Value bias of verbal memory

Sucheta Chakravarty[1], Esther Fujiwara[3,2], Christopher R. Madan[1,4],

Sara E. Tomlinson[3], Isha Ober[2], and Jeremy B. Caplan[1,2]

[1]Department of Psychology, [2]Neuroscience and Mental Health Institute,

and [3]Department of Psychiatry, University of Alberta, Edmonton, AB, Canada;

[4]School of Psychology, University of Nottingham, Nottingham, UK.

Author Note

Corresponding author: Jeremy Caplan. jcaplan@ualberta.ca. Department of Psychology, Biological Sciences Building, University of Alberta, Edmonton, AB, T6G 2E9, Canada, Tel: +1.780.492.5265, Fax: +1.780.492.1768.

We thank Jeff Pisklak and Brenna Schuldhaus for helping with participant testing. Supported in part by the Alberta Gaming Research Institute, Natural Sciences and Engineering Research Council of Canada (NSERC), and Canadian Institutes of Health Research (CIHR).

Abstract

A common finding is that items associated with higher reward value are subsequently remembered better than items associated with lower value. A confounding factor is that when a higher value stimuli is presented, this typically signals to participants that it is now a particularly important time to engage in the task. When this was controlled, Madan et al. (2012) still found a large value-bias of memory. Their value-learning procedure, however, explicitly pitted high- against low-value words. Our novel value-learning procedure trained words one at a time, avoiding direct competition between words, but with no difference in words signalling participants to engage in the task. Results converged on null effects of value on subsequent free recall accuracy. Re-analyses attributed Madan et al.'s value-bias to competition between choice items that were paired during learning. Value may not bias memory if it does not signal task importance or induce inter-item competition.

Keywords: reward value; free recall; choice; lexical decision

Value bias of verbal memory

Introduction

Consider a situation in which a person sees items one at a time, and can bet on, or against, each item. Ample research has investigated how people learn, via feedback, which items are associated with a high or a low payout. Here we are interested in a possible downstream phenomenon. In some learning situations, it may be adaptive to let reward-value influence learning about the *items themselves*. Depending on their payout, high-value items may continue to lead to high reward in the future (Anderson, 2013). In Estes' (1972) restaurant metaphor, a customer learns, by trial-and-error, which menu items are associated with high- versus low reward-outcome: If choosing the Alberta bison steak consistently yields reward (a tasty meal) compared to the Ontario rainbow trout (seasoning was less than inspired), the customer may not only want to remember the item–reward relationship, but anything more there is to know about the highly rewarding item— how big the steak was, which farmer raised it, which sides went well with it; even recognizing the bison steak faster on a different menu may be advantageous. Further information about the low-reward item (trout, in this example) may be less useful.

Thus, a bias to learn more about items previously associated with high-value than those associated with low-value may be strategic, and this ability is necessary for effective decision-making, as has been suggested (Shohamy & Adcock, 2010). Alternatively, the inability to overcome a value-bias of memory might lead to maladaptive reward-seeking behaviour such as in gambling situations. For example, previous knowledge about high-value items can lead to memory illusions like overweighting wins and the near-miss effect (e.g., Clark, Crooks, Clarke, Aitken, & Dunn, 2012; Madan, Ludvig, & Spetch, 2014, 2017; Kassinove & Schare, 2001; Winstanley, Cocker, & Rogers, 2011). In the earlier metaphor, the customer might continue to choose a poorly cooked bison steak, even though it was delicious only one time, and miss out on the tasty rainbow trout.

There is plenty of evidence that information accompanied by high reward is remembered better than information accompanied by low reward (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Castel, Benjamin, Craik, & Watkins, 2002; Watkins & Bloom, 1999; Wittmann et al., 2005; Wittmann, Dolan, & Düzel, 2011). For example, when directly instructed to learn words associated with different point values, participants selectively biased their encoding towards higher-value words (e.g., Adcock et al., 2006; Castel et al., 2002; Castel, Farb, & Craik, 2007; Castel, Balota, & McCabe, 2009; Castel, Lee, Moore, & Humphreys, 2011; Watkins & Bloom, 1999; Weiner & Walker, 1966), even when they either know, at encoding, that values will play no role later in retrieval; or are told, at retrieval, to ignore values (e.g., Cuvo, 1974; Gruber & Otten, 2010; Tarpy & Glucksberg, 1966; Weiner, 1966).

Similar effects are found for paired-associate memory (e.g., Harley, 1965; Kuhl, Shah, DuBrow, & Wagner, 2010; Loftus & Wickens, 1970; Soderstrom & McCabe, 2011; Wolosin, Zeithamova, & Preston, 2012). Reward-predicting cues are associated with activation in the subcortical dopaminergic regions of the midbrain such as the ventral tegmental area (VTA) and ventral striatum. Neuroimaging studies (e.g., Adcock et al., 2006; Knutson, Westdorp, Kaiser, & Hommer, 2000) suggest that this value-based memory enhancement is mediated by structural and functional connectivity from the midbrain dopamine systems to the medial temporal lobe (MTL) including the hippocampus, which is related to declarative memory formation. If dopaminergic modulation enhances the function of MTL regions, this may result in better encoding of high- over low-rewarding items.

Importantly, as pointed out by Madan, Fujiwara, Gerson, and Caplan (2012), the value of an item can also influence its usefulness for the participant in earning reward, completely distinct from any subjective preference. Many existing value-learning paradigms are designed such that the objective is to maximize total amount of reward and remembering high-value items at the expense of low-value items may provide the optimal way of achieving this objective. Thus, each high- and low-value item not only differs by the value associated with it, but is also weighted differently in its usefulness towards performing optimally in the task. Indeed, it is important to test if humans possess the capability to adopt such a rational strategy as per task demand but this makes it difficult to determine whether results are due to participants having a memory-bias, or simply follow from a rational strategy. Put differently, under such circumstances, value could be telling the participant which information during the task demand more versus less engagement, thus manipulating engagement and thereby having a secondary effect on memory. In contrast, if value influences memory beyond simply signalling (rational) usefulness of the item within the task, value could still bias memory, even when task-usefulness has been controlled. In other words, we ask, when it is neither a task-demand nor a rational strategy to prioritise high- over low-value items, will the participants still elicit a bias in favour of the high-value items?

The Monetary Incentive Encoding (MIE) task, introduced by Adcock et al. (2006), adapting a procedure by Knutson et al. (2000), is an example of a paradigm for which usefulness varies alongside value. In the MIE, prior to studying each of a set of images, one of two reward symbols appears, indicating that the upcoming image will lead to either a high or a low reward if they remember the image on a later memory test. Thus, images are associated with different reward-values, but the participant has a clear incentive to apply more effort and resources into encoding high-value than low-value images because high-value images have greater usefulness (can lead to more reward) than low-value images. Indeed, it is precisely the goal of such studies to examine different resource allocation to encoding of high- versus low-value information. We reason that such a task may very well be testing how participants optimally prioritise their encoding (and retrieval)— as would be rational for the task, and as they were directly told to do. Thus, the confound regarding usefulness is important in understanding the precise interpretation of such studies; for example, implicit in the dopamine hypothesis is that stimulus-value increases dopamine levels, and dopamine, in turn, modulates episodic encoding. Alternatively, however, stimulus-value may influence usefulness and usefulness may in turn independently influence both dopamine levels and engagement in the task, with spillover effects on encoding that are not directly caused by dopaminergic modulation. We are therefore interested in whether value will continue to bias memory, even when both high- and low-value items are equally useful in earning the same amount of reward.

Madan et al. (2012) found large-magnitude reward-value memory biases, even while controlling for usefulness, with unrewarded memory tests. Participants had to learn the values of a large set (36) of words through trial-and-error, with feedback. Half the words were designated high-value and half were designated low-value. During value-learning, participants chose between a random pair consisting of one high-value word and one low-value word. Because the response was always forced-choice between one high- and one low-value word, high- and low-value words had equal usefulness. That is, knowing that the high-value word was 'high' was theoretically just as useful in collecting points (and subsequent bonus payment) as knowing that the low-value word was 'low': Choosing the former had the same high-reward outcome as *not* choosing the latter. Because of this symmetry, any effect of value on memory for the individual words later could be interpreted as a bias that the participants brought to the task, rather than simply optimal encoding that might have been demanded by the task. After the value-learning task, participants were presented with each item individually for a lexical decision test (judging whether each item was a word or a non-word; response time was the dependent measure) and then given a free recall test (attempt to recall all the words from the experiment, in any order). Two important features of the procedure should be noted: first, value-learning was rewarded but the subsequent memory tests were not. Second, during value-learning, participants were not asked to remember the words themselves, only to learn to make optimal choices between pairs of items. Thus, the item-memory tests followed incidental item-learning but intentional value-learning. Results showed that previously learnt high-value words had much higher free-recall probabilities and faster lexical-decision times than previously learnt low-value words. Interestingly, the effect of value on these two tasks was slightly (but significantly) negatively correlated, suggesting that value can influence memory distinctly from its influence on access speed.

