experiments are summarized in Table 3. Participants completed a four-part task over the course of a single online session designed to measure whether people access individual episodes

Table 3 | Summary of experimental differences

	Experiment 1	Experiment 2	Experiment 3	Experiment 4	Experiment 5
Number of rounds	5	7	8	9	1
Number of stimulus features	2: colour and category	4: texture, location, animacy and size	2: colour and category	2: colour and category	30: unspecified
Reward range	-9 to 9	-2 to 2	-2 to 2	-2 to 2	-1 to 1
Decisions per round	Sample A: 6; sample B: 3	4	Sample A: 3; sample B: 3	1 (6 in final round)	30
Feature uncertainty manipulation	No	No	Yes	Yes	No
Surprise choice manipulation	No	No	No	Yes	No

to make decisions based on multiple features of past experiences (Fig. 1a). Completing all four parts (a round) took approximately 5 min, and participants completed five (experiment 1), seven (experiment 2) or eight (experiments 3 and 4) rounds in total. In experiment 5, participants completed a single round that took approximately 20 min.

Stimuli. In experiments 1, 3 and 4, stimuli varied across two features: colour (red, yellow, blue or green) and category (animal, object, food or scene). In experiment 2, stimuli instead varied across four binary features: texture (with or without a pattern), location (aquatic or not aquatic), animacy (animal or object) and size (larger or smaller than a microwave). In experiments 1–4, a total of 16 possible items were used, with a subset of 6 pseudo-randomly sampled to be used in each round. All two-feature items and an example subset (highlighted in red) are shown in Fig. 1b. As demonstrated in this figure, the six items used in each round were selected such that there were always three items with a specific instantiation of each feature (for example, three green items and three animals), two of another type (for example, two blue items and two objects) and one of another type (for example, one food and one yellow). Items could repeat across rounds, and the number of repetitions per item was pseudo-randomly balanced within each session. Items repeated between one and four times in a single session depending on the experiment. Each time an item appeared, it was associated with a different reward. In experiment 5, stimuli were naturalistic images with many possible features and 21 images were shown a single time each.

Encoding phase. In the first part of a round, participants completed a task designed to allow them to encode individual items and an associated reward (which we refer to throughout as an episode). Each item was presented on the screen for 1 s, after which its reward appeared alongside it for another 6 s. An item's associated reward was a pseudo-randomly sampled integer (excluding 0) between either -9 and 9 (experiment 1), between -2 and 2 (experiments 2, 3 and 4) or between -1 and 1 (experiment 5). In experiment 5, to aid future recall, the reward was presented using a coloured border around the presented image (green for 1 and red for -1). To create a balance of future offer values in each round, the reward assignment algorithm ensured that (1) an equal balance of positive and negative reward was used in each round, (2) no feature dimension (for example, blue or object) had a summed value of exactly zero and that (3) at least one feature dimension had a positive summed value and at least one had a negative total value. Immediately after viewing the episode, participants completed an attention check consisting of the item alongside two options, either the associated reward that was just shown or another randomly selected reward. They had 3 s to respond. Each episode was viewed only once, for a total of six trials per round.

Distractor phase. Immediately following the encoding task, participants completed a 90-s distractor task to prevent active rehearsal of the

episodes. This distractor consisted of a two-back working memory task in which participants were shown one of several letters in sequence. Participants were asked to identify whether the current letter matched the one presented two steps earlier.

Decision phase. Immediately following the distractor task, participants then made several decisions based on the features of each item (six decisions: experiment 1 (sample A); four decisions: experiment 2; three decisions: experiments 1 (sample B) and 3; one decision: experiment 4; 30 decisions: experiment 5). Each decision consisted of an offer in which a single feature (for example, animal) was displayed on the screen, and participants were asked to either take or leave this offer. Participants were informed that the value of each offer consisted of the sum of each episode that was described by the offer (for example, the value of the animal offer would be the sum of the rewards associated with all animals seen during encoding) and that they should take positive offers and leave negative offers. Participants had 7.5 s to make each decision.

Memory phase. Finally, immediately after the decision task, we assessed participants' memory for the episodes in two ways. In experiments 1–4, participants were asked to freely recall the items that they saw in each round. They were provided with six empty text boxes and were told to enter the items in the order in which they remembered them. Participants were further told to halt their recall and move on to the next task if they could no longer remember any items. Following the free recall portion, participants were shown each item and were asked to provide their memory for the reward that was associated with each item. In experiment 5, free recall proceeded similarly to the previous experiments except that participants were given a single large box to enter their responses. Value recall in experiment 5 consisted of a two-alternative forced choice between each image with either a red or green border, and participants were instructed to choose the option that matched what they saw earlier in the task.

Feature uncertainty manipulation (experiments 3 and 4). To determine whether episodic memory is used preferentially when it is unclear which features should be prioritized during encoding, we manipulated whether information was provided to participants about upcoming choices in experiments 3 and 4. In these experiments, participants were told either before the encoding phase (four rounds) or after the distractor phase (four rounds) that they would be shown offers of only one feature type (either colour or category) during the decision phase. The order in which participants were shown each condition was counterbalanced.

