Memory-based and online strategies in retrospective evaluations.

Alice Mason^{1 2}, Gordon D. A. Brown², Geoff Ward³, & Simon Farrell¹

² University of Western Australia

¹ University of Warwick

³ University of Essex

Abstract

The “construction” view of preference holds that we sample information from memory in the moment to make judgments and decisions. Building on previous work examining the relationship between memory performance and judgment, we used incentive-compatible methods to investigate both the extent to which memory supports evaluation and the potential trade-off between online and memory-based retrospective evaluation. In a series of pre-registered experiments, we presented people with a sequence of numerical values and asked them to complete either a free recall task, a willingness to-pay (WTP) task, or both. Using Bayesian mixed effects modeling, we predicted WTP (for the values in the remembered sequence) from both recalled items and presented items, and examined how the position of a value in a sequence determined its contribution to overall evaluation. Uniquely, we compared models of overall performance in memory across trials (e.g., primacy and recency) with those that predict evaluation from specific items recalled on individual trials. Previous studies are likely to have underestimated the strength of the relationship by focusing only on general trends in memory performances. Our findings indicated that across multiple incentivization and prioritization manipulations an averaging model is the best predictor of evaluation, even when we account for memory at the level of individual people and individual sequences. We replicate the findings that when participants anticipate a judgment task, they show less accurate recall from memory and show a preference for online updating. We demonstrated that even under these conditions, relying on memory is an effective strategy for performing evaluation, and discuss implications for understanding how people resolve the trade-off between online and retrospective evaluation.

Keywords: free recall; sampling; retrospective evaluation; judgment; episodic memory

Memory-based and online strategies in retrospective evaluations.

Introduction

Our evaluations of past experiences are powerful predictors of our future intentions and choices across a variety of domains, ranging from where we choose to go on holiday (Wirtz, Kruger, Scollon, & Diener, 2003) to whether we book in for a painful medical procedure (Redelmeier, Katz, & Kahneman, 2003). Our subsequent evaluations fail to capture, and often fail to predict, our moment-by-moment enjoyment of these experiences (Fredrickson, Kahneman, Hayes, & Knopoff, 1993; Wirtz et al., 2003), which may explain why we find ourselves repeatedly camping in the rain or hosting family parties (Mitchell, Thompson, Peterson, & Cronk, 1997). This discrepancy has been explained by a memory-based framework in which the evaluations that inform our decision-making require relevant information to be retrieved from memory (Gigerenzer & Goldstein, 1996; Slovic, 1995; Stewart, Chater, & Brown, 2006; Vlaev, Chater, Stewart, & Brown, 2011) and hence may not accurately reflect the value experienced in the past.

Previous work has established that—in some situations—people’s evaluations and judgments of events can indeed be predicted from their memory for those events. What has also been shown, however, is that people will rely less on memory—and more on online updating of evaluations of experiences as they unfold—if they expect to be asked to report such evaluations in the future. In this paper, we use *incentive-compatible* methods to examine people’s use of memory-based strategies in retrospective evaluation. Using *incentive-compatible* mechanisms means that we can examine the costs and benefits to using online updating versus memory-based evaluation, and examine how people trade off those costs and benefits when they could use either online or memory-based evaluation (or both).

Memory-based versus online evaluations

Challenging the idea that preferences are fixed and stable (Kahneman & Tversky, 1981; Slovic, 1995), it has been shown that preferences are both sensitive to context (Lichtenstein &

Slovic, 1971), and are often constructed during the decision process (Gigerenzer & Goldstein, 1996; Slovic, 1995; Stewart et al., 2006; Vlaev et al., 2011; Warren, McGraw, & Van Boven, 2011). Despite the obvious importance of encoding and retrieval to preference construction, it is only recently that memory-based processes have begun to be incorporated into subsequent models of judgment and decisions from experience (Bornstein, Khaw, Shohamy, & Daw, 2017; Lieder, Griffiths, & Hsu, 2018; Shohamy & Daw, 2015; Stewart et al., 2006; Weber & Johnson, 2006). For example, by extending the MINERVA model of memory, Dougherty, Gettys, and Ogden (1999) demonstrated that memory processes can account for probability judgments, and explain cognitive biases in those judgements. One of the most compelling pieces of evidence for a memory-based account of retrospective evaluations is the *peak-end* rule. Across a range of decision contexts, individuals tend to overweight the most intense, and the most recent, events in their overall evaluation of an experience. This parallels findings in the memory literature that emotional or value-laden moments are better recalled than neutral events, and that recent items are more likely to be recalled (Aldrovandi, Poirier, Kusev, & Ayton, 2015). The peak-end rule has been used to explain evaluations such as people's enjoyment of a meal (Robinson, Blissett, & Higgs, 2011), ratings of film clips (Fredrickson et al., 1993), evaluations of actual holidays (Kemp, Burt, & Furneaux, 2008), ratings of painful medical procedures (Redelmeier & Kahneman, 1996) and people's financial habits, ranging from purchasing entertainment tickets (Dixon & Verma, 2013) to describing preferences for payment sequences (Langer, Sarin, & Weber, 2005) and loan repayments (Hoelzl, Kamleitner, & Kirchler, 2011). More generally, it is well established that certain events in a sequence are "weighted" more heavily in an individual's subsequent evaluation (Aldrovandi et al., 2015; Fredrickson, 2000; Fredrickson et al., 1993; Montgomery & Unnava, 2009; Robinson et al., 2011).

There have been some attempts to directly link these patterns in judgment to known characteristics of memory, including the better recall of more recent and more extreme events (Aldrovandi et al., 2015; Kemp et al., 2008; Montgomery & Unnava, 2009). Montgomery and Unnava (2009) suggested that summary judgments of past experiences rely heavily on

memories formed during that experience and that over-weighting of events is linked to their accessibility in memory. Recency effects are well-established in the free-recall literature (Farrell, 2012; Murdock, 1962; Ward, Tan, & Grenfell-Essam, 2010), but an important exception is that introducing a delay between learning and testing (whilst keeping presentation rate otherwise the same) reduces or eliminates the recency effect and produces a shift to primacy (Bjork & Whitten, 1974; Glanzer & Cunitz, 1966; Postman & Phillips, 1965). Montgomery and Unnava (2009) tested whether this pattern replicated in global evaluations of sequences. They presented participants with sequences with either increasing or decreasing utility profiles. Previous studies have consistently shown that people prefer increasing to decreasing sequences of rewards (Loewenstein & Prelec, 1993; but see Frederick & Loewenstein, 2008) when preferences are elicited immediately. However, if there is a delay, then this pattern is reversed so that people show a greater preferences for decreasing sequences (Aldrovandi et al., 2015; Montgomery & Unnava, 2009). In this sense, evaluations after a delay can be seen as primacy driven which is consistent with the way in which information is recalled from memory after a delay (Glanzer & Cunitz, 1966; Postman & Phillips, 1965).

Another test of the memory-evaluation relationship has examined the extent to which recall of a particular experience (among a number of experiences) predicts overall evaluation of that experience. In other words, if a person happens to recall some more positive aspects of an experience, does this produce a more positive retrospective evaluation? Montgomery and Unnava (2009) measured participants' free recall of individual items in each presented sequence. Participants were presented with sentences describing events in a holiday. The events ranged from negative to positive in valence. Using a mediation analysis, the authors found evidence that the final evaluation of the sequence was mediated by the accessibility of the final three items in memory. Aldrovandi et al. (2015) replicated this effect using ratings of emotional materials; when tested after a delay, participants had worse recall of the final items and overall evaluation was less influenced by a negative item at the end of a list than when there was no delay.

