Cognition 231 (2023) 105315

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Value restructures the organization of free recall

Elizabeth A. Horwath [a], Nina Rouhani [b], Sarah DuBrow [c], Vishnu P. Murty [a,*]

[a] Department of Psychology, Temple University, Philadelphia, PA, United States of America
[b] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA, United States of America
[c] Department of Psychology, University of Oregon, Eugene, OR, United States of America

ARTICLE INFO

Keywords:
Episodic memory
Free recall
Reward
Motivation
Context

ABSTRACT

A large body of research illustrates the prioritization of goal-relevant information in memory; however, it is unclear how reward-related memories are organized. Using a rewarded free recall paradigm, we investigated how reward motivation structures the organization of memory around temporal and higher-order contexts. To better understand these processes, we simulated our findings using a reward-modulated variant of the Context Maintenance and Retrieval Model (CMR; Polyn et al., 2009). In the first study, we found that reward did not influence temporal clustering, but instead shifted the organization of memory based on reward category. Further, we showed that a reward-modulated learning rate and source features of CMR most accurately depict reward's enhancement on memory and clustering by value. In a second study, we showed that reward-memory effects can exist in both extended periods of sustained motivation and frequent changes in motivation, by showing equivalent reward effects using mixed- and pure-list motivation manipulations. However, we showed that a reward-modulated learning rate in isolation can support reward's enhancement of memory in pure-list contexts. Overall, we conclude that reward-related memories are adaptively organized by higher-order value information, and contextual binding to value contexts may only be necessary when rewards are intermittent versus sustained.

1. Introduction

Imagine going to the grocery store to buy ingredients for an upcoming dinner as well as a few other essentials. If you were asked to recall what you purchased afterwards, it is possible that you could start by describing the first things you picked up and move forward in time with each item you put in the cart. But, we often buy many things in one trip, which makes it unlikely that you would remember every single item. Instead, individuals prioritize information that is relevant to their goals (Murty & Adcock, 2017), such that you would be more likely to recall items for your upcoming dinner as opposed to more mundane, basic items. This type of prioritization may bias the organization of your recall away from the order in which you selected items but rather towards the importance of items to your goal state while shopping (e.g., grouping ingredients you need for dinner). Prior research has detailed the role of reward and goal relevance in influencing which items are stored in memory (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Miendlarzewska, Bavelier, & Schwartz, 2016; Murty & Adcock, 2017; Shohamy & Adcock, 2010), but it remains unknown how reward-related memories are organized.

A large body of research has characterized memory organization and free recall dynamics for neutral memoranda. While information is bound to the temporal context of encoding, it is most common for items to be recalled in the order they were learned (Howard & Kahana, 2002b; Polyn, Norman, & Kahana, 2009). During goal-directed behavior, however, this may not be the most adaptive form of memory organization. Organizing memories based on temporal order may lead to an inability to discern important shared relationships among items that span across temporal context, which could hinder the achievement of goals. Rather, organizing memories by a value-related context would lead to the clustering of goal-relevant information to support more adaptive behaviors like generalization, inference, and insight (Cowan, Schapiro, Dunsmoor, & Murty, 2021).

In prominent models of free recall, items are thought to be bound to different contextual features during encoding which are reinstated during retrieval to organize recall (Howard & Kahana, 2002b; Kahana, 1996; Polyn et al., 2009). The Context Maintenance and Retrieval Model (CMR; Polyn et al., 2009) describes how items are learned and later recalled based on associations between items and their context. During encoding, as an item is presented, it is stamped into a continuously

* Corresponding author at: 1701 N 13th St, Philadelphia, PA 19122, United States of America.
E-mail address: vishnu.murty@temple.edu (V.P. Murty).

https://doi.org/10.1016/j.cognition.2022.105315
Received 3 March 2022; Received in revised form 12 October 2022; Accepted 22 October 2022
Available online 15 November 2022

drifting context (Polyn et al., 2009). The further away two items are presented in time, the less they are associated with each other since their temporal contexts differ as a function of time. After learning, items compete to be recalled and are cued by the context of the most recently recalled item. One hypothesis regarding how value may influence memory organization would suggest that reward modulates how strongly items are bound to a temporal context, resulting in greater temporal contiguity at recall. The binding of information to its temporal context has been shown to be supported by the hippocampus (HPC; (Davachi, 2006; Manning, Polyn, Baltuch, Litt, & Kahana, 2011), and a large literature has shown that the HPC supports reward-motivated memory in conjunction with involvement of the ventral tegmental area (VTA; Adcock et al., 2006; Murty & Adcock, 2017). Given the hippocampus's role in supporting both temporal memory and reward-related memory, it may be hypothesized that reward biases memory organization towards temporal context, a process which could be mediated by the HPC.

However, some evidence against this hypothesis suggests that reward disrupts temporal recall dynamics. For example, Murphy and colleagues have shown that reward motivation biases recalltowards strategic value-related organization and away from typical serial organization which would be structured by temporal context (Murphy & Castel, 2022; Murphy, Schwartz, & Castel, 2022). Additionally, Stefanidi, Ellis, and Brewer (2018) have suggested that reward disrupts typical primacy and recency effects by showing that recall is more likely to begin with a high-value word, regardless of serial position. Together, this suggests that reward may disrupt temporal organization and instead drive recall towards a value-related structure.

In the CMR framework, there are additional contextual features, such as task demands or the modality in which information is delivered (i.e., source), that influence the likelihood of an item being recalled. Items that belong to the same source and share task contexts are often recalled contiguously (Murdock & Walker, 1969; Polyn et al., 2009). Extending on this prior work, the source context has been used to differentiate between high- and low-salience items (Talmi, Lohnas, & Daw, 2019, 2021). Here, we are interested in how these source features may influence memory organization for goal-relevant items. A competing hypothesis therefore predicts that high-reward items will become more associated with each other, and therefore more likely to be recalled contiguously, since they share a 'value context'. Through this reward context, during recall, we hypothesize that transitions between words will be to ones that share the same reward context rather than ones that were encoded closer in time.

A few additional pieces of neural evidence lend to the second hypothesis that reward may target higher-order features of memory. While it has been widely accepted that the HPC supports episodic memory (Davachi, 2006), recent work has shown that the HPC is also involved in supporting memory for higher-order concepts, such as famous people and places (Morton, Zippi, Noh, & Preston, 2021). Given that the hippocampus can support conceptual information, as well as reward-related memory, it may also facilitate memory organization through higher-order categories such as reward value (Adcock et al., 2006; Murty & Adcock, 2017). Second, there is evidence that salient information facilitates interactions between VTA and anterior temporal networks (Cowan, Fain, O'Shea, Ellman, & Murty, 2021), which are known to support value-related and semantic information, respectively. Given this prior work, we propose that rather than influencing memory through temporal features of an experience, reward may bias memory organization towards de novo reward contexts. Behaviorally, we probe this hypothesis by characterizing clustering around reward value, rather than temporal context.

A related, yet parallel, question concerns the mechanisms by which the salience of reward motivation influences encoding. Prior research proposes that affect only influences memory when salience is triggered by the local environment (i.e., an emotional item surrounded by neutral information) and does not extend to time periods that involve a sustained state of affective information (i.e., a list of all emotional words; Talmi, Luk, McGarry, & Moscovitch, 2007, 2019). This theoretical framework was developed under the context of emotional valence, however, has been extended to the domain of reward motivation in studies demonstrating that reward's influence on memory only occurs in the context of mixed lists (Talmi et al., 2007, 2019). In these studies, high- and low-salience items (i.e., high and low reward value) were learned intermixed within the same list, driving attentional salience, rather than pure lists, where high- and low-salience items were presented and tested separately (Talmi et al., 2007, 2019). However, these findings challenge neurobiological models of reward-motivated memory which suggest that motivation can be sustained over extended periods of time (Murty & Adcock, 2017; Shohamy & Adcock, 2010), resulting in more tonic increases in memory. Thus, comparing reward's benefit on memory across mixed and pure lists, and testing memory after learning words in both categories rather than separately in the pure list condition, would help disambiguate among these two mechanisms.

Across two studies, we use a rewarded free recall task in which items are learned under the context of either high or low reward. We measure recall organization to determine how value influences the extent to which items are organized by the order in which they were encoded, compared to a reward context, which we consider higher-order, value information. We then extend the interpretation of our findings with a computational modeling approach by using a reward-motivated variant of CMR to simulate features of the encoding and retrieval processes to understand the underlying mechanisms of reward-related memory organization.

