Investigating how reward retroactively improves memory for associations and items

Hannah Tarder-Stoll[1], Margaret L. Schlichting[2,*], Katherine D. Duncan[2,*]

[1] Rotman Research Institute, Baycrest Hospital, Toronto, Ontario, Canada

[2] Department of Psychology, University of Toronto, Toronto, Ontario, Canada

*these authors contributed equally to this work

Please address correspondence to: Hannah Tarder-Stoll (hrt2116@columbia.edu)

ORCiD:
Hannah Tarder-Stoll: https://orcid.org/0009-0007-0957-4499
Margaret L. Schlichting: https://orcid.org/0000-0002-5433-8870
Katherine D. Duncan: https://orcid.org/0000-0003-0456-851X

Abstract

Given we cannot remember all our experiences, it seems intuitive that we should prioritize events of consequence, like those that lead to a reward. However, rewards have had inconsistent effects on people's memory for preceding neutral experiences in experiments, raising questions as to whether our minds are equipped to remember adaptively. Here, we will use online testing in a registered report to investigate whether rewards tend to retroactively enhance certain types of memories. Consistent with prior research and theories involving how dopamine modulates the hippocampus, we hypothesize that reward retroactively enhances memory in a delay-dependent manner, emerging only after a period of consolidation. Crucially, we further predict that this long-delay modulation of memory will be stronger for associative memory than for item recognition—a factor not previously considered despite its clear theoretical grounding. This work will answer important open questions about the mechanisms by which reward retroactively enhances memory, elucidating the conditions under which this key aspect of adaptive memory emerges.

Introduction

Prioritizing useful experiences in our memory is the cornerstone of adaptive behaviour[1,2]. For instance, experiences that allow us to anticipate and seek out rewards are more important to remember than those that were less consequential. Complicating this prioritization, many events are seemingly insignificant in the moment but later become meaningful through their association with a positive outcome[3]. How can we enhance memories for these events when their outcomes only later reveal themselves?

Despite the centrality of this problem to understanding adaptive learning and memory[3–5], whether reward does in fact retroactively enhance human memory is unclear due to mixed findings. While some previous studies have found evidence for such an effect[6–8], others have not[9,10], raising important questions about the particular circumstances under which it manifests. Identifying the boundary conditions of retroactive, reward-related memory enhancement is crucial to characterizing when different adaptive memory functions operate. Here, we explore a neurobiologically-grounded factor that could reconcile mixed results, namely distinctions between how we remember associations and recognize the items comprising them. Using a well-powered, registered design we will test the hypothesis that rewards enhance how we form rich, relational memories of preceding experiences, but are less beneficial for forming cruder memories sufficient to support item recognition.

Several previous studies of human behaviour have elegantly shown that reward can retroactively enhance memory for preceding neutral information[6–8]. In these studies, participants typically view individual items (like words or pictures of objects) and then have a chance to win a reward in an unrelated game. Using this structure, for example, Murayama and

Kitagami[6] had participants view lists of pictures and then play one of two games: one that offered monetary rewards and one that did not. The researchers found that people were better at later recognizing items that preceded the reward compared to the no reward game[6], hereafter referred to as the retroactive reward effect. This finding has been echoed across other similar studies[7,8]. Notably, these studies also show that the retroactive reward effect is delay-dependent, emerging only when memories are tested after a period of consolidation but not immediately[6–8].

The striking delay-dependence of the retroactive reward effect was impressively predicted from earlier neurobiological research in rodents. Reward is associated with dopamine release[11–13], and in turn dopamine strengthens connections between neurons in the hippocampus[14–16], which are critical for forming memories[17]. Even recently activated connections are strengthened by dopamine[18–20], explaining how reward can retroactively enhance memory[21–26]. But, because it takes time to remodel neural connections, dopamine's effect on both the hippocampus[15,27–29] and memory performance[24] are typically seen only after several hours—mirroring the effect of reward on human memory. Rodent research also shows that dopamine may further retroactively enhance our memory by fostering the replay of hippocampal memories formed prior to either a reward[30,31] or dopamine stimulation[32]. Together, these effects of dopamine in the hippocampus could mechanistically explain how rewards can reach back in time to enhance our memory for previously unimportant events.

However, this seemingly straightforward story is complicated by other studies of human behaviour which show that rewards do not retroactively enhance memory[9,10], raising questions about the reliability and boundary conditions of this phenomenon. Why are retroactive reward

effects on human memory so variable? One explanation may be that the types of memory assessed in work on this topic are not especially reliant on the hippocampus, despite this region playing a central role in the underlying theories. Specifically, the hippocampus is thought to rapidly bind the disparate aspects of an experience into a memory so that it can later be retrieved with rich contextual details—that is, support associative memory[33–35]. Underscoring the importance of this distinction, past work has shown that different motivational states engaged in the moments prior to memory formation benefit associative versus item memory[42]. While this supports the notion that a similar dissociation may also exist in the context of retroactive reward effects, to our knowledge no past work on this topic has probed associative aspects of memory. Instead, all studies have required participants to simply recognize individual stimuli (i.e., assessing item memory)[6–10], a task that could certainly benefit from, but does not require, retrieval of contextually rich memories[36–41].

Here, we explore whether the variability in how rewards retroactively enhance memory may stem, in part, from using assessments that do not consistently rely on the hippocampus. Certain design differences between experiments that have found positive versus null retroactive reward effects hint that this may be the case. For example, in Braun et al.[7] participants learned items while navigating through a maze, yielding a reliable reward modulation effect; such a spatial memory task could have engaged the hippocampus more than viewing objects on the screen alone[43,44]. Further, Patil et al.[8] found a reliable retroactive reward effect when limiting their analyses to high confidence recognition responses, which may indicate retrieval of contextually rich, hippocampal dependent memories despite the item recognition task[45]. In contrast, Kalbe & Schwabe[10] and Oyarzun et al.[9] collapsed across all trials,

regardless of confidence, in their analyses and reported null effects of reward. Another study, however, used recognition memory tests without confidence ratings or a clear link to hippocampal memory and still showed retroactive reward effects[6]. Given this mixed picture emerging from the literature, we need a systematic test of the prediction that reward retroactively enhances hippocampally dependent memories. As such, we plan to directly compare the retroactive reward effect on item recognition and associative memory to determine if retroactive enhancements are significantly more robust for associative aspects of memory, which have been closely linked to hippocampal function[33,34].