Surprise choice manipulation (experiment 4). We aimed to further test whether participants had access to features that were made irrelevant by the feature uncertainty manipulation, as we predicted that one hallmark of using episodic memory for decision-making would

be the availability of all of the features of an event at decision time. To accomplish this, we added a final round to Experiment 4 in which participants were informed—unexpectedly—that, regardless of whether they had previously been told they would receive offers concerning only a single feature, they would now be required to make decisions about all possible offers.

Task instructions and practice. Across experiments 1–4, participants were told a version of the following at the beginning of the experiment, with slight variations depending on the experimental condition. For brevity, we have left out aspects of the instructions related to button presses and task timing:

Today you will be playing a game where your job is to earn as many points as possible. You will first see several images and will learn how many points each image is worth. You will then be asked to make choices based on what you have learned in order to earn points. The game will unfold over several different phases, and you will complete multiple rounds of these phases. Each round is independent of the others, so what you learn in one round does not influence the other rounds at all.

In the first phase, you will be shown several images, one at a time, alongside how many points it is worth. Each image can be worth either a positive or negative number of points. After viewing an image and its value, you will be asked to identify the number of points it was worth. You will be shown two numbers on either the left or right side of the screen. One of these numbers will be the image's value, while the other will not. You will see each image and its point value only once in a round.

You will need to use the value of each image to make choices. You will be given an offer on the screen and will be asked whether you would like to take or leave the offer. You should take an offer if you think it will allow you to earn more points, and you should leave an offer if you think it will cause you to lose points. All offers will consist of different features of the images you saw. Each offer is worth the sum of the point values of all images that match that feature.

In between the two phases you just learned about, you will complete another brief task. During this phase, you will see a sequence of letters presented one at a time. Your job is to determine if the letter on the screen matches the letter that appeared two letters before.

The last phase is a brief memory test where you will be asked to recall information about the items and point values that you learned about in the first phase.

The instructions for experiment 5 used this same language but did not include any information about rounds or the memory test.

During the instructions, participants were provided with two practice trials for each phase except the memory phase. Participants were required to achieve 100% on a multiple choice comprehension test with ten questions about the instructions in order to proceed with the study. They were further required to repeat all instructions related to any missed questions until they answered correctly. If they missed more than three questions, they were required to repeat the entire sequence of instructions, including practice trials.

Participants. All experiments were approved by the New York University Institutional Review Board, and all participants provided informed consent before their participation. Participants with normal

or corrected-to-normal vision were recruited from the New York University participant pool. Compensation was provided in the form of course credit. Participants were excluded if they indicated on a post-task questionnaire that they wrote any information down during the study or if they either failed to answer or provided nonsensical responses to the post-task questions. To incentivize honest responses on this questionnaire, participants were told that they would receive compensation regardless of their answers.

Experiment 1 (sample A), 83 participants were recruited and 16 were excluded, leading to a final sample of 67 participants ($M_{age} = 19.05 \pm 0.15$; 23 males, 42 females, 2 declined to say). For experiment 1 (sample B), 58 participants were recruited and 8 were excluded, leading to a final sample of 50 participants ($M_{age} = 18.76 \pm 0.2$; 14 males, 36 females). For experiment 2, 90 participants were recruited and 15 were excluded, leading to a final sample of 75 participants ($M_{age} = 18.89 \pm 0.17$; 17 males, 56 females, 2 declined to say). For experiment 3 (sample A), 61 participants were recruited and 11 were excluded, leading to a final sample of 50 participants ($M_{age} = 18.93 \pm 0.18$; 13 males, 37 females). For experiment 3 (sample B), 97 participants were recruited and 13 were excluded, leading to a final sample of 84 participants ($M_{age} = 18.81 \pm 0.14$; 25 males, 58 females, 1 declined to say). For experiment 4, 139 participants were recruited and 14 were excluded, leading to a final sample of 125 ($M_{age} = 19.14 \pm 0.12$; 42 males, 83 females). We recruited a larger sample for experiment 4 because each participant could complete only one of the before/after conditions during the final 'surprise' round, and we aimed to have roughly comparable numbers for each condition to our prior experiments. For experiment 5, 72 participants were recruited and 9 were excluded, leading to a final sample of 63 participants ($M_{age} = 18.78 \pm 0.16$; 17 males, 46 females). We determined our sample sizes based on effects measured in an initial pilot study, which is reported in Supplementary Fig. 5.