An alternative to the memory-based account is that evaluations are formed as information

is encountered. Hogarth et al. (1992) suggested that as a sequence unfolds, information is integrated in a step-by-step manner. Accordingly, rather than evaluation being done at the point at which an evaluation is requested, evaluation is continuously conducted throughout the sequence, so that a final evaluation is in place when the sequence ends. A critical and open question is, therefore: what determines the use of memory-based versus online strategies in retrospective evaluations? In their belief updating model, Hogarth et al. (1992) highlight the possibility that, under certain conditions, information will be processed retrospectively at the end of the sequence. Retrospective evaluation places a higher load on memory, as all (or at least a representative set) of sequence elements need to be retained in order to form a reliable evaluation. In contrast, online updating requires memory for only the running total or average. Hogarth et al. (1992) suggest that people are more likely to adopt a step-by-step online updating approach, and rely less on memory, when the cognitive demands of a task are high. This step-by-step approach is more likely to result in recency in judgments, whereas an end-of-sequence model is associated with primacy.

A related key suggestion, and the focus of previous research, has been around task expectancy (Aldrovandi et al., 2015; Hastie & Park, 1986; Hogarth et al., 1992; Lichtenstein & Srull, 1987). When there is no expectation that an evaluation is required, or when people expect to complete a memory-based task (e.g., justifying the evaluation; Hogarth et al., 1992), people are thought to compute the evaluation after acquiring information. In order to carry out such retrospective evaluation, an individual is likely to retrieve recent samples from memory, leading to recency in judgment (Lichtenstein & Srull, 1987). The retrieval of samples from memory to form the judgment leads to a tight coupling between the items recalled and the evaluation. Hastie and Park (1986) suggested that if a judgment task is anticipated, participants will engage in online updating (or step-by-step) strategies to calculate evaluations; this suggestion was based on the assumption that memory-based strategies are less accurate and more effortful. Furthermore, people are likely to retrieve, as opposed to construct, underlying preferences in familiar situations (Feldman & Lynch, 1988; Hastie & Park, 1986). Aldrovandi et al. (2015)

conducted a series of experiments in which participants could evaluate sequences online and there was no expectation of a memory test. They provided evidence demonstrating that the memory-judgment relationship holds even when participants expect to make a judgment and have no awareness of a memory test, and therefore no reason to remember items (Aldrovandi et al., 2015). Hogarth et al. (1992) indicated that other factors including series length, response mode and stimulus complexity can also impact on the strategies people adopt. In the current experiments, we control these additional factors across experiments, allowing us to better examine the role of task expectancy in the memory-evaluation relationship.

Overview of experiments

Here we extend the research in memory-based retrospective evaluations by introducing incentive-compatible methods to examine the conditions under which people engage in memory-based strategies, and how the serial position of an item influences memory and evaluation of temporal sequences. A key distinction between our experiments and many retrospective evaluation studies is that we use numbers as the stimuli [but see (Langer et al., 2005; Varey & Kahneman, 1992)], allowing us to collect fine-grain measures of memory performance. Furthermore, participants in our experiments were told that the numbers represent values in a fictitious currency. The use of numbers allows us to have precise control over the valence of stimuli and also means that the stimuli have a consistent value across participants. Previous research has used words varying in valence, descriptions of people or holiday experiences or subjective measures such as pain (Fredrickson, 2000; Kemp et al., 2008; Lichtenstein & Srull, 1987), which can make it difficult to compare evaluations across participants. In order to elicit accurate preferences, we use a bidding procedure to determine participant's Willingness-To-Pay (WTP) for each sequence of numbers (Becker, DeGroot, & Marschak, 1964). We also collect and analyze fine-grained trial-level free recall data from participants, allowing us to develop a detailed account of what information participants are sampling from memory. In some experiments we incentivize recall. This use of controlled

stimuli and incentive-compatible methodology gives us a basis for examining the trade-off between different strategies for judgment.

Across three experiments, we tested the conditions under which participants adopt memory-based strategies in retrospective evaluations. Using model comparisons, we examine the ability of theoretically-driven models to predict evaluation under a range of experimental conditions (see Table 1 for a summary of the models). The first model, referred to as the “Expected Value” model, predicts the WTP estimate from the average of all presented items—the true value of the sequence. If participants based their evaluations on the expected value of the sequence they would maximize their earnings in the bidding task. The model makes no assumptions about how the average is calculated or whether it is done in a step-by-step online manner or at the end of the sequence. The second, the “Free” model, allows each of the serial positions in the list to take on its own parameter and can therefore capture both primacy and recency. We also include a value updating model. This “Temporal Difference Learning” updates the evaluation at each time step, and effectively produces a recency weighted function across serial positions. Finally we test two memory-based models. First, the “Serial Position (SP) Memory” model uses individual participants’ average accuracy at each serial position, across all trials, to weight the items in the list. Each participant will have a different shape to their serial position curve and therefore a different weighting will be applied to the items. Second, in Experiments 1 and 2 where we collect both memory and evaluation data on a single trial, we also test a “Recall-based” model, where memory for each item is used to predict evaluation. This is essentially means that there is a participant-specific, trial-specific shape to the serial position curve for both memory and evaluation.

Across the experiments, we focus on whether patterns known to dominate recall such as serial position, primacy, and recency effects, also influence evaluation. In Experiment 1 we collected recall and evaluation data from each participant on each trial, enabling us to assess the contribution of memory for individual items to evaluation. After viewing each sequence of numbers participants were either asked to recall as many items as possible or to complete an

incentivized willingness-to-pay task. Experiment 2 accounted for an asymmetry in incentivization that might explain these results, and introduced an explicit incentive for both memory and judgment. Experiment 3 allowed us to examine the potential trade-off between online evaluation and memory encoding by manipulating task expectancy. That experiment used a between subjects design to compare the effects of pre- and post-cueing each of the tasks. The implications of the rich data for the role of memory in evaluation are unpacked as the experiments are presented.

General Method

Participant Recruitment

For all experiments, participants were recruited via Prolific Academic to participate in the experiment online. We set initial sample sizes for each experiment, which although arbitrary, were based on money considerations and precedent from previous related literature. Our use of Bayesian statistics mean standard power analysis is inappropriate. The recruitment platform allows a specified number of participants to sign up to the study but some of these participants will not complete the experiment. If they return their submission they are automatically replaced but there are some errors in the process (participants may not complete the experiment for reasons outside their control, in which case we would still credit them), meaning that we did not have precise control over how many people would participate and complete the experiment.

To be eligible to take part in the experiments participants needed to be aged 18-65, have English as their first language (self-reported), be a resident of the UK, USA, Ireland, Australia, New Zealand or Canada, and have a Prolific Academic approval rating of over 90%. In addition, participants could only take part in one experiment in the series. Participants were reimbursed for their time according to the standard rates on Prolific Academic (£5-£6/hour). In all of the experiments, participants could earn an additional performance related bonus.

Task Design

In all experiments participants completed 32 trials. On each trial, participants were presented with a sequence of two-digit numbers. Participants were told that numbers represented amounts in a fictitious currency, Galactic Credits (GC). Accordingly, each sequence had a currency value that was simply the average of the numbers in that sequence. Each number was presented one at a time for 1500 ms, with an inter-stimulus interval of 1000 ms. After each sequence of numbers, participants completed either a recall task, an evaluation task, or both.

Sequences. The sequences were uniformly sampled, and the minimum and maximum values varied across each list but were always between 11 and 99. When generating the sequences, the variability within sequences was made large relative to between-sequence variability so that sampling more numbers would be most obviously beneficial to accurately estimate the mean. To achieve this, the mean of each sequence was randomly sampled from the numbers 40 to 70 (a uniform distribution). We then calculated the highest and lowest value for the distribution of each sequence by adding or subtracting 24 from the mean (i.e., each sequence distribution had a range of 58 but different minimum and maximum values). Finally, 7 numbers were uniformly sampled from this sequence distribution.

Free recall Task. If a recall cue appeared, participants had 15 seconds to recall as many items as possible in any order. Participants were instructed to type their responses, pressing enter after each item, into a response box at the centre of the screen. The screen was cleared each time the participant pressed enter so that previous responses were not visible. We chose a short recall period of 15 seconds because Aldrovandi et al. (2015) suggested that when a longer period of two minutes is used participants engage in exhaustive recall, whereas we were interested in the immediately accessible items that would plausibly contribute to evaluation in this setting and in everyday life (Kitayama & Burnstein, 1989). Miller, Weidemann, and Kahana (2012) highlight that recall is more likely to be terminated as output position (and therefore recall time) increases.