2. Methods

2.1. Participants

Two separate cohorts, without any history of psychiatric or neurological conditions, were recruited from Temple University to participate in two studies. Informed consent was obtained from each participant in a manner approved by Temple University's Institutional Review Board. Sixty-three people participated in Study 1; three were removed due to failure to return for the 24-h test and four because of experimental error, resulting in a final sample of 56 (47 females, 9 males, ages 17–39, median age: 19). Sixty-two people were recruited to participate in Study 2; two were removed because of failure to return for the 24-h test and one due to failure to follow task instructions, resulting in a final sample of 59 participants (48 females, 10 males, 1 non-binary, ages 18–25, median age: 20). Sample sizes were chosen based on standards in the literature of free recall (for example, Kahana, Howard, Zaromb, & Wingfield, 2002; Talmi et al., 2007) and motivation (for example, Murayama & Kitagami, 2014; Murayama & Kuhbandner, 2011; Patil, Murty, Dunsmoor, Phelps, & Davachi, 2017) and are sufficient to show both a medium effect size of 0.5 and a strong effect size of 0.9 with a power level of 0.8 (Patil et al., 2017).

2.2. Encoding

After consent was obtained, participants were given instructions for the task. All participants performed encoding on day 1, and recall followed by recognition on day 2. Half of the participants in each study also completed the recall test immediately after encoding on day 1. During encoding, participants learned 70 words in total that were randomly chosen from the Toronto Noun Word Pool (Friendly, Franklin, Hoffman, & Rubin, 1982), only 50 of which were used in our analyses. The first and last 10 words, which were associated with a low reward value, were removed from analyses to preclude recency and primacy effects (which we refer to as "buffer" words). The remaining 50 target words were split into two lists, 25 in each. The words included in each list were randomly selected for each participant. Each list was presented three times in the same order (i.e., all three repetitions of one list followed by three repetitions of the other list) to ensure accuracy would be greater than floor. The buffer lists were only presented one time each. Each word was preceded by a green or gray star and was surrounded by the same-colored box to ensure that participants fully associated the reward category of each word. Participants were told that they could earn a flat fee of up to $15 as a bonus for remembering the words. Participants were directly instructed that "words preceded by a green star are associated with a high reward and words preceded by a gray star are associated with a low reward. Remembering green star words will work towards a $14 bonus at the end of the study and remembering gray star words will work towards a $1 bonus at the end of the study. The more green words you remember, the more green stars you will accumulate and the closer you will be to earning the $14 bonus. The more gray words you remember, the more gray stars you will accumulate and the closer you will be to earning the $1 bonus. All together, you can earn up to a $15 bonus by the end of the study for remembering the words". Rather than incentivizing each individual word or the entire list, we purposefully instructed participants that remembering each word would contribute to earning the bonus of the associated category (green or gray star). Our prior work has shown that this vague threshold is more effective in raising motivation than incentivizing single words with a lower monetary value (Murty, Tompary, Adcock, & Davachi, 2017; Patil et al., 2017). We intentionally did not provide details about the minimum number of recalled words that were necessary to earn the bonus as we wanted focus to be given to all words rather than prioritizing only a certain amount.

In Study 1 (Fig. 1, top), participants learned mixed lists of words, where the high- and low-reward words were intermixed throughout the two lists, with reward order randomized for each participant. In this study, the star cue was presented before each word for 1 s, followed by the presentation of the colored box for 1 s, followed by the word for 3 s with a 1.5 s ITI. Participants in Study 2 learned the words in pure lists (see Fig. 1, bottom), such that all the high-reward words were learned in one list and low-reward words in the other, with list order counterbalanced across participants. Participants were provided details about both reward conditions, exactly as in the instructions above, before learning the first list. Here, the star cue was presented for 2 s before the entire list, followed by the presentation of each word for 3 s with a 1.5 s ITI. Notably, across both studies, the reward value associated with a word was the same upon all 3 repetitions.

2.3. Test

The test phase for both studies was identical. All participants completed a free recall test 24 h after encoding. They were given 5 min to verbally recite as many words as they could remember learning with no penalty for false recalls. Half of the participants in each study also completed this task on day 1 immediately after encoding. The test was completed after learning all words rather than a traditional study-test design in which many short lists would be learned with a test between each. In addition to not being amenable to delay tests, study-test allows participants to adopt different strategies across lists to maximize reward, and we wanted to avoid changes in strategy interfering with the reward manipulation. After the 24-h recall test, participants were given a recognition test in which they were asked to give an old/new judgement to which they could respond "definitely old", "maybe old", "maybe new", or "definitely new", again with no penalty for false alarms. Participants completed 140 trials during the recognition test, 70 of which were old (i.e., presented during encoding) and 70 were new (i.e., novel foils). The test was self-paced.

The unbalanced nature of our design, in which only half of the participants completed the immediate test, was used in case a reward effect would only be captured at one time point, or if recall on day 1 would influence recall on day 2. This design allowed us the ability to look at subsets of the data by group and day to determine the timepoints at which reward influences memory.

The data is available for public access on Open Science Framework (https://osf.io/tf4rd/).

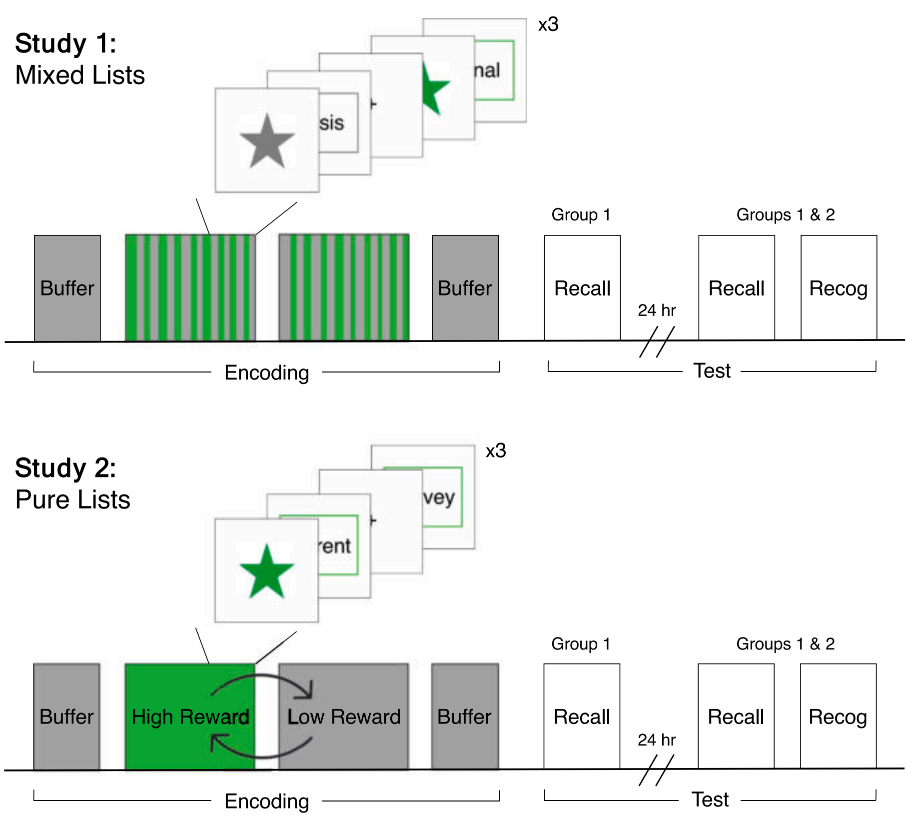

Fig. 1. Task design. Participants in Study 1 learned 2 mixed lists of 25 words each, with high- and low-reward words mixed among both lists. Participants in Study 2 learned 2 pure lists of 25 words each (1 high-reward and 1 low-reward). Buffer words, which were removed from analyses to avoid primacy and recency effects, were learned once, while the lists containing the high- and low-reward target words were repeated three times. Recall was tested for all participants (groups 1 and 2) after 24 h, with a recognition test following. Half of the participants (group 1) also completed the recall test immediately after encoding.

2.4. Analysis

R (version 4.1.2) was used for all analyses, and the lmer package for all multilevel models. Plots and statistical tables were generated with sjPlot in which error bars represent 95% confidence intervals.

First, we wanted to ensure that our findings replicate previous work, which showed enhanced recognition memory for high-reward information. To do this, we calculated corrected recognition scores for each participant by measuring the number of hits (saying "old" to an "old" word) and subtracting the number of false alarms (saying "old" to a "new" word) for each reward condition. We then compared corrected recognition across high- and low-reward conditions using multilevel models with a subject-level random intercept.

We then measured free recall accuracy on a word-level basis by determining whether each encoded word was recalled. The first and last 10 (buffer) words were excluded from this analysis to remove primacy and recency effects. Multilevel models were used to capture the influence of reward and day on recall accuracy. In addition to these main effects of interest, the models included a fixed effect for group and random intercepts describing subject identity and word. We used multilevel models to account for the unbalanced design in which not all participants completed both recall tests, and for differences in memorability across encoded words independent of condition. Effects of reward on recall were captured with the following model comparison: recalled ~ 1 + day + group + (1|subject) + (1|word) versus recalled ~ reward + day + group + (1|subject) + (1|word). Similarly, the effect of delay was measured with: recalled ~ 1 + reward + group + (1|subject) + (1|word) versus recalled ~ day + reward + group + (1|subject) + (1|word). Finally, reward x delay interactions were determined with: recalled ~ reward + day + group + (1|subject) + (1|word) versus recalled ~ reward * day + group + (1|subject) + (1|word).