Model simulations

We formalized the episodic-based decision strategy using a toy process model. This model makes decisions by sequentially sampling individual memories and their associated rewards without replacement. Memories each consist of a binary feature vector and reward. When given an offer, the model first samples a memory i with probability computed using a softmax function with sampling bias parameter τ . The logit is defined as

$$\text{logit}_i = \begin{cases} \tau & \text{if memory } i \text{ is relevant to current decision} \\ 0 & \text{if memory } i \text{ is irrelevant} \end{cases}$$

The sampling probability for each memory is then

$$p_i = \frac{\exp(\text{logit}_i)}{\sum_{j=1}^N \exp(\text{logit}_j)}.$$

This creates a sampling bias where relevant memories have probability $\frac{\exp(\tau)}{\sum_j \exp(\text{logit}_j)}$ while irrelevant memories have probability $\frac{1}{\sum_j \exp(\text{logit}_j)}$. When $\tau = 0$, all memories are sampled uniformly. As τ increases, the model increasingly favours sampling memories that share features with the current offer.

After sampling a memory, the model adds its reward to a running sum and increments the elapsed decision time by a fixed recall time. The model then decides whether to continue sampling using a geometric stopping rule: at each step, sampling terminates with probability ϕ . If sampling continues, the model selects another memory from the remaining unsampled memories, with probabilities renormalized over the available set. The recalled value for memory i includes Gaussian noise to mimic imprecise recall of reward:

$$v_i = r_i + \epsilon_i,$$

where r_i is the true reward and $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. The total recalled value is then

$$V_{\text{recalled}} = \sum_{i \in \text{sampled}} v_i.$$

Sampling terminates when either (1) the stopping criterion is met, (2) all memories have been sampled or (3) maximum decision time is reached. The total decision time is then

$$RT = n_{\text{sampled}} \times t_{\text{recall}},$$

where t_{recall} is a constant. Finally, once sampling has terminated, the model makes a binary choice based on whether the total recalled value is positive:

$$\text{Choice} = \begin{cases} 1 & \text{if } V_{\text{recalled}} > 0 \text{ and at least} \\ & \text{1 relevant memory was sampled} \\ \text{Bernoulli}(0.5) & \text{if no relevant memories} \\ & \text{were sampled} \end{cases}.$$

We also formalized the feature-based strategy using a model that precomputes the value of each feature during encoding by summing the rewards associated with all items sharing that feature. During the decision phase, the model simply retrieves the cached value for the offered feature and makes a choice according to a logistic choice rule with inverse temperature β . This strategy predicts constant decision times based on only on a non-retrieval time ρ because only a single cached value must be referenced on each choice.

To first illustrate the behaviours that allowed us to distinguish between episodic and feature-based decision strategies in Fig. 1c, we simulated each model on the same basic structure used with human participants in experiments 1–4 ($N = 6$ memories). We ran 5,000 simulations for each model with the following values for each parameter. For the episodic model these were: recall time $t_{\text{recall}} = 1$, reward recall noise $\sigma = 0.5$, stopping probability $\phi = 0.1$, sampling bias $\tau = 0$ and a maximum decision time of 7.5 s. In these simulations, t_{recall} exerts no influence over recall because its product with the number of possible sampled memories cannot exceed the maximum decision time. For the feature-based model this was: non-retrieval time $\rho = 1.5$ s and inverse temperature $\beta = 5.0$. We chose these parameters purely to illustrate (1) the expected positive relationship between decision response times and the number of sampled memories and (2) that recalled offer value should differ from true offer value as a function of the sampled memories and distortions in reward recall in the episodic memory model. In general, these properties are expected under any values of ϕ and σ large enough to create variation in recalls across decisions.

We also simulated episodic memory agents completing experiment 5 (Fig. 5e) by generating $N = 21$ memories with features chosen to create a variety of offers with different numbers of relevant memories. Performance was then simulated across several parameter combinations (sampling bias $\tau \in \{0, 1, 2.5, 5, 10\}$ and probability of stopping $p_{\text{stop}} \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$), with 5,000 independent agents for each combination. We used an upper bound of 7.5 s for each decision. The reward recall noise σ was fixed to a negligible value (0.001) to more clearly assess the impacts of sampling bias and the probability of stopping on choice accuracy. We also fixed the recall time t_{recall} to 1 in order to ensure that agents with minimal stopping tendencies (low ϕ) could not exhaustively search all available memories, thereby preserving the variance in retrievals necessary to create individual differences in performance. Because decision response times are determined by the overall number of memories recalled, we focused on assessing how each parameter combination impacts

the relationship between the number of sampled memories that were irrelevant to each offer and accuracy. We calculated the amount of variance in response times explained by irrelevant memories across all of the agents with each combination of parameters.

Statistical analysis

All data were analysed with regression models estimated using hierarchical Bayesian inference such that group-level priors were used to regularize participant-level estimates unless otherwise specified. Predictors were specified as fixed effects alongside random slopes and intercepts that were allowed to vary across participants. In experiments 3A, 3B and 4, parameter estimates for rounds completed in either the before or after conditions were fitted separately. The joint posterior was approximated using No-U-Turn Sampling as implemented in Stan²⁰. Four chains with 2,000 samples (1,000 discarded as burn-in) were run for a total of 4,000 posterior samples per model. Chain convergence was determined by ensuring that the Gelman–Rubin statistic R was close to 1. Default weakly informative priors implemented in the brms package were used for each regression model²¹. For all models, fixed effects are reported in the text as the mean of each parameter’s marginal posterior distribution alongside 95% HDIs. In figures, 95%, 80% and 50% HDIs are shown, each indicating the range that contains the corresponding percentage of the posterior density. Parameter values outside of these ranges are unlikely given the model, data and priors.

