In for a pound, in for a penny: How the opportunity to gain reward influences the competition for memory resources

Deborah Talmi[a], Deimante Kavaliauskaite[a], Nathaniel D. Daw[b]

[a]Division of Neuroscience and experimental psychology, University of Manchester

[b]Princeton Neuroscience Institute and Department of Psychology, Princeton University

Address for correspondence

Deborah Talmi, Division of neuroscience and experimental psychology, School of Biological Sciences, University of Manchester, Manchester, UK, M139PL.

Telephone: 0161 2751968 Email: Deborah.Talmi@manchester.ac.uk

Abstract

When people encounter items that they believe will help them gain reward, they later remember them better than those that do not. While it is adaptive to preferentially remember experiences that will be useful later, it is unknown how the competition for memory resources is implemented in time, through the processes of encoding, consolidation, and retrieval. In two experiments we promised participants £1 for remembering some pictures, but only 10 pence for remembering others. Their ability to describe the pictures was tested after one minute and after 24 hours. Memory at immediate test showed effects of list composition, suggesting local competition at encoding and/or retrieval. These results are consistent with our recently-proposed emotional Context Maintenance and Retrieval model, supporting it as a general account of motivated memory. In contrast, relative to this baseline, more valuable memories were not preferentially retained following delay, suggesting no detectable role of competition for consolidation.

Introduction

It is paramount to remember information that can help us attain our goals and increase our subjective utility. Because memory resources are limited, it is adaptive for the brain to prioritise experiences that can be leveraged to maximise utility over less-important experiences (Gershman & Daw, 2017). This intuition was expressed by Sherlock Holmes, Arthur Conan Doyle's timeless protagonist, when (justifying his ignorance of astronomy) he remarked that "It is of the highest importance… not to have useless facts elbowing out the useful ones". Holmes' intuition agrees with rational considerations, which suggest that if memory resources are limited, it is best to hold on to the most valuable information, compared globally against all other information (Anderson & Schooler, 1991). But how such competition plays out in practice is not so simple, for at least two reasons. First, experiences are encountered sequentially rather than all at once, meaning that a rational agent must decide which memories to retain without knowing the full set of competitors. Second, Holmes' picture of the mind as a storeroom of fixed, limited capacity is grossly oversimplified: memories can compete, interfere with each other and promote each other at many stages of encoding, consolidation, and retrieval. For a rational analysis it's not clear to what extent each of these stages represents a meaningful resource bottleneck. Both of these considerations suggest that competition may occur locally, only among items that are close to each other on some contextual dimension, such as time (the preceding or following items) or space (various objects in the same scene). Our research question, therefore, concerns the scope of the competition of reward-predicting information for memory resources.

To answer this question it is useful to consider how the opportunity to maximise utility influences the cognitive mechanisms that underlie successful recall. In an experimental setting where remembering some items earn more money or ‘points’, high-value items are better attended and remembered (e.g. Ariel & Castel, 2014; Lee, Greening, & Mather, 2015; Wittmann et al., 2005), although some observe memory facilitation only in delayed tests (e.g. Montagrin, Brosch, & Sander, 2013; Spaniol, Schain, & Bowen, 2014). The mechanistic relationship between reward anticipation, attention, and memory has not been worked out (Lisman, Grace, & Duzel, 2011). Because a signal that reward may be available is goal-relevant and goal-congruent is “emotional” by definition (Sander, Grandjean, & Scherer, 2005), it may be useful to refer to the emotional memory literature to aid this investigation. The emotional Context Maintenance and Retrieval model (eCMR), a variant of retrieved-context models that describes the influence of emotional value on free recall (Talmi, Lohnas, & Daw, in press), suggests that the extra attention participants allocate to high-value items binds items more tightly to their encoding context. When asked to freely recall these items later this tighter binding enhances the competitive advantage of high-value items against low-value items, but makes it difficult to link them to a new context later on (Madan, Fujiwara, Gerson, & Caplan, 2012). This account assumes that the temporal context narrows the memory search during recall to focus solely on the study set (Lohnas, Polyn, & Kahana, 2015), such that high-reward items only compete with items that were studied in the same set as them. The implication is that the competition for retrieval resources is local (intra-list) rather than global (inter-list).

There is evidence for local competition for recall from experiments with emotionally-arousing words, sentences, and pictures: all items that are known to capture attention preferentially, and which might heuristically be interpreted as fitness-relevant. In these experiments memory for emotionally-arousing items is typically increased compared to memory for neutral ones when emotional and neutral items are presented and recalled together, in 'mixed' lists (Barnacle, Tsivilis, Schaefer, & Talmi, 2017; Hadley & MacKay, 2006; MacKay et al., 2004; Talmi, Luk, McGarry, & Moscovitch, 2007), or when the emotional item is the 'oddball' among neutral ones (Hurlemann et al., 2005; Mather & Knight, 2009). The recall advantage of emotional items is less pronounced and sometimes disappears altogether when emotional and neutral items are presented and recalled separately, in 'pure' lists. eCMR simulates these situations successfully (Talmi, Lohnas, & Daw, in press). There is also evidence that recalling items that are not personally valued, but attract attention because they are bizarre, is enhanced because of local completion during retrieval (McDaniel, Dornburg, & Guynn, 2005; Schmidt, 1991). It is unknown, at present, whether these accounts generalise to information that predicts reward directly and explicitly. To the extent that competition for memory resources is global in nature, high-value items should out-compete low-value items regardless of the local context of their study set. Indeed, it would be irrational for a fully informed observer to allocate limited resources to low-value items in detriment to high-value items.