Critical to understanding Madan et al.'s results, reward value, without controlling for usefulness, has previously been found to prioritise information in attention, even when high-value information is not otherwise more salient, or even task-irrelevant (Anderson, 2013). Because a high-value and a low-value word were always pitted against one another during value-learning in Madan et al.'s experiment, value could have biased attention competitively, in favour of high-value words, directly at the level of individual trials. Competition between simultaneously presented words should be all the more pronounced, since the participants' goal was to optimize reward during learning. Estes (1976) suggested participants might rehearse high-value words more than low-value words. Given that rehearsal acts much like additional encoding (e.g., Brodie & Murdock, 1977; Ward & Tan, 2004), this is in line with Cuvo (1974), who found that explicit incentive value did indeed increase overt rehearsal and free-recall probability. Thus, it is possible that the value biases Madan et al. (2012) had observed in free-recall (and in lexical decision) may have been produced by such competitive effects during value-learning. Our main goal was to find out whether the value-bias of memory, as measured by free recall, is entirely due to such direct competition during value-learning; if not, the value-bias should remain even if items were not directly pitted against one another during value-learning.

To this end, in the current experiments we used a one choice value-learning task that also controlled for usefulness of high-value and low-value words (Figure 1). On each trial of value-learning, the participant had to choose the word, or choose the fixed standard letter-string, HHHHH. Participants were explicitly instructed that on each trial, the word would either be high value or low value. The phrasing was that a high word is "worth" 10 points if chosen, and a low word is "worth" 1 point if chosen. High-value words were to be chosen and low-value words not-chosen and the letter string HHHHH was to be chosen instead. If the participant made the correct response (choosing a high-value word or not-choosing a low-value word and choosing the letter string instead), ten points would be earned; otherwise, one point would be earned. Choosing an item can increase participants' preference for the item, whereas not-choosing reduces preference (Leotti & Delgado, 2011; Sharot, De Martino, & Dolan, 2009) and affective value (Ferrey, Frischen, & Fenske, 2012). The act of choosing also increases activity in the ventral tegmental area, a major source of dopaminergic modulation of memory-relevant brain regions (Leotti & Delgado, 2011). Thus, when responded correctly for a number of trials, high- and low-value words could presumably differ on average in subjective preference and affective value, perhaps with a concurrent increase in dopamine response during the trial (although this was not measured). Notably, choosing versus not-choosing has been reported by Coverdale, Pandeirada, and Nairne (in press) to bias free recall, with no manipulation of value, with Coverdale and Nairne (in press) ruling out a task-congruity account. However, unlike Madan et al. (2012) or Coverdale et al. (in press) or Coverdale and Nairne (in press), in our procedure, the high- and low-value words were not directly pitted against each other on a trial-by-trial basis, which could attenuate the choice bias. Further, in our procedure, correct responses to high-value words entailed making a spatially congruent response, whereas correct responses to low-value words entailed making a spatially incongruent response. Although we had no "no-go" condition, it is possible that the incongruent response is related to no-go behaviour, in which case, learnt, low-value words may have reduced affective value compared to learnt, high-value words (Frischen, Ferrey, Burt, Pistchik, & Fenske, 2012).

One published procedure that avoids the usefulness confound was conducted by Murayama and Kitagami (2014). Their procedure eliminated the direct association between the reward and the stimuli in an incidental learning paradigm by presenting them in different contexts, but still within the same trial. Therefore, the participants would have no extra incentive to encode the stimuli from the reward trials better than the control and accordingly, in an item recognition test that immediately followed the task, memory for stimuli from the reward trials were not found to be different than that for control trials (unrewarded).[1] Thus, it is plausible that the effect of value in the Madan et al. (2012) study will not hold when high- and low-value items are not placed in direct competition with one another during study.

If value can bias memory even when there is no rational reason to do so, high-value words

[1] Our calculations based on their reported statistical tests produced a Bayes Factor of 0.97, thus inconclusive. Murayama and Kitagami's data suggest that the picture might be different in a one-week delayed test, for which they did find an effect of rewarded versus unrewarded trials, but this effect was just under-threshold; by our calculations, the Bayes Factor was 2.7. When combined with their internal replication, the Bayes Factor reduces to 0.57, considered inconclusive by convention. In an additional experiment manipulating reward, but in the negative direction, unrewarded versus reduction in reward trials, there was strong evidence for the null, $BF_{10} = 0.16$.

would draw more processing and encoding than low-value words. Consequently, we predicted that free recall would be facilitated for high-value compared to low-value words. Alternatively, if value, with usefulness controlled, did not produce a bias in free recall, it may be argued that usefulness, not value, invokes the dopaminergic signal thereby modulating cognitive resources and subsequent memory outcome.

In Experiments 1 and 2, value-learning (rewarded) was followed by a test of free recall of all trained words, unrewarded. An unrewarded lexical-decision test was also included, for two reasons: first, for continuity with Madan et al. (2012), to simply find out whether the value-bias of lexical decision remained after direct competition during value-learning was eliminated; and second, in case a value-bias of free recall were found, to be able to specify whether it might have the same underlying cause as the lexical-decision effect, in that case, potentially related to priming, or distinct, as found by Madan and colleagues. Experiment 3 followed up on a small value-bias of free recall found in Experiment 2. This experiment included a rewarded form of the free-recall test, to determine if the value-bias of free recall might have been a side effect of value biasing output-order of recalled words. In a final section, we conduct additional analyses of the data reported by Madan et al. (2012), to test possible mechanisms of the value bias in two-word value-learning procedures.

Experiment 1

Methods

Participants. A total of 148 introductory psychology students at the University of Alberta participated for partial fulfillment of course credit. Five participants did not complete the experiment due to either machine error ($N = 4$) or because they fell asleep ($N = 1$), yielding a total of 143 participants who completed the experiment. All participants had learnt English before the age of six, had normal or corrected-to-normal vision, and were comfortable typing. Participants gave written informed consent prior to the study, which was approved by a University of Alberta research ethics board.

Materials. A total of 72 words were selected from the list of 160 words previously used in Madan et al. (2012). Briefly, the words were selected such that: (a) they were six to seven letters in length and had exactly two syllables; (b) word imageability and frequency were held at mid-levels based on norms from the MRC Psycholinguistic Database (Wilson, 1988); and (c) they did not have emotional connotations (based on normative ratings from the ANEW database Bradley & Lang, 1999). Of these, 48 were selected at random for each participant, to be used in the value-learning task. In turn, half (24) were randomly designated high-value, the other half (24) designated low-value.

For the lexical decision task, 72 pronounceable non-words were selected from the same set as used in Madan et al. (2012) and Madan and Spetch (2012). These non-words were generated using LINGUA (Westbury, Hollis, & Shaoul, 2007), using a pre-compiled word-frequency dictionary (Shaoul & Westbury, 2006), and were matched in length to the words. Non-words were generated using a Markov chaining length of three (Westbury et al., 2007), as we have done previously (Madan, Shafer, Chan, & Singhal, 2017; Madan & Singhal, 2012).

Procedure. Participants were informed that they would perform a “word choice” task, hereafter referred to as the value-learning task. At the beginning of the experiment, they were informed that they would receive a monetary bonus of up to $5 based on their performance in this task, in addition to their partial course credit. The value-learning task was followed by unexpected and unrewarded tests of lexical decision and free recall. Since the value learning phase did not require participants to learn the items themselves, but only to learn to make reward-maximizing choices, free recall can be considered to measure incidental item-learning.

Between value-learning and lexical decision, as well as between lexical decision and free recall, participants completed a brief distractor task consisting of 10 randomly selected math questions. Each question was presented in the form of $A+B+C=$ __, where A, B, and C were randomly selected integers from two to eight, inclusive. Each equation was presented at the centre of the screen and remained until the participant typed a number and pressed the ‘ENTER’ key. Each question was followed by a 500-ms inter-trial interval.