Response time analysis. To examine how the number of recalled memories impacted the amount of time it took participants to make their choices in all experiments, we used a linear mixed-effects model to predict trial-wise response times. Response times were modelled using a shifted lognormal distribution, which accounts for the positive skew typical of response times. We included both random intercepts and slopes for participants to account for individual differences in both baseline response speed and sensitivity to the number of memories recalled. The model can be written as

$$RT_{ij} \sim \text{ShiftedLogNormal}(\mu_{ij}, \sigma, \theta) \\ \mu_{ij} = \beta_0 + \beta_1 n_{\text{Memories}}_{ij} + u_{0j} + u_{1j} n_{\text{Memories}}_{ij},$$

where i indexes trials and j indexes participants, and the tilde (\sim) indicates how the outcome variable is distributed. The parameters u_{0j} and u_{1j} represent the random intercepts and slopes for each participant, σ is the scale parameter and θ is the shift parameter of the distribution. The number of memories participants recalled in each round (n_{Memories}_{ij}) was included as a continuous predictor. We further conducted a separate analysis of the number of offer-relevant memories that participants recalled in each round, which consisted of an identical model but with this predictor instead.

In experiment 5, the mixed-effects model of the total number of memories recalled by participants included only a random intercept because this measure did not vary within participants. We also assessed how individual differences in task performance related to the impact of recalled memories on response times in two ways. First, using the same mixed-effects modelling framework, we modelled the relationship between the number of offer-irrelevant memories participants had and their performance as

$$\mu_{ij} = \beta_0 + \beta_1 n_{\text{IrrelMemories}}_{ij} + \beta_2 \text{Accuracy}_j + \beta_3 n_{\text{IrrelMemories}}_{ij} \\ \times \text{Accuracy}_j + u_{0j} + u_{1j} n_{\text{IrrelMemories}}_{ij},$$

where Accuracy_j is each participant’s average choice accuracy. Second, we determined whether the total number of memories or the number of offer-relevant memories better predicted response times as a function of performance. Specifically, we split participants into two groups based on the median task performance, refit the original mixed-effects

models to each group and then assessed model fit by separating the data into ten folds and cross-validating, as described in detail in our decision analyses.

Decision analysis. To analyse the extent to which decision-relevant variables influence participants' choices in all experiments, we fitted two mixed-effects logistic regression models. Binary choices (1 = offer taken, 0 = offer rejected) were predicted by either an offer's true value or its recalled value. The true value of an offer (TrueValue_{ij}) was just the sum of all offer-relevant items that were shown to the participants. By contrast, we computed the recalled offer value ($\text{RecalledValue}_{ij}$) as the sum over the reward that was remembered for each offer-relevant item that was also recalled during the free recall phase. These predictors were z-scored before model fitting to facilitate comparison between them. Each model included random intercepts and slopes to account for individual differences in both baseline offer acceptance rates and value sensitivity. These models can be written as

$$\text{Choice}_{ij} \sim \text{Bernoulli}(p_{ij})$$

$$\text{logit}(p_{ij}) = \beta_0 + \beta_1 \text{TrueValue}_{ij} + u_{0j} + u_{1j} \text{TrueValue}_{ij}$$

$$\text{logit}(p_{ij}) = \beta_0 + \beta_1 \text{RecalledValue}_{ij} + u_{0j} + u_{1j} \text{RecalledValue}_{ij}$$

To determine the extent to which either recalled offer value or true offer value best predicted choices, we compared model fit using either predictor. Specifically, model fit was assessed by separating the data into ten folds and cross-validating. The ELPD was then computed by summing the log likelihood for each held out datapoint and then used as a measure of out-of-sample predictive fit for each model, where higher ELPD values suggest better model fit, as they indicate a higher likelihood of accurately predicting new data. To compare models, we then subtracted the pointwise ELPD estimates and calculated the standard error of this difference to quantify uncertainty of the comparison^{72,73}.