Evaluation Task. An auction procedure was used to obtain participants' Willingness-To-Pay (WTP) for each sequence of numbers, elicited using the bidding procedure developed by Becker et al. (1964). For each trial, and therefore for each bid, participants were given an endowment of 100 Galactic Credits (GCs). Participants placed a bid for each sequence (also in GCs) and a selling price was randomly drawn from a uniform distribution of prices (the range used in the experiment). If the bid was below the selling price, the participant did not purchase the item and kept only the endowment. If the bid was above the randomly drawn selling price, then the participant was automatically used their endowment to buy the sequence at the selling price, and kept the remainder of the endowment. In this latter case, the participants earning was then the remainder of the endowment, plus the currency value (the average) of the sequence that had been purchased. The optimal strategy—placing a bid equal to the estimated true value of the sequence—was explained to participants. This method is incentive compatible as only the participant's bid determines whether or not they buy the item, and the greatest payoff would be obtained by bidding the believed true value of the sequence (see <https://osf.io/92x4w/> for full WTP instructions shown to participants).

Data analysis

The inferential framework used was Bayesian estimation and model comparison. For the analyses where we are comparing mean recall or WTP across conditions, we used the BayesFactor package (Morey & Rouder, 2018) to estimate Bayes Factors. For these Bayes Factor ANOVA analyses, the scale of the effect size for fixed effects (r scale) was set to .707, labeled the “medium” prior in the package. Where we are testing multiple effects, we use the “top” method to estimate Bayes Factors, otherwise the default is to compare against a null model that only has a subject-specific intercept. For t -tests, the analyses used an uninformative Jeffrey's prior on the variance, and a standard Cauchy prior of $\sqrt{2}/2$ on the r scale value. For a discussion on priors, see Rouder, Morey, Verhagen, Province, and Wagenmakers (2016).

The value of the Bayes factors quantifies the strength of evidence in favor of one model

with respect to another given the data obtained. It informs us how much our prior beliefs should shift in response to the data obtained. Although there are no strict cut-offs, we apply the verbal labels used by Wagenmakers, Wetzels, Borsboom, and Maas (2011) to describe our results. For reference these are summarized in Table 2.

Across the experiments we present the following analyses.

Free recall task. We use three different measures of free recall. First, we examine overall accuracy at each serial position to produce a serial position curve (SPC): the mean probability of correct recall at each serial input position. Second, we report the first recall probability (FRP) function, which plots the probability that recall was initiated with an item from each serial position in the original list. Third, as per our pre-registration, we also computed the lag conditional response probability function (lag-CRPs). We report these in the Supplementary Materials, as they did not shed light on the main research question, the relationship between memory and evaluation. These measures are presented as benchmarks of free recall, and as data often used to evaluate or compare theories of free recall (Farrell, 2012; Howard & Kahana, 1999; Ward et al., 2010).

Evaluation task. The true value of the sequence was defined as the average of all the values shown. To measure how accurate participants' judgements were relative to the true mean, we calculated the root mean square error between participants' bid amounts and the actual sequence values.

Predicting evaluation. A primary interest was the extent to which different theoretically-motivated predictors were related to people's evaluations (auction bids). Table 1 lists the five models that were compared across the five experiments, and as described conceptually above. The normative *Expected Value* model represents the predictions of an integrator model that averages all values in the sequence, with equal weighting for the values. The *Free* model allows each position in the sequence to be given its own weighting in predicting evaluation, and allows the model to capture primacy and recency in the weighting of

information at evaluation. The *Temporal Difference Learning* model is a simple value updating model. Each of the stimuli is used to update the estimated value by the difference between the new stimulus and the current estimate. This temporally weighted average following the last presented number in the sequence is then used as the predictor. The *Serial Position* model represents a memory model that uses information about people's memory performance to predict their evaluation responses. This predictor was constructed for each sequence by using a person's overall memory serial position function to calculate a weighted average of the values in the sequence. Finally, the *Recall-based* model was a trial-based memory model. In experiments and conditions where memory and evaluation performance were available for individual trials, the predicted evaluation for a trial was the average of the values recalled for that trial, including both true and false recollections (i.e., intrusions).

We used Bayesian multilevel modeling to predict participants' WTP estimates using the R package *brms* (Bürkner, 2017). For the *Temporal Difference Learning* model we used the R package *R2jags* (Su & Yajima, 2015). We specified weakly informative priors on the mean of each of the population level effects, included in the models. The "Expected Value" model had a single population level parameter which was the average value of the sequence. The "Serial Position" model had as many population level coefficients as the list length. In all cases we used a normally distributed prior. In the case where there was a single predictor of WTP, e.g., "Expected Value", the prior on the slope was set to $N(1, 0.5)$. When we were predicting WTP from individual items, the prior was defined as $N(1/\text{list length}, 0.5)$. For all models we also included "participant" as a group-level effect on intercept. The prior on the mean intercept was $N(0, 10)$. The standard error on the group level effect (participant) was a half student $t(3, 0, 10)$, which is the default setting in the *brms* package (Bürkner, 2017). For all of the Bayesian analyses using the *brms* package the warm-up or burn-in period was 1000, with an additional 4000 iterations to estimate the posterior distribution of each parameter; four chains were run with these values. All $\hat{R}$ values were 1.05 or below indicating good convergence of the chains. For the "Temporal Difference Learning" model the initial burn-in period was 5000, with

an additional 10000 iterations and 4 chains.

To compare each of the models in our experiments we selected the model with the highest marginal likelihood. Then, using the Bayes Factor function in `brms` we obtained a Bayes Factor for the relative evidence in favor of each of the models with respect to the best fitting one.

The “Temporal Difference Learning” required that participants had completed all trials in the experiment, therefore as the models were being compared, for all the models only participants who had completed all trials were included in these analyses.

Experiment 1

In Experiment 1, participants were presented with sequences of seven numbers and were successively cued to perform both tasks after seeing each sequence. The order of the tasks was counterbalanced across trials within participants.

Results

In our preregistration we set our sample size 60 (see <https://osf.io/se2gg/> for details). We excluded 5 participants as they failed to complete at least 80% of all trials in the experiment. We also had exclusion criteria for the memory and evaluation tasks. Five participants were excluded for poor performance in the memory task (on average less than 1 item correct per trial). For the evaluation task, 9 participants were excluded for not completing at least 80% of the trials. A total of 19 participants were excluded, and 72 were included in the analysis.

Memory performance. For both conditions (RE: recall followed by evaluation; ER: evaluation followed by recall) the accuracy serial position functions show a small amount of extended primacy, and recency across the last two serial positions (Figure 1). There is extreme evidence of an effect of task order on recall accuracy ($BF_{10} = 5.15e+04$), with accuracy being greater in the RE compared to the ER condition. For each task order, a mixed effects logistic regression predicting recall probability from serial position was conducted. There is an extreme

effect of serial position on accuracy for each of the task order conditions (RE: $BF_{10RE} = 6.45e+206$); ER: $BF_{10ER} = 1.77e+211$). This second condition is effectively delayed recall and as expected we see reduced recency compared to the RE condition where recall is immediate. The first recall probability functions complement the above findings, with recency dominated recall in the ER (recall last) condition and greater primacy in the RE (recall first) condition.

Accuracy of evaluation. Participants’ accuracy in the WTP task is shown in Figure 3. One participant was excluded from all the evaluation analyses as their WTP estimates were outside of the required range (0–100). A Bayes Factor t -test indicates anecdotal evidence against the effect of task order on WTP estimates ($BF_{order} = 0.33$).