Here we test this prediction by examining memory for stimuli whose later recall explicitly predicts high and low reward, encoded in pure and mixed lists. In two experiments, our participants viewed lists of neutral pictures that were presented framed or unframed. They knew that they could gain larger monetary reward by recalling framed items on later test. The critical manipulation was the arrangement of framed and unframed items in the list, such that some lists included both item types (mixed) and some only one (pure). The frame provides a clear signal of which individual stimulus should be prioritised (Mather, Clewett, Sakaki, & Harley, 2016; Mather & Sutherland, 2011). Importantly, in the global context of the experiment – and, indeed, the participant's day outside the laboratory – remembering framed pictures promised more reward than average, certainly more reward than unframed pictures. If the competition for memory resources is a global, inter-list phenomenon, rational participants should attend the framed pictures preferentially regardless of the list in which they were presented and recall should mirror this order of priority. By contrast, if the competition for recall is local, as predicted by eCMR (Simulation 2 in Talmi et al., 2018), the recall advantage for high- over low-value items will be greater in mixed lists, a pattern called a 'list-composition effect'. Confirming this prediction would support eCMR as a general model of motivated memory, across very different operationalisations of value.

Theories of modulated consolidation ascribe enhanced delayed memory for items that are highly emotional or that are associated with reward to competition for resources during memory consolidation (for review, see Lisman et al., 2011; see also Patil, Murty, Dunsmoor, Phelps, & Davachi, 2017). Crucially, modulated-consolidation theories hold that high-value items will be forgotten more slowly holds irrespective of the list in which they were presented. Despite decade-long interest in list composition effects, to our knowledge this effect has never been tested with a prolonged retention interval.

In summary, experiments 1-2 tested memory immediately to assess whether the list-composition effect will be observed. Experiment 2 added a delayed test to test the hypothesis that reward attenuates the forgetting rate regardless of the local context.

Experiment 1

Methods

Participants

The sample size was selected using Gpower to give us 95% power to detect a large effect size, as observed in our previous work with the emotional list composition task (e.g. Barnacle et al., 2017), in a two-tailed paired t-test. Participants in Experiment 1 (N=29, 28 females, 1 male, age range 18-21) were all students, and received course credits for participation. To be eligible participants were required to have normal or corrected-to-normal vision. Participants could not take part if they were between 18-35 years old, had past or current neurological or psychiatric problem, taking centrally-acting medications, or any non-prescription recreational drugs. The project gained ethical approval from the University of Manchester Research Ethics Committee.

Materials

Stimuli consisted of pictures retrieved from freely-available databases, adjusted in size to 280x210 pixels. In addition, the experiments used 9 printed sheets with 36 easy arithmetic subtraction and addition problems for the distractor task, with one sheet assigned randomly to each block. The experiment was implemented using Cogent200 on a MATLAB platform. 48 pictures of clothing and 48 pictures of stationary were used. Pictures were retrieved from the Bank of Standardized Stimuli (Brodeur, Dionne-Dostie, Montreuil, & Lepage, 2010; Brodeur, Guérard, & Bouras, 2014) and from Food-Pics (Blechert, Meule, Busch, & Ohla, 2014; scores for pictures drawn from the Food-pic database were converted from a 1-100 scale to the BOSS 1-5 scale). Pictures from the two categories were matched for familiarity (clothing: M=4.30, SD=.32; stationary: M=4.39, SD=.35) and complexity (clothing: M=2.29, SD=.37; stationary: M=2.30, SD=.49). Each list included 8 pictures from each category. In the pure list condition they were all either framed with a grey square, or unframed. In the mixed list condition, 4 pictures of each type were framed and 4 unframed. The assignment of pictures to conditions was randomised for every participant.

Procedure

Participants were tested individually, in a quiet room, in the presence of a single experimenter. After providing informed consent, the experiment began with instructions and practice and ended with reward delivery. It comprised of 6 blocks. Each block included a picture task, a distractor task, and a free recall test. These are described below. The allocation of picture lists to the picture task in each block, the order of pictures in each list, and the order of blocks were randomised for each participant. The experiment included 2 blocks of pure lists where all pictures were framed (pure lists, high reward), 2 blocks where none of the pictures were framed (pure lists, low reward), and 2 blocks where half the pictures were framed and the other half unframed (mixed lists).

Picture task. In each block participants viewed 16 pictures which could all be framed, all unframed, or half framed and half unframed. Each picture was individually presented at the centre of the screen on a white background. It was presented for 2s followed by a blank screen for 3.5, 4 or 4.5s (randomised ISI) before the next picture was presented.

Distractor task. Participants were given 60s to complete as many arithmetic questions as possible.

Free recall test. Participants were given a few minutes to describe in writing each of the pictures they remembered (3 minutes in Experiment 1, 4 minutes in Experiment 2). They were prompted to keep trying if they stopped before the time. A beep signalled the start and end of the memory task.

Instructions. Before the experiment began participants were given instructions about each of the tasks. For the picture task, they were asked to view the pictures and try to commit them to memory. They were told that they will earn £1 for recalling framed pictures, and 10 pence for recalling unframed pictures, but not given any information about the proportion of framed pictures. For the Distractor task, they were informed that to qualify for reward they must (1) complete at least eight questions (2) those two that are randomly selected for checking must be completed accurately. For the free recall test, participants were asked to describe each of the pictures they could recall by writing three details about the picture, separated by a comma; e.g. “A plant, spikey leaves, in round pot”. Participants were explicitly told that their goal was to maximise their monetary reward.

Practice. After the instructions were delivered, participants practiced each of the tasks using a set of 16 additional pictures (3 framed). During the practice block a lower reward rate was in effect (10 pence vs. 1 penny). At the end of the practice block participants were given feedback about their performance and paid any money they earned.

Reward delivery. At the end of the session, one block was selected randomly, and participants who qualified for reward, based on their performance on the distractor task, were paid according to their free recall performance. They were debriefed and thanked.

Results

Picture recall was scored as 'correct' if participants recalled at least 3 correct details about a picture. The proportion of correct recalls was calculated for each condition (out of 16 in the two pure list conditions, and out of 8 in the mixed list condition).