Lastly, we assessed overall task performance using a separate mixed-effects logistic regression model. The model included only fixed and random intercepts to assess whether accuracy (1 = correct, 0 = incorrect) was from different from chance-level performance (0.5). The model was simply

$$\text{Correct}_{ij} \sim \text{Bernoulli}(p_{ij})$$

$$\text{logit}(p_{ij}) = \beta_0 + u_{0j}$$

Surprise round analysis. In experiment 4, we also assessed performance on the final 'surprise' round where participants were required to make decisions about the offers that they expected (in the before condition) as well as for offers that were unexpected (in both conditions). To do so, we calculated a Δ Performance score for each participant, which consisted of taking their difference in accuracy on the final round relative to their average accuracy on previous rounds (previous round accuracy – last round accuracy). We calculated this score separately for expected and unexpected offers in each condition and then used Δ Performance as the outcome variable in separate simple linear regression models:

$$\Delta\text{Performance}_{ij} \sim \text{Normal}(\mu, \sigma)$$

$$\mu = \beta_0$$

Memory analysis. We assessed performance on each part of the memory phase of experiments 1–4 using two complementary models. First, we examined participants' overall recall rates, which we defined as the proportion of items they accurately recalled relative to all of the items they were shown. We compared these recall rates with chance performance, which we defined as the probability of correctly guessing items when randomly selecting 6 items from the pool of 16 possible

items that could be shown on a given round, which was 37.5% (or 6/16). For each participant, we computed the difference between their mean recall rate and chance (Recall_i) and then fitted a simple linear regression model to these difference scores:

$$\text{Recall}_i \sim \text{Normal}(\beta_0, \sigma)$$

Next, to assess the accuracy of participants' memory for the reward associated with successfully recalled items, we fitted a mixed-effects linear regression model predicting participants' remembered rewards from the true associated rewards. Both remembered ($\text{RememberedReward}_{ij}$) and true (TrueReward_{ij}) rewards were normalized by dividing by the maximum absolute value in the dataset. The model included random intercepts and slopes for participants to account for individual differences in both baseline memory and value sensitivity:

$$\text{RememberedReward}_{ij} \sim \text{Normal}(\mu_{ij}, \sigma)$$

$$\mu_{ij} = \beta_0 + \beta_1 \text{TrueReward}_{ij} + u_{0j} + u_{1j} \text{TrueReward}_{ij}$$

Because experiment 5 assessed participants' memory for associated reward using a two-alternative forced choice task, we fitted a simple mixed-effects model to determine the extent to which their responses differed from chance. This model was identical to the model used to assess overall task performance.

Replication analysis for experiments 1–4. To assess the extent to which our effects replicated and/or differed across experiments sharing similar designs, we also ran our primary analyses with an identifier for each sample included as a fixed interaction effect. Data were pooled across comparable experiments (experiments 1 and 2; experiments 3 and 4), and we used sum-to-zero (effect) coding for the experiment factor. This approach allows each experiment's coefficient to represent its deviation from the pooled mean across all experiments, rather than a comparison with an arbitrary reference experiment. We performed these pooled analyses for both the response time and decision models. In experiments 3 and 4, we additionally assessed interaction effects of condition on decision response times in a single mixed-effects model and assessed the difference in ELPD differences between conditions from pooled choice models. The results of these analyses are reported in Supplementary Tables 4 and 5.

Controlling for individual differences in effort. Across all experiments, we also assessed possible contributions of individual differences in effort to the relationship between the number of recalled memories and response times in three ways. First, we looked for evidence of a speed–accuracy trade-off during the decision phase of each experiment by fitting mixed-effects models predicting choice accuracy from response times:

$$\text{Correct}_{ij} \sim \text{Bernoulli}(p_{ij})$$

$$\text{logit}(p_{ij}) = \beta_0 + \text{RT}_{ij} + u_{0j} + u_{1j} \text{RT}_{ij}$$

Second, we included z-scored performance on the two-back distractor task as a fixed-effect covariate in the response time models. Third, to distinguish between-participant and within-participant effects of the number of memories presented on response times, we decomposed this predictor into two orthogonal components: a between-participant component representing each participant's average number of memories recalled, and a within-participant component representing the deviation from each participant's personal average. The results of these analyses are reported in Supplementary Tables 6–9.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data are available via GitHub at https://github.com/jonathanicholas/nm2025_emdm.

Code availability

The code is available via GitHub at https://github.com/jonathanicholas/nm2025_emdm.