Joint analysis. In this experiment we have memory and evaluation data from each participant for each trial. This allows us to test whether memory for a specific set of items predicts evaluation using the “Recall-based” model. 33 participants completed all trials and their data was included in the following analysis. In this scenario, we see extreme evidence in favour of the “Expected Value” over all the other models (see Table 4 for details). We pre-registered a secondary analysis to examine whether the memory-evaluation relationship is modulated by task order (see Table 5 for full model fits). We see a small effect of task order in that analysis: When the evaluation task is completed first the “Expected Value” model is the best predictor of evaluation; when the recall task is first there is substantial evidence that the “Recall-based” model is a better predictor.

For this and other model-based analyses reported in the paper, one question is how accurate the Bayes Factor comparisons are. In other words, what is the likelihood of correctly recovering a model when the true generating model is known? To address this, we report a confusion matrix for the models. We simulated data from each of the models and then fitted each of the models to the simulated data. We noted which model had the lowest marginal likelihood and counted the number of times (out of 50) that each model won. The results are summarized in Table 3. There was perfect recovery for all models except for the “Free” model, which is consistently identified as the “Expected Value” model. This is not surprising as they are

functionally very similar but the “Free” model is punished for complexity, due to its greater number of parameters.

Discussion

After each sequence of numbers, participants were asked to complete both the recall and evaluation tasks, the order of tasks being counterbalanced across trials. The recall patterns were as expected for seven item lists; we observed one-item primacy and graded recency. There was a main effect of task order on recall; in the ER condition we effectively have a delayed recall condition and therefore see increased primacy (Postman & Phillips, 1965). < More critically, by collecting recall and evaluation data on each trial we were able to test a model that predicted evaluation on the basis of recall on individual trials. Overall, the normative model “Expected Value” that takes the average of the true sequence was the best predictor of the data.

When we account for task order we see that the “Recall-based” model provided a convincingly superior account of evaluation than any of the other models. These order effects support evidence from previous studies that have found memory for individual items on trial-by-trial basis predicts evaluations (Aldrovandi et al., 2015; Montgomery & Unnava, 2009). Montgomery and Unnava (2009) tested this relationship by examining the difference between improving and declining sequences, whereas Aldrovandi et al. (2015) altered the position of a negative item within a neutral list. In our experiments, we have used lists that vary randomly in their value profile. This enabled us to also examine the overall patterns of serial position on recall across trials. However, in the RE condition the model comparisons suggest that trial-by-trial recall provides a better account of the evaluation data. To our knowledge, this is the first study to relate recall and evaluation in this way. In addition, and in contrast to the experiments of Montgomery and Unnava (2009) and Aldrovandi et al. (2015), we used incentive-compatible methodology to elicit evaluations.

Nevertheless, the “Expected value” model was a better overall indicator of evaluation.

Given that only the WTP task was incentivized, it is possible that people focused more on the WTP task at list presentation. If people can use different strategies during encoding to form evaluations—specifically, memory-based evaluation and online updating—there will likely be some trade-off between these tasks. Indeed, one way to conceptualize our paradigm is as a dual-task procedure where participants must divide attention between memory encoding and a subsidiary task (Baddeley, Scott, Drynan, & Smith, 1969; Craik, Naveh-Benjamin, Ishaik, & Anderson, 2000; Murdock, 1965). The interesting complications in our paradigm are that participants are free to give more or less weight to memory encoding (vs online updating) during sequence presentation, as there is no explicit instruction to perform online updating; and that memory updating is potentially an effective strategy for performing the task to which online updating is directed (i.e., evaluation). Potentially, people are using the incentivization of the WTP task as a cue to perform online updating, at the expense of memory performance.

This idea of a trade-off between online updating and recall is also recognized in discussions around expectancy effects on the evaluation-recall relationship. Previous studies have suggested that expectancy influence the evaluation strategy people adopt (Aldrovandi et al., 2015; Hastie & Park, 1986; Hogarth et al., 1992; Lichtenstein & Srull, 1987), with online updating being preferred when an evaluation is expected, and participants relying more on recall when either a) evaluation is not expected; or b) recall is potentially expected in addition to evaluation. Nonetheless, a memory-evaluation relationship has been observed even when recall is unexpected (Aldrovandi et al., 2015). A key feature of our incentive-compatible paradigm is that we can measure the cost—in terms of the discrepancy between the evaluation of a sequence and its normative value—to following different strategies.

Experiment 2 addresses this issue, by more explicitly controlling the prioritizing of evaluation and recall.

Experiment 2

Previous research has focused on the influence of task expectancy on the evaluation strategies that people use (Aldrovandi et al., 2015; Hastie & Park, 1986; Hogarth et al., 1992; Lichtenstein & Srull, 1987). One issue is that none of these studies examining test expectancy used incentive-compatible methods to elicit preferences. In addition, those studies examine different mechanisms for forming an evaluation, while the more specific trade-off we examine here is that between online updating and memory encoding during list presentation.

Experiment 2 aimed to test whether people would use different evaluation strategies depending on which task was incentivized and therefore (by assumption) prioritized. Incentivizing both tasks allows us to address the trade-off quantitatively using the incentive-compatible auction method. Participants were presented with sequences of seven numbers and were retrospectively cued to perform both the free recall and WTP tasks. Task incentivization was manipulated across two experimental blocks. In each block, 75% of trials were incentivized for a given task. Participants were told that each sequence of numbers had been pre-labeled as either WTP or free recall. We then selected a random trial from each block to pay out on. It was possible to earn a bonus for a recall trial in a WTP block. The order of the tasks was post-cued and counterbalanced across participants and conditions (for pre-registration see <https://osf.io/beak/>).

Results

Our pre-registration was for 100 participants. Seven participants were excluded from analysis for not completing at least 80% of the experiment, and 9 participants were excluded for not completing 80% of the evaluation trials. 5 were excluded for poor performance (less than 1 item correct per trial on average in the memory task). A total of 20 participants were excluded, and 94 were included in the analysis.

Memory performance. We conducted a mixed effects logistic regression predicting recall probability from serial position, prioritized task, and task order. The accuracy serial position functions plotted in Figure 1 show similar performance across the incentivization conditions and task orders. There was extreme evidence of an effect of task order ($BF_{10} = 2.09e + 07$) and substantial evidence of an interaction between serial position and task order ($BF_{10} = 10.24$). There is evidence of primacy for both task orders and there is increased recency when the recall task is completed first, consistent with the longer effective retention interval in the ER condition. This is consistent with the findings from the previous experiment indicating more recency for immediate versus delayed recall.

There is extreme evidence against an effect of prioritized task on overall accuracy ($BF_{10} = 0.02$). Mirroring the accuracy serial position functions, the first recall probability functions plotted in Figure 2 also indicate primacy across both conditions, and increased recency in the RE (i.e., immediate recall) condition.

Incentivization of recall and evaluation tasks. Accuracy in the WTP tasks is plotted in Figure 3. A Bayesian t -test found only anecdotal evidence for an effect of task incentivization on WTP accuracy ($BF_{10} = 0.42$). In the previous experiment we looked at recall accuracy in terms of the proportion of items correctly recalled. In this experiment we have incentivized the recall in Galactic Credits so in addition to looking proportion of correct recalls we can also examine recall in terms of the incentives [i.e. earning per recall] when either the evaluation or the recall task was incentivized (see Figure 4). A Bayesian t -test demonstrated substantial evidence against an effect of task incentivization on recall performance ($BF_{10} = 0.16$).

Joint analysis. 45 participants completed all trials and were included in the analysis. First, we analyzed the data collapsing across task incentivization and order (see Table 4 for all model details). As in Experiment 1 we are able to fit memory data on a trial by trial basis. The Bayes Factors indicate extreme evidence in favour of the “Expected Value” model compared to all other models.

We then analyzed each of the order conditions separately (see Table 5). The results here are very similar to the order effects in Experiment 1. For the ER condition, there is extreme evidence in favour of the “Expected Value” model compared to the “Recall-based”, “Free”, “Temporal Difference Learning” and “Serial Position models”. In the RE condition, the “Recall-based” model is the best fitting model, but note that it differs little in fit from the “Expected Value” model. There is extreme evidence in favour of the “Recall-based” model compared to the “Temporal Difference Learning”, “Serial Position” and “Free” models.