The statistical analysis of average recall follows exactly on the analytic approach we have used in previous work. That work led to the prediction that we would find an interaction between list type and reward. Therefore, the four comparisons that ensue – within lists (high vs. low reward, in pure and mixed lists separately) and across lists (high and low reward in mixed vs. pure lists) are all planned contrasts. To be conservative we corrected these for familywise error by using a Bonferroni-corrected p-value threshold of $p<.0125$.

The average free recall data are depicted in Figure 1. They were analysed with a 2 (list: pure, mixed) x 2 (reward: high, low) repeated-measures ANOVA. The effect of list type was not significant, but the interaction with reward was significant and large, $F(1,28)=20.09$, $p<.001$, $\eta p2 =.42$, qualifying the significant main effect of reward, $F(1,28)=10.71$, $p<.001$, $\eta p2 =.28$. These results reveal a reward-dependent list-composition effect. Planned paired t-tests showed that the effect of reward was only significant in the mixed list condition, t(28)=4.38, p<.001, Cohen's d = .81. There was no significant difference between recall of pure lists, t<1 or between memory for high-rewarded items in mixed and pure lists, t(28)=1.69, p>.1. Low-rewarded items were recalled less well in mixed lists, t(28)=4.72, p<.001. The average recall of mixed lists, collapsing across reward (M=41.27, SD=12.48) was numerically lower than the average recall of unrewarded pure lists (M=43.86, SD=13.79).

Plotting the recall as a function of output position suggests that participants recalled high-rewarded items earlier than low-rewarded items (Figure 1), but the interaction between reward and output position did not reach significance, F(15,855=1.49, p=.10). Analysis of contiguity effects in mixed lists (see Polyn, Norman, & Kahana, 2009 for the calculation of these scores) for participants who recalled at least 2 items in each of these lists suggested that recall clustered based on temporal contiguity, t(28)=3.62, p=0.001, but not based on reward, t=1, or semantic distance, t<1.

Experiment 2

Experiment 2 was conducted as a conceptual replication of Experiment 1. It used a very similar methodology in the test of immediate memory, but added a delayed memory test after 24 hours.

Methods

Participants

Participants in Experiment 2 were recruited from the local community through advertising. They took part in two experimental sessions, and received £14 for their time and effort. We told participants explicitly that they can only qualify for reward if they do not make mistakes in the distractor task, and have excluded 2 participants who made more than 2 errors on average across all lists (they were statistical outliers, with error numbers that exceeded the inter-quartile range by more than 3 SDs). Of the remaining sample, 9 participants indicated they expected a memory test in session 2, and three of these 9 stated they have rehearsed the pictures in preparation for this (see "Session 2", below). These 3 participants were excluded from the analysis to avoid contaminating the forgetting rate data, but their exclusion did not alter any of the statistical results. The final sample included 22 females and 8 males who were 19-35 years old.

Materials

Stimuli consisted of 16*9 pictures from 9 different categories: clothing, animals, cityscapes, food, household objects (non-kitchen), kitchen objects, landscapes, office items, and people. They were drawn from the same sources as in Experiment 1, as well as the Nencki Affective Picture System (Marchewka, Zurawski, Jednoróg, & Grabowska, 2014), ImageNet (image-net.org), and internet sources. None had an overt emotional tone. While categories and pictures differed on multiple attributes, they were randomly assigned to conditions for each participant, eliminating the danger of systematic error. Each list included 16 pictures, which all belonged to a single semantic category. The Clothing list was always used in the practice block;

otherwise, allocation of list to condition was randomised for each participant. In the pure list condition all pictures were either framed or unframed; in the mixed list condition 8 pictures were framed and 8 unframed.

Procedure

Participants in Experiment 2 took part in two Sessions, 24 hours apart. The procedure of Session 1 was almost identical to the procedure of Experiment 1 other than the differences noted below. Session 2 included a delayed free recall test of each of the blocks encoded in Session 1.

Session 1 included a total of 8 blocks (instead of 6 in Experiment 1). As in Experiment 1, four included pure lists (half high-reward and half low-reward), but here there were four blocks with mixed lists. Each block began by describing the semantic category of the pictures in that block, e.g. “landscapes”. While in Experiment 1 all blocks ended with a free recall test, here this was true only for half of the blocks in each of the three conditions. The free recall test was omitted in the other half. Participants received the same instructions as participants in Experiment 1, but additional instructions were appended at the end. The purpose of the additional instructions was to equate the nature of encoding across blocks allocated to the immediate and delayed recall condition, and to decrease the likelihood of rehearsal between Session1 and Session 2. Participants were told that in some blocks, the immediate free recall test will be omitted. Furthermore, they were told that one of these blocks would be selected for a delayed free recall test at the end of Session 1. We hoped that the inclusion of a delayed memory test in Session 1 would discourage participants from expecting another one in Session 2. Finally, participants were told that Session 2 will include different tasks in Session 2. To avoid deception and maintain credibility, Session 1 always began with a bogus, additional practice block, which appeared to participants to be part of the experiment proper, and ended with a delayed free recall test of pictures from that block. The practice block always used the same Clothing list, and included only the picture and the distractor task. In the delayed test, participants were reminded they saw a list of "clothing" and asked to recall them. Data from this test were not analysed. The session ended with reward delivery as in Experiment 1.

Session 2. The session began with a delayed free recall test of each of the 8 blocks presented in Session 1. Participants were shown the title of the category of a list from one block, and given 4 minutes to recall the pictures from that block, before the next title was presented. Note that this was the first test of half of the blocks, but the second test of the other half, those that were tested immediately. Participants were informed that they would be rewarded for their recall according to the same schedule and rate. Session 2 also included additional picture rating tasks that are not reported here. At the end of the session participants were asked whether they expected a memory test, and if they answered in the affirmative, whether they made any notes or rehearsed pictures in their head. Then they were payed, debriefed and thanked.