Value learning. Participants were presented with two strings of letters, a word (which varied from trial to trial) and the letter string, 'HHHHH' (illustrated in Figure 1). Each word occurred equally often on the left or the right side of the screen, with HHHHH on the other side. Participants were instructed to choose the word or the HHHHH string using the 'Q' and 'P' keys of a computer keyboard, for the left and right options, respectively. Depending on the chosen option, participants were presented with visual and auditory feedback indicating high or low reward. Visual feedback was in form of photos of a single coin or a pile of coins shown at the center of the screen for 1,500 ms. The message, 'Points Accumulated,' was displayed, followed by the number of points earned in the current trial (i.e., low reward: 1 point, single coin picture; high reward: 10 points, pile of coins picture), as well as their cumulative sum across all trials. Auditory feedback were respective sound clips, either many coins being dispensed or a single coin hitting a table, played simultaneously to the visual feedback. Auditory feedback was provided via headphones, which participants were asked to wear during the value-learning task. Feedback was followed by a blank screen for a duration drawn from a uniform random distribution between 500 and 700 ms. Within each cycle, this procedure was repeated for a total of 48 words. For the 24 of these words randomly assigned to be 'high value,' choosing the word (responding on the same side as the word) deterministically yielded a 10-point reward, whereas choosing HHHHH yielded a 1-point reward. For the remaining 24 words, assigned to be 'low value,' choosing the word yielded a 1-point reward, and choosing HHHHH yielded a 10-point reward. That is, *every* trial could result in a 10-point reward: if the participant bet on high-value words and if they bet against low-value words by choosing HHHHH. Participants completed 16 cycles of 48 trials, with the words presented in a new random sequence in each cycle.

There was no time-limit to select the word or HHHHH. To prevent participants from responding before processing the word, and to reduce accidental keypress responses, there was a minimum response time (RT) of 200 ms; if a participant made a choice in less than 200 ms, they were presented with the text 'Response too fast. Invalid trial.' for 5 s and earned no points on that trial. This removed any incentive to speed through the task. The task was preceded by two practice trials (one low-value trial and one high-value trial), where participants were explicitly instructed that choices can result in either 10 points or 1 point. The two words from the practice trials were not used in the main experiment.

Lexical decision. Participants were instructed that they would be presented with many letter strings, one at a time; some of the letter strings would be words, others would not, and that they had to decide whether the letter string is a word or not. If a letter string was a word, participants were instructed to press ‘P’; if the letter string was not a word, to press ‘Q’. Participants were explicitly told ‘You will NOT receive any bonus payment for this task.’ Participants were also told ‘Note: The letters will be shown very quickly; if you miss the letters please make your best guess.’ The first trials were practice examples (3 words, 3 non-words).

On each trial, participants were shown a fixation cross for 1000 ms. The letter string was then presented for 200 ms. The screen was then blank until a key was pressed. The lexical decision task consisted of 144 trials: 72 words and 72 non-words (described in Materials). The 72 words were comprised of 24 words each of three types: previously assigned to be high value (in the value-learning task), previously assigned to be low value, and ‘new’ words not previously presented in the experiment session.

Free recall. Participants were next given 5 minutes to recall all of the words they could from the experiment session, in any order. Words were recorded when the participant typed the word and then pressed the ‘ENTER’ key. After each response and ‘ENTER’ keypress, participants were presented with a yellow square to confirm that their response was recorded but gave no indication of accuracy. Participants were explicitly told ‘You will NOT receive any bonus payment for this task.’

At the end of the experiment session, participants received an honorarium of up to $5, based on their performance in the initial value-learning task. The honorarium was calculated as the total number of points that the participant earned multiplied by 0.000625. The result was rounded up to the nearest quarter dollar.

Data analysis

Free recall. Repetitions were ignored: a word was considered recalled if typed correctly. Proportions of all and learnt high- and low-value words recalled were measured for subsequent analysis.

Lexical decision. Responses were eliminated from analyses if they exceeded the participant's mean plus three standard deviations, and only response times from correct trials were analyzed. Median response time in each condition (high, low, and new) was computed for each participant before entering them into ANOVAs.

Results and discussion

Value learning. Figure 2a plots the value learning curve. Accuracy on the first cycle was greater for high-value than low-value words. Because no values could be known during the first cycle, this indicates that when guessing, participants had a tendency to prefer to choose the word as opposed to the 'HHHHH' option. Over the course of learning, this word-choice preference greatly reduced, and was no longer significant in the last cycle (16^{th}), $t(142) = 1.20,\ p = 0.23$. Figure 2e plots the distribution of mean accuracy over the last four cycles of value-learning.

On the first cycle, response times on correct trials (all of which must have been guesses) were no different for high-value (choose response) and low-value (not-choose) words, $t(140) = -0.45,\ p = 0.65$ (Figure 2c). As the values were learnt, correct responses to high-value words became significantly faster than correct responses to low-value words (last cycle: $t(142) = -2.08,\ p = 0.039$).

Lexical decision. Accuracy in the lexical decision task was high, particularly for words that were in the value-learning task (Mean ± SEM: high = 0.9388 ± 0.0090, low $= 0.9289 \pm 0.0085$, New$= 0.8409 \pm 0.0117$, Nonwords $= 0.8564 \pm 0.0111$). The effect of value on accuracy (difference between high- and low-value words) was a non-significant trend ($p = 0.07$). Correct lexical decisions (Figure 3a) were nominally faster in response to high- than to low-value words, but this was not significant ($p = 0.30$). We wondered if an underlying value bias might have been washed out by the inclusion of words whose values were not known (or not very well) to the participant. To test this, we asked if this trend might be significant for the subset of words whose values were learnt during value-learning. We measured this subset for each participant, by finding the words that were responded correctly on all four final value-learning cycles. Indeed, we found that the lexical decision time difference for value-learnt words was significant (Figure 3c), $t(142) = -2.50,\ p < 0.05$ (difference $\simeq$20 ms), correct responses to high-value words being faster than those to low-value words.

Free recall. Free-recall rates did not differ by value, $t(135) = 1.14,\ p = 0.26$. When restricting the analysis to words whose values were learnt, there was still no significant effect of value on free-recall, $t(133) = 0.36,\ p = 0.72$ (Figure 4a). We followed up this analysis with Bayesian *t* tests (conducted with a MATLAB function written by Schwarzkopf, 2015), which produces a Bayes Factor. The advantage is that Bayesian model comparison techniques provide us with support for one model over another, in contrast to classical hypothesis testing, which looks for evidence against only one model (the null hypothesis). The Bayes Factor, which quantifies this, is the ratio of Bayesian probabilities for the alternative and null hypotheses; $BF_{10} = \frac{p(H_1)}{p(H_0)}$. For smaller values of this ratio, the null is more strongly supported by the data under consideration than the alternate hypothesis, whereas a greater value of this ratio would indicate otherwise. By convention (Kass & Raftery, 1995), there is “some” evidence for the null when $BF_{10} < 0.3$, and correspondingly, “some” evidence for the alternate hypothesis when $BF_{10} > 3$. “Strong” evidence is infered when $BF_{10} < 0.1$ or > 10. Accordingly, our results provided *some* evidence favouring the null with a $BF_{10} = 0.18$ for all value-trained words and strong evidence favouring the null ($BF_{10} = 0.10$) when restricted to learnt words only. Thus, our data are consistent with no effect of value on free recall.

Figure 4c plots the distribution of total number of free-recalls of high- and low-value words. We note that 22 participants had $\leq$5 recalls (high- and low-value words combined). It is possible that these participants lacked the motivation to produce higher number of free recalls. We explore this possibility in Experiment 3.

In sum, training the value of only one word at a time, and controlling for usefulness, many participants reached a near-perfect value-learning criterion. A value-bias of lexical decision times was significant when restricted to words whose values had clearly been learnt. Thus, weaker but in some sense, we found an extension of the boundary conditions on the value-bias of lexical decision found by Madan et al. (2012). However, our expected effect on free recall was not found; there was no significant effect of value on free-recall probability, and the Bayesian t test favoured a no-difference interpretation.

Experiment 2

Participants learned the values of the words very quickly in Experiment 1. We wondered if the lack of influence of value on free recall may have been because value-learning was too explicit or because values were over-learned. Perhaps if participants had less ready access to word-values, a memory bias would emerge. In Experiment 2, we switched to probabilistic reward-values, with 90:10 ratios. The probabilistic value-learning procedure made learning more difficult. However, in our earlier attempts with the 90:10 procedure we also found a number of participants with virtually no value learning. Thus, to partly offset this, we decreased the set size from 48 to 36 words. This prevented elicitation of ceiling accuracies among the participants without too many participants with no value learning. We also manipulated task-order; one group of participants performed lexical decision prior to free recall (Group LD-FR), and the other group performed free recall prior to lexical decision (Group FR-LD). This enabled us to test if a value-bias of free recall might emerge without the intervening lexical decision task.