In this study, we will investigate the retroactive effect of reward on item and associative memory using a well-powered, registered, within-subjects design. We will adapt a task from Murayama and Kitagami[6] in which participants are exposed to short lists of images which are either followed by an unrelated reward ($1) or no reward ($0) game. We will include a filler task after each encoding block to ensure that reward effects to not spill over into the subsequent set of images, as in Murayama and Kitagami[6]. Unlike Murayama and Kitagami[6], who only investigated item memory, we will additionally consider memory for trial-unique object-object associations. Specifically, following the incidental memory formation phase, we will immediately test participants' memory for (1) half of the objects (item memory) and (2) the corresponding object-object associations (associative memory). Participants will be tested again one day later, at which point we will test their memory for the other half of the objects and object-object associations.

We will investigate three pre-registered, *a priori* hypotheses. Firstly, we hypothesize a retroactive reward effect on memory for object-object associations, but only when tested at a

delay. Specifically, we hypothesize that rewards will not influence associative memory in the immediate test (Hypothesis 1a, see **Table 1**), but will reliably boost associative memory at the long delay (Hypothesis 1b, see **Table 1**). We will also test the interaction between reward and delay, but we have not registered the hypothesis because our power analyses revealed that we will not achieve 95% power to detect a null effect even at our maximum sample size of 240 participants (see **Supplementary Information**).

Next, replicating past work[6–8], we hypothesize that subsequent rewards will not influence item recognition on the immediate test (see **Table 1**). We do not have strong predictions about reward modulation at the long delay (Hypothesis 2b, see **Table 1**). On the one hand, the item test is a high-powered conceptual replication of Murayama & Kitagami[6] who did show retroactive reward enhancement; as such, the same might be expected here. On the other hand from a theoretical perspective we do not expect a strong role for hippocampus in item memory, nor empirically do we see any evidence for this effect in our pilot data (see **Supplementary Information**). We will again test for a reward by delay interaction but have not preregistered this hypothesis due to power reasons.

Finally, we will directly compare retroactive reward modulation across memory test types (item vs. associative) using normalized scores. We hypothesize that the retroactive reward effect will be larger for associative than item memory when tested at a delay (Hypothesis 3, see **Table 1**). For all of our hypotheses, we will interpret significant positive effects from null hypothesis significance testing in the predicted or opposite direction, and in the case of null results will perform Bayesian equivalence testing using the ROPE procedure[46] to assess evidence for null hypotheses as detailed in **Table 1**.

Methods

Ethics Information.

All procedures will be in accordance with a study protocol approved by the Research Ethics Board at the University of Toronto. All participants will provide informed consent and be compensated $8 CAD per hour for participating. They will also have the opportunity to earn an additional monetary bonus of up to $8 depending on their performance in the reward section of the task (see *Design*).

Design.

Materials.

Stimuli will consist of 288 color images of emotionally neutral objects (e.g., lamp, pencil, mug) collected from the BOSS stimulus set[47]. Objects will be grouped into pairs chosen to have low semantic and visual similarity to each other. Semantic similarity between objects in a pair will be assessed using WordNet::Similarity[48], an online lexical database assessing semantic relatedness between words. We will use a similarity score of 0.2 as our cut off for semantic relatedness within our triads. Visual similarity will be assessed by two independent raters, who will be told to give binary responses about the colour and shape similarity of object pairs and then regroup pairs if both raters agree that objects were visually similar. Ninety-six object pairs will be presented during encoding, and the additional 48 object pairs (96 objects) will be reserved and used as lures in the item recognition test (see *Procedure)*. Object pairs will be counterbalanced across participants such that each will be equally likely to appear during

encoding or as lures during the item recognition test, in the rewarded versus not rewarded condition, and in the immediate versus long delay condition.

Procedure.

We will conduct a 2 (reward, no reward) x 2 (immediate, long delay), within-subjects experiment adapted from the procedure used by Murayama and Kitagami[6]. The experiment will be completed across two sessions (**Figure 1A**). The first session of the experiment will consist of three phases: an encoding phase, an item recognition test, and an associative memory test. Participants will then complete the second session one day later, which will consist of another item recognition and associative memory test. The experiment will be programmed in and hosted through Inquisit 5.0 (Millisecond Software). Participants will perform the task on their personal computers and will not be eligible to participate with a phone or tablet. Data collection and analysis will not be performed blind to the conditions of the experiments, but participants will not be aware of the experimental manipulations or hypotheses and will not interact with the experimenter during online data collection.

Encoding Phase. The incidental encoding phase will consist of a series of 16 blocks. Each block will contain a list of six object pairs followed by a reward modulation event and then six flanker trials.

Each block will begin with the presentation of pairs of objects on a computer screen, with each pair visible for 4 seconds with a 500 millisecond ISI in between trials (hereafter referred to as encoding trials; **Figure 1B**). To encourage incidental memory encoding, participants will be instructed to imagine the two depicted objects interacting (no response required) but will not be told that their memory for these objects will be tested later.

Participants will view six object-object associations in one block, for a total of 96 object-object associations across the entire experiment.

During the reward modulation event (**Figure 1B**), participants will play the "clock game" which will be associated with a monetary reward half of the time. First, participants will see a 1 second cue that says "reward clock" on reward trials or "no reward clock" on no reward trials. They will then see a fixation cross for a variable duration of 1.5-2 seconds. Participants will then see a circle appear on the screen and, on both reward and no reward trials, they will have to press the spacebar within 500ms. The circle will be a different colour (green or blue) to signify reward versus no reward trials, with the colour-condition mapping counterbalanced across participants. On reward trials, there could be one of two outcomes: if participants press the spacebar fast enough, they will be rewarded with $1, and will see a message that says "Congratulations! You won $1" for 2 seconds. If they are too slow on reward trials, they will fail to get a reward and will see a message that says "Too slow! The next trial will begin shortly" for 2 seconds. In comparison, on no reward trials, regardless of how quickly participants press the spacebar, they will not get a reward and will see a message that says "The next trial will begin shortly" for 2 seconds. Half of the blocks will be from the reward condition and half from the no reward condition. The order of reward and no reward blocks will be pseudorandomized across participants to ensure that (1) rewards are equally distributed between the first and second half of the experiment and (2) participants receive no more than two reward blocks in a row.