Results

The average free recall data were analysed with a 2 (delay: immediate, delayed) x 2 (list: pure, mixed) x 2 (reward: high, low) repeated-measures ANOVA. This analysis only includes lists that were tested once – either immediately (session 1) or only after a delay (only tested in session 2). The three-way interaction was significant and large, $F(1,29)=13.83$, $p<.001$, $\eta p2=.32$. We unpacked this interaction with two separate repeated-measures ANOVAs. The first was a list-by-reward analysis which focused on average recall in the immediate condition; these data are depicted in Figure 1. Its aim was to examine whether this experiment replicated Experiment 1. The effect of list type was not significant, $F<1$, but the interaction with reward was significant and large, $F(1,29)=30.05$, $p<.001$, $\eta p2=.51$, qualifying the significant main effect of reward, $F(1,29)=20.06$, $p<.001$, $\eta p2=.41$, which was also observed in Experiment 1. Planned paired t-tests showed that as in Experiment 1, the effect of reward was only significant in the mixed list condition, $t(29)=6.60$, $p<.001$, Cohen's $d= 1.20$, where participants again recalled high-rewarded items earlier than low-rewarded items (Figure 1). By contrast, there was no difference between the recall of the two pure lists, $t<1$. As in Experiment 1, low-rewarded items were recalled less well in mixed compared to pure lists, $t(29)=4.44$, $p<.001$. However, while in Experiment 1 high reward did not increase memory, here recall of high-rewarded items was better than their recall in pure lists $t(29) = 4.32$, $p<.001$. The average recall of mixed lists, collapsing across reward ($M=26.45$, $SD=10.10$) was equivalent to the average recall of unrewarded pure lists ($M=26.25$, $SD=9.34$), $t<1$.

Further analysis of the mixed list condition showed that in this experiment, the interaction between reward and output position was significant, $F(15,1035)=13.90$, $p<.001$ $\eta p2 =.17$. Recall did not cluster based on temporal contiguity ($t<1$), but did cluster based on reward, $t(28)=2.08$, $p=.047$, such that participants had a greater-than-chance propensity to recall high-rewarded (and low-rewarded) items together. Clustering based on reward predicted the recall advantage of high-reward over low-reward items, $r=.74$, $p<.001$.

The list composition effect was less pronounced in the delayed test. Indeed, after a delay the interaction between list type and reward was small in size and not statistically significant, $F(1,29)=2.27$, $p=.14$, $\eta p2 =.07$, while the effect of reward was significant and large, $F(1,29)=18.95$, $p<.001$, $\eta p2 =.39$. Nevertheless, planned paired t-tests showed that the effect of reward was only significant in the mixed list condition, $t(29)=3.44$, $p<.001$, but the same trend was not significant in the pure list condition, $t(29)=1.61$, $p=0.12$.

To understand whether reward modulated forgetting we computed proportional forgetting scores [(average immediate recall – average delayed recall)/ (average immediate recall)], depicted in Figure 2. While a variety of forgetting scores have been proposed in the literature (Wixted, 1990), it was imperative here to take immediate performance into account in the calculation of forgetting because it differed between conditions. Proportional forgetting scores were computed for each list-by-reward condition, and therefore necessitated excluding 6 participants who did not recall any items immediately in one of these four conditions. This calculation was carried out twice. First, we examined lists that were only tested once - immediately or after a delay, i.e. in session 1 or 2, but not both; and second, we examined proportional forgetting of the same lists. Forgetting of items that were already recalled once in Session 1 (ranging from 15% to 31%) was lower than forgetting of items that were not tested in Session 1 (ranging from 62% to 74%). Forgetting scores were analysed with two list-by-reward repeated-measures ANOVA. None of the factors were significant (across lists: list type: $F<1$, Reward: $F(1,29)=1.61$, $p=.22$; interaction: $F<1$; within list: all $F<1$).

General discussion

In two experiments, participants only recalled high-value items more than low-value items when they were studied and recalled together, in a mixed list, but not when they were studied and recalled separately, in pure lists. This result supports our hypothesis that the competition for memory resources is local, constrained by the context specified by the task set - here, the temporal context of the previous list. It aligns with a core prediction of eCMR, supporting it as a general model of motivated memory.

Our immediate mixed-list results, where the effect of reward was large, corroborate previous reports (e.g. Ariel & Castel, 2014) and reward robustly enhances immediate free recall. Reward failed to increase immediate free recall and source memory in some previous work (Ngaosuvan & Mantila, 2005; Nilsson, 1987). Our findings suggest that these failures are due to the fact that they manipulated reward anticipation between-groups, a manipulation akin to our pure-list condition. The list-composition paradigm, which was here used with a reward manipulation for the first time, allowed us to document a large, replicable influence of reward anticipation on immediate memory (in mixed lists) that nevertheless failed to enhance memory when competition was abolished (in pure lists). The null effect of value in pure lists does not mean that there is no effect of value on memory in that condition, but the significant interaction demonstrates that it is smaller in pure than mixed lists. This pattern extends the results of Ngaosuvan and Mantila and of

Nilsson, where the evidence for the success of the manipulation of reward was only documented in subjective ratings of motivation.

The remarkable similarity between the list-composition effects obtained with emotional pictures in our previous work (summarised in Table 1 in Talmi et al., 2018) and high-reward pictures here suggests that they result from the same mechanism. eCMR explains enhanced recall for emotional information as resulting from the interplay of encoding, storage, and retrieval processes (Talmi et al., 2018). In eCMR, items that are encoded in the same way – through similar encoding operations - share a 'source' context. Emotional items trigger unique cognitive, affective and physiological changes, which distinguish their source context from that of neutral items. During encoding, preferential attention to emotional items binds them more strongly to their source context. During retrieval, when the context is used to probe memory, increased connection strength gives emotional items an advantage over neutral ones, allowing them to win the competition for recall.