To establish whether free recall accuracy differed by group at the 24-h delay test, a model comparison with and without an interaction between reward and group (recalled ~ reward + group + (1|subject) + (1|word) versus recalled ~ reward * group + (1|subject) + (1|word)) was conducted. In Study 1, we found a significantly worse model fit by including the reward x group interaction ($\chi^2(1) = 6.26$, $p < .05$), and in Study 2, there was not a significant interaction ($\chi^2(1) = 0.87$, $p = .35$). This suggests that there were no group differences across reward at the 24-h delay test in either study.

For Study 1, we conducted analyses inspired by the Context Maintenance and Retrieval Model (Polyn et al., 2009) to calculate temporal clustering and reward-category clustering. To compute temporal clustering, we calculated a percentile score between the temporal distance of the current and next recall and the temporal distance of the current recall and all words that could have been recalled next. This allows for the comparison of the distance of the current transition versus any transition that could have been made. This measure was calculated across reward condition such that for high-reward clustering, transitions from high-reward words to either a high- or low-reward word were included, and vice versa for low-reward transitions to either high or low. Any trial that included a transition *to* or *from* a buffer word was not included in this analysis.

In a second set of clustering analyses, we asked whether recall was organized by reward category (i.e., transitioning from a high-reward word to another high-reward word, or transitioning from a low-reward word to another low-reward word). To calculate reward clustering, we separately measured the proportion of transitions that stayed within reward condition (i.e., high to high, or low to low) versus switched across reward condition (i.e., high to low, or low to high). The proportion of stay transitions was calculated as the number of transitions that stayed within reward condition (separately for high- and low-reward) divided by the number of transitions in which the participant *could* have stayed within condition. Again, any transition that started or ended with a buffer word was taken out of this analysis.

It is possible that any reward clustering effects could be due to differences in overall recall accuracy across conditions, such that greater clustering of high-reward words appears to occur because they are more likely to be recalled. In other words, greater high-reward clustering would emerge by chance that there are more high-reward items available in memory to transition to. Standard methods of clustering that use shuffling procedures to quantify clustering by chance across conditions (Howard, Youker, & Venkatadass, 2008; Polyn, Erlikhman, & Kahana, 2011) do not account for differences in accuracy by condition, which was confirmed via simulation analyses. Therefore, we used two methods to account for this possibility to ensure that any clustering results were not due to accuracy differences by reward. First, the denominator for the transition probabilities was normalized, such that rather than using the total number of transitions made, we added the total number of *all possible* stay (or switch) transitions that *could* have been made across each trial so that there would be no bias towards recalling a high-reward word next (Fig. 2: Normalized). Second, we estimated dynamic, idiosyncratic levels of chance performance for each participant to account for the baseline levels of clustering that may exist given any accuracy differences by reward (Fig. 2: Corrected). To do this, we first calculated the real transition probabilities in our data by capturing the number of stay (or switch) transitions from each reward condition and divided that by the total number of transitions made from that reward condition (i.e., non-normalized). Then for each transition, we measured the chance of staying within (or switching across) reward condition by calculating the total number of recalled words remaining from each reward condition and divided that by the total number of transitions remaining in the recall. We then took the difference between the real and chance probabilities for each subject to gain a corrected transition probability score.

3. Results

3.1. Study 1

First, we measured recognition accuracy at the subject-level to confirm that our study replicated prior work on reward-memory enhancements (Adcock et al., 2006; Murty & Adcock, 2017). Using the following model comparison: recognition ~ 1 + group + (1|subject) versus recognition ~ reward + group + (1|subject), we found a significant main effect of reward in recognition memory ($\chi^2(1) = 13.52$, $p < .001$) suggesting that people were more likely to correctly recognize high- than low-reward words.

Then, we asked whether free recall accuracy differed by reward and delay. Here, the results showed a significantly better model fit with reward ($\chi^2(1) = 113.75$, $p < .001$; Fig. 3; Table 1) suggesting better memory for items of high versus low value. We found a significantly worse model fit with a term for delay ($\chi^2(1) = 7.79$, $p < .01$), suggesting no difference in memory immediately versus after 24 h. We did not find a significant interaction between reward and delay ($\chi^2(1) = 1.36$, $p = .24$), meaning that the effect of reward on memory does not change whether tested immediately or after 24 h. Finally, we asked whether the probability of recalling an item on day 2 was influenced by the probability of it being recalled on day 1 for participants that completed both tests. This comparison (recalled ~1 + reward + (1|subject) + (1|word) versus recalled ~ previously recalled + reward + (1|subject) + (1|word)) revealed that previously recalling an item at the immediate test significantly influenced its likelihood of being recalling at the 24-h delay test ($\chi^2(1) = 1220.7$, $p < .001$).

Next, we tested how reward influences the use of temporal context to organize memory. Interestingly, we did not find any significant effects of reward ($\chi^2(1) = 0.26$, $p = .61$; Fig. 4; Table 2) or delay ($\chi^2(1) = 0.01$, $p = .94$). There was also not a significant reward x delay interaction ($\chi^2(1) = 0.01$, $p = .93$). For full clarity, although there was no reward x delay interaction, the effect of reward is consistent at both time points (both p's < 0.68). This suggests that the amount one organizes their memory based on temporal context is not influenced by reward value, nor does it change after a delay.

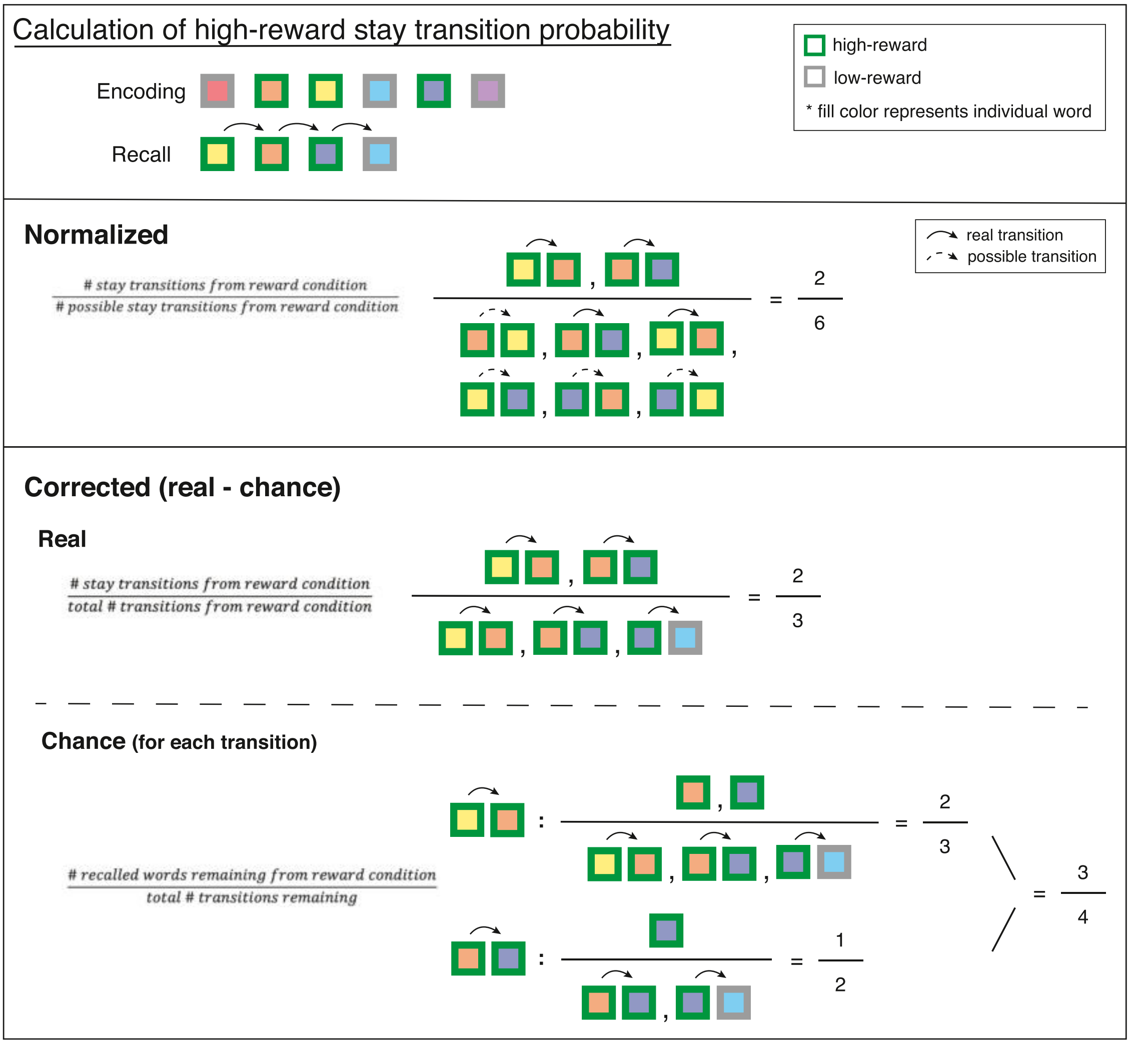

Fig. 2. Example calculation of reward clustering for high-reward stay transitions. Low-reward clustering and switch transition probabilities were calculated in a similar manner. Each colored square represents an individual word, while the outline reflects its reward value. We calculated clustering in two ways to ensure that clustering by reward was not due to the greater availability of high-reward words: a normalized and a corrected score. The normalized transition probability was calculated as the number of stay transitions from the given reward condition divided by the number of *all possible* transitions that could have been made from that reward condition. Second, the corrected clustering score was defined as the difference between real and chance transition probabilities. The real transition probability is non-normalized, where the number of stay transitions from the reward condition was divided by total number of *actual* transitions made from that reward condition. Then the chance transition probability was taken for each transition, such that the number of recalled words remaining from the reward condition that hadn't been recalled yet was divided by the total number of transitions that were remaining.