In addition, we looked at the effect of prioritized task (see Table 5). When participants were incentivized for the evaluation task the “Expected value” model was best, there was little evidence over the “Recall-based” model. There was extreme evidence that the “Expected value” model was better than the other models. When recall was prioritized the “Expected Value” model was the best model and there was extreme evidence it was better than the “Recall-based”, “Serial Position”, “Free” and “Temporal Difference Learning” models.

It was also possible to perform this analysis based on the task that participants reported prioritizing. At the end of the experiment participants were shown three multiple choice options on screen and asked to select the option that best described the strategy they adopted when completing the tasks. The options were: a) “I focused more on the task (recall vs evaluation) that was more likely to be rewarded”, b) “I tried equally hard on both tasks regardless of the reward”, c) “Neither of the above options best describes the strategy I used”. We excluded the 3 participants who reported using a different strategy than those listed. There were 71 participants who reported focusing more on the rewarded task and 9 who reported focusing equally on both tasks. We ran the above analysis separately for the two identified groups of participants (see Table 5). In both cases there was extreme evidence in favour of the “Expected Value” model compared to all other models.

Discussion

The results from Experiment 2 suggest that when collapsed across the incentivization manipulations “Expected Value” is the best predictor of choice. This supports the findings from Experiment 1. As in Experiment 1, there was also a suggestion that the “Recall-based” model did relatively better (though not better than the “Expected Value” model), when the recall task was performed prior to the evaluation task. Critically, there was no evidence that the incentivization task resulted in people adopting different strategies to perform the evaluation task, even when participants reported giving greater priority to the more heavily incentivized task. Neither recall accuracy (Figure 4) or WTP accuracy (Figure 3) changed as a function of task incentivization.

The free recall data indicate that people are equally accurate and show very similar serial position curves irrespective of which task is incentivized. The experiment also counterbalanced for task order across trials and we see greater recency when the recall task is immediate as opposed to delayed (as seen in Experiment 1). The fact that the incentivization does not influence recall patterns suggests that in this and the previous experiment, participants engage in similar averaging strategies regardless of which task is prioritized.

In both Experiments 1 and 2, the memory-based model performs relatively better in the recall first condition, although not necessarily giving the best account of the data. One could argue that the memory-based models would perform better when the recall task is completed first, as recall has been performed anyway and so basing judgment on the recalled items is straightforward. However, in this experiment we have the added issue that participants were also asked to prioritize one of the tasks which may have interacted with the order effects and influenced strategies used. The combined evidence from Experiments 1 and 2 is inconclusive and future experiments could better test this hypothesis by using sequences that would produce different evaluations depending on which items were recalled.

The lack of effect of prioritization potentially implies there is not a trade-off between online updating and memory encoding, in contrast to the expectations for this experiment.

However, it may also be that participants were unable or unwilling to respond to the incentivization structure. Specifically, participants might be completing both tasks at encoding, and are unable (or choose not) to give priority to any one task. While previous work has shown that substantially varying the amount of practice given to the concurrent task (so as to emphasize that task) can determine the extent of dual-task interference (Naveh-Benjamin, Kilb, & Fisher, 2006), the instructed incentives here may not have been sufficient to move participants' trade-off between the two tasks. Accordingly, Experiment 3 implemented a stronger, and more standard, manipulation by using a single- vs dual-task comparison.

Experiment 3

In Experiment 3, we compared conditions where participants complete only one task (either evaluation or recall) to when people complete both. This is more akin to the standard dual-task design used in studies of divided attention (Baddeley et al., 1969; Craik et al., 2000; Murdock, 1965). Participants were presented with sequences of seven values, and we manipulated both task (recall vs evaluation) and task expectancy. In the sole task condition, participants were pre-cued to complete either the evaluation or the free recall on all trials. In the post-cue condition, participants were retrospectively cued to complete either WTP or recall on a trial by trial basis. Both tasks were incentivized such that the incentive amounts were roughly equated across the two tasks. The pre-registration can be found at <https://osf.io/ean6u/>.

Results

Our pre-registration was for 150 participants. Five participants were excluded for not completing at least 80% of the experiment and 6 were excluded for not completing 80% of the evaluation trials. A total of 11 participants were excluded, and 139 were included in the analysis.

Memory performance. The accuracy serial position functions (Figure 1) indicate better recall performance in the recall only compared to the post-cued condition, the effect of

condition having substantial support from the Bayes Factors, ($BF_{10} = 35.72$). There is extreme an extreme effect of serial position ($BF_{10} = >10^{+99}$). There is substantial evidence against a condition $\times$ serial position interaction ($BF_{10} = 0.15$). The first recall probability functions (Figure 2) show similar graded recency and one-item primacy across the two conditions. It is worth noting here that the standard error bars shown in Figure 1 are within-subject error bars so can only be used to compare the within-subject variables.

Accuracy of evaluation. There was substantial evidence that accuracy in the WTP task, shown in Figure 3, was the same in the two conditions ($BF_{10} = 0.22$), such that evaluations were no closer to the “true” sequence values in the pre-cue condition than in the post-cue condition.¹

Joint analysis. 33 participants completed all trials and were included in the analysis. For the post-cued condition where participants completed both the evaluation and the recall tasks we were able run model comparisons. However, given that we did not have memory data for every trial we could not include the “Recall-based” model. As shown in Table 4, there was strong evidence that the “Expected Value” model was a better predictor of WTP.

Memory as a strategy. If participants are trading off memory encoding against online updating, we should be able to quantify this trade-off as both memory and evaluation can be measured on the same scale. In particular, we can ask about the consequences of turning the task into a dual-task—and attempting to jointly carry out online updating and memory encoding—for performance in terms of earnings on the WTP task. In other words, was the dual-task cost participants appear to have imposed on themselves in Experiment 3 worth it in terms of increased earnings, compared to the alternative strategy of simply performing memory encoding, and then performing memory-based evaluation?

To answer this, we conducted an exploratory analysis across experiments to examine how accurate participants would have been in the evaluation task if they had based their WTP

¹ We pre-registered an additional analysis regarding the speed of WTP evaluations. There is extreme evidence that participants were faster to make evaluations in the sole task condition, $BF_{10} = 3.20 \times 10^4$.

estimates on memory recall. For each participant in the post-cued condition we used the average of the items they recalled to calculate a WTP estimate. We then looked at how accurate this estimate was compared to the true value of the sequence. Surprisingly, the results (see Figure 5) indicate that using this memory-based strategy would have led to more accurate WTP estimates than the ones participants actually provided (those shown in Figure 3). This means that participants were apparently choosing to turn the post-cued task into a dual task (as reflected by the cost of the post-cueing to memory accuracy) with no benefit to WTP accuracy, when relying just on a memory-based strategy (and thus not incurring a cost to memory performance) would also have benefited WTP performance.

Discussion

This experiment enabled us to probe the conditions under which participants adopt memory-based evaluation strategies, and the accuracy of such strategies. In the condition in which only recall was required, participants showed greater accuracy and more primacy in their recall compared to the post-cued condition. There was no cost to accuracy of WTP in the post-cued condition in comparison to the evaluation-only condition.

We found evidence that when participants are asked to complete both the evaluation task and the recall task for a given sequence, the “Expected Value” model is the best predictor of evaluation. The “Serial Position” model was not predictive of the evaluation data. However, our additional analysis showed that the cost imposed in the post-cued task was not a necessary one. Had participants treated the post-cued task simply as a memory encoding task—followed by either memory-based recall or memory-based evaluation—they would have incurred no cost to memory performance, and would actually have produced more accurate WTP estimates (relative to the “true” value of the sequences). We defer discussion of this surprising result to the General Discussion.