Methods

Participants. A total of 201 introductory psychology students at the University of Alberta participated for partial fulfillment of course credit. Twelve participants did not complete the experiment due to either machine error ($N = 10$) or because they fell asleep ($N = 2$), yielding a total of 189 participants who completed the experiment. All participants had learnt English before the age of six, had normal or corrected-to-normal vision, and were comfortable typing.

Participants gave written informed consent prior to the study, which was approved by a University of Alberta research ethics board. None of the participants in Experiment 2 were included in Experiment 1.

Materials. A total of 54 words and 54 non-words were selected from those used in Experiment 1.

Procedure. The basic procedure was the same as Experiment 1. The value-learning task was followed by unexpected and unrewarded tests of lexical decision and free recall, the order of which were now counterbalanced across participants. The lexical-decision task was identical to Experiment 1, except that there were only 108 trials (after the initial six practice trials): 54 words and 54 non-words. The 54 words included the 36 value-trained words (18 previously high-value and 18 previously low-value) and 18 'new' words not part of the value-learning. The free-recall task was identical to Experiment 1, apart from the smaller set size.

Value learning. The value-learning procedure was identical to Experiment 1 (i.e., response keys, position of words/images on the screen, timing). As before, every trial could result in the participant earning a 10-point reward. Of 36 words, 18 were randomly assigned to be 'high' value, and the remaining, 'low' value. Participants completed 20 cycles of 36 trials each, with the words presented in a new random sequence in each cycle.

Outcomes occurred probabilistically, with a ratio of 90:10. Thus, choosing a specific high-value word every time over the first ten cycles, rather than HHHHH, resulted in a 10-point reward nine times and a 1-point reward one time. If HHHHH was chosen, this would yield the opposite outcome, as in Experiment 1. For low-value words, the reward contingency was reversed, such that choosing the word led to a 1-point reward for nine out of ten trials, but resulted in a 10-point reward on one trial out of ten.

The honorarium was calculated as the total number of points that the participant earned multiplied by 0.0007, rounded up to the nearest quarter dollar.

Results and discussion

Value learning. Figure 2b plots the value learning curve. Note that because rewards were probabilistic, a response is scored as "accurate" if it would be the best bet given full knowledge of values— betting on a high-value word and betting against a low-value word— regardless of actual reward outcome. As in Experiment 1, there was a word-choice bias which decreased over the course of learning, but here, remained significant; last (20$^{\text{th}}$) cycle, $t(188) = 3.85,\ p < 0.001$.

As in Experiment 1, on the first cycle, response times on correct trials did not differ, $t(188) = -0.80,\ p = 0.42$, suggesting no difference in making spatially congruent versus incongruent responses. After the first few cycles, response times on correct trials (Figure 2d) were slower for low- than for high-value words; last (20$^{\text{th}}$) cycle, $t(186) = -2.00,\ p = 0.047$. Figure 2f plots the distribution of mean accuracy over the last four cycles of value-learning, illustrating that value-learning in Experiment 2 was more difficult than in Experiment 1.

Lexical decision. Lexical decision accuracy did not differ significantly between high-value (0.869 ± 0.016) and low-value (0.868 ± 0.016) words, $t(188) = 0.047,\ p = 0.96$, nor when broken down by task order (Group LD-FR: $t(96) = 1.68,\ p = 0.10$; Group FR-LD: $t(91) = -1.47,\ p = 0.15$). Figure 3b plots the mean lexical decision times (correct responses only) for High, Low and New probes. Lexical decision times were not significantly different in response to high-value (601.0 ± 8.6 ms) than low-value (605.1 ± 8.0 ms) words, $t(180) = -0.90,\ p = 0.37$, nor when each group was analyzed individually (Group LD–FR: $t(90) = -0.83,\ p = 0.41$; Group FR–LD: $t(89) = -0.46,\ p = 0.64$).

When restricted to words whose values had been learnt in value-learning (correct responses on the last four training cycles), the effect of value was still non-significant (high: 622.5 ± 13.2 ms, low: 632.9 ± 10.5 ms, $t(148) = -0.76,\ p = 0.45$). Broken down by task-order group (Figure 3d), the effect was significant for Group LD-FR (high: 592.6 ± 12.5 ms, low: 615.1 ± 13.6 ms, $t(75) = -2.29,\ p < 0.05$), but not for Group FR-LD ($t(72) = 0.08, p = 0.94$). Thus, value facilitated lexical decision, but only reliably when the lexical-decision task immediately followed value-learning. Note that lexical decision also preceded free recall in

Experiment 1, where we did find a significant value-bias of lexical decision for words whose value had been learnt, and of similar magnitude ($\sim$20 ms in Experiment 1 and $\sim$22 ms in Experiment 2, Group LD-FR).

Free recall. As in Experiment 1, free-recall rates did not differ by value, $t(182) = 0.12,\ p = 0.90$, and the comparison remained non-significant when broken down by task order ($p > 0.8$). However, when restricted to words whose values were learnt by the end of value-learning, there was a significant modulation of free recall by value. For the full sample, 28.3% of high-value words learnt in value-learning were recalled in free-recall, compared to 25.0% of low-value words, $t(187) = 2.88,\ p < 0.01,\ \text{Cohen's}\ d = 0.15$ (Figure 4b). The effect was in the same direction for both task orders, and was significant for Group FR-LD ($M_{high} = 29.9\%,\ M_{low} = 25.7\%, t(90) = 2.74,\ p < 0.01,\ \text{Cohen's}\ d = 0.18$), but not Group LD-FR ($M_{high} = 26.8\%, M_{low} = 24.3\%, t(96) = 1.46,\ p = 0.15\ \text{Cohen's}\ d = 0.12$).

Thus, when values were learnt, they did appear to influence probability of free recall with small effect size, but this effect was attenuated when there was an intervening lexical-decision task. Similar to Experiment 1, we followed up this analysis with Bayesian t test. It produced a Bayes Factor, $BF_{10} = 3.89$, indicating the difference was favoured more than 3:1 over the null effect, considered "some" evidence for the effect. This suggests that the difference due to value may not be simply due to our large N.[2]

In sum, Experiment 2 replicated the value-bias of lexical-decision effect from Experiment 1 and Madan et al. (2012), but again, was only reliable when restricted to words whose values had been learnt. Findings also revealed that value influenced free-recall probability, with small effect size, when free-recall immediately followed value learning (even across words, within-subjects). However, these effects were weakened beyond significance when there was an intervening lexical decision task.

---

[2] We carried out additional tests looking at the relation between the number of accurate responses to an item during value-learning to its free recall probability, as detailed in the Supplementary Materials section.

Experiment 3

Only Experiment 2 produced an effect of value on free recall, and only when the analyses were restricted to learnt items and when free recall preceded lexical decision. In Experiment 3, we sought to replicate this value-bias of free recall. Because both the value-bias of lexical decision and of free recall seemed to be attenuated when another task intervened, we eliminated lexical decision from this experiment and tested free recall right after the value-learning phase, using the same probabilistic value-learning procedure as in Experiment 2. Here, our objective was to obtain a less noisy measure of the putative value-bias of free recall. Inspection of the distributions of the number of free-recall responses given in Experiments 1 and 2 (Figure 4c,d) shows that many participants produced fewer than 10 recalls. Although it is possible that those participants could not recall more words, it is also possible that they lacked the will or interest to do so, or gave up recall prematurely. The low response counts for many participants may have made the free-recall analyses somewhat noisy. Thus, we made one further change to the procedure: whereas in Experiment 1 and 2, free recall was unrewarded, in Experiment 3, recall of each value-trained word was rewarded. The relationship between the reward in free recall to previously learnt value differed across three experimental groups as follows.

For the first group, Neutral, previously high- and previously low-value words were rewarded equally. We reasoned that if the value-bias of free recall was truly due to differences in the strength of items in memory, then an incentive to produce more free-recall responses should strengthen the value-bias, due to increased number of data points per participant. We expected the Neutral group to provide the best estimate of a putative value-bias of free recall because reward did not differ for recalled high-value than recalled low-value words. If we were to find a value-bias for the Neutral group, we wondered if participants could willingly amplify the bias, deliberately prioritizing high- over low-value words in recall. This was tested in the second group, Congruent, in which participants were rewarded more for high- than for low-value words. If we failed to replicate the value-bias of free recall in the Neutral group (which, as it turned out, was the outcome for the Neutral group), the Congruent group would instead give us an additional test of whether a value bias could emerge if participants were given an explicit incentive to do so at time of test. Finally, whether or not a value-bias of free recall were found for the Neutral or Congruent groups, we were curious as to whether participants could overcome or reverse a putative value-bias. To test this, we included a third group, Incongruent, in which participants were rewarded more for low- than for high-value words.