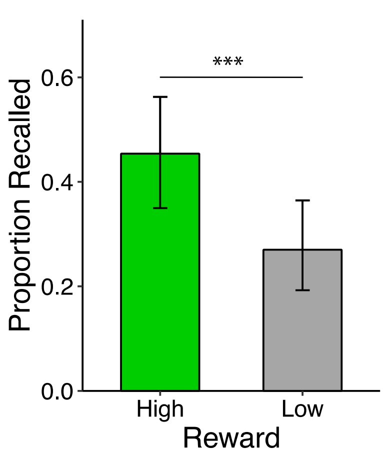

Fig. 3. Free recall accuracy difference by reward. Error bars represent 95% CI. *** $p < .001$.

Finally, our main question of interest was to understand if reward may be considered a higher-order category that can be used to organize memory. To answer this question, we calculated the normalized probability of transitions in which people either stayed or switched reward category (i.e., high-high or low-low versus high-low or low-high, respectively). For both stay and switch transition probabilities, we found significant main effects of reward (stay: $\chi^2(1) = 34.31$, $p < .001$; Fig. 5 and Table 3, left, switch: $\chi^2(1) = 34.30$, $p < .001$; Fig. 5 and Table 3 right), with greater stay among high- and greater switch among low-reward words. This suggests that when a high-reward word had just been recalled, the next word was more likely to be high-reward. Similarly, when a low-reward word had just been recalled, the next recall was again more likely to be high-reward. We found non-significant effects of delay (stay: $\chi^2(1) = 0.09$, $p = .77$, switch: $\chi^2(1) = 0.50$, $p = .48$) and reward x delay interactions (stay: $\chi^2(1) = 1.74$, $p = .19$, switch: $\chi^2(1) = 0.80$, $p = .37$). Again, the stay and switch effects remain significant at both time points individually (all p's < 0.05). Since these calculations were normalized as a function of all words that could have possibly been recalled next, not just the words that were recalled, this

Table 1
Multilevel model comparison measuring reward, delay, and group influences on word-level recall accuracy. Bold *p*-value represents significant effect.

		Recall		
Predictors	*M*	*Odds ratios*	*CI*	*p*
(Intercept)		1.16	0.46–2.97	0.751
Reward		2.25	1.93–2.61	**<0.001**
High	0.42			
Low	0.28			
Day		0.78	0.65–0.93	**0.005**
Immediate	0.42			
Delay	0.31			
Group		0.48	0.27–0.87	**0.015**
1	0.40			
2	0.25			
Random effects				
σ^2	3.29			
τ_{00} subject	1.12			
τ_{00} word	0.29			
ICC	0.30			
$N_{subject}$	56			
N_{word}	50			
Observations	4200			
Marginal R^2 / Conditional R^2	0.067 / 0.347			

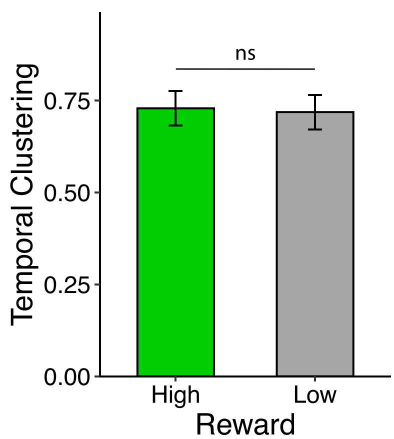

Fig. 4. Temporal clustering by reward. Error bars represent 95% CI. ns $p > .05$.

Table 2
Multilevel model comparison measuring reward, delay, and group influences on temporal clustering. Bold *p*-value represents significant effect.

		Temporal clustering		
Predictors	*M*	*Estimates*	*CI*	*p*
(Intercept)		0.76	0.65–0.87	**<0.001**
Reward		−0.01	−0.05–0.03	0.609
High	0.72			
Low	0.71			
Day		0.00	−0.05–0.05	0.940
Immediate	0.72			
Delay	0.71			
Group		−0.04	−0.11–0.03	0.304
1	0.72			
2	0.69			
Random effects				
σ^2	0.02			
τ_{00} subject	0.01			
ICC	0.33			
$N_{subject}$	54			
Observations	158			
Marginal R^2 / Conditional R^2	0.012 / 0.342			

reward clustering effect supersedes the possibility of clustering simply due to greater availability of high-reward words in memory.

To further confirm that these findings were not due to the reward difference in free recall accuracy, we compared clustering scores against chance performance (calculated individually and dynamically for each participant). We show fewer stay ($\chi^2(1) = 2.94$, $p = .09$) and switch transitions ($\chi^2(1) = 16.35$, $p < .001$) from high-reward words. Although the effect of reward on stay transitions is not as strong here as in the normalized analysis, the pattern remains consistent. There was not a significant effect of delay or reward x delay interaction for stay or switch transitions (all p's > 0.34). While only significant for switch transitions across condition, this is confirmation that value creates a context used to organize memory, above and beyond the explanation that recalling high-reward words together may result from the greater availability of them in memory.

3.2. Study 2

In Study 1, we showed that reward structures memory organization around higher-order contexts rather than temporal features. Prior work has suggested that motivationally-relevant influences on memory may only exist in the context of mixed lists in study-test designs where salient items (e.g., high reward) have increased transient attentional demands since they are intermixed with less important items (Talmi et al., 2019; Talmi, Kavaliauskaite, & Daw, 2021). It is possible that the memory enhancement and categorical organization shown in Study 1 could be a result of either increased local attentional demands due to a “pop-out” effect of salient items, or reward could lead to increased sustained states of motivated encoding. To differentiate between these two possibilities, in Study 2 we characterized reward enhancements on recall accuracy using pure lists, with a test only after all learning was completed, where there would be no bias in transient attentional demands for high-reward items given that all items in each list have the same reward value. Therefore, if we see the same reward accuracy effects as Study 1, this would suggest that the reward-memory enhancement can be sustained for extended periods of time, rather than relying on salience within the local temporal environment.

Similar to recognition memory in Study 1, here, we again found a significant effect of reward ($\chi^2(1) = 3.85$, $p = .05$), suggesting better recognition of high-value items. In the free recall domain, results showed significant effects of reward ($\chi^2(1) = 43.02$, $p < .001$; Fig. 6; Table 4) and delay ($\chi^2(1) = 10.97$, $p < .001$). Together, this suggests better recall memory for items of high value and overall worse memory after a 24-h delay. Finally, we found a non-significant interaction between reward and delay ($\chi^2(1) = 1.95$, $p = .16$), meaning that reward's influence on memory is not delay-dependent. Given that we saw significant effects of reward both in the context of mixed and pure lists, we conclude that this reward-memory benefit can exist in a state of sustained high salience across a list, and not solely when the local environment induces attentional salience changes.