References

1. Tulving, E. in *Organization of Memory* (eds Tulving, E. & Donaldson, W.) 381–403 (Academic Press, 1972).
2. Eichenbaum, H. The hippocampus and mechanisms of declarative memory. *Behav. Brain Res.* **103**, 123–133 (1999).
3. Bartlett, F. C. *Remembering: A Study in Experimental and Social Psychology* (Cambridge Univ. Press, 1932).
4. Anderson, J. R. & Milson, R. Human memory: an adaptive perspective. *Psychol. Rev.* **96**, 703–719 (1989).
5. Schacter, D. L., Guerin, S. A. & Jacques, P. L. S. Memory distortion: an adaptive perspective. *Trends Cogn. Sci.* **15**, 467–474 (2011).
6. Shohamy, D. & Adcock, R. A. Dopamine and adaptive memory. *Trends Cogn. Sci.* **14**, 464–472 (2010).
7. Lengyel, M. & Dayan, P. Hippocampal contributions to control: the third way. In *Proc. Advances in Neural Information Processing Systems* (eds Platt, J. C. et al.) 889–896 (Curran Associates, 2008).
8. Gershman, S. J. & Daw, N. D. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annu. Rev. Psycho.* **68**, 101–128 (2017).
9. Nagy, D. G., Orbán, G. & Wu, C. M. Adaptive compression as a unifying framework for episodic and semantic memory. *Nat. Rev. Psychol.* **4**, 484–498 (2025).
10. Rescorla, R. A. & Wagner, A. R. in *Classical Conditioning II: Current Research and Theory* (eds Black, A. H. & Prokasy, W. F.) 64–99 (Appleton-Century-Crofts, 1972).
11. Shohamy, D. & Daw, N. D. in *The Cognitive Neurosciences* 5th edn (eds Gazzaniga, M. S. & Mangun, G. R.) 591–603 (Boston Review, 2014).
12. Knowlton, B. J., Mangels, J. A. & Squire, L. R. A neostriatal habit learning system in humans. *Science* **273**, 1399–1402 (1996).
13. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
14. Bayer, H. M. & Glimcher, P. W. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* **47**, 129–141 (2005).
15. Shohamy, D., Myers, C. E., Grossman, S., Sage, J. & Gluck, M. A. The role of dopamine in cognitive sequence learning: evidence from Parkinson's disease. *Behav. Brain Res.* **156**, 191–199 (2005).
16. Zaghloul, K. A. et al. Human substantia nigra neurons encode unexpected financial rewards. *Science* **323**, 1496–1499 (2009).
17. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* 2nd edn (MIT Press, 2018).
18. Wilson, R. C. & Niv, Y. Inferring relevance in a changing world. *Front. Hum. Neurosci.* **5**, 189 (2012).
19. Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V. & Niv, Y. Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron* **93**, 451–463 (2017).
20. Niv, Y. et al. Reinforcement learning in multidimensional environments relies on attention mechanisms. *J. Neurosci.* **35**, 8145–8157 (2015).
21. Miller, G. A. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.* **63**, 81–97 (1956).
22. 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).
23. Duncan, K., Semmler, A. & Shohamy, D. Modulating the use of multiple memory systems in value-based decisions with contextual novelty. *J. Cogn. Neurosci.* **31**, 1455–1467 (2019).
24. Mason, A., Madan, C. R., Simonsen, N., Spetch, M. L. & Ludvig, E. A. Biased confabulation in risky choice. *Cognition* **229**, 105245 (2022).
25. Nicholas, J., Daw, N. D. & Shohamy, D. Uncertainty alters the balance between incremental learning and episodic memory. *eLife* **11**, e81679 (2022).
26. Montaser-Kouhsari, L., Nicholas, J., Gerraty, R. T. & Shohamy, D. Differentiating reinforcement learning and episodic memory in value-based decisions in Parkinson's disease. *J. Neurosci.* **45**, e0911242025 (2024).
27. Plonsky, O., Teodorescu, K. & Erev, I. Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychol. Rev.* **122**, 621–647 (2015).
28. Daw, N. D. & Dayan, P. The algorithmic anatomy of model-based evaluation. *Philos. Trans. R. Soc. B* **369**, 20130478 (2014).
29. Kool, W., Cushman, F. A. & Gershman, S. J. When does model-based control pay off? *PLoS Comput. Biol.* **12**, e1005090 (2016).
30. Kahana, M. J. *Foundations of Human Memory* (Oxford Univ. Press, 2012).
31. Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
32. Sederberg, P. B., Howard, M. W. & Kahana, M. J. A context-based theory of recency and contiguity in free recall. *Psychol. Rev.* **115**, 893–912 (2008).
33. O'Reilly, R. C. & McClelland, J. L. Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus* **4**, 661–682 (1994).
34. O'Reilly, R. C. & Rudy, J. W. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol. Rev.* **108**, 311–345 (2001).
35. Rudy, J. W. & O'Reilly, R. C. Conjunctive representations, the hippocampus, and contextual fear conditioning. *Cogn. Affect. Behav. Neurosci.* **1**, 66–82 (2001).
36. Lee, S. W., O'Doherty, J. P. & Shimojo, S. Neural computations mediating one-shot learning in the human brain. *PLoS Biol.* **13**, e1002137 (2015).
37. Packard, M. G. & McGaugh, J. L. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol. Learn. Mem.* **65**, 65–72 (1996).
38. Poldrack, R. A. & Packard, M. G. Competition among multiple memory systems: converging evidence from animal and human brain studies. *Neuropsychologia* **41**, 245–251 (2003).
39. Stewart, N., Chater, N. & Brown, G. D. Decision by sampling. *Cogn. Psychol.* **53**, 1–26 (2006).
40. Biele, G., Erev, I. & Ert, E. Learning, risk attitude and hot stoves in restless bandit problems. *J. Math. Psychol.* **53**, 155–167 (2009).
41. Shadlen, M. N. & Shohamy, D. Decision making and sequential sampling from memory. *Neuron* **90**, 927–939 (2016).
42. Wang, S., Feng, S. F. & Bornstein, A. M. Mixing memory and desire: how memory reactivation supports deliberative decision-making. *WIREs Cogn. Sci.* **13**, e1581 (2022).
43. Hertwig, R. & Erev, I. The description–experience gap in risky choice. *Trends Cogn. Sci.* **13**, 517–523 (2009).
44. Schlichting, M. L. & Preston, A. R. in *The Hippocampus from Cells to Systems: Structure, Connectivity, and Functional Contributions to Memory and Flexible Cognition* (eds Hannula, D. E. & Duff, M. C.) 405–437 (Springer, 2017).
45. Whittington, J. C. et al. The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249–1263 (2020).