Methods

Participants

A total of 139 introductory psychology students at the University of Alberta participated for partial fulfillment of course credit. Data from 3 participants were lost due to machine error. All participants had learnt English before the age of six, had normal or corrected-to-normal vision, and were comfortable typing.[3]

All participants gave written informed consent prior to the study, which was approved by a University of Alberta research ethics board. Participants were assigned at random, based on order of arrival to the testing room, to one of the three experimental groups: Congruent ($N = 45$), Neutral ($N = 49$) and Incongruent ($N = 42$).

Materials

Stimuli were identical to Experiment 2.

Procedure

The basic procedure consisted of value-learning followed by unanticipated rewarded free recall test. The value-learning and the post-value-learning distractor task were identical to Experiment 2, with three changes: for the value learning task, generic coin images were used as reward feedback stimuli, no auditory feedback was given to the participants and for the math distractor task, correct responses turned green and incorrect responses turned red.

[3] In this experiment, we restricted participation to participants without gambling tendencies, details of which can be found in the Supplementary Materials section.

Free recall. Free recall was identical to the procedure used in Experiments 1 and 2, with the following modifications. Participants in the Congruent group were rewarded 300 points for each high- and 100 for each low-value word. Participants in the Neutral group were given 200 points for each correctly recalled word, high or low. Participants in the Incongruent group were rewarded 300 points for each low- and 100 for each high-value word. Each participant was presented with the corresponding instructions about their rewarding scheme just prior to free-recall. When the participant typed a word and pressed the 'ENTER' key, the response was recorded and at the bottom of the screen they were shown the text 'You won: ' which was followed by the points earned for that particular word. In addition, 'Points accumulated: ' appeared at the top of the screen, followed by the cumulative sum of points earned in free recall. Spelling mistakes and repeats were not rewarded. Participants were informed of repeated responses with the text message 'Already typed!', displayed for 2 s. Points earned through value-learning and free recall were summed and multiplied by 0.0007, rounded up to the nearest quarter dollar to calculate the honorarium for each participant. This, combined with the points per word, meant that the free-recall and value-learning tasks were weighted approximately equally: maximum possible reward was $5 for value-learning, plus $5 for free-recall, which was also expressed in the instructions. The feedback during the recall phase served to reinforce participants' understanding of the reward mapping, and increase the salience of the reward, with the goal to increase the number of recalled words.

Results and discussion

Value learning

A 2×3 ANOVA with the within subject factor Value (high and low) and the between subject factor Group (Congruent, Incongruent and Neutral) on learning accuracy (of cycle 1) revealed a significant main effect of Value, $F(1,133) = 54.16,\ MSE = 0.03,\ p < 0.001,\ \eta_p^2 = 0.29$ (high: 0.53 ± 0.01, low: 0.36 ± 0.01). No significant effect of Group, $F(2,133) = 0.87,\ MSE = 0.01,\ p = 0.42,\ \eta_p^2 = 0.01$, or interaction Value×Group,

$F(2,133) = 0.25,\ MSE = 0.03,\ p = 0.78,\ \eta_p^2 < 0.01$, was found. The main effect of Value remained significant till the end of value-learning, $F(1,133) = 21.41,\ MSE = 0.02,\ p < 0.001,\ \eta_p^2 = 0.14$ for cycle 20 (Figure 5a–c).

Over the course of value-learning, response times on correct trials (Figure 5d–f) were faster for high- than low-value words; on the last (20^{th}) cycle, the main effect of Value was significant, $F(1,132) = 5.92,\ MSE = 0.08\ s^2,\ p < 0.05,\ \eta_p^2 = 0.04$, while the main effect of Group was not, $F(2,132) = 2.09,\ MSE = 0.30\ s^2,\ p = 0.13,\ \eta_p^2 = 0.03$, nor was their interaction, $F(2,132) = 0.53,\ MSE = 0.08\ s^2,\ p = 0.59,\ \eta_p^2 = 0.01$. Thus, there was no difference in value-learning across all three groups (see Figure 5a–f). Within each group, value-learning performance was similar and showed main effects of Value, with higher accuracy (Figure 5a–c) and faster acquisition (Figure 5d–f) of high-value than low-value words. Figure 5g–i plots the distributions of average accuracy over the last four cycles (17–20) for all participants in each group. Comparison with Figure 2f suggests that difficulty was similar to Experiment 2, which used the same value-learning procedure.

Free recall

Participants tended to recall more items in Experiment 3 (Figure 6a–c) than in Experiments 1 and 2 (Figure 4c,d), suggesting that rewarding free recall and providing accuracy feedback during recall increased motivation to recall more words, as intended. Free recall probability for high-value words was calculated by the proportion of high words recalled out of the total number of correct recalls and similarly for low-value words. Thus, for each participant, these two probabilities sum to one, making it incompatible to include Value (high and low) as a two-level factor in a repeated-measures ANOVA design.

We conducted a paired t-test for high and low free recall probabilities for all participants across the three groups. This showed that free recall rates did not differ by Value, $t(135) = -1.11,\ p = 0.27$. Further, we conducted a univariate ANOVA with Group as the factor for the recall rate (high or low) which revealed no significant effect,

$F(2,133) = 1.79,\ MSE = 0.02,\ p = 0.17,\ \eta_p^2 = 0.03$ (Figure 7a). Next, we carried out analysis for learnt items only (Figure 7b). Of the 136 participants, 21 (Congruent: $N = 5$, Neutral: $N = 10$, Incongruent: $N = 6$) failed to recall any learnt word and could not be included in this analysis. For the remaining participants, free recall probability for high words were calculated by the proportion of high words that were learnt and correctly recalled out of the total number of learnt high words, and similarly for low-value words. Thus, the free recall probabilities for high- and low-value words were not linearly dependent and Value could be used as a two-level factor in a repeated-measures ANOVA design. A 2×3 ANOVA with Value (high and low) as a within-subjects factor and Group (Congruent, Incongruent and Neutral) as a between-subjects factor revealed no main effect of Value, $F(1,112) = 0.32,\ MSE = 0.05,\ p = 0.57,\ \eta_p^2 < 0.01$, nor of Group, $F(2,112) = 2.47,\ MSE = 0.08,\ p = 0.06,\ \eta_p^2 = 0.04$. The interaction Value×Group was far from significant, $F(2,112) = 1.25,\ MSE = 0.05,\ p = 0.29,\ \eta_p^2 = 0.02$.

To test the stability/robustness of these null effects, we conducted a Bayesian ANOVA using JASP (JASP Team, 2016). Assuming uniform prior probabilities for all models, the Bayesian ANOVA produces a quantification of model support, known as the $BF_{inclusion}$. Similar to the Bayes Factor, BF_{10}, as mentioned in Experiments 1 and 2, $BF_{inclusion}$ indicates if a model fits better with a particular main effect or interaction included versus excluded. By convention (Kass & Raftery, 1995), when $BF_{inclusion} > 3$, the non-null effect is considered to have "some" support, and when $BF_{inclusion} < 0.3$, the null effect is considered to have "some" support— i.e., the model fits better without the effect. Paralleling the conventional ANOVA, $BF_{inclusion}$ values were tested for each main effect and the interaction. For Value, this supported the null ($BF_{inclusion} = 0.12$). For Group, this was not conclusive ($BF_{inclusion} = 0.35$). More importantly, for the interaction Value×Group, there was a strong support for the null ($BF_{inclusion} = 0.04$). Thus, the absence of modulation of free recall by value alone or differential modulation by value per group was supported by the Bayesian ANOVAs.[4]

Thus, Experiment 3 failed to replicate the effect of value on free recall found for Group FR-LD in Experiment 2. Not only did the Neutral group failed to show a value-bias, despite recalling more words than in Experiment 2, the Congruent group also failed to show a value-bias of free recall. Considering that the congruent group in Experiment 3 received more reward for recalling high-value words, this group could have uncovered even larger reward-based memory advantages acting at the recall stage. However, this was not the case (a relative 5.4% recall advantage of high- over low-value words in Experiment 3 congruent; versus 16% in Experiment 2). The failure to replicate the value-bias of free recall in all three groups, but especially, the Neutral and Congruent groups in Experiment 3, raises the possibility that the result in Experiment 2, Group FR-LD was a *Type I* error.