4. CMR simulations

4.1. Simulation approach

Next, we simulated our data from both studies using the Context Maintenance and Retrieval model (Polyn et al., 2009) to understand the potential mechanisms that support our accuracy and reward clustering effects. CMR is a neural network model of memory search that describes how items are encoded and compete to be retrieved. Each item is associated with a context state, which changes over time as more items are encoded or retrieved. The model contains a feature, or item, layer and a temporal context layer. When a new item is presented, the context is updated through a feature-to-context matrix during both encoding and retrieval. During retrieval, a context-to-feature matrix guides memory search in which the current state of the context cues retrieval of

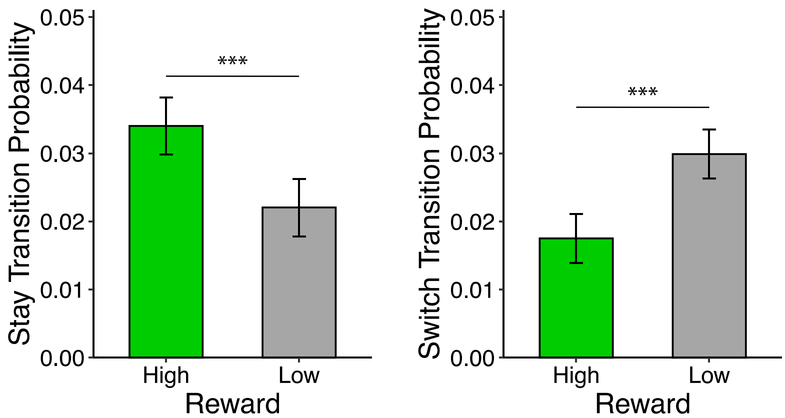

Fig. 5. Stay and switch transition probability by reward. Probability of transitioning within (stay; left) or across (switch; right) reward condition. Error bars represent 95% CI. *** $p < .001$.

Table 3
Multilevel model comparison measuring reward, delay, and group influences on stay and switch transition probabilities. Bold *p*-value represents significant effect.

	Stay transitions				Switch transitions			
Predictors	*M*	*Estimates*	*CI*	*p*	*M*	*Estimates*	*CI*	*p*
(Intercept)		0.04	0.03–0.05	**<0.001**		0.04	0.03–0.05	**<0.001**
Reward		−0.01	−0.02 to −0.01	**<0.001**		-0.01	−0.02 to −0.01	**<0.001**
High	0.03				0.03			
Low	0.02				0.02			
Day		−0.00	−0.01–0.00	0.768		−0.00	−0.01–0.00	0.768
Immediate	0.03				0.03			
Delay	0.03				0.03			
Group		−0.00	−0.01–0.00	0.144		−0.00	−0.01–0.00	0.144
1	0.03				0.03			
2	0.02				0.02			
Random effects								
σ^2	0.00				0.00			
$\tau_{00\ subject}$	0.00				0.00			
ICC	0.30				0.30			
$N_{subject}$	54				54			
Observations	160				160			
Marginal R^2 / Conditional R^2	0.167 / 0.415				0.167 / 0.415			

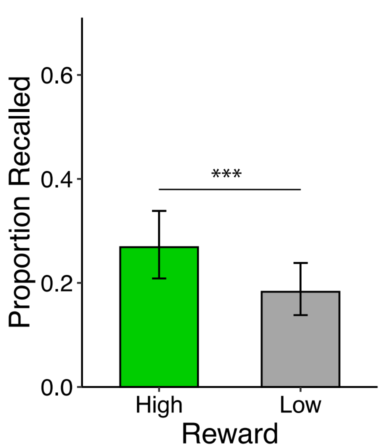

Fig. 6. Free recall accuracy difference by reward. Error bars represent 95% CI. *** $p < .001$.

the next word as items compete to be recalled. The item with the resulting highest activation given the current context will be retrieved. CMR also contains a source layer, which indicates contextual, task, or

Table 4
Multilevel model comparison measuring reward, delay, and group influences on word-level recall accuracy. Bold *p*-value represents significant effect.

	Recall			
Predictors	*M*	*Odds ratios*	*CI*	*p*
(Intercept)		1.37	0.55–3.45	0.499
Reward		0.61	0.53–0.71	**<0.001**
High	0.33			
Low	0.24			
Day		0.75	0.64–0.89	**0.001**
Immediate	0.34			
Delay	0.26			
Group		0.56	0.31–1.01	0.054
1	0.31			
2	0.22			
Random effects				
σ^2	3.29			
τ_{00} subject	1.15			
τ_{00} word	0.23			
ICC	0.30			
$N_{subject}$	59			
N_{word}	50			
Observations	4450			
Marginal R^2 / Conditional R^2	0.039 / 0.323			

state shifts. We used a source layer here to discriminate between high- versus low-reward items (Talmi et al., 2019, 2021).

Model simulations tested the effect that three parameters had on the effects of interest: drift rate, learning rate, and the source layer. Notably, these three measures have been manipulated to account for multiple findings in free recall in prior work (Howard & Kahana, 2002a; Lohnas, Polyn, & Kahana, 2015; Polyn et al., 2009; Rouhani, Norman, Niv, & Bornstein, 2020). Drift rate is defined as the distance by which the temporal context changes, or "drifts", when a new item is presented. An item that enters temporal context with a high drift rate, compared to a low drift rate, triggers greater representational change to the temporal context at encoding. Learning rate determines how strongly an item is bound or associated with the context, including both temporal and source contexts when relevant. A higher learning rate means an item is stamped more strongly into the context. Typically, a variable learning rate would be applied to both the temporal and source layers, but our results showed a similar pattern when learning rate was only modulated by reward within the source layer and kept static in the temporal layer. Therefore, our results will be presented with the latter, simpler model. The source layer, which contains two units, one for each reward category, determines the reward context to which each item belongs.

We evaluated whether manipulating these parameters as contingent on reward (i.e., a higher drift/learning rate for high-reward items and having a source layer or not) would exhibit the same behavioral effects presented above. Defining two separate source categories for the reward conditions is identical to the approach used by Talmi et al. (2019, 2021) in the context of both reward- and emotion-related effects on memory. Given the relatively minimal amount of data available from each participant (i.e., without a study-test design), we were not able to conduct model fitting in which the value of parameters would be determined based on a given participant's data. Instead, the values for the baseline of all parameters (i.e., for low-reward items) were taken from Polyn et al., 2009; see Table 5). The values for the high-reward learning and drift rates were moderately increased from baseline, but not drawn from any previously fit values. Our recall data was collected immediately after encoding as well as after a 24-h delay. To simulate the delay, we decreased the drift rate at recall for the delay test (Talmi et al., 2019). We ran 10,000 simulations for each parameter, as well as all possible combinations, for the immediate and delay tests separately. For each simulation, a random participant's data was selected to inform the model of the reward category for each encoded word. We qualitatively compared each individual parameter and their combinations to our behavioral data to determine which model most closely represented our accuracy and reward clustering findings (Talmi et al., 2019). For each study, we will describe the most plausible model simulation along with the pattern exhibited by each parameter on its own (see Polyn et al., 2009 for model equations).

Table 5
CMR parameters including drift (β) and learning (γ) rates used across all models (left) when the parameters were static or contingent on reward, and when there was a source layer. Encoding parameters were consistent across both delays. Recall parameters were defined separately for the immediate ('imm') and delay ('del') tests. Parameters used in the most plausible models for Study 1 (middle) and Study 2 (right).

Parameter	Value	Study 1 most plausible model	Study 2 most plausible model
		Static	
β_{enc}	0.6	✓	✓
γ_{enc}^{CF}	1		
γ_{enc}^{FC}	0.581	✓	✓
$\beta_{rec,\ imm}$	0.36	✓	✓
$\beta_{rec,\ del}$	0.25	✓	✓
γ_{rec}^{CF}	0	✓	✓
γ_{rec}^{FC}	0	✓	✓
		Reward-contingent	
$\beta_{enc,\ low}$	0.6		
$\beta_{enc,\ high}$	0.8		
$\gamma_{enc,\ low}^{CF}$	1	✓	✓
$\gamma_{enc,\ high}^{CF}$	2	✓	✓
$\beta_{enc,\ low}^{source}$	0.5	✓	
$\beta_{enc,\ high}^{source}$	0.7	✓	
$\gamma_{enc,\ low}^{CF\ source}$	0.1	✓	
$\gamma_{enc,\ high}^{CF\ source}$	0.3	✓	
		Source layer features	
γ_{enc}^{FC}	0.898	✓	
$\beta_{rec,\ imm}^{temp}$	0.51	✓	
$\beta_{rec,\ del}^{temp}$	0.4	✓	
$\beta_{rec,\ imm}^{source}$	0.59	✓	
$\beta_{rec,\ del}^{source}$	0.59	✓	

4.2. Study 1

Since we tested memory immediately and after a 24-h delay, we simulated our data at each time point separately. However, the pattern of effects in the immediate and delay tests were similar, so for simplicity, we will be discussing model simulations across delay conditions, but visualize them separately in our figures.

The model that resulted in a pattern most like our behavioral findings included a reward-dependent learning rate and a source layer that differentiated the two reward categories. This set of parameters showed accuracy (Fig. 7) and reward clustering (Fig. 8) patterns that mirrored the greater probability of recall for high-reward words and greater stay compared to switch transition probability for high-reward words, respectively. The role of the source layer is to create two separate reward contexts through which items of the same value category are associated. These distinct categorical contexts may enhance memory for items in both reward categories because they all have stronger contextual features without distinguishing one category as more goal-relevant than the other. However, the learning rate creates stronger associations between items and their contexts, more so for high-reward words, which serves memory for the goal-relevant items. Since the high-reward items have stronger contextual associations, they are more likely to win the recall competition, leading to greater free recall and clustering rates. Combining learning rate with source creates an additive benefit where the source associations lead to enhanced memory and clustering for each reward category, and the learning rate leads to stronger item-context associations for high-reward words. Overall, we conclude this model to show the strongest and most plausible overlap to our findings above and beyond any parameter by itself or any possible combination of the three parameters with each other. While our manipulation of the parameter space was limited to these three, it is possible that another sub-model could provide a better account of the data. Fitting the model, rather than simulating it as we did here, could reveal a more precise picture; however, this approach was not possible given our design and limited amount of data available.