46. Biderman, N., Bakkour, A. & Shohamy, D. What are memories for? The hippocampus bridges past experience with future decisions. *Trends Cogn. Sci.* **24**, 542–556 (2020).

47. Squire, L. R. & Zola, S. M. Structure and function of declarative and nondeclarative memory systems. *Proc. Natl Acad. Sci. USA* **93**, 13515–13522 (1996).

48. Poldrack, R. A. et al. Interactive memory systems in the human brain. *Nature* **414**, 546 (2001).

49. 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).

50. 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).

51. Greve, A., Cooper, E., Tibon, R. & Henson, R. N. Knowledge is power: prior knowledge aids memory for both congruent and incongruent events, but in different ways. *J. Exp. Psychol. Gen.* **148**, 325–341 (2019).

52. van Kesteren, M. T. R., Ruiter, D. J., Fernández, G. & Henson, R. N. How schema and novelty augment memory formation. *Trends Neurosci.* **35**, 211–219 (2012).

53. Rouhani, N. & Niv, Y. Signed and unsigned reward prediction errors dynamically enhance learning and memory. *eLife* **10**, e61077 (2021).

54. Jang, A. I., Nassar, M. R., Dillon, D. G. & Frank, M. J. Positive reward prediction errors during decision-making strengthen memory encoding. *Nat. Hum. Behav.* **3**, 719–732 (2019).

55. Howard, M. W. & Kahana, M. J. A distributed representation of temporal context. *J. Math. Psychol.* **46**, 269–299 (2002).

56. Polyn, S. M., Norman, K. A. & Kahana, M. J. A context maintenance and retrieval model of organizational processes in free recall. *Psychol. Rev.* **116**, 129–156 (2009).

57. Kahana, M. J. Computational models of memory search. *Annu. Rev. Psychol.* **71**, 107–138 (2020).

58. Bornstein, A. M. & Norman, K. A. Reinstated episodic context guides sampling-based decisions for reward. *Nat. Neurosci.* **20**, 997–1003 (2017).

59. Aka, A. & Bhatia, S. What I like is what I remember: memory modulation and preferential choice. *J. Exp. Psychol. Gen.* **150**, 2175–2184 (2021).

60. Michelmann, S., Staresina, B. P., Bowman, H. & Hanslmayr, S. Speed of time-compressed forward replay flexibly changes in human episodic memory. *Nat. Hum. Behav.* **3**, 143–154 (2019).

61. Wimmer, G. E., Liu, Y., Vehar, N., Behrens, T. E. J. & Dolan, R. J. Episodic memory retrieval success is associated with rapid replay of episode content. *Nat. Neurosci.* **23**, 1025–1033 (2020).

62. Liu, Y., Mattar, M. G., Behrens, T. E. J., Daw, N. D. & Dolan, R. J. Experience replay is associated with efficient nonlocal learning. *Science* **372**, eabf1357 (2021).

63. McFadyen, J., Liu, Y. & Dolan, R. J. Differential replay of reward and punishment paths predicts approach and avoidance. *Nat. Neurosci.* **26**, 627–637 (2023).

64. Wimmer, G. E., Liu, Y., McNamee, D. C. & Dolan, R. J. Distinct replay signatures for prospective decision-making and memory preservation. *Proc. Natl Acad. Sci. USA* **120**, e2205211120 (2023).

65. Carmichael, L., Hogan, H. P. & Walter, A. A. An experimental study of the effect of language on the reproduction of visually perceived form. *J. Exp. Psychol.* **15**, 73–86 (1932).

66. Nagy, D. G., Török, B. & Orbán, G. Semantic compression of episodic memories. Preprint at <https://arxiv.org/abs/1806.07990> (2018).

67. Zhou, D. et al. A compressed code for memory discrimination. Preprint at *bioRxiv* <https://doi.org/10.1101/2025.10.12.681901> (2025).

68. Kerrén, C., Reznik, D., Doeller, C. F. & Griffiths, B. J. Exploring the role of dimensionality transformation in episodic memory. *Trends Cogn. Sci.* **29**, 614–626 (2025).

69. Nicholas, J., Daw, N. D. & Shohamy, D. Proactive and reactive construction of memory-based preferences. *Nat. Commun.* **16**, 1618 (2025).