A second possibility is that value may have led participants to produce high-value words earlier in the response sequence than low-value words. Then, lacking an incentive to produce a large number of responses, participants in Experiment 2 may have terminated recall before they had exhausted their memory search. In this way, an output-order effect could produce a difference in probability of recall, even if the strength of high- and low-value words in memory were equivalent. In Experiment 3, participants may not have produced high-value items earlier in recall. Alternatively, given an incentive to produce more recalls, the earlier-recalled high-value words might have been offset by eventual, later recalls of low-value words. Following the approach of Madan et al. (2012), after ignoring responses that were not among the 36 trained words (neither high, nor low-value words), we computed Mann-Whitney U tests on the response-position within the free-recall sequence of high- versus low-value words. We did so for each participant, and then tested whether the z-transformed U values were significantly different than zero. A negative value indicates high-value words tending to be recalled earlier in output. For Experiment 2, the t test

4 A crossed-random mixed effects analysis was performed for Experiment 3 which included word-effects as a random factor, our findings persisted with this mixed model; details can be found in the Supplementary Materials section.

indeed produced a significantly negative value, $t(167) = -2.10,\ p < 0.05$. When restricted to learnt words only, the effect remained significant, $t(163) = -2.03,\ p = 0.044$. In Experiment 3, participants recalled high-value words significantly earlier than low-value words, $t(131) = -2.31,\ p < 0.05$. When broken down by group, this effect was significant for Congruent, $t(43) = -2.34,\ p < 0.05$, but not for Neutral, $t(46) = -1.37,\ p = 0.18$, or Incongruent, $t(40) = -0.41, p = 0.68$. When restricted to learnt words only, the effect was overall non-significant, $t(103) = -1.42,\ p = 0.15$, nor was it significant when broken down by group: Congruent, $t(34) = -1.42,\ p = 0.16$; Neutral, $t(36) = -0.23,\ p = 0.82$; Incongruent, $t(31) = -0.81,\ p = 0.42$. In sum, support was found for a small tendency for high-value words to be recalled earlier, which may explain the positive result found in Experiment 2.[5]

In sum, Experiment 3 failed to replicate the value-bias of free recall, even when participants were explicitly incentivized to produce such a bias (Congruent group). These findings suggest that the value-bias of free-recall observed for the FR-LD group in Experiment 2 may have been a consequence of output-order prioritization combined with premature termination of recall, rather than a difference in memory strength between high- and low-value words.

Analyses of the two-choice value learning paradigm

In stark contrast to the overall pattern of results with the present, one-word, value-learning procedure, Madan et al. (2012) found a robust effect of value (high or low) on free-recall probability when the high- and low-value words were presented together during value-learning. To reconcile these findings, we sought to understand why a value bias is present when high and low items are pitted against each other during training (“two-word” procedure), but is absent when items are presented alone (“one-word” procedure). We revisited the Madan et al. (2012) data set and tested two hypotheses about the origin of the value bias with the two-word procedure. Methods are reported in detail in Madan et al. (2012); in brief, in the value-learning phase, participants were shown a high-value word paired with a low-value word, and learned, through

5 Similar to Experiment 2, we tested the relation between the number of accurate responses to an item during value-learning to its free recall probability, as detailed in the Supplementary Materials section.

trial and error, to choose high-value, and not choose low-value words. Choosing the high-value word earned the participant 10 points and choosing the low-value word earned the participant 1 point. Each of the 36 words was used in exactly one trial in each of 13 cycles. This was followed by lexical decision of each trained word (presented individually), and free recall of all words from the experiment.

Hypotheses and overview

Two distinct mechanisms present themselves. First, the Set-level Differentiation hypothesis, is that the two-choice paradigm simply accentuated the distinction between the set of high-value words and the set of low-value words. Alternatively, the Choice-based Competition hypothesis is that the act of *choosing* the high-value word on a given trial while *not-choosing* (avoiding) the low-value word, draws more attention towards the chosen word and away from the non-chosen word. To the extent that Choice-based Competition is at play, recall of high- and low-value words should be negatively correlated across training trials; if training trials are indexed by i, then the correlation Q, between recall of H_i and L_i, the high- and low-value words presented on trial i, respectively, should be less than zero. Viewed from an individual-differences perspective, the more negative the value of Q, the greater the bias to recall high- over low-value words. In contrast, Set-level Differentiation predicts no such effect. With the one-word procedure, there is no such pairing of H_i and L_i; thus, if Set-level Differentiation were favoured, this would imply that the distinction between high- and low-value words is too subtle to influence memory with the one-choice procedure. If Choice-based Competition were instead supported, this would imply that the high- and low-value words compete directly for encoding processing only (noticeably) when they are presented simultaneously.

Data analyses and results

Measure of the value-bias of free recall. To quantify the value-bias of free recall, we calculated, for each participant in the Madan et al. (2012) data set, the log-odds ratio of the free recall probabilities, $B = log \frac{P_{high}}{(1-P_{high})}$, where P_{high} is the probability of recall of high value words as a proportion of words recalled from the value-learning phase ($P_{high} + P_{low} = 1$). A Bayesian t test confirmed very strong support for the value bias, $t(91) = 4.10,\ p < 0.001,\ BF_{10} = 209$.

Set-level Differentiation versus Choice-based Competition. To test whether the act of choosing one word while simultaneously not choosing the other word during value-learning induces direct competition in memory between the pair of words, we computed the correlation between free-recall probability (recalled or not recalled) and between the chosen and not-chosen words in a pair, across value-learning pairs, for each participant. Because these were dichotomous data, Yule's Q was the measure of correlation, defined as $Q = \frac{ad-bc}{ad+bc}$, where a, b, c, d are tallies of the number of trials for which both words were recalled, only the chosen word was recalled, only the non-chosen word was recalled, and neither word was recalled, respectively. Besides acting on dichotomous data, Q can be interpreted similar to Pearson correlation. Thus, $Q > 0$ would indicate that chosen and non-chosen words of a pair were likely to be both recalled or both not recalled. $Q < 0$ would indicate some level of mutual exclusivity, such that if the chosen word is recalled, the non-chosen word is less likely to be recalled, or vice versa (controlling for the marginal probabilities), which would support the Choice-based Competition hypothesis. Finally, $Q = 0$ indicates independence, which would be consistent with the Set-level Differentiation hypothesis. To better satisfy the assumption of normality, Q values were log-odds transformed. Further, to keep the analyses relatively simple, we report Q values collapsed across all value learning cycles.

Q was not significantly negative for the sample as a whole, $t(91) = -0.76$, $p = 0.45$, $BF_{10} = 0.15$. However, our specific question is not whether the effect is present on average, but whether it could explain the cause of the value-bias of free recall. By computing Q for each participant, we could ask whether individuals who have more Choice-based Competition are those who produce more free recalls of high- than low-value words. Supporting the Choice-based Competition hypothesis, there was a strong, significant negative correlation between B and Q, $r(90) = -0.34,\ p < 0.001,\ BF_{10} = 19.87$, suggesting that participants with more negative Q values had a greater value bias, B (Figure 9 in Supplementary Materials).

To test whether this relationship could explain the presence of the value-bias in the aggregate, we subdivided participants into two subsets: (i) $Q < 0$ and (ii) $Q \geq 0$ (figure 9). Whereas the Bayes Factor BF_{10} still indicated strong support for the value bias for participants with $Q < 0$, $t(45) = 3.59,\ p < 0.001,\ BF_{10} = 35.26$, for participants with $Q >= 0$, the Bayes Factor was inconclusive, $t(44) = 2.09,\ p = 0.042,\ BF_{10} = 1.16$ (although statistically reliable according to classical statistics, at the $\alpha = 0.05$ level).[6] Given the similar sample sizes of these sub-groups, this suggests the value-bias of free recall in Madan et al. (2012) was not a general effect but influenced by individual differences, and when the two-word procedure induces participants to remember words competitively, that competition favours the chosen over the non-chosen word.

Choice versus value. We wondered whether the Choice-based Competition effect might actually be better viewed as Value-based Competition. That is, we wondered whether words compete based on which word was chosen versus not chosen, or based on high versus low value. Thus, we ran a parallel set of analyses to those reported in the previous section, computing Q based on high- versus low-value rather than choice. Similar to the previous analyses, Q values were not significantly different from zero, $t(91) = 1.46,\ p = 0.15,\ BF_{10} = 0.32$. But in this case, no significant correlation was found between Q and B, $r(90) = 0.03,\ p = 0.77,\ BF_{10} = 0.09$. In fact, the Bayes Factor suggests very strong support for a null relationship between Q and B. This suggests that the act of choosing, not the ultimate value of the word, is what produced the value-bias of free recall in Madan et al. (2012).