Next, we detail how manipulating each parameter on its own influenced accuracy and clustering simulations. First, by turning on a dynamic drift rate for rewarded items, we can determine whether increasing the representational distance between learned items in the temporal context influences the degree of recalling and clustering rewarded items. Drift rate appeared to have the least impact on both effects such that greater context change does not lead to a greater likelihood of high-reward items being recalled or clustered. Second, we use the source layer to measure the extent to which creating categorical contexts, which links items of the same reward condition, influences the likelihood of words being retrieved and recalled contiguously. Creating separate reward contexts led to no difference between high- and low-reward accuracy or clustering; however, it equally boosted memory and clustering for both conditions beyond any of the other parameters. Third, we manipulate the learning rate to determine how creating stronger item-context associations influences recall dynamics. The learning rate exhibited a reward difference in accuracy and reward

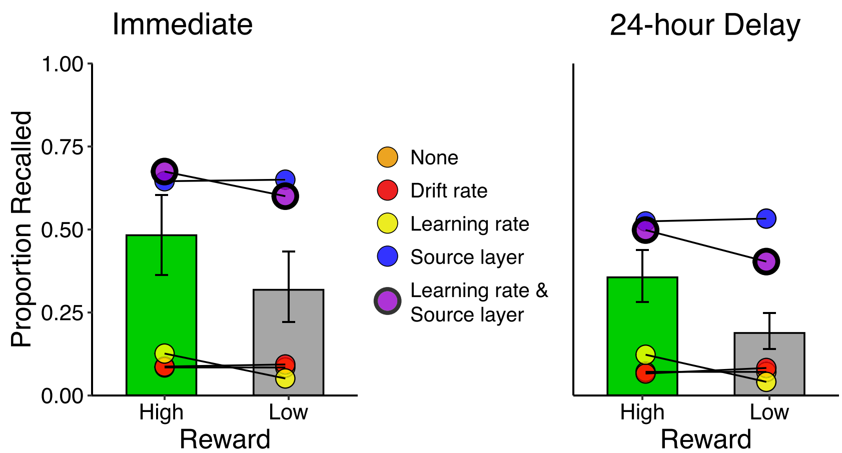

Fig. 7. Free recall accuracy difference by reward simulated with CMR at the immediate (left) and 24-h delay (right) tests. The simulations are overlaid on the behavioral data at each time point. The most plausible model, learning rate and source, is represented by the purple dot with thick black outline. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

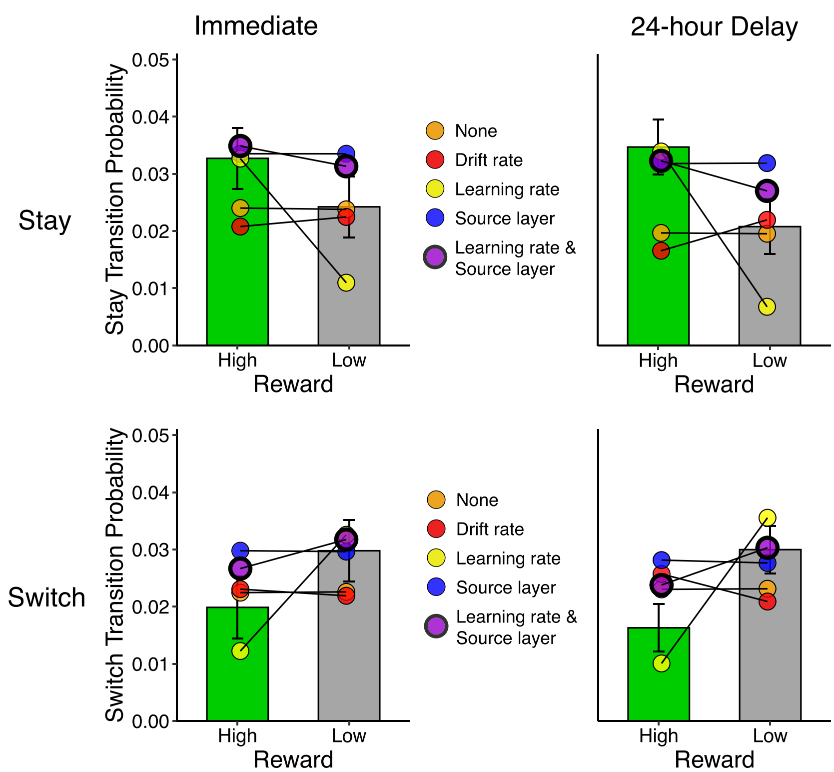

Fig. 8. Stay (top) and switch (bottom) transition probability differences by reward simulated with CMR at the immediate (left) and 24-h delay (right) tests. The simulations are overlaid on the behavioral data at each time point. The most plausible model, learning rate and source, is represented by the purple dot with thick black outline. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

clustering similar to our behavioral findings, but at a lower rate than our data. Creating stronger item-context links for goal-relevant items may cue other high-reward items that also share strong item-context associations, leading to greater probability of recall and clustering. Finally, in testing each combination of parameters, with a main interest in combining a dynamic learning rate and source layer for rewarded items, we conclude that learning rate and source together are the most reasonable features to explain our data by more strongly associating the high-reward items with their context and creating a categorical context for high- and low-reward words.

4.3. Study 2

We simulated CMR in Study 2 in the same manner as Study 1 at each time point separately. Each parameter on its own showed a similar pattern to Study 1 (Fig. 9). First, the drift rate again did not result in better memory for high-value items, and in fact showed the opposite effect, potentially because representing low-reward items closer in time when learned in one list increases their associations with one another. The source layer also enhanced memory, but specifically for low-reward words. This may again be likely because of the lower drift rate for low-reward words making them more similar within category. Similar to Study 1, learning rate was crucial in supporting reward's enhancement

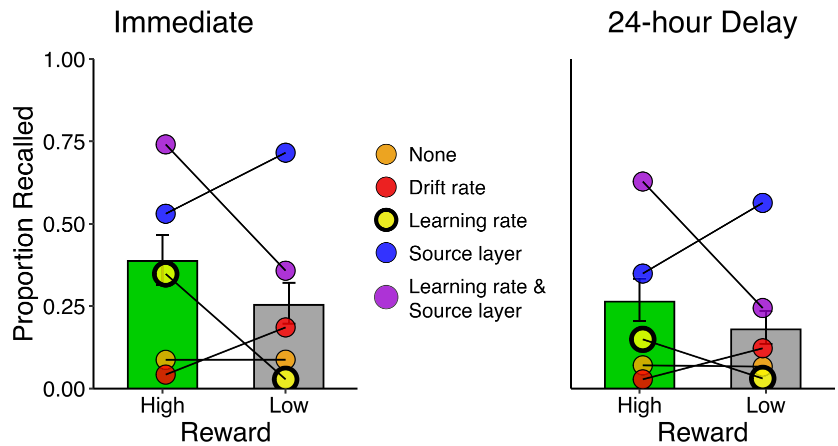

Fig. 9. Free recall accuracy difference by reward simulated with CMR at the immediate (left) and 24-h delay (right) tests. The simulations are overlaid on the behavioral data at each time point. The most plausible model, learning rate, is represented by the yellow dot with thick black outline. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

on recall accuracy, but interestingly was enough to show this effect and did not need the additional support of the source layer. The nature of the design in Study 2 intrinsically groups words by reward category and it is therefore unnecessary for participants to use this higher-order type of organization when the design already does so. Thus, creating stronger item-context associations for valuable information is enough to support memory when the learning environment already contains a higher-order category structure. Again, since our simulations tested a constrained search of the parameter space, it is possible that a better sub-model could be revealed with model fitting.

5. Discussion

Across two studies, we investigated how reward influences the organization of memories, delineating roles for temporal context and de novo higher-order reward contexts, as well as putative mechanistic underpinnings using model simulations. In Study 1, we found that reward did not influence the extent to which participants clustered their recall around temporal context, but instead created a category for high-reward items where the probability of transitioning within category was greater for high- than low-reward words. Given that there was greater clustering by high reward, we posit that although a reward category is created for both high- and low-value information, the stronger binding of high-reward words to their source context leads to the difference in reward-related organization. In Study 2, we showed that reward enhances memory in pure lists as it does for mixed lists, suggesting that sustained salience, rather than just attentional changes, supports reward-memory enhancements.