The opportunity to gain is thought to be associated with dopamine release, and many believe that this underlies enhanced memory for high-value items after a delay (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Bunzeck, Dayan, Dolan, & Duzel, 2010; Lisman et al., 2011; Mason, Farrell, Howard-Jones, & Ludwig, 2017; Mason, Ludwig, & Farrell, 2017; Patil et al., 2017; Spaniol et al., 2014; Wittmann et al., 2005). Here the forgetting rate was equivalent regardless of whether items predicted high or low reward. The non-significant effect of reward on forgetting aligns with the general finding in cognitive psychology that manipulations that increase immediate memory do not attenuate proportional forgetting rates (Loftus, 1985; Slamecka, 1985). Of note, we powered the study to detect large effects, based on our previous results with the list composition task. Experiment 2 should be repeated with a larger sample size to determine whether there is evidence for smaller-size effects of reward-modulated forgetting.

In summary, we found that memory for reward-predicting information was not any better than memory for any other information - unless the two competed during recall. Reward also did not make items more resistant to forgetting. While it is certainly important not to have “useless facts” elbow out the useful ones, our results suggest that Sherlock Holmes may have been wrong to attribute enhanced memory for useful facts to what “the skilful workman…takes into his brain-attic” – namely, to biased encoding. eCMR (Talmi et al., 2018) predicted the pattern that we obtained here, suggesting that it could be a result of competitive advantage of well-attended, high-reward items during recall, rather than encoding. The main caveat of the work reported here is that it did not measure attention allocation during encoding. This limits our ability to illuminate the mechanism that underlies the novel behaviours we observed. Further research should examine attention to high- and low-reward items during encoding, and test the prediction that removing the competition at recall could salvage “useless facts” even when they are encoded in mixed lists.

Our experiments show that in some situations, people recall just as much for a penny as they do for a pound. These findings could have implications for the design of reward schemes in educational settings, because it suggests that motivating learners to commit a subset of material to memory by promising a prize is unlikely to enhance memory for that material – unless an assessments include material that learners did not prioritise, in which case the memory advantage for former will be at the cost of the latter. If further research supports eCMR as a model of motivated memory the model can be used to predict how reward schemes might be best employed to support learning.

Acknowledgements

DT was supported by the Royal Society [IE160027]. DK was supported by the University of Manchester Learning Through Research initiative. ND was supported by grant DA038891 from NIDA, part of the CRCNS program, and grant 57876 from the John Templeton Foundation. We thank Charlotte Ho, Emma Kavanagh and Rebecca Lawless for collecting the data for Experiment 1.

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), 507–17. https://doi.org/10.1016/j.neuron.2006.03.036

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*(6), 396–408. https://doi.org/10.1111/j.1467-9280.1991.tb00174.x

Ariel, R., & Castel, A. D. (2014). Eyes wide open: enhanced pupil dilation when selectively studying important information. *Experimental Brain Research*, *232*(1), 337–44. https://doi.org/10.1007/s00221-013-3744-5

Barnacle, G. E., Tsivilis, D., Schaefer, A., & Talmi, D. (2017). 1 Local context influences memory for emotional stimuli but not electrophysiological markers of emotion-dependent attention Gemma E. Barnacle, (0161), 1–35. https://doi.org/10.1111/psyp.13014

Blechert, J., Meule, A., Busch, N. A., & Ohla, K. (2014). Food-pics: an image database for experimental research on eating and appetite. *Frontiers in Psychology*, *5*, 617. https://doi.org/10.3389/fpsyg.2014.00617

Brodeur, M. B., Dionne-Dostie, E., Montreuil, T., & Lepage, M. (2010). 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*. https://doi.org/10.1371/journal.pone.0010773

Brodeur, M. B., Guérard, K., & Bouras, M. (2014). Bank of Standardized Stimuli (BOSS) phase ii: 930 new normative photos. *PLoS ONE*. https://doi.org/10.1371/journal.pone.0106953

Bunzeck, N., Dayan, P., Dolan, R. J., & Duzel, E. (2010). A common mechanism for adaptive scaling of reward and novelty. *Human Brain Mapping*, *31*(9), 1380–94. https://doi.org/10.1002/hbm.20939

Gershman, S. J., & Daw, N. D. (2017). Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework, *68*(5), 101–128. https://doi.org/10.1146/annurev-psych-122414-033625

Hadley, C. B., & MacKay, D. G. (2006). Does emotion help or hinder immediate memory? Arousal versus priority-binding mechanisms. *Journal of Experimental Psychology: Learning Memory and Cognition*, *32*(1), 79–88. https://doi.org/10.1037/0278-7393.32.1.79

Hurlemann, R., Hawellek, B., Matusch, A., Kolsch, H., Wollersen, H., Madea, B., … Dolan, R. J. (2005). Noradrenergic Modulation of Emotion-Induced Forgetting and Remembering. *Journal of Neuroscience*, *25*(27), 6343–6349. https://doi.org/10.1523/JNEUROSCI.0228-05.2005

Lee, T.-H., Greening, S. G., & Mather, M. (2015). Encoding of goal-relevant stimuli is strengthened by emotional arousal in memory. *Frontiers in Psychology*, *6*, 1173. https://doi.org/10.3389/fpsyg.2015.01173

Lisman, J., Grace, A. A., & Duzel, E. (2011). A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, *34*(10), 536–47. https://doi.org/10.1016/j.tins.2011.07.006

Loftus, G. R. (1985). Evaluating Forgetting Curves. *Journal of Experimental Psychology*, *11*(2), 397–406. https://doi.org/10.1007/BF02645546