Following the reward event, participants will complete six trials of a flanker task (**Figure 1B**) to reduce the carryover of the reward modulation to memories formed in the next block. Participants will see three arrows on the screen and will be told to indicate the direction that

the middle arrow is pointing, with the outer arrows randomly selected to point in the same or different direction as the middle arrow. The direction of the middle arrow will point left and right each half of the time. They will have 2 seconds to make their response with a 500 millisecond ISI between trials. If participants fail to respond they will automatically progress to the next trial. In total, the filler task will be 15 seconds long. This filler block duration is longer than in Murayama and Kitagami[6] to further reduce the chance of reward carryover.

Each block will take approximately 50 seconds to complete, and the entire encoding phase will take approximately 13.5 minutes. Participants will have the opportunity to take one self-timed break in the middle of the encoding phase.

Immediate Memory Tests. Immediately following the encoding phase, participants will complete an item recognition test (**Figure 1C**) for half of the individual objects presented in the encoding phase. They will see 96 old objects (48 objects from the reward condition and 48 objects from the no reward condition) and 48 lures. They will be instructed to respond "old" if they believe they saw the object during the encoding session or "new" if they believe that the object was not previously presented. Participants will have 1.5 seconds to respond to each item, immediately after which a 1 second ISI will begin, regardless of when the participants respond. The item recognition test will take approximately 6 minutes to complete.

Following the item recognition test, participants will complete a 3 alternative forced choice associative memory test (**Figure 1D**) for the object-object associations presented during the encoding phase. The associative memory test will always happen after the item recognition test because while the associative memory test re-exposes participants to old objects and therefore may facilitate their memory, the recognition test does not re-expose participants to

associations. In the associative memory test, participants will be shown one object at the top of the screen and will be asked to indicate which of the three objects shown on the bottom was associated with it. Participants will be tested on 48 associations, comprising the same objects tested in the item recognition test (24 associations from the reward condition and 24 from the no reward condition). Additionally, the lures in the 3 alternative forced choice associative memory test will include one object from a different association encoded in the same block, and one encoded in a different block. All lures will be from the same reward and delay condition as the correct answer on a given trial (e.g., if the correct answer on a trial is from the reward, immediate condition, both lures will also be from the reward, immediate condition), and therefore will have been tested in the preceding recognition test. Participants will have 4 seconds to respond, immediately after which a 1 second ISI will begin, regardless of when the participants respond. The associative memory test will take approximately 4 minutes to complete.

Long Delay Memory Tests. Participants will be invited to participate in the second session one calendar day after completing session one, with a minimum delay of 12 hours between sessions. They will be prohibited from participating in session two if the delay is more than one calendar day. In the second session, participants will then complete item recognition and associative memory tests in the same format as described above. In the item recognition test, they will be tested on the 96 previously untested objects from encoding (48 objects from the reward condition and 48 objects from the no reward condition) along with 48 novel lures. In the associative memory test, they will be tested on the previously untested 48 associations (24 associations from the reward condition and 24 associations from the no reward condition) from

encoding. Finally, participants will complete a short post-task questionnaire in which they are asked about their strategies during the experiment. They will be paid their monetary bonus of up to $8, corresponding to their performance on reward trials during incidental encoding in the first session, upon completion of the long delay memory tests.

Quality Checks.

We will perform quality checks on the data to ensure the absence of floor or ceiling effects. This analysis will be performed after compliance-related exclusion but before memory-related exclusions (see Participants). We will report the proportion of participants with perfect performance on either memory test (hit rate of 1 and false alarm rate of 0 on the recognition test, 100% accuracy on the associative memory test) and with below chance performance (d' of 0, 33% accuracy on associative memory test). We will monitor performance of the first 40 participants. If more than 10% of participants are at floor or ceiling, we will stop data collection and adjust the task to make it easier or more difficult, depending on whether participants tend to be at floor or ceiling (e.g., we will increase or decrease the allotted time to respond on the memory tests). However, given the lack of floor of ceiling effects in our pilot data (see below), this is unlikely to be necessary. To validate our reward manipulation, we will report the speeding and accuracy differences between reward and no reward trials on the clock game. If there are no significant speeding or accuracy differences between reward conditions (assessed using null hypothesis significance testing) in the first 40 participants, we will increase the reward to $1.50. Again, given the robust effect of reward on clock game behavior in our pilot data (see below), this is unlikely to be necessary.

Sampling Plan.

Participants.

Participants will be recruited through CloudResearch, an online participant management tool associated with Mechanical Turk[49]. To be eligible to participate, they must report that they are between the ages of 18 and 30, are a native English speaker, and have normal or corrected to normal vision. Participants will be excluded after participation if they report that they take drugs with well-documented direct effects on the dopamine system. A participant's data will also be excluded from analyses if they fail to complete the second session of the experiment one calendar day after the first session, if they stop partway through a session, or if their data are not properly recorded or transferred to Inquisit.

Participants will also be excluded based on task compliance and performance on the memory tests using a multistep approach. First, we will exclude participants who are not compliant (a) in the clock game, defined as less than 75% success rate across both reward and no reward conditions, and/or (b) in the flanker task, defined as less than 50% response rate. Within this sample of compliant participants, we will the calculate the proportion of participants whose performance is at floor or ceiling (e.g., to determine whether any design changes are required, as described above). Next, we will exclude participants if they do not perform numerically above chance (d' of 0 in the item recognition test, 33% in the associative memory test) in one or more of the four memory tests. Exclusion-related memory analyses will be performed across all trials, irrespective of clock game performance. Data needing to be excluded after it has been collected will be replaced by recruiting new participants until we reach our target sample size (see below).