To bolster and extend our empirical findings, we simulated these results using a neural network model of free recall, a reward-mediated variant of CMR. This approach allowed us to determine that in Study 1, the most plausible account of our accuracy effects came from a model that included a biased learning rate for high-reward items and a separate source layer for each reward category. This combination of parameters modeled value-based contextual associations through the source layer by enhancing memory for both high- and low-value information, while learning rate enhanced memory specifically for high-reward words by creating stronger item-context associations. These findings dovetail with a recently developed variant of CMR (eCMR) tailored to model emotional memory, and has been replicated in the reward domain (Talmi et al., 2019, 2021). eCMR demonstrated that emotional-memory enhancements are driven by an emotion-contingent learning rate and source layer, suggesting parallelism between the mechanistic underpinnings of value- and emotion-related memory. Simulations in Study 2 showed that when a higher-order categorical structure is already established in the learning environment, learning rate alone is enough to support increased memory for high-value information. These two studies reveal different potential mechanistic underpinnings of reward-memory benefits, but this is not particularly surprising given their differing goals and designs. In Study 2, temporal and reward contexts were confounded so that we could examine memory accuracy in the context of pure lists as opposed to mixed lists in Study 1. Moreover, given that we only manipulated a small portion of the parameter space in CMR, it is possible that with more data available to fit the model as opposed to simulating it, a different sub-model could provide a more accurate depiction of the data. However, our simulations provide a proof of concept that item-context and contextual associations underly reward-related memory enhancements and organization.

Our CMR simulations showed very similar patterns among simulations of reward clustering. By creating stronger item-context associations through the learning rate, and defining separate reward contexts through the source layer, high-reward items were more likely to be recalled contiguously. Within the source layer, there is a separate learning rate for the temporal and source contexts. It is typical to manipulate the learning rate in both contexts, but we interestingly found the same patterns when only modulating the source context (presented above). Although we cannot directly equate the temporal and source contexts because they contain a different number of units, it is interesting to note that this model with only a source learning rate was sufficient to replicate the effects we saw in the full model. While we are not able to make direct comparisons, this presents some evidence that categorical features of memory may be targeted by reward to drive accuracy and clustering enhancements.

Together, these patterns support the notion that reward may not influence memory solely by strengthening the binding of items to their temporal features as events unfold. Rather, our findings suggest that individuals construct de novo categories around value, and that episodic information is bound to these contexts to support more adaptive forms of organization. While in neutral contexts it is typical to recall information in a similar temporal order to that which it was learned, under a motivational context such as reward, this may not be the most adaptive feature to target. Instead, it may be more adaptive to note important relationships that exist between items, such as how related they are to one another based on their relative importance to a current or future goal. However, future work tapping into how memory organization by reward value relates to future adaptive decision-making is necessary to test these hypotheses.

The creation of de novo categories around which memories can be organized is not a novel idea, and in fact has been discussed in length by Barsalou (for example, see Barsalou, 1983). Barsalou discussed how new categories of information can be created (e.g., "items to sell at a garage sale" or "things to take on a camping trip") but that these differ from common categories (e.g., "fruits" or "mammals") due to their necessity in certain contexts. Common categories, which are well established in memory, allow for associations to be formed during encoding and retrieval (Barsalou, 1983). When existing common categories are not specific enough to help achieve a goal, it would be necessary to create a new one (Barsalou, 1991), such as determining what to bring when preparing for a camping trip, or when instructed to remember certain items to earn a monetary bonus. Creating categories that link items of like value allows them to be better remembered, leading to the receipt of the goal.

These behavioral findings provide some insight into the underlying neural mechanisms driving differences in memory organization, when considering them in conjunction with prior neuroimaging work. Reward's influence on memory is known to be largely supported by hippocampal involvement in concordance with the VTA (Adcock et al., 2006; Murty & Adcock, 2017). It is widely accepted that the HPC supports episodic memory and item-context binding (Davachi, 2006). However, recent evidence suggests that both of these regions may be more involved in conceptual knowledge than previously thought (Cowan, Fain, et al., 2021; Morton et al., 2021). Given the VTA and HPC's involvement in reward and conceptual aspects of memory, we posit that reward may facilitate higher-order memory, the functions of which could extend to the domain of reward and value generalization. In fact, in other domains, hippocampal engagement has been shown to support the generalization of reward information across related items (Kahnt, Park, Burke, & Tobler, 2012; Wimmer, Daw, & Shohamy, 2012; Wimmer & Shohamy, 2012), a process which has also been shown behaviorally where reward generalizes across semantic categories to other category members that were retroactively tagged as important (Patil et al., 2017). Relatedly, HPC-VTA interactions have been shown to support retroactive generalization of valuable information, however in the context of threat (Clewett, Dunsmoor, Bachman, Phelps, & Davachi, 2020). Given this work, in addition to our findings here, we propose that reward targets higher-order features of memory in service of providing a more adaptive form of memory organization.

In a parallel question, we explored the nature of the mechanisms which may be driving reward-memory effects, namely whether these effects exist in states of both high- and low- transition frequencies between salient and non-salient items (Talmi et al., 2019). We distinguish between these two mechanisms by measuring reward's influence on memory in pure and mixed lists. Within this framework, an enhancement of memory only in mixed lists would suggest reward-memory effects may be due to high transitions between salient items, whereas enhanced memory in pure lists would suggest a mechanism present in states of sustained salience. In both Study 1, which used a mixed list paradigm, and Study 2, which used pure lists, we found significant reward accuracy effects in which high-reward words were more likely to be later recognized and recalled. This suggests that reward-memory effects may not only be due to salience in local temporal environments, but rather may extend across periods of time, which dovetails well with models of VTA-HPC circuits guiding periods of tonic motivational engagement (Murty & Adcock, 2017; Shohamy & Adcock, 2010). We take this to indicate that in salient states, regardless of the transition frequency, reward-memory enhancements and organization of memory by reward category are likely supported by creating stronger item-context and categorical associations for valuable information.

While our goal was to understand how reward motivation shapes memory organization, there are several factors that are important to disentangle from our effects. First, it is possible that these effects of reward could be explained by lower-level processes such as allocating greater attention, rehearsal, or strategy use to high-value information. Previous work has shown that similar attentional processes are dedicated to both high- and low-value information measured by equivalent fixations to both sets of information (Ariel & Castel, 2014). Similarly, individual differences in working memory capacity and strategy use (Elliott, McClure, & Brewer, 2020) as well as more rehearsal of high-reward information (Stefanidi et al., 2018) are not related to value-directed remembering. All of this is taken to suggest that low-level processing does not drive reward's effect on memory. Instead, the encoding of valuable stimuli seem to engage deeper encoding strategies (Elliott & Brewer, 2019), relying on frontotemporal semantic networks (Cohen, Rissman, Suthana, Castel, & Knowlton, 2014), suggesting the necessity of higher-level processing. Further, another potential explanation of these effects is that participants explicitly remember the value associated with the given words. However, much evidence has shown that people have low source memory for reward value, which suggests that this is an automatic encoding process and does not rely on explicit knowledge of the words' reward category. While source memory would likely rely on higher-order processes, this evidence is contradictory and therefore does not reveal a clear mechanism. Thus, further work is necessary to fully understand the potential processes at play.

Additionally, given the large literature on consolidation-dependent reward memory (for example, see Murayama & Kitagami, 2014; Murayama & Kuhbandner, 2011; Wang & Morris, 2010), it was expected that consolidation may play a role in reward-related memory organization such that this higher-order form of organization would only appear after a delay. However, evidence from a recent study showing that some features of memory (i.e., schema use) are constructed during encoding but may not express themselves until after a delay (Tompary, Zhou, & Davachi, 2020), supports our findings that an organizational structure may occur immediately. It would be important for future research to tease apart the circumstances under which reward's influence on memory is and is not consolidation-dependent.

In summary, we have shown enhanced memory for items of high value, effects which can exist in sustained states of reward. We found that reward does not influence memory through temporal features during reward-motivated encoding, but rather organizes memory by de novo reward information that provides an associative link in memory for items deemed valuable. While we have investigated these aspects of memory under states of reward motivation, these effects may also be true for goal-relevant information or in the context of agency, which engages similar neural circuitry (Murty, DuBrow, & Davachi, 2015). Further research is needed to understand whether these features extend into other motivational domains. Taken together, as opposed to remembering all of the details of an event in order, it may be more adaptive to target and link together important information so that it may be more accessible when working towards a goal.

Data availability

The data repository is included in the manuscript.

Acknowledgement

This work was supported by grant R21 DA043568 from the National Institute of Health (V.P.M.).

References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron, 50*(3), 507–517. https://doi.org/10.1016/j.neuron.2006.03.036

Ariel, R., & Castel, A. D. (2014). Eyes wide open: Enhanced pupil dilation when selectively studying important information. *Experimental Brain Research, 232*(1), 337–344. https://doi.org/10.1007/s00221-013-3744-5

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition, 11*(3), 211–227. https://doi.org/10.3758/BF03196968

Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (pp. 1–64). San Diego: Academic Press.