General Discussion

The experiments presented here provide a detailed analysis of the relationship between recall of sequences from memory and global evaluations of these sequences. Using incentive-compatible methods we have been able to assess and quantify the relationship between online and memory-based strategies in evaluation. Across experiments and conditions the findings suggest that people do not rely on memory to evaluate sequences. However, in Experiments 1 and 2 there was some evidence that this is moderated by the order of the recall and evaluation tasks. In Experiment 1 we collected both memory and evaluation data on each trial. This enabled us to directly assess the influence of memory on evaluation. We found strong evidence that overall online strategies predict evaluation. However, when the recall task is completed first then we see at the trial-level that recalls predict evaluation. This builds on previous studies that have found memory for items appearing at the start, middle or end of a list mediates evaluation of the list (Montgomery & Unnava, 2009). Earlier studies focused on the serial position function across participants by using lists with improving or declining utility profiles (Aldrovandi et al., 2015; Montgomery & Unnava, 2009) but have not examined the position of each item in the list. The order results from Experiment 2 were less conclusive and overall the evidence from Experiment 2 is more difficult to interpret as the incentivization manipulation was unsuccessful. We therefore conducted Experiment 3 to better understand how evaluation strategies compete.

In Experiment 3 participants where we compared the different strategies, participants appear to pit these two strategies against each other and engage in both online updating and memory-based strategies, but this is not done in an optimal manner and comes at a cost to both. Previous work has suggested that task expectancy is likely to influence the particular strategies people adopted. Hastie and Park (1986) claimed that memory will only predict evaluation when a memory task is anticipated. Contrary to this, Aldrovandi et al. (2015) demonstrated that when people anticipate an evaluation task they do use memory-based strategies. Previous studies

have partly by-passed this issue by collecting a single memory and evaluation estimate from participants. With the exception of the evaluation-only condition in Experiment 3, participants were aware that there could be a memory task. Arguably, this may have encouraged use of memory-based strategies. In order to understand the task demands that might be governing people's use of memory-based strategies we introduced incentivization to both tasks (Experiment 2) to try to shift people's use of strategies. In fact, people performed equally well in both tasks, regardless of which task was incentivized. Therefore, we introduced a task expectancy manipulation in Experiment 3 and examined whether there is a trade-off between memory-based and online updating strategies. This enabled us to address the tension between being able to model the memory data accurately and using a task design that prevents participants from adopting experiment specific strategies; for example, relying more on memory given they know memory will be tested.

The results of Experiment 3 highlight the fact that task expectancy by itself does not determine the strategy employed. In the post-cued condition, where either task might be cued, there was a cost to memory performance (compared to the single-task condition), and no effect on evaluation accuracy. One interpretation of this result, as presented above, is that participants imposed on themselves a trade-off between online updating and memory encoding when people were unsure of the upcoming task. Accordingly, having the expectation of a memory test is not sufficient for participants to rely solely on memory-based evaluation, and people may continue to rely on online evaluation as long as a judgment is expected (irrespective of anticipation of a memory probe) (Hastie & Park, 1986).

One argument for people relying partially on online updating even when a memory task is expected is that memory-based evaluation will be slower and less reliable (Hastie & Park, 1986). The use of incentive compatible methods adds to this picture by allowing us to quantify the trade-off to performance when using different strategies. We have convincing evidence across all our experiments that contradicts the claim that memory-based strategies lead to less accurate evaluations (Hastie & Park, 1986) and our results highlights that the opposite is true; a

memory-based strategy would be more effective and accurate than online strategies. The cost to memory-performance suggests that participants are attempting to use online-based evaluation strategies in this condition, this is in line with Hastie and Park (1986) predictions for conditions when evaluation is expected, but in a way that is detrimental to performance on both the evaluation task and the memory task.

Why is this? One suggestion is that participants assume that online updating will be more effective than memory-based judgment, even though it is not for the stimuli and evaluation task used here. Another is that although encoding for effective recall-based evaluation would lead to as (or more) accurate judgments, it does require more effort, as suggested by Hastie and Park (1986). People prefer to put less—rather than more—cognitive effort into tasks (Kool, McGuire, Rosen, & Botvinick, 2010), and rewards appear to be discounted by the amount of effort required to obtain those rewards (Botvinick, Niv, & Barto, 2009). Accordingly, although the strategy adopted by our participants may not appear optimal from a purely monetary perspective, it may well be that our participants adopted a more sophisticated cost-benefit analysis that also acknowledged the cognitive effort associated with memory encoding and memory-based evaluation.

The results from these experiments also add to our understanding of the recall of value-based information and its relationship to evaluation. Few studies have examined patterns of free recall of numbers (Dale & Baddeley, 1966), and by using numbers as stimuli we have been able to provide detailed analyses of the recall of such stimuli. Using standard manipulations from the free-recall literature our experiments indicate that recall patterns of numbers follow well-established patterns (Bjork & Whitten, 1974; Farrell, 2012; Grenfell-Essam & Ward, 2015; Postman & Phillips, 1965; Spurgeon, Ward, & Matthews, 2014; Stefanidi, Ellis, & Brewer, 2018; Ward et al., 2010), though with some quantitative exceptions. In Experiments 1 and 2, we manipulated task order and effectively created an immediate and delayed recall task; in both experiments we see evidence of more primacy in immediate compared to delayed recall (Bjork & Whitten, 1974). Our manipulation in Experiment 3 created a dual-task scenario and we

saw a cost to memory recall accuracy when both tasks were completed. Taken together, the findings from the recall data are in line with our expectations from previous free-recall studies. There were however some cases where we observed flatter serial position functions or reduced primacy, which raises the question of whether there is something “special” about the recall of digits as opposed to words. This was examined to some extent by Dale and Baddeley (1966), who found that people are more likely to recall digits in ascending and descending order. Although our results do not suggest that recall of digits is notably different from recall of words, organizational differences may modulate the expression of recall mechanisms in patterns of recall.

Another interesting issue is whether the value of an item affects how much it influences evaluation. According to the peak-end rule, the highest and lowest values will influence evaluation to a greater degree than more central items (Fredrickson, 2000). This effect has been replicated in the decisions-from-experience literature whereby the best and worst outcomes in a decision context will be overweighted in both memory and risky choice (Ludvig, Madan, McMillan, Xu, & Spetch, 2018; Madan, Ludvig, & Spetch, 2014). Although the causality of this relationship is still under investigation, these findings are tied to the fact that the highest and lowest values are often better remembered (Madan & Spetch, 2012). The current experiments were not designed to address whether the value of the items made them more or less likely to influence evaluation and as such the items were sampled from a uniform distribution with a narrow range. It is possible that varying the distribution of items may shift people’s use of memory-based strategies.

We have provided a comprehensive examination of both the conditions under which people use, and the accuracy of, memory-based strategies in retrospective evaluations. Taken together, our results show strong evidence that online evaluation strategies are favoured under a wide range of conditions. However, there is some suggestion that preferences are sometimes determined at the point of evaluation, based on the recent recall of items from the sequence. The importance of memory to evaluation has been stressed in models of decision-making and

judgment in which samples of individual episodes or pieces of information are the fodder for evaluation (Bornstein et al., 2017; Gershman & Daw, 2017; Lieder et al., 2018; Stewart et al., 2006). The challenge presented by these data and others is understanding under what circumstances memory will be involved, which in turn will depend on task expectancy, effort discounting, and the goals of the participant.