70. Hoffman, M. D. & Gelman, A. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593–1623 (2014).

71. Bürkner, P.-C. *Brms*: an R package for Bayesian multilevel models using Stan. *J. Stat. Softw.* **80**, 1–28 (2017).

72. Vehtari, A., Gelman, A. & Gabry, J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Stat. Comput.* **27**, 1413–1432 (2017).

73. Sivula, T., Magnusson, M., Matamoros, A. A. & Vehtari, A. Uncertainty in Bayesian leave-one-out cross-validation based model comparison. Preprint at <https://arxiv.org/abs/2008.10296> (2023).

Acknowledgements

We thank members of the Mattar Lab as well as K. Jensen, Q. Lu, N. Biderman and D. Shohamy for their helpful comments and feedback on the project. This work was supported by the National Science Foundation SBE Postdoctoral Research Fellowship under grant no. 2507527.

Author contributions

J.N. and M.G.M. conceived the research and study design. J.N. conducted the data collection and analysis and drafted the paper. J.N. and M.G.M. reviewed and edited the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41562-025-02383-3>.

Correspondence and requests for materials should be addressed to Jonathan Nicholas.

Peer review information *Nature Human Behaviour* thanks Silvy H. P. Collin, Richard Henson and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2026

Corresponding author(s): Jonathan Nicholas

Last updated by author(s): Nov 9, 2025

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection Custom code written in jspsych (version 6.0.1) was used to run the experiment. The experiment was hosted on Heroku.

Data analysis Custom code written in python (version 3.8.3) and R (version 4.3.3) was used to analyze the data. All code are available on github: https://github.com/jonathanicholas/nm2025_emdm. Exact package versions for all dependencies are available in the repo environment.yml file.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

All data collected in the manuscript are available here: https://github.com/jonathanicholas/nm2025_emdm

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	Participants were asked about their biological sex, and this information is reported in the manuscript. Consent was not obtained to provide individual data linked to sex.
Reporting on race, ethnicity, or other socially relevant groupings	No socially constructed variables were used in the manuscript.
Population characteristics	Age and sex information are reported in the manuscript.
Recruitment	Participants were recruited through the NYU Subject Pool.
Ethics oversight	NYU Institutional Review Board

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	The study is a quantitative experimental design.
Research sample	New York University undergraduates completed the study for course credit. Experiment 1 (sample A), 83 participants were recruited and 16 were excluded, leading to a final sample of 67 participants (Mean age = 19.05; 23 males, 42 females, 2 declined to say). For Experiment 1 (sample B), 58 participants were recruited and 8 were excluded, leading to a final sample of 50 participants (Mean age = 18.76; 14 males, 36 females). For Experiment 2, 90 participants were recruited and 15 were excluded, leading to a final sample of 75 participants (Mean age = 18.89; 17 males, 56 females, 2 declined to say). For Experiment 3 (sample A), 61 participants were recruited and 11 were excluded, leading to a final sample of 50 participants (Mean age = 18.93; 13 males, 37 females). For Experiment 3 (sample B), 97 participants were recruited and 13 were excluded, leading to a final sample of 84 participants (Mean age = 18.81; 25 males, 58 females, 1 declined to say). For Experiment 4, 139 participants were recruited and 14 were excluded, leading to a final sample of 125 (Mean age = 19.14; 42 males, 83 females). We recruited a larger sample for Experiment 4 because each participant could complete only one of the before/after conditions during the final "surprise" round, and we aimed to have roughly comparable numbers for each condition to our prior experiments. For Experiment 5, 72 participants were recruited and 9 were excluded, leading to a final sample of 63 participants (Mean age = 18.78; 17 males, 46 females).
Sampling strategy	Convenience sampling of NYU undergraduate psychology students was used. Sample sizes were chosen based on comprehensive behavioral piloting (with the results of this pilot experiment reported in Supplementary Information) and, for later experiments, effects measured in earlier experiments. This approach is consistent with similar published studies.
Data collection	Data collection occurred online, on each participant's personal computer with a screen size of at least 1350 x 900.
Timing	All data was collected between Nov. 7 2023 and June 31 2025.
Data exclusions	Participants were excluded if they indicated on a post-task questionnaire that they wrote any information down during the study or if they either failed to answer or provided nonsensical responses to the post-task questions. These exclusion criteria were established in advance.
Non-participation	None.
Randomization	Counterbalancing of condition in experiments 3 and 4 was conducted randomly.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	Antibodies
<input checked="" type="checkbox"/>	Eukaryotic cell lines
<input checked="" type="checkbox"/>	Palaeontology and archaeology
<input checked="" type="checkbox"/>	Animals and other organisms
<input checked="" type="checkbox"/>	Clinical data
<input checked="" type="checkbox"/>	Dual use research of concern
<input checked="" type="checkbox"/>	Plants

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	ChIP-seq
<input checked="" type="checkbox"/>	Flow cytometry
<input checked="" type="checkbox"/>	MRI-based neuroimaging

Plants

Seed stocks

n/a

Novel plant genotypes

n/a

Authentication

n/a