Power Analyses.

The effect size of retroactive reward modulation at a long delay is uncertain, with reports in the literature ranging from $d = 0.09$[10] to $d = 0.54$[6] for item recognition, and there is no published data from associative memory tests. Given this uncertainty, we will employ a sequential analysis procedure[50] with predefined checkpoints, including the lower and upper bounds for our sample size. Conservatively, we will select n = 90 participants who meet all inclusion criteria (including performance requirements) as our *minimum* sample size, which will provide 95% power to detect and effect size of $d = 0.39$ and 99.9% power to detect the effect size reported by Murayama and Kitagami ($d = 0.54$)[6]. Due to resource and time constraints, we will select n = 240 as our maximum sample size. With this sample size, we will have 95% power to detect a small effect size of $d = 0.24$ or higher. We will have one additional checkpoint at n = 165. We will use the GroupSeq package in R[51] to control the type I error rate at 0.05 by adjusting the alpha boundary at each checkpoint (Checkpoint 1: n = 90, alpha boundary = 0.0031; Checkpoint 2: n = 165, alpha boundary = 0.0183; Checkpoint 3: n = 240, alpha boundary = 0.044). At each checkpoint, we will assess evidence for the null for all hypotheses that do not reach significance using the Bayesian Region of Practical Equivalence (ROPE)[46] procedure. First, we will define the ROPE for each analysis as zero ± 0.17 times the standard deviation of our dependent variable for each analysis (boundaries determined to achieve 95% power through simulations; see *Supplement* for simulation results). Next, we will run Bayesian models and determine whether 95% of the Highest Density Interval (HDI) falls inside the ROPE[46]. At a given checkpoint, if all hypotheses have informative answers—either reaching the alpha boundary or yielding sufficient evidence for the null—we will stop data collection. Otherwise, we will continue data collection until we reach our maximum sample size.

Analysis Plan.

We will conduct all analyses using a mixture of frequentist and Bayesian models in the R programming language. For all analyses related to our hypothesis tests, we will restrict our consideration of reward blocks to only those in which the participant successfully obtained the reward. Missed responses on the memory tests will be coded as incorrect. To conduct our primary frequentist analyses, we will use linear mixed effects models (LMMs, lmer function in the *lme4* package)[52]. Because our analyses model summary statistics (i.e. d'), all models will include random intercepts, but not random slopes, grouped by participant. Confidence intervals for models will be calculated with the confint function in R (*lme4 package*)[52]. If the models fail to converge, we will use bobyqa optimization and increase the number of iterations.

If our frequentist models do not reveal significant results, we will use Bayesian generalized non-linear multivariate multilevel models (*brms* package)[53] to determine evidence for the null hypothesis for each of our *a priori* hypotheses. Bayesian models will be fit with the default weakly informative priors, with a student's t distribution centered at 0 with 1 degree of freedom and a scale parameter of .25. Using the *bayestestR* package, we will then determine the percentage of the HDI that falls within the ROPE (0 ± 0.17; see Supplement Results for simulations)[54]. If 95% of the HDI for the Bayesian model parameters of interest fall inside the ROPE, this will indicate evidence in favour of the null hypothesis. See *Supplementary Information* for all frequentist and Bayesian model specifications.

Associative Memory. We will assess whether reward influences associative memory in the immediate and delayed tests (H1a and H1b). We will first calculate average accuracy on the associative memory test for each participant in each reward and delay condition. We opted to

model average performance rather than trial-wise accuracy to keep this model as similar as possible to our item recognition model (see below). Next, we will predict average accuracy as a function of reward (-0.5 = no reward, 0.5 = reward) for the immediate and delayed tests separately using LMMs. In these models, a positive, significant effect of reward would suggest that reward retroactively enhances associative memory in either the immediate test (H1a) or the delayed test (H1b). Negative, significant effects would suggest that reward retroactively impairs associative memory, contrary to our hypotheses. If there is not a significant effect of reward, we will assess evidence for the null using a brms model with the same structure and the ROPE procedure described above.

Item Recognition. Next, we will calculate d' (z(hit rate) – z(false alarm rate)) for each combination of reward (reward vs. no reward) and delay (immediate vs. long) conditions for each participant. To account for 1's and 0's in the data, we will add 0.5 to the numerator and 1 to the denominator of the hit and false alarm rates[55]. To determine whether reward retroactively enhances item memory recognition (H2a and H2b), we will predict d' as a function of reward (-0.5 = no reward, 0.5 = reward) for the immediate and delayed tests separately using LMMs. In these models, a positive, significant main effect of reward would suggest that reward enhances item recognition in either the immediate test (H2a) or the delayed test (H2b). Negative, significant effects would suggest that reward retroactively impairs item recognition. If there is not a significant main effect of reward, we will assess evidence for the null using a brms model with the same structure and the ROPE procedure described above.

Comparing Item Recognition and Associative Memory. To assess our final *a priori* hypothesis (H3) that the retroactive effect of reward will be stronger for associative memory as

compared to item recognition at the long delay, we will put both dependent variables on the same scale. Specifically, we will separately z-score the distributions of delayed item and associative memory performance, both of which contain delayed memory performance in each reward condition for each participant. We will then run an LMM comparing the normalized scores as a function of memory test (-0.5 = item recognition, 0.5 = associative memory), reward (-0.5 = no reward, 0.5 = reward), and their interaction at the long delay only. A significant interaction, such that associative memory exhibits a greater difference between reward and no reward conditions than item memory, would be consistent with our hypothesis that reward enhances associative more than item recognition memory after a delay (H3). Other significant patterns would be inconsistent with our hypothesis. If there is not a significant interaction between reward and memory test, we will assess evidence for the null using a brms model with the same structure and the ROPE procedure described above.