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

MacKay, D. G., Shafto, M., Taylor, J. K., Marian, D. E., Abrams, L., & Dyer, J. R. (2004). Relations between emotion, memory, and attention: evidence from taboo stroop, lexical decision, and immediate memory tasks. *Memory & Cognition*, *32*(3), 474–488.

https://doi.org/10.3758/BF03195840

Madan, C. R., Fujiwara, E., Gerson, B. C., & Caplan, J. B. (2012). High reward makes items easier to remember , but harder to bind to a new temporal context, *6*(August), 1–15. https://doi.org/10.3389/fnint.2012.00061

Marchewka, A., Zurawski, Ł., Jednoróg, K., & Grabowska, A. (2014). The Nencki Affective Picture System (NAPS): introduction to a novel, standardized, wide-range, high-quality, realistic picture database. *Behavior Research Methods*, *46*(2), 596–610. https://doi.org/10.3758/s13428-013-0379-1

Mason, A., Farrell, S., Howard-Jones, P., & Ludwig, C. J. H. (2017). The role of reward and reward uncertainty in episodic memory. *Journal of Memory and Language*, *96*, 62–77. https://doi.org/10.1016/J.JML.2017.05.003

Mason, A., Ludwig, C., & Farrell, S. (2017). Adaptive scaling of reward in episodic memory: a replication study. *Quarterly Journal of Experimental Psychology*, *70*(11), 2306–2318. https://doi.org/10.1080/17470218.2016.1233439

Mather, M., Clewett, D., Sakaki, M., & Harley, C. W. (2016). Norepinephrine ignites local hotspots of neuronal excitation: How arousal amplifies selectivity in perception and memory. *Behavioural Brain Research*, *39*, e200. https://doi.org/10.1017/S0140525X15000667

Mather, M., & Knight, M. (2009). Reconciling findings of emotion-induced memory enhancement and impairment of preceding items. *Emotion*, *9*(6), 763–781. https://doi.org/10.1037/a0017281.Reconciling

Mather, M., & Sutherland, M. R. (2011). Arousal-Biased Competition in Perception and Memory. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, *6*(2), 114–33. https://doi.org/10.1177/1745691611400234

McDaniel, M. A., Dornburg, C. C., & Guynn, M. J. (2005). Disentangling encoding versus retrieval explanations of the bizarreness effect: Implications for distinctiveness. *Memory & Cognition*, *33*(2), 270–279. https://doi.org/10.3758/BF03195316

Montagrin, A., Brosch, T., & Sander, D. (2013). Goal conduciveness as a key determinant of memory facilitation. *Emotion*, *13*(4), 622–628. https://doi.org/10.1037/a0033066

Ngaosuvan, L., & Mantila, T. (2005). Rewarded remembering: Dissociations between self-rated motivation and memory performance. *Scandinavian Journal of Psychology*, *46*(4), 323–330. https://doi.org/10.1111/j.1467-9450.2005.00462.x

Nilsson, L.-G. (1987). Motivated memory: Dissociation between performance data and subjective reports. *Psychological Research*, *49*(2–3), 183–188. https://doi.org/10.1007/BF00308685

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

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

Sander, D., Grandjean, D., & Scherer, K. R. (2005). A systems approach to appraisal mechanisms in emotion. *Neural Networks*. https://doi.org/10.1016/j.neunet.2005.03.001

Schmidt, S. R. (1991). Can we have a distinctive theory of memory? *Memory & Cognition*, *19*(6), 523–542. https://doi.org/10.3758/BF03197149

Slamecka, N. J. (1985). On Comparing Rates of Forgetting. Comment on Loftus (1985). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*(4), 812–816. https://doi.org/10.1037/0278-7393.11.1-4.812

Spaniol, J., Schain, C., & Bowen, H. J. (2014). Reward-Enhanced Memory in Younger and Older

Adults. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, *69*(5), 730–740. https://doi.org/10.1093/geronb/gbt044

Talmi, D., Lohnas, L. J., & Daw, N. D. (2018). A retrieved context model of the emotional modulation of memory. *Psychological Review*.

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

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H. J., & Düzel, E. (2005). Reward-related fMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron*. https://doi.org/10.1016/j.neuron.2005.01.010

Wixted, J. T. (1990). Analyzing the empirical course of forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(5), 927–935. https://doi.org/10.1037/0278-7393.16.5.927

Figure captions

Figure 1. Average immediate free recall as a function of reward and list composition. High and low reward are presented in dark and light grey, respectively.

Top: immediate memory test in Experiments 1.

Middle: immediate memory test in Experiment 2.

Bottom: Recall probability as a function of output order in Experiment 1(left) and 2 (right). * $p<.001$. ns $p>0.05$.

Figure 2. Proportional forgetting scores of high and low-reward items over 24 hours.

Proportional forgetting scores for items only tested once, after 24 hours, or twice, both immediately and after a delay.

Figure S1. Recall probability in Experiment 2, as a function of serial position during encoding.

Figure 1.