Discussion

In sum, our further analyses of the Madan et al. (2012) data support our view that the value-bias was caused by a factor unique to procedures in which high- and low-value words are presented together. Specifically, for some participants (but not all), the act of choosing one word and not choosing the other word produced competition in memory between the chosen and

6 When re-run as a median split on Q, the same basic pattern was found: for $Q < Q_{median}$, $t(45) = 3.44,\ p < 0.005,\ BF_{10} = 23.83$; for $Q \geq Q_{median}$, $t(45) = 2.29,\ p < 0.05,\ BF_{10} = 1.68$.

not-chosen word on a given trial. Because participants were supposed to be learning to choose high-value words over low-value words, this tended to bring out a value-bias of free recall for such participants. This mechanism is unavailable in the present, one-word procedure, and thus cannot produce a value-bias of memory.

General Discussion

The three experiments presented here were motivated by the idea that information associated with high value should be remembered better than information associated with low value. Although the information had either high or low value, the usefulness of that value as a learning signal was equated: learning to choose high-value words could earn participants the same amount of points (and bonus payment) as learning *not* to choose low-value words. This way, there was no rational reason to expend more cognitive resources on studying high-value words initially. Any effect of value on later memory for trained words could not simply reflect an optimal study strategy. Under these circumstances, value did not bias later memory in any substantial way.

When an advantage for high-value words emerged, with the sizeable samples reported here (Experiment 2, group FR-LD), it dissipated when lexical decision was tested first (Experiment 1 as well as the LD-FR group of Experiment 2). Experiment 3 suggested that the one instance of a small-magnitude (relative difference of $\sim$16%) value-bias of free-recall in Experiment 2 may have been due to participants recalling high-value words early, and then terminating recall prematurely.

Incentivizing recall in Experiment 3 increased overall recall rates. However, again, value produced no net bias in recall. Even when recall incentives were higher for high- than low-value words (Congruent group, Experiment 3), participants were unsuccessful at prioritizing their recall toward high-value words. This suggests no underlying difference in the quality of memory for high- versus low-value words. Thus, value may bias free recall within very narrow boundary conditions; when it does, it may act by biasing output order.

Madan et al. (2012) found a robust and large-magnitude value-bias of free recall ($\sim$25% benefit to recall-probability), plainly visible without restricting the analysis to words whose values were learnt. This suggests the quality of memory for words may be substantially affected when high- and low-value words are pitted against one another during value-learning, even when usefulness is equated. Pinpointing the origin of the value-bias in the two-word procedure, our re-analyses of the Madan et al. (2012) data suggested that the two-word procedure induced a competitive relationship in memory between items within each pairing. Without any thought to the question of how value might influence memory, choosing can enhance memory relative to not choosing (Coverdale et al., in press; Coverdale & Nairne, in press). For example, Coverdale and Nairne (in press) had no manipulation nor mention of value; participants were asked to make choices between pairs of words based on their usefulness (more or less useful) or based on how much they represent a specific category (more or less representative); such as, in one experiment, participants decided whether "bowl" or "ruler", was either more or less useful in a survival situation. The memory component was incidental, similar to Madan et al. (2012). In both their experiments, participants remembered the chosen items significantly better than the non-chosen items, even when taking into account whether the chosen item was congruent ("more" wording) or incongruent ("less" wording) with the instruction. Thus, the act of choosing versus not-choosing can influence free recall, without any notion of value. Our analyses of the Madan et al. (2012) data suggest that paradigms similar to Coverdale and Nairne (in press) might also find that choosing/not-choosing induces a negative correlation, indicating direct competition between paired items, rather than effects of choice acting on the stimulus sets (chosen/not-chosen) as a whole. Moreover, if the findings of Coverdale and Nairne (in press) and Madan et al. (2012) are to be understood along with the present findings, the implication is that the act of choosing, and not value *per se*, induces competition between words judged simultaneously, but not when words are judged individually. Thus, we predict that the choice bias of free recall reported by Coverdale et al. (in press) or Coverdale and Nairne (in press) would be eliminated in a follow-up experiment in which words were chosen or not-chosen individually, akin to our one-word procedure here.

As to why the effect of choice is absent in the one-word value-learning procedure used here, consider that attentional competition is only between a trial-unique word and a fixed nonsense string (HHHHH). After choosing the nonsense string, the nonsense string itself may have not been as salient as the non-chosen, trial-unique word. For this reason, such an attentional bias might be attenuated here.

Unlike free recall, the value-bias of lexical decision was clearly present, when restricted to words whose values had been learnt. Such an access-speed bias could, in principle, bias attention during other forms of behaviour involving the trained stimuli, and might eventually be able to bias memory for the words themselves. Although the cause of the value-bias of lexical decision is unknown, we explore two potential explanations.

First, participants, tended to prefer to choose a word than not-choose it (the word-choice bias); this is evident in that accuracy on Cycle 1 was greater for high- than for low-value words (Figures 2a,b). Although this must reflect a guessing strategy (no values are known on Cycle 1), this word-oriented response behaviour might have differentially influenced memory for high- than low-value words. To check this, we plotted the value-bias of lexical decision (difference in lexical decision time for high minus that for low-value words, restricted to words whose values were learnt) as a function of a measure of the word-choice bias (accuracy during value-learning, Cycle 1, for high-value words minus that for low-value words). Figure 10a,b shows that negative values of the value-bias of lexical decision were observed at similar rates for participants with positive and negative word-choice biases. Second, over the course of learning, participants took longer to respond to low- than to high-value words (Figures 2c,d). It is possible that this difference in response time during value-learning directly translated into the value-bias of lexical decision time. However, a similar plot, depicting the value-bias of lexical decision against the difference in value-learning response time (last cycle) also shows no regular relationship between these two behavioural effects. Although a non-significant correlation does not confirm a true null correlation, it does render it less plausible that the relationship could have alone produced the effect of interest (value-bias of lexical decision time).

One study disentangled reward-value from memory by functionally separating the rewards from the to-be-remembered stimuli themselves (Murty & Adcock, 2014). In this study, participants monitored repeated presentations of an object for a visual change from colour to greyscale. Series of trials were pre-cued as to whether performance— a button press indicating colour change— would lead to either a high or a low reward. In some trial series, a similar distractor object was shown among the repetitions of the original target object (e.g., if the target object had been a photo of balloons, the distractor object was a different balloon picture). Participants were instructed that this change, creating an expectancy violation, was irrelevant to their task, which was to indicate a colour change in repeats of the original picture. Distractor objects were remembered better in a later surprise memory test if they had been shown during high-reward trial series than low-reward trial series. Again, this result is important to understanding the basis of reward influences on memory, but it is still possible that the usefulness difference between trial-values explains these results. Expecting a high reward may mobilize visual attention and vigilance more than expecting low reward. Thus, a surprising stimulus may in fact have been processed better if it appeared during a high- than low-value trial.

Different from this procedure, our goal for the current studies was not to separate value from the items, but rather, to tie value to items while controlling their usefulness in the value-learning task. Our procedure is similar to studies that have found that prior reward value can bias visual attention (summarized by Anderson, 2013). In those paradigms, a stimulus feature (e.g., a designated colour such as red) signals a high-reward trial and a different feature (e.g., blue) signals a low-reward trial. On subsequent tests such as visual search (e.g. Anderson & Yantis, 2013; Della Libera & Chelazzi, 2009) or a Flanker task (Anderson, Laurent, & Yantis, 2012), stimuli with the high-value feature capture attention whereas stimuli with the low-value feature may even repel attention. Our finding that prior value-learning facilitated lexical decision for high-value words more than for low-value words (given that values were learnt) might be an extension of these phenomena, but acting at the level of objects (here, words) rather than features (anticipated by Anderson, 2013). Our stimuli are words. Thus, their “features” must be either letters, or, arguably, line segments, all of the same colour, intensity, etc. Thus, “features” of our stimuli are mixed at random between high- and low-value words.