If we do not find a significant effect of reward on either recognition or associative memory with our final sample size, we will try running analyses on a subset of participants who perform statistically above chance on all trials (i.e., exhibit both above-chance item recognition d', determined by bootstrapping the null for each participant; as well as above-chance associative memory performance of over 47.91% accuracy, determined using a binomial test). However, this will not influence our sampling plan, which is based on memory performance being numerically above chance (see *Sampling Plan)*.

Pilot Data.

We collected 36 pilot participants on the task described above (see *Procedure)* to validate our design and provide preliminary evidence for each of our three preregistered

hypotheses (see *Table 1*). Briefly, piloting confirmed that the task produced reasonable levels of performance across participants: There were no exclusions necessary related to task compliance; no observations were at ceiling, and only one participant had below floor memory performance (item d' below 0). Therefore, only one participant was excluded out of 36. Our pilot data further confirmed the effectiveness of our reward manipulation: participants were significantly faster and more accurate on reward vs no reward trials of the clock game (see Supplementary Information). Please see *Supplementary Information* for a full description and statistical analysis of the pilot data.

Data Availability.

All pilot data are publicly available on Open Science Framework (https://osf.io/35fcs/). Behavioural data from the full experiment will be anonymized and made publicly available.

Code Availability.

All randomization, experiment, and analysis scripts, along with resulting simulations from the pilot data are publicly available on Open Science Framework (https://osf.io/35fcs/). Analysis scripts from the full experiment will also be made publicly available.

References

1. Shohamy, D. & Adcock, R. A. Dopamine and adaptive memory. *Trends Cogn. Sci.* **14**, 464–472 (2010).

2. Richards, B. A. & Frankland, P. W. The Persistence and Transience of Memory. *Neuron* **94**, 1071–1084 (2017).
3. Hull, C. L. *Principles of behavior: an introduction to behavior theory*. x, 422 (Appleton-Century, 1943).
4. Izhikevich, E. M. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex N. Y. N 1991* **17**, 2443–2452 (2007).
5. Brzosko, Z., Schultz, W. & Paulsen, O. Retroactive modulation of spike timing-dependent plasticity by dopamine. *eLife* **4**, e09685 (2015).
6. Murayama, K. & Kitagami, S. Consolidation power of extrinsic rewards: Reward cues enhance long-term memory for irrelevant past events. *J. Exp. Psychol. Gen.* **143**, 15–20 (2014).
7. Braun, E. K., Wimmer, G. E. & Shohamy, D. Retroactive and graded prioritization of memory by reward. *Nat. Commun.* **9**, 4886 (2018).
8. Patil, A., Murty, V. P., Dunsmoor, J. E., Phelps, E. A. & Davachi, L. Reward retroactively enhances memory consolidation for related items. *Learn. Mem. Cold Spring Harb. N* **24**, 65–69 (2017).
9. Oyarzún, J. P., Packard, P. A., de Diego-Balaguer, R. & Fuentemilla, L. Motivated encoding selectively promotes memory for future inconsequential semantically-related events. *Neurobiol. Learn. Mem.* **133**, 1–6 (2016).
10. Kalbe, F. & Schwabe, L. On the search for a selective and retroactive strengthening of memory: repeated failure to find category-specific behavioral tagging. Preprint at https://doi.org/10.31234/osf.io/ed3z8 (2020).

11. Schultz, W. Predictive reward signal of dopamine neurons. *J. Neurophysiol.* **80**, 1–27 (1998).
12. Schultz, W. Getting formal with dopamine and reward. *Neuron* **36**, 241–263 (2002).
13. Montague, P. R., Hyman, S. E. & Cohen, J. D. Computational roles for dopamine in behavioural control. *Nature* **431**, 760–767 (2004).
14. Lisman, J. E. & Grace, A. A. The Hippocampal-VTA Loop: Controlling the Entry of Information into Long-Term Memory. *Neuron* **46**, 703–713 (2005).
15. Lisman, J., Grace, A. A. & Duzel, E. A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends Neurosci.* **34**, 536–547 (2011).
16. Tsetsenis, T. *et al.* Midbrain dopaminergic innervation of the hippocampus is sufficient to modulate formation of aversive memories. *Proc. Natl. Acad. Sci.* **118**, e2111069118 (2021).
17. Scoville, W. B. & Milner, B. LOSS OF RECENT MEMORY AFTER BILATERAL HIPPOCAMPAL LESIONS. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).
18. Frey, U. & Morris, R. G. Synaptic tagging and long-term potentiation. *Nature* **385**, 533–536 (1997).
19. Frey, U. & Morris, R. G. Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci.* **21**, 181–188 (1998).
20. Redondo, R. L. & Morris, R. G. M. Making memories last: the synaptic tagging and capture hypothesis. *Nat. Rev. Neurosci.* **12**, 17–30 (2011).
21. Morris, R. G. M. Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. *Eur. J. Neurosci.* **23**, 2829–2846 (2006).

22. Dunsmoor, J. E., Murty, V. P., Clewett, D., Phelps, E. A. & Davachi, L. Tag and capture: how salient experiences target and rescue nearby events in memory. *Trends Cogn. Sci.* **26**, 782–795 (2022).
23. Moncada, D. & Viola, H. Induction of Long-Term Memory by Exposure to Novelty Requires Protein Synthesis: Evidence for a Behavioral Tagging. *J. Neurosci.* **27**, 7476–7481 (2007).
24. Wang, S.-H., Redondo, R. L. & Morris, R. G. M. Relevance of synaptic tagging and capture to the persistence of long-term potentiation and everyday spatial memory. *Proc. Natl. Acad. Sci. U. S. A.* **107**, 19537–19542 (2010).
25. de Carvalho Myskiw, J., Benetti, F. & Izquierdo, I. Behavioral tagging of extinction learning. *PNAS Proc. Natl. Acad. Sci. U. S. Am.* **110**, 1071–1076 (2013).
26. Salvetti, B., Morris, R. G. M. & Wang, S.-H. The role of rewarding and novel events in facilitating memory persistence in a separate spatial memory task. *Learn. Mem. Cold Spring Harb. N* **21**, 61–72 (2014).
27. Rossato, J. I., Bevilaqua, L. R. M., Izquierdo, I., Medina, J. H. & Cammarota, M. Dopamine controls persistence of long-term memory storage. *Science* **325**, 1017–1020 (2009).
28. Bethus, I., Tse, D. & Morris, R. G. M. Dopamine and Memory: Modulation of the Persistence of Memory for Novel Hippocampal NMDA Receptor-Dependent Paired Associates. *J. Neurosci.* **30**, 1610–1618 (2010).
29. O'Carroll, C. M., Martin, S. J., Sandin, J., Frenguelli, B. & Morris, R. G. M. Dopaminergic modulation of the persistence of one-trial hippocampus-dependent memory. *Learn. Mem.* **13**, 760–769 (2006).