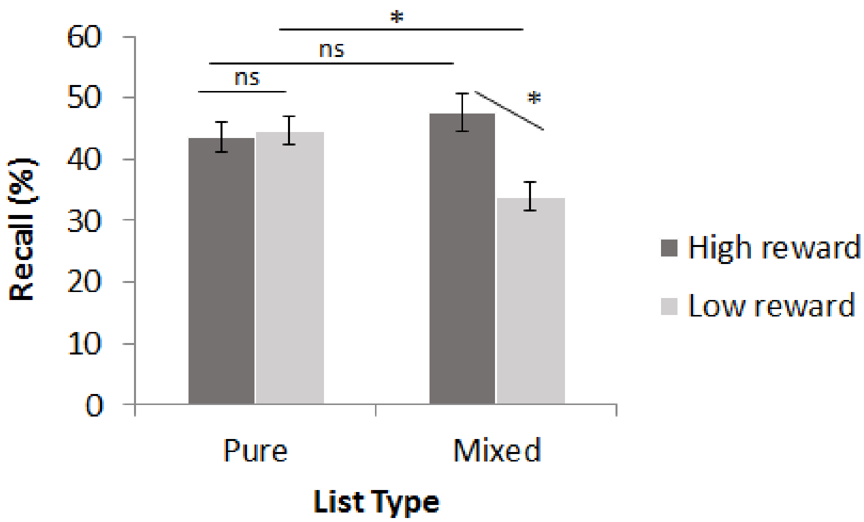
60
50
40
30
20
10
0
Recall (%)
ns
ns
*
*
Pure
Mixed
List Type
High reward
Low reward

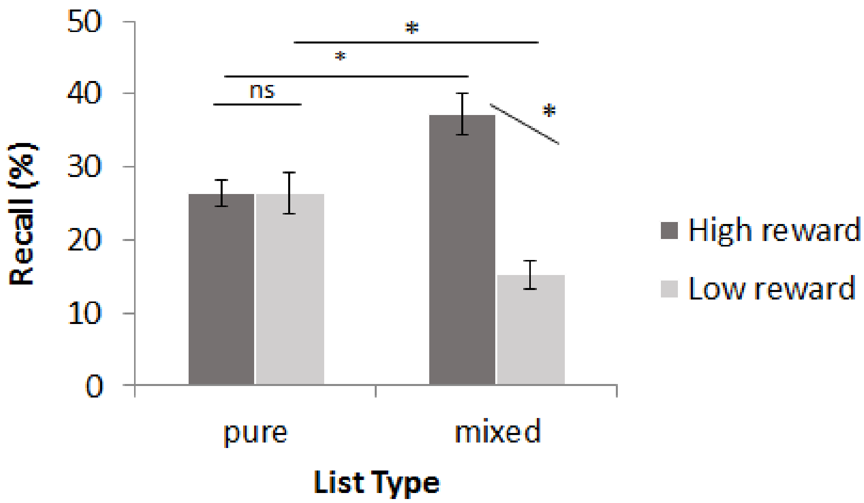
50
40
30
20
10
0
Recall (%)
ns
*
*
*
pure
mixed
List Type
High reward
Low reward

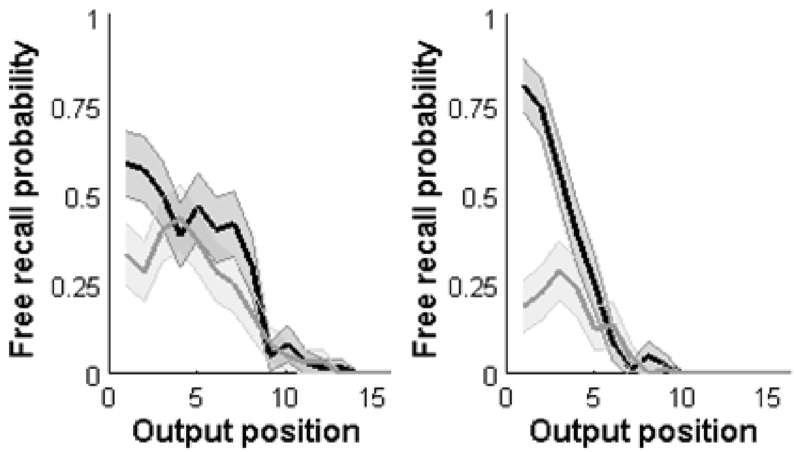
Free recall probability
Output position
Free recall probability
Output position

Figure 2.

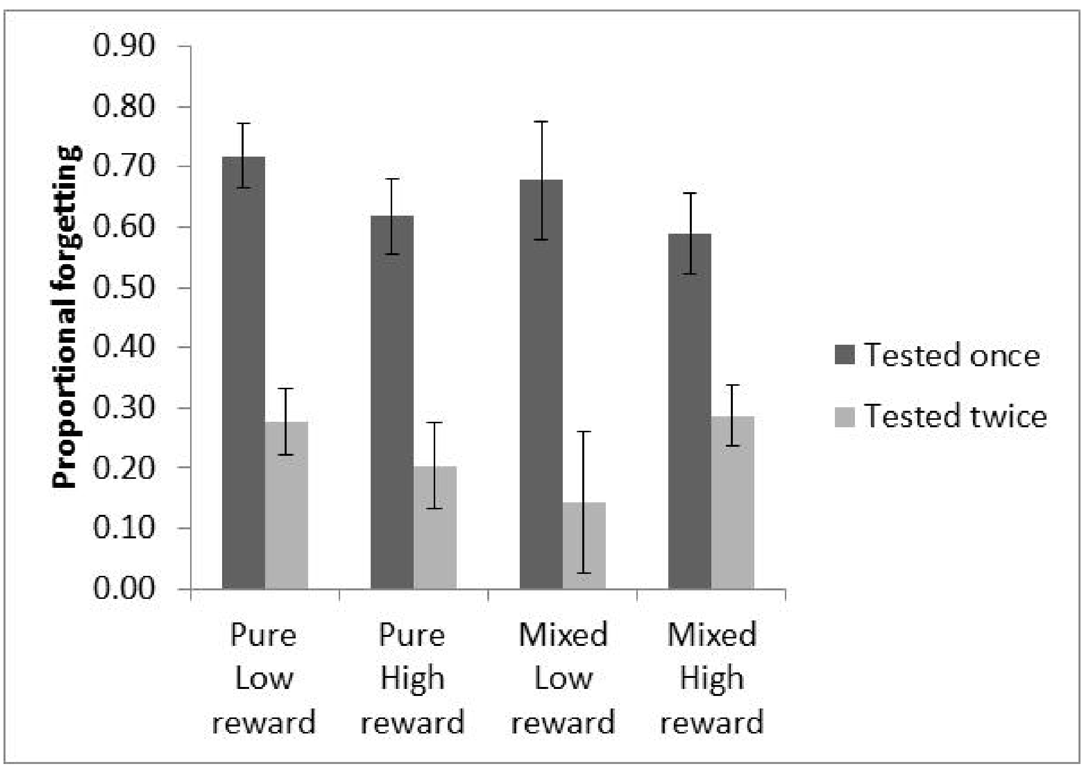
Proportional forgetting
Tested once
Tested twice
Pure Low reward
Pure High reward
Mixed Low reward
Mixed High reward