References

- Aldrovandi, S., Poirier, M., Kusev, P., & Ayton, P. (2015). Retrospective Evaluations of Sequences. *Experimental Psychology*, 62(5), 320–334. <https://doi.org/10.1027/1618-3169/a000301>
- Baddeley, A. D., Scott, D., Drynan, R., & Smith, J. C. (1969). Short-term memory and the limited capacity hypothesis. *British Journal of Psychology*, 60(1), 51–55. <https://doi.org/10.1111/j.2044-8295.1969.tb01175.x>
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a Single Response Sequential Method. *Systems Research and Behavioural Science*, 9(3), 226–232. <https://doi.org/10.1002/bs.3830090304>
- Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, 6(2), 173–189. [https://doi.org/10.1016/0010-0285\(74\)90009-7](https://doi.org/10.1016/0010-0285(74)90009-7)
- Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, 8(1), 15958. <https://doi.org/10.1038/ncomms15958>
- Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition*, 113(3), 262–280. <https://doi.org/10.1016/j.cognition.2008.08.011>
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1). <https://doi.org/10.18637/jss.v080.i01>
- Craik, F. I. M., Naveh-Benjamin, M., Ishaik, G., & Anderson, N. D. (2000). Divided Attention during Encoding and Retrieval: Differential Control Effects? *Journal of Experimental Psychology: Learning Memory and Cognition*, 26(6), 1744–1749. <https://doi.org/10.1037/0278-7393.26.6.1744>

- Dale, H. C., & Baddeley, A. D. (1966). Remembering a list of two-digit numbers. *The Quarterly Journal of Experimental Psychology*, 18(3), 212–219.
<https://doi.org/10.1080/14640746608400032>
- Dixon, M., & Verma, R. (2013). Sequence effects in service bundles: Implications for service design and scheduling. *Journal of Operations Management*, 31(3), 138–152.
<https://doi.org/10.1016/j.jom.2012.12.002>
- Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review*, 106(1), 180–209.
<https://doi.org/10.1037/0033-295X.106.1.180>
- Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic memory. *Psychological Review*, 119, 223–271. <https://doi.org/10.1037/a0027371>
- Feldman, J. M., & Lynch, J. G. (1988). Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of Applied Psychology*, 73(3), 421–435.
<https://doi.org/10.1037/0021-9010.73.3.421>
- Frederick, S., & Loewenstein, G. (2008). Conflicting motives in evaluations of sequences. *Journal of Risk and Uncertainty*, 37(2-3), 221–235. <https://doi.org/10.1007/s11166-008-9051-z>
- Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition and Emotion*, 14(4), 577–606.
<https://doi.org/10.1080/026999300402808>
- Fredrickson, B. L., Kahneman, D., Hayes, A., & Knopoff, K. (1993). Duration Neglect in Retrospective Evaluations of Affective Episodes, 65(1).
- Gershman, S. J., & Daw, N. D. (2017). Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework. *Annual Review of Psychology*, 68(1), 101–128.
<https://doi.org/10.1146/annurev-psych-122414-033625>
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the Fast and Frugal Way. *Psychological*

- Review*, 103(4), 650–669. <https://doi.org/10.1037/0033-295X.103.4.650>
- Glanzer, M., & Cunitz, A. R. (1966). Two storage mechanisms in free recall. *Journal of Verbal Learning and Verbal Behavior*, 5(4), 351–360.
- Grenfell-Essam, R., & Ward, G. (2015). The effect of selective attention and a stimulus prefix on the output order of immediate free recall of short and long lists. *Canadian Journal of Experimental Psychology*, 69(1), 1–16. <https://doi.org/10.1037/cep0000045>
- Hastie, R., & Park, B. (1986). The Relationship Between Memory and Judgment Depends on Whether the Judgment Task is Memory-Based or On-Line. *Psychological Review*, 93(3), 258–268.
- Hoelzl, E., Kamleitner, B., & Kirchler, E. (2011). Loan repayment plans as sequences of instalments. *Journal of Economic Psychology*, 32(4), 621–631. <https://doi.org/10.1016/j.joep.2011.02.002>
- Hogarth, R. M., Einhorn, H. J., Levin, I., Judd, C., Loewenstein, G., Messier, W., & Tubbs, R. (1992). Order Effects in Belief Updating : The Belief-Adjustment Model, 55.
- Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 25(4), 923–941. <https://doi.org/10.1037/0278-7393.25.4.923>
- Kahneman, D., & Tversky, A. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458. <https://doi.org/10.1126/science.7455683>
- Kemp, S., Burt, C. D. B., & Furneaux, L. (2008). A test of the peak – end rule with extended autobiographical events. *Memory & Cognition*, 36(1), 132–138. <https://doi.org/10.3758/MC.36.1.132>
- Kitayama, S., & Burnstein, E. (1989). The relation between opinion and memory: Distinguishing between associative density and structural centrality. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 91–122). Lawrence Erlbaum Associates, Inc.

- Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, 139(4), 665–682. <https://doi.org/10.1037/a0020198>
- Langer, T., Sarin, R., & Weber, M. (2005). The retrospective evaluation of payment sequences: Duration neglect and peak-and-end effects. *Journal of Economic Behavior and Organization*, 58(1), 157–175. <https://doi.org/10.1016/j.jebo.2004.01.001>
- Lichtenstein, M., & Srull, T. K. (1987). Processing objectives as a determinant of the relationship between recall and judgment. *Journal of Experimental Social Psychology*, 23(2), 93–118. [https://doi.org/10.1016/0022-1031\(87\)90027-8](https://doi.org/10.1016/0022-1031(87)90027-8)
- Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology*, 89(1), 46–55. <https://doi.org/10.1037/h0031207>
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of Extreme Events in Decision Making Reflects Rational Use of Cognitive Resources. *Psychological Review*, 1, 1.
- Loewenstein, G., & Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review*, 100(1), 91–108. <https://doi.org/10.1037/0033-295X.100.1.91>
- Ludvig, E. A., Madan, C. R., McMillan, N., Xu, Y., & Spetch, M. L. (2018). Living near the edge: How extreme outcomes and their neighbors drive risky choice. *Journal of Experimental Psychology: General*, 147(12), 1905–1918. <https://doi.org/10.1037/xge0000414>
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, 21, 629–636. <https://doi.org/10.3758/s13423-013-0542-9>
- Madan, C. R., & Spetch, M. L. (2012). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica*, 139, 343–349. <https://doi.org/10.1016/j.actpsy.2011.12.010>

- Miller, J. F., Weidemann, C. T., & Kahana, M. J. (2012). Recall termination in free recall. *Memory & Cognition*, 40(4), 540–550. <https://doi.org/10.3758/s13421-011-0178-9>
- Mitchell, T. R., Thompson, L., Peterson, E., & Cronk, R. (1997). Temporal Adjustments in the Evaluation of Events: The "Rosy View". *Journal of Experimental Social Psychology*, 33(33), 421–448. <https://doi.org/10.1006/jesp.1997.1333>
- Montgomery, N. V., & Unnava, H. R. A. O. (2009). Temporal Sequence Effects : A Memory Framework. *Journal of Consumer Research*, 36(1), 83–92. <https://doi.org/10.1086/595278>
- Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of Bayes Factors for Common Designs. Retrieved from <https://cran.r-project.org/package=BayesFactor>
- Murdock. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64(5), 1–4. Retrieved from <http://psycnet.apa.org/journals/xge/64/5/482/>
- Murdock, B. B. (1965). Effects of a subsidiary task on short-term memory. *British Journal of Psychology*, 56(4), 413–419. <https://doi.org/10.1111/j.2044-8295.1965.tb00983.x>
- Naveh-Benjamin, M., Kilb, A., & Fisher, T. (2006). Concurrent task effects on memory encoding and retrieval: Further support for an asymmetry. *Memory and Cognition*, 34(1), 90–101. <https://doi.org/10.3758/BF03193389>
- Postman, L., & Phillips, L. W. (1965). Short-term temporal changes in free recall. *Quarterly Journal of Experimental Psychology*, 17(2), 132–138. <https://doi.org/10.1080/17470216508416422>
- Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, 66(1), 3–8. [https://doi.org/10.1016/0304-3959\(96\)02994-6](https://doi.org/10.1016/0304-3959(96)02994-6)
- Redelmeier, D. A., Katz, J., & Kahneman, D. (2003). Memories of colonoscopy : a randomized trial. *Pain*, 104, 187–194. [https://doi.org/10.1016/S0304-3959\(03\)00003-4](https://doi.org/10.1016/S0304-3959(03)00003-4)
- Robinson, E., Blissett, J., & Higgs, S. (2011). Peak and end effects on remembered enjoyment of