30. Singer, A. C. & Frank, L. M. Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* **64**, 910–921 (2009).
31. Ambrose, R. E., Pfeiffer, B. E. & Foster, D. J. Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron* **91**, 1124–1136 (2016).
32. McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N. & Dupret, D. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat. Neurosci.* **17**, 1658–1660 (2014).
33. Eichenbaum, H. A cortical–hippocampal system for declarative memory. *Nat. Rev. Neurosci.* **1**, 41–50 (2000).
34. Davachi, L. Item, context and relational episodic encoding in humans. *Curr. Opin. Neurobiol.* **16**, 693–700 (2006).
35. Mayes, A., Montaldi, D. & Migo, E. Associative memory and the medial temporal lobes. *Trends Cogn. Sci.* **11**, 126–135 (2007).
36. Wan, H., Aggleton, J. P. & Brown, M. W. Different contributions of the hippocampus and perirhinal cortex to recognition memory. *J. Neurosci. Off. J. Soc. Neurosci.* **19**, 1142–1148 (1999).
37. Eldridge, L. L., Knowlton, B. J., Furmanski, C. S., Bookheimer, S. Y. & Engel, S. A. Remembering episodes: a selective role for the hippocampus during retrieval. *Nat. Neurosci.* **3**, 1149–1152 (2000).
38. Brown, M. W. & Aggleton, J. P. Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nat. Rev. Neurosci.* **2**, 51–61 (2001).

39. Wais, P. E. Hippocampal signals for strong memory when associative memory is available and when it is not. *Hippocampus* **21**, 9–21 (2011).

40. Davachi, L. & Wagner, A. D. Hippocampal contributions to episodic encoding: insights from relational and item-based learning. *J. Neurophysiol.* **88**, 982–990 (2002).

41. Davachi, L., Mitchell, J. P. & Wagner, A. D. Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proc. Natl. Acad. Sci. U. S. A.* **100**, 2157–2162 (2003).

42. Murty, V. P. & Alison Adcock, R. Distinct Medial Temporal Lobe Network States as Neural Contexts for Motivated Memory Formation. in *The Hippocampus from Cells to Systems: Structure, Connectivity, and Functional Contributions to Memory and Flexible Cognition* (eds. Hannula, D. E. & Duff, M. C.) 467–501 (Springer International Publishing, 2017). doi:10.1007/978-3-319-50406-3_15.

43. Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. & Tanila, H. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* **23**, 209–226 (1999).

44. Burgess, N., Maguire, E. A. & O'Keefe, J. The human hippocampus and spatial and episodic memory. *Neuron* **35**, 625–641 (2002).

45. Rugg, M. D. *et al.* Item memory, context memory and the hippocampus: fMRI evidence. *Neuropsychologia* **50**, 3070–3079 (2012).

46. Kruschke, J. K. Rejecting or Accepting Parameter Values in Bayesian Estimation. *Adv. Methods Pract. Psychol. Sci.* **1**, 270–280 (2018).

47. Brodeur, M. B., Dionne-Dostie, E., Montreuil, T. & Lepage, M. The Bank of Standardized Stimuli (BOSS), a new set of 480 normative photos of objects to be used as visual stimuli in cognitive research. *PloS One* **5**, e10773 (2010).
48. Princeton University. WordNet. (2010).
49. Litman, L., Robinson, J. & Abberbock, T. TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behav. Res. Methods* **49**, 433–442 (2017).
50. Lakens, D. Performing high-powered studies efficiently with sequential analyses. *Eur. J. Soc. Psychol.* **44**, 701–710 (2014).
51. Pahl, R. GroupSeq: Group Sequential Design Probabilities - With Graphical User Interface. (2022).
52. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using **lme4**. *J. Stat. Softw.* **67**, 1–48 (2015).
53. Bürkner, P.-C. brms: An R Package for Bayesian Multilevel Models Using Stan. *J. Stat. Softw.* **80**, 1–28 (2017).
54. Makowski (@Dom_Makowski), D. *et al.* bayestestR: Understand and Describe Bayesian Models and Posterior Distributions. (2023).
55. Snodgrass, J. G. & Corwin, J. Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *J. Exp. Psychol. Gen.* **117**, 34–50 (1988).

Acknowledgments

We gratefully acknowledge Lana Duan and Carissa DeMarinis for assistance with data collection in earlier versions of this study. Thanks to Matti Vuorre and Monica Thieu for advice on the Bayesian analysis approach, and to Anuya Patil for helpful conversations. This work was supported by a Natural Sciences and Engineering Research Council (NSERC) Discovery Grant (PJT-159492) to KDD and MLS; Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant (RGPIN-2023-04231) to KDD; Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant (RGPIN-04933-2018) to MLS; Canada Foundation for Innovation JELF to MLS and KDD; and Ontario Research Fund (36876 to MLS; 34479 to KDD). The funders have no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

HTS, MLS, and KDD conceptualized the research question, experiment design, and analytic approach. HTS collected the data and performed the analysis. MLK and KDD acquired funding and supervised. HTS wrote the original draft and MLK and KDD revised and edited the draft.

Competing interests

The authors declare no competing interests.