Figure S1

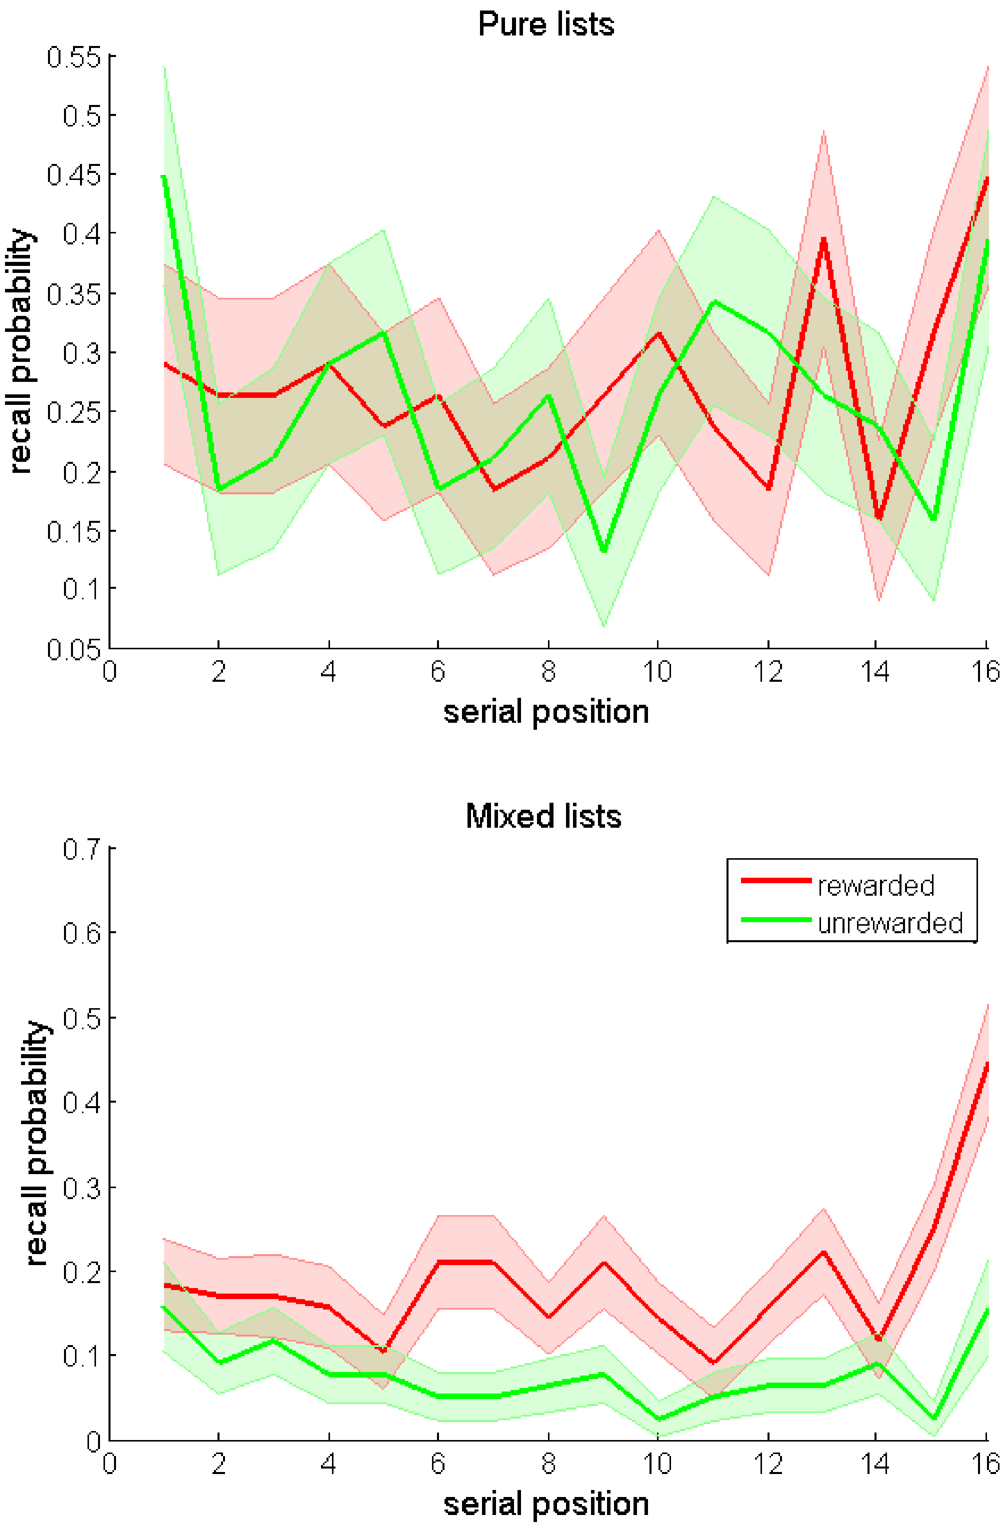
Pure lists
recall probability
serial position
0.55
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
2
4
6
8
10
12
14
16
Mixed lists
rewarded
unrewarded
recall probability
serial position
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0