Clewett, D., Dunsmoor, J., Bachman, S., Phelps, E., & Davachi, L. (2020). *Survival of the salient: Emotion rescues otherwise forgettable memories via neural reactivation and post-*

encoding hippocampal connectivity [Preprint]. *Neuroscience.*. https://doi.org/10.1101/2020.07.07.192252

Cohen, M. S., Rissman, J., Suthana, N. A., Castel, A. D., & Knowlton, B. J. (2014). Value-based modulation of memory encoding involves strategic engagement of fronto-temporal semantic processing regions. *Cognitive, Affective, & Behavioral Neuroscience, 14*(2), 578–592. https://doi.org/10.3758/s13415-014-0275-x

Cowan, E. T., Fain, M., O'Shea, I., Ellman, L. M., & Murty, V. P. (2021). VTA and anterior Hippocampus target dissociable neocortical networks for post-novelty enhancements. *The Journal of Neuroscience, 41*(38), 8040–8050. https://doi.org/10.1523/JNEUROSCI.0316-21.2021

Cowan, E. T., Schapiro, A. C., Dunsmoor, J. E., & Murty, V. P. (2021). Memory consolidation as an adaptive process. *Psychonomic Bulletin & Review, 28*(6), 1796–1810. https://doi.org/10.3758/s13423-021-01978-x

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology, 16*(6), 693–700. https://doi.org/10.1016/j.conb.2006.10.012

Elliott, B. L., & Brewer, G. A. (2019). Divided attention selectively impairs value-directed encoding. *Collabra: Psychology, 5*(1), 4. https://doi.org/10.1525/collabra.156

Elliott, B. L., McClure, S. M., & Brewer, G. A. (2020). Individual differences in value-directed remembering. *Cognition, 201*, Article 104275. https://doi.org/10.1016/j.cognition.2020.104275

Friendly, M., Franklin, P. E., Hoffman, D., & Rubin, D. C. (1982). The Toronto word Pool: Norms for imagery, concreteness, orthographic variables, and grammatical usage for 1,080 words. *Behavior Research Methods & Instrumentation, 14*(4), 375–399. https://doi.org/10.3758/BF03203275

Howard, M. W., & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of Mathematical Psychology, 46*(3), 269–299. https://doi.org/10.1006/jmps.2001.1388

Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language, 46*(1), 85–98. https://doi.org/10.1006/jmla.2001.2798

Howard, M. W., Youker, T. E., & Venkatadass, V. S. (2008). The persistence of memory: Contiguity effects across hundreds of seconds. *Psychonomic Bulletin & Review, 15*(1), 58–63. https://doi.org/10.3758/PBR.15.1.58

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition, 24*(1), 103–109. https://doi.org/10.3758/BF03197276

Kahana, M. J., Howard, M. W., Zaromb, F., & Wingfield, A. (2002). Age dissociates recency and lag recency effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(3), 530–540. https://doi.org/10.1037/0278-7393.28.3.530

Kahnt, T., Park, S. Q., Burke, C. J., & Tobler, P. N. (2012). How glitter relates to gold: Similarity-dependent reward prediction errors in the human striatum. *Journal of Neuroscience, 32*(46), 16521–16529. https://doi.org/10.1523/JNEUROSCI.2383-12.2012

Lohnas, L. J., Polyn, S. M., & Kahana, M. J. (2015). Expanding the scope of memory search: Modeling intralist and interlist effects in free recall. *Psychological Review, 122*(2), 337–363. https://doi.org/10.1037/a0039036

Manning, J. R., Polyn, S. M., Baltuch, G. H., Litt, B., & Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences of the United States of America, 108*(31), 12893–12897. https://doi.org/10.1073/pnas.1015174108

Miendlarzewska, E. A., Bavelier, D., & Schwartz, S. (2016). Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews, 61*, 156–176. https://doi.org/10.1016/j.neubiorev.2015.11.015

Morton, N. W., Zippi, E. L., Noh, S., & Preston, A. R. (2021). Semantic knowledge of famous people and places is represented in hippocampus and distinct cortical networks. *The Journal of Neuroscience, 41*(12), 2762–2779.

Murayama, K., & Kitagami, S. (2014). Consolidation power of extrinsic rewards: Reward cues enhance long-term memory for irrelevant past events. *Journal of Experimental Psychology: General, 143*(1), 15–20. https://doi.org/10.1037/a0031992

Murayama, K., & Kuhbandner, C. (2011). Money enhances memory consolidation—But only for boring material. *Cognition, 119*(1), 120–124. https://doi.org/10.1016/j.cognition.2011.01.001

Murdock, B. B., & Walker, K. D. (1969). Modality effects in free recall. *Journal of Verbal Learning and Verbal Behavior, 8*(5), 665–676. https://doi.org/10.1016/S0022-5371(69)80120-9

Murphy, D. H., & Castel, A. D. (2022). The role of attention and ageing in the retrieval dynamics of value-directed remembering. *Quarterly Journal of Experimental Psychology, 75*(5), 954–968. https://doi.org/10.1177/17470218211046612

Murphy, D. H., Schwartz, S. T., & Castel, A. D. (2022). Serial and strategic memory processes in goal-directed selective remembering. *Cognition, 225*, Article 105178. https://doi.org/10.1016/j.cognition.2022.105178

Murty, V. P., & Adcock, R. A. (2017). Motivated memory: Anticipated reward and punishment shape encoding via differential medial temporal network recruitment. In D. E. Hannula, & M. C. Duff (Eds.), *The Hippocampus from cells to systems: Structure, connectivity, and functional contributions to memory and flexible cognition* (pp. 467–501). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-50406-3.

Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *Journal of Neuroscience, 35*(16), 6255–6264. https://doi.org/10.1523/JNEUROSCI.4181-14.2015

Murty, V. P., Tompary, A., Adcock, R. A., & Davachi, L. (2017). Selectivity in Postencoding connectivity with high-level visual cortex is associated with reward-motivated memory. *The Journal of Neuroscience, 37*(3), 537–545. https://doi.org/10.1523/JNEUROSCI.4032-15.2016

Patil, A., Murty, V. P., Dunsmoor, J. E., Phelps, E. A., & Davachi, L. (2017). Reward retroactively enhances memory consolidation for related items. *Learning and Memory, 24*(1), 65–69. https://doi.org/10.1101/lm.042978.116

Polyn, S. M., Erlikhman, G., & Kahana, M. J. (2011). Semantic cuing and the scale insensitivity of recency and contiguity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 37*(3), 766–775. https://doi.org/10.1037/a0022475

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organization. *Psychology Review, 116*(1), 129–156. https://doi.org/10.1037/a0014420.A

Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create event boundaries in memory. *Cognition, 203*(104269).

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences, 14*(10), 464–472. https://doi.org/10.1016/j.tics.2010.08.002

Stefanidi, A., Ellis, D. M., & Brewer, G. A. (2018). Free recall dynamics in value-directed remembering. *Journal of Memory and Language, 100*, 18–31. https://doi.org/10.1016/j.jml.2017.11.004

Talmi, D., Kavaliauskaite, D., & Daw, N. D. (2021). In for a penny, in for a pound: Examining motivated memory through the lens of retrieved context models. *Learning & Memory, 28*(12), 445–456. https://doi.org/10.1101/lm.053470.121

Talmi, D., Lohnas, L. J., & Daw, N. D. (2019). A retrieved context model of the emotional modulation of memory. *Psychological Review, 126*(4), 455–485. https://doi.org/10.1037/rev0000132

Talmi, D., Luk, B. T. C., McGarry, L. M., & Moscovitch, M. (2007). The contribution of relatedness and distinctiveness to emotionally-enhanced memory. *Journal of Memory and Language, 56*(4), 555–574. https://doi.org/10.1016/j.jml.2007.01.002

Tompary, A., Zhou, W. X., & Davachi, L. (2020). Schematic memories develop quickly, but are not expressed unless necessary. *Scientific Reports, 10*(1), 1–17. https://doi.org/10.1038/s41598-020-73952-x

Wang, S. H., & Morris, R. G. M. (2010). Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. *Annual Review of Psychology, 61*(February), 49–79. https://doi.org/10.1146/annurev.psych.093008.100523

Wimmer, G. E., Daw, N. D., & Shohamy, D. (2012). Generalization of value in reinforcement learning by humans: Generalization of value. *European Journal of Neuroscience, 35*(7), 1092–1104. https://doi.org/10.1111/j.1460-9568.2012.08017.x

Wimmer, G. E., & Shohamy, D. (2012). Preference by association: How memory mechanisms in the Hippocampus Bias decisions. *Science, 338*(6104), 270–273. https://doi.org/10.1126/science.1223252