Table 1 | Design Table

Question	Hypothesis	Sampling plan (e.g. power analysis)	Analysis Plan	Interpretation given to different outcomes
1A Is there a retroactive effect of reward on associative memory in the immediate test?	Associative memory will show no meaningful differences between the reward and no reward conditions in the immediate test.	Given uncertainty surrounding retroactive reward effect sizes in the literature, we will use an adaptive design with three checkpoints (n = 90, n = 165, n = 240), including participants that meet all inclusion criteria (see *Methods*). With 240 participants, we would be 95% powered to detect an effect size of d = 0.24, suggesting that we are well-situated to detect even small effect sizes. Bootstrapped power simulations based on our pilot data suggested that the maximum sample size in our adaptive design (240 participants) will achieve 99% power to detect a null effect of reward on associative memory at the short delay (see *Supplementary Information*).	We will conduct a frequentist linear mixed effects model predicting average associative memory accuracy in the reward vs. no reward conditions in the immediate test, with random intercepts grouped by participant. If there is no significant coefficient for reward, we will conduct a Bayesian mixed effects model predicting average associative memory accuracy in the reward vs. no reward conditions in the immediate test, with random intercepts grouped by participant. We will then determine the percentage of the HDI within a ROPE of 0 ± 0.17 times the standard deviation of the associative memory score (see *Methods*).	A significant positive coefficient for reward will be interpreted as evidence that memory was enhanced immediately for associations that preceded a reward vs no reward game during encoding. A significant negative coefficient for reward will be interpreted as evidence that memory was impaired for associations that preceded a reward vs. no reward game during encoding. If 95% of the HDI is inside the ROPE, we will consider this as evidence for a null effect of reward. If less than 95%, but more than 5% of HDI is in the ROPE, we will consider the results inconclusive with respect to the null.
1B Is there an effect of reward on associative memory in the delayed test?	Associative memory accuracy will be significantly higher in the reward compared to the no reward condition in the delayed test.	Given uncertainty surrounding retroactive reward effect sizes in the literature, we will use an adaptive design with three checkpoints (n = 90, n = 165, n = 240), including participants that meet	We will conduct a frequentist linear mixed effects model predicting average associative memory accuracy in the reward vs. no reward conditions in the delayed test, with random intercepts grouped by participant.	A significant positive coefficient for reward will be interpreted as evidence that memory was enhanced at a delay for associations that preceded a reward vs. no reward game during encoding. A significant negative coefficient for reward will be interpreted as evidence that memory was impaired for

		all inclusion criteria (see *Methods*). With 240 participants, we would be 95% powered to detect an effect size of *d* = 0.24, suggesting that we are well-situated to detect even small effect sizes. A power analysis based on the effect size in our pilot data for hypothesis 1b (d = 0.57) suggested that the maximum sample size in our adaptive design (240 participants) will achieve 99% power to detect an effect of reward on associative memory at the long delay.	If there is no significant coefficient for reward, we will conduct a Bayesian mixed effects model predicting average associative memory accuracy in the reward vs. no reward conditions in the delayed test, with random intercepts grouped by participant. We will then determine the percentage of the HDI within a ROPE of 0 ± 0.17 times the standard deviation of the associative memory score (see *Methods*).	associations that preceded a reward vs. no reward game during encoding. If 95% of the HDI is inside the ROPE, we will consider this as evidence for a null effect of reward. If less than 95%, but more than 5% of HDI is in the ROPE, we will consider the results inconclusive with respect to the null.
2A Is there a retroactive effect of reward on d' in the item recognition test in the immediate test?	Item recognition will show no meaningful differences in the reward and no reward conditions in the immediate test.	Given uncertainty surrounding retroactive reward effect sizes in the literature, we will use an adaptive design with three checkpoints (n = 90, n = 165, n = 240), including participants that meet all inclusion criteria (see *Methods*). With 240 participants, we would be 95% powered to detect an effect size of *d* = 0.24, suggesting that we are well-situated to detect even small effect sizes. Bootstrapped power simulations based on our pilot data suggested that the maximum sample size in our adaptive design	We will conduct a frequentist linear mixed effects model predicting d' in the reward vs. no reward conditions in the immediate test, with random intercepts grouped by participant. If there is no significant coefficient for reward, we will conduct a Bayesian mixed effects model predicting d' in the reward vs. no reward conditions in the immediate test, with random intercepts grouped by participant. We will then determine the percentage of the HDI within a ROPE of 0 ± 0.17 times the standard deviation of the d' score (see *Methods*).	A significant positive coefficient for reward will be interpreted as evidence that recognition memory was enhanced immediately for items that preceded a reward vs. no reward game during encoding. A significant negative coefficient for reward will be interpreted as evidence that recognition memory was impaired for associations that preceded a reward vs. no reward game during encoding. If 95% of the HDI is inside the ROPE, we will consider this as evidence for a null effect of reward. If less than 95%, but more than 5% of HDI is in the ROPE, we will consider the results inconclusive with respect to the null.

		(240 participants) will achieve 99% power to detect a null effect of reward on item memory at the short delay (see *Supplementary Information)*.		
2B Is there a retroactive effect of reward on item recognition in the delayed test?	We are agnostic as to whether item recognition will be modulated by retroactive reward. Replicating past work[6–8], we may find that item recognition is higher in the delayed test when followed by reward vs no reward. Alternatively, in line with our pilot data, item recognition may show no meaningful differences in the reward and no reward conditions in the delayed test.	Given uncertainty surrounding retroactive reward effect sizes in the literature, we will use an adaptive design with three checkpoints (n = 90, n = 165, n = 240), including participants that meet all inclusion criteria (see *Methods*). With 240 participants, we would be 95% powered to detect an effect size of d = 0.24, suggesting that we are well-situated to detect even small effect sizes. Bootstrapped power simulations based on our pilot data suggested that the maximum sample size in our adaptive design (240 participants) will achieve 99% power to detect a null effect of reward on item memory at the long delay (see *Supplementary Information*).	We will conduct a frequentist linear mixed effect model predicting d' in the reward vs. no reward conditions in the delayed test, with random intercepts grouped by participant. If there is no significant coefficient for reward, we will conduct a Bayesian mixed effects model predicting d' in the reward vs. no reward conditions in the delayed test, with random intercepts grouped by participant. We will then determine the percentage of the HDI within a ROPE of 0 ± 0.17 times the standard deviation of the d' score (see *Methods*).	A significant positive coefficient for reward will be interpreted as evidence that recognition memory was enhanced at a delay for items that preceded a reward vs no reward game during encoding. A significant negative coefficient for reward will be interpreted as evidence that recognition memory was impaired for associations that preceded a reward vs. no reward game during encoding. If 95% of the HDI is inside the ROPE, we will consider this as evidence for a null effect of reward. If less than 95%, but more than 5% of HDI is in the ROPE, we will consider the results inconclusive with respect to the null.
3 Is the effect of reward greater for associative than object memory in the delayed test?	Z scores will be higher for associative compared to item memory in the reward compared to no reward condition.	Given uncertainty surrounding retroactive reward effect sizes in the literature, we will use an adaptive design with three checkpoints (n = 90, n = 165, n = 240), including	We will z score average accuracy in the associative memory test and d' for the delayed tests. We will then conduct a frequentist linear mixed effect model predicting z score as a function of the	A significant positive coefficient for the reward X memory test interaction will be interpreted as evidence that, after a delay, memory is enhanced in the reward compared to no reward condition to a greater degree for associative compared to item memory tests. A significant