- eating in low and high restrained eaters §. *Appetite*, 57(1), 207–212.
<https://doi.org/10.1016/j.appet.2011.04.022>
- Rouder, J. N., Morey, R. D., Verhagen, J., Province, J. M., & Wagenmakers, E.-J. (2016). Is There a Free Lunch in Inference? *Topics in Cognitive Science*, 8(3), 520–547.
<https://doi.org/10.1111/tops.12214>
- Shohamy, D., & Daw, N. D. (2015). Integrating memories to guide decisions. *Current Opinion in Behavioral Sciences*, 5, 85–90. <https://doi.org/10.1016/j.cobeha.2015.08.010>
- Slovic, P. (1995). The construction of preference. *Decision Research*, 50(5), 371. Retrieved from <http://search.ebscohost.com/login.aspx?direct=true%7B/&%7Ddb=pdh%7B/&%7DAN=1995-35778-001%7B/&%7Dsite=ehost-live>
- Spurgeon, J., Ward, G., & Matthews, W. J. (2014). Why do participants initiate free recall of short lists of words with the first list item?: Toward a general episodic memory explanation. *Journal of Experimental Psychology: Learning Memory and Cognition*, 40(6), 1551–1567.
<https://doi.org/10.1037/xlm0000028>
- 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>
- Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1–26. <https://doi.org/10.1016/j.cogpsych.2005.10.003>
- Su, Y. S., & Yajima, M. (2015). Package ‘R2jags’. R package version 0.03-08.
- Varey, C., & Kahneman, D. (1992). Experiences extended across time: Evaluation of moments and episodes. *Journal of Behavioral Decision Making*, 5(3), 169–185.
<https://doi.org/10.1002/bdm.3960050303>
- Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. A. (2011). Does the brain calculate value ? *Trends in Cognitive Sciences*, 15(11), 546–554. <https://doi.org/10.1016/j.tics.2011.09.008>

- Wagenmakers, E. J., Wetzels, R., Borsboom, D., & Maas, H. L. J. van der. (2011). Why Psychologists Must Change the Way They Analyze Their Data: The Case of Psi: Comment on Bem (2011). *Journal of Personality and Social Psychology*, 100(3), 426–432. <https://doi.org/10.1037/a0022790>
- Ward, G., Tan, L., & Grenfell-Essam, R. (2010). Examining the relationship between free recall and immediate serial recall: the effects of list length and output order. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 36, 1207–1241. <https://doi.org/10.1037/a0020122>
- Warren, C., McGraw, A. P., & Van Boven, L. (2011). Values and preferences: Defining preference construction. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(2), 193–205. <https://doi.org/10.1002/wcs.98>
- Weber, E. U., & Johnson, E. J. (2006). Constructing Preferences from Memory. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1301075>
- Wirtz, D., Kruger, J., Scollon, C. N., & Diener, E. (2003). What to do on Spring Break? The Role of Predicted, on-Line, and Remembered Experience in Future Choice. *Psychological Science*, 14(5), 520–524. <https://doi.org/10.1111/1467-9280.03455>

Table 1

List of models tested across the experiments

Predictor	Description
Expected Value	<i>Average of all values in the sequence</i>
Free	<i>Each item in the list can take on its own parameter in the model</i>
Serial Position	<i>For each individual each value in the sequence is weighted by the participants' average memory accuracy at that position</i>
Recall-based	<i>The average of all the values recalled for each trial, for each individual</i>
Temporal Difference	<i>A recency weighted learning rule</i>

Table 2

Bayes Factor Interpretation Guidelines according to Wagenmakers et al., (2011), where H_1 is the model in the numerator, and H_0 is the model in the denominator (usually the null hypothesis).

>100	Extreme evidence for H_1
30–100	Very strong evidence for H_1
10–30	Strong evidence for H_1
3–10	Substantial evidence for H_1
1–3	Anecdotal evidence for H_1
1	No evidence
1/3–1	Anecdotal evidence for H_0
1/10–1/3	Substantial evidence for H_0
1/30–1/10	Strong evidence for H_0
1/100–1/30	Very strong evidence for H_0
<1/100	Extreme evidence for H_0

Table 3

Confusion matrix for model simulations for Experiment 1

Generating Model	Recall Based	Serial Position	Expected Value	Free	Temporal Difference
Recall Based	50	0	0	0	0
Serial Position	0	50	0	0	
Expected Value	0	0	50	0	0
Free	0	0	50	0	0
Temporal Difference	0	0	0	0	50

Table 4

*Bayes Factor model comparisons for main analysis in each experiment. The first column lists each of the models run in the analysis for the experiments. The model marked with a * has the highest Marginal Likelihood and the Bayes Factors are calculated with respect to that model. The higher the Bayes Factor the more evidence there is in favour of the best fitting model.*

Model Name	Bayes Factor
Experiment 1	
Recall-based	1.55e+13
Serial Position	1.74e+10
Expected Value	*
Free	4.99e+08
Temporal Difference	1.37e+15
Experiment 2	
Recall-based	1.67e+79
Serial Position	7.90e+74
Expected Value	*
Free	3.56e+75
Temporal Difference	5.78e+75
Experiment 3	
Serial Position	787.35
Expected Value	*
Free	244.76
Temporal Difference	68.25

Table 5

Bayes Factor model comparisons for each additional analyses

Model Name	Bayes Factor
Experiment 1 Order	
Evaluation - Recall-based	9.32e+10
Evaluation - Serial Position	1.26e+03
Evaluation - Expected Value	*
Evaluation - Free	5.25e+04
Evaluation - Temporal Difference	3.92e+01
Recall- Recall-based	*
Recall- Serial Position	1.56e+07
Recall - Expected Value	3.16e+00
Recall - Free	6.73e+06
Recall - Temporal Difference	3.61e+11
Experiment 2 Order	
Evaluation - Recall-based	4.48e+61
Evaluation - Serial Position	1.91e+47
Evaluation - Expected Value	*
Evaluation - Free	1.33e+48
Evaluation - Temporal Difference	1.21e+53
Recall- Recall-based	*
Recall- Serial Position	4.40e+52
Recall - Expected Value	1.00e+00
Recall - Free	3.13e+54
Recall - Temporal Difference	5.73e+54
Experiment 2 Prioritized Task	
Evaluation - Recall-based	1.04e+00
Evaluation - Serial Position	4.56e+42
Evaluation - Expected Value	*
Evaluation - Free	4.56e+42
Evaluation - Temporal Difference	3.53e+43
Recall- Recall-based	1.79e+37
Recall- Serial Position	2.66e+38
Recall - Expected Value	*
Recall - Free	1.90e+44
Recall - Temporal Difference	3.99e+44
Experiment 2 Reported Strategy	
Prioritized - Recall-based	9.84e+22
Prioritized - Serial Position	3.28e+06
Prioritized - Expected Value	*
Prioritized - Free	7.27e+14
Prioritized - Temporal Difference	1.86e+51
Both- Recall-based	6.44e+44
Both- Serial Position	1.91e+51
Both - Expected Value	*
Both - Free	1.67e+53
Both - Temporal Difference	1.86e+51

Experiment 1

Experiment 2

Experiment 3

Figure 1. Serial Position Curves for each experiment showing the mean proportion of items recalled from each list length as a function of serial position

Experiment 1

Experiment 2

Experiment 3

Figure 2. First Recall Probability curves showing the probability of first recall as a function of serial position

Experiment 1

Experiment 2

Experiment 3

Figure 3. Evaluation task accuracy: Root mean squared error of WTP estimate in Galactic Credits

Experiment 2

Figure 4. Experiment 2 Recall Accuracy in Galactic Credits - 1 item is equal to 10 GCs

Experiment 3

Figure 5. Predicted Evaluation Task Accuracy from memory recall in GC - Root mean squared error of WTP estimate