		participants that meet all inclusion criteria (see *Methods*). With 240 participants, we would be 95% powered to detect an effect size of *d* = 0.24, suggesting that we are well-situated to detect even small effect sizes. A power analysis based on the effect size in our pilot data for hypothesis 3 (d = 0.56) suggested that the maximum sample size in our adaptive design (240 participants) will achieve 99% power to detect an interaction between reward and memory test at the long delay.	interaction between reward (reward vs. no reward) and memory test (item vs. associative test) in the delayed test, with random intercepts grouped by participant. If there is no significant coefficient for the interaction between reward and memory test, we will also conduct a Bayesian mixed effects model predicting z score as a function of the interaction between reward (reward vs. no reward) and memory test (item vs. associative test) in the delayed test, with random intercepts grouped by participant. We will then determine the percentage of the HDI within a ROPE of 0 ± 0.17 times the standard deviation of the z scores.	negative coefficient for the reward X memory test interaction will be interpreted as evidence that, after a delay, memory is enhanced in the reward compared to no reward condition to a greater degree for item compared to associative memory tests. If 95% of the HDI is inside the ROPE, we will consider this as evidence for a null effect. If less than 95%, but more than 5% of HDI is in the ROPE, we will consider the results inconclusive with respect to the null.

Figure 1 | Task Schematic

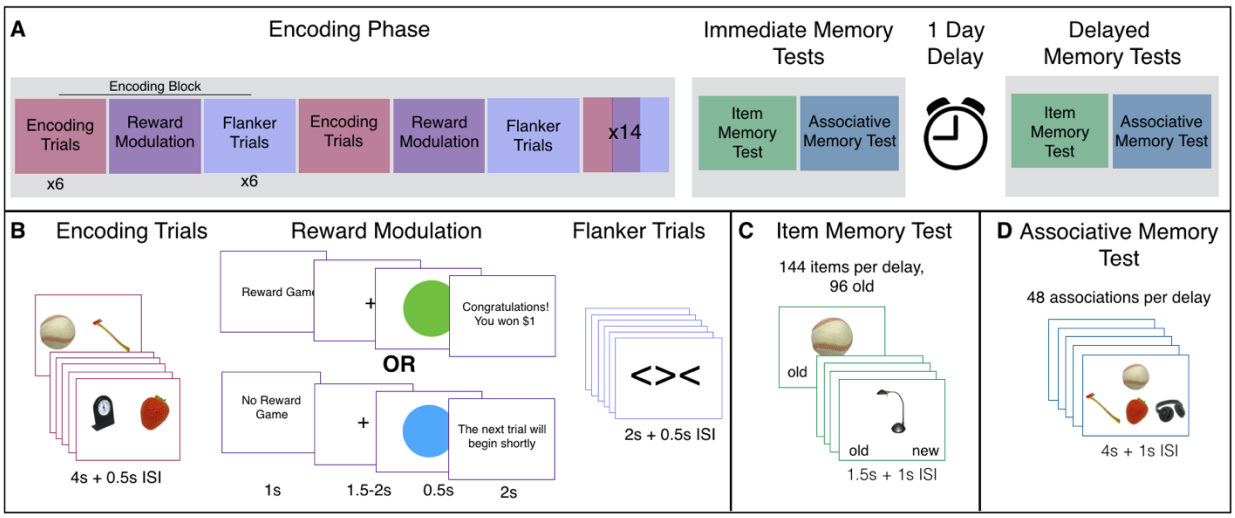

(**A**) Task Overview. In the first session, participants will complete the encoding phase and the immediate memory tests. The encoding phase will consist of encoding trials, a reward modulation event, and flanker trials that are completed sequentially in 16 blocks. Participants will then complete item recognition and associative memory tests. Participants will return one day later for a second session in which they will complete additional object and associative memory tests. (**B**) Schematic of the Encoding Phase. In each encoding trial, participants will encode an object-object association. Participants will then see a circle that is animated to look like a clock on the screen for 3 seconds (reward modulation event). They will be told to press the space bar within 500ms of seeing the circle. On reward trials, if they press the space bar quickly enough, they will be given $1. On no reward trials, they will not earn any money. Participants will then complete flanker trials in which they see three arrows on the screen and are told to indicate which direction the middle arrow is pointing. The flanker trials act as a reward washout period before beginning the next block of encoding trials. (**C**) Item Recognition Test. Participants will complete an item recognition test for the individual objects learned during the encoding phase. Participants will see an object on the screen and must indicate whether it was old or new. Half of the objects will be tested immediately, and the other half will be tested after the delay. (**D**) Associative Memory Test. Participants will then be tested on the object-object associations learned during the encoding phase. They will see an object on the top of the screen and will have to indicate which of the three objects on the bottom it had been paired with during encoding. Half of the associations will be tested immediately, and the other half of the associations will be tested after the delay.