Our findings have important implications for basic memory processes. In many existing reward-memory procedures, reward is used as a shaping cue. Reward value in these procedures leads participants to find an optimal strategy. Reward is thus— by design— confounded with usefulness. Because of this, the reward signal may inform participants that they have very good reasons to devote more of their precious cognitive resources to certain (high-value) stimuli at the expense of other (low-value) stimuli. Effects of reward-value on memory, then, may be measuring participants' ability to prioritise their study and retrieval processes to maximize their reward (points, money, etc.). This is made explicit in prioritisation procedures (Watkins & Bloom, 1999) and observations with such procedures have shown that there is considerable individual-variability in this ability (e.g., Castel et al., 2002, 2007). Our findings show that, when the usefulness of reward is controlled, leaving no rational incentive to prioritise some items over others, words previously designated as having a high value are accessed faster (lexical decision) and may sometimes be recalled earlier, than words designated as low value. However, even these effects were small in magnitude, and quite fragile, and value failed to produce a bias of free-recall probability that could not be explained as a prioritization of high-value words within the recall sequence. Thus, prior findings that value influences memory may largely be due to usefulness rather than any other characteristic related to value, itself. It remains to be tested if increased activity in the ventral tegmental area, suggesting increased dopaminergic modulation of memory-encoding regions, responds to usefulness, or to value even when usefulness is controlled (Clewett & Mather, 2014).

Although the three experiments reported here converge on the idea that value, when usefulness was not a confounding factor, does not influence quality of memory, and exerts only a small and transient influence on accessibility of information in memory or attentional orienting bias, these conclusions are clearly limited to the procedures we used. Most notably, free recall is only one way to test memory; it is possible that with different memory tests, clear value-biases of memory would emerge. With a longer delay (one week) between value-learning and the memory test, an effect of value might emerge, as Murayama and Kitagami (2014) suggested.

In sum, we found no support for the idea that value biases the quality of memory for the trained items themselves. If items are pitted against each other (Madan et al., 2012), the forced choice procedure can bias memory by inducing a competitive relationship between simultaneously presented items that differ in value, but such effects also occur in the absence of any manipulation of value. Thus, although a small value-bias of memory may influence behaviour via accessibility or orienting attention, effects on the quality of memory reported previously are likely attributable to differences in usefulness, not value, when prioritizing high- over low-value materials is, in fact, a rational strategy. Even when usefulness is equated, forced-choice procedures can induce value-biases in memory that are better understood as choice, not value, effects.

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*, 507-517.

Anderson, B. A. (2013). A value-driven mechanism of attentional selection. *Journal of Vision*, *13*(3), 1-16.

Anderson, B. A., Laurent, P. A., & Yantis, S. (2012). Generalization of value-based attentional priority. *Visual Cognition*, *20*(6), 647-658.

Anderson, B. A., & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(1), 6-9.

Bates, D. (2007). *Linear mixed model implementation in lme4* (Vol. 15). University of Wisconsin.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): instruction manual and affective ratings* (Tech. Rep.). The Center for Research in Psychophysiology, University of Florida.

Brodie, D. A., & Murdock, B. B. (1977). Effect of presentation time on nominal and functional serial-position curves of free recall. *Journal of Verbal Learning and Verbal Behavior*, *16*, 185-200.

Castel, A. D., Balota, D. A., & McCabe, D. P. (2009). Memory efficiency and the strategic control of attention at encoding: impairments of value-directed remembering in Alzheimer's Disease. *Neuropsychology*, *23*(3), 297-306.

Castel, A. D., Benjamin, A. S., Craik, F. I. M., & Watkins, M. J. (2002). The effects of aging on selectivity and control in short-term recall. *Memory & Cognition*, *30*(7), 1078-1085.

Castel, A. D., Farb, N. A. S., & Craik, F. I. M. (2007). Memory for general and specific value information in younger and older adults: measuring the limits of strategic control. *Memory & Cognition*, *35*(4), 689-700.

Castel, A. D., Lee, S. S., Moore, A. N., & Humphreys, K. L. (2011). Memory capacity, selective control, and value-directed remembering in children with and without attention-deficit/hyperactivity disorder (ADHD). *Neuropsychology*, *25*(1), 15-24.

Clark, L., Crooks, B., Clarke, R., Aitken, M. R. F., & Dunn, B. D. (2012). Physiological responses to near-miss outcomes and personal control during simulated gambling. *Journal of Gambling Studies*, *28*(1), 123-137.

Clewett, D. V., & Mather, M. (2014). Not all that glittered is gold: neural mechanisms that determine when reward will enhance or impair memory. *Frontiers in Neuroscience*, *8*(194).

Coverdale, M. E., & Nairne, J. S. (in press). The mnemonic effect of choice. *Psychonomic Bulletin & Review*.

Coverdale, M. E., Pandeirada, J. N. S., & Nairne, J. S. (in press). Survival processing in a novel choice procedure. *American Journal of Psychology*.

Cuvo, A. J. (1974). Incentive level influence on overt rehearsal and free recall as a function of age. *Journal of Experimental Child Psychology*, *18*(1), 167-181.

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*(6), 778-784.

Estes, W. K. (1972). Reinforcement in human behavior: reward and punishment influence human actions via informational and cybernetic processes. *American Scientist*, *60*(6), 723-729.

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, *83*(1), 37-64.

Ferrey, A. E., Frischen, A., & Fenske, M. J. (2012). Hot or not: response inhibition reduces the hedonic value and motivational incentive of sexual stimuli. *Frontiers in Psychology*, *3*(575), 1-7.

Ferris, J., & Wynne, H. (2001). *The Canadian problem gambling index: final report* (Tech. Rep.). Ottawa, ON: Canadian Centre on Substance Abuse.

Frischen, A., Ferrey, A. E., Burt, D. H. R., Pistchik, M., & Fenske, M. J. (2012, 169-179). The affective consequences of cognitive inhibition: devaluation or neutralization? *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 1.

Gruber, M. J., & Otten, L. J. (2010). Voluntary control over prestimulus activity related to encoding. *Journal of Neuroscience*, *30*(29), 9793-9800.

Harley, J., William F. (1965). The effect of monetary incentive in paired associate learning using a differential method. *Psychonomic Science*, *2*(1), 377-378.

JASP Team. (2016). *JASP (Version 0.8.1.1)[computer software].* Retrieved from `https://jasp-stats.org[jasp-stats.org]`

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Society*, *90*(430), 773-795.

Kassinove, J. I., & Schare, M. L. (2001). Effects of the "near miss" and the "big win" on persistence at slot machine gambling. *Psychology of Addictive Behaviors*, *15*(2), 155-158.

Knutson, B., Westdorp, A., Kaiser, E., & Hommer, D. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *NeuroImage*, *12*(1), 20-27.

Kuhl, B. A., Shah, A. T., DuBrow, S., & Wagner, A. D. (2010). Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *NatureNeuroscience*, *13*(4), 501-506.

Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, *22*(10), 1310-1318.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, *1*(4), 476-490.

Loftus, G. R., & Wickens, T. D. (1970). Effect of incentive on storage and retrieval processes. *Journal of Experimental Psychology*, *85*(1), 141-147.

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. *Frontiers in Integrative Neuroscience*, *6*(61), 1-15.

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.

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *Quarterly Journal of Experimental Psychology*, *70*(10), 2048-2059.

Madan, C. R., Shafer, A. T., Chan, M., & Singhal, A. (2017). Shock and awe: distinct effects of taboo words on lexical decision and free recall. *Quarterly Journal of Experimental Psychology*, *70*(4), 793-810.

Madan, C. R., & Singhal, A. (2012). Encoding the world around us: motor-related processing influences verbal memory. *Consciousness and Cognition*, *21*(3), 1563-1570.

Madan, C. R., & Spetch, M. L. (2012). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica*, *139*(2), 343-349.

Murayama, K., & Kitagami, S. (2014). Consolidation power of extrinsic rewards: reward cues enhance long-term memory for irrelevant past events. *Journal of Experimental Psychology: General*, *143*(1), 15-20.

Murty, V. P., & Adcock, R. A. (2014). Enriched encoding: reward motivation organizes cortical networks for hippocampal detection of unexpected events. *Cerebral Cortex*, *24*(8), 2160-2168.

Schwarzkopf, S. (2015, March). *Data from SamPenDu: Bayes Factors Matlab functions.* https://figshare.com/articles/Bayes_Factors_Matlab_functions/1357917. doi: 10.6084/m9.figshare.1357917.v1

Shaoul, C., & Westbury, C. (2006). *USENET orthographic frequencies for the 40,481 words in the English Lexicon Project (2005-2006).* Edmonton, AB: University of Alberta (downloaded from http://www.psych.ualberta.ca/~westburylab/downloads/elp.download.html).

Sharot, T., De Martino, B., & Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *Journal of Neuroscience*, *29*(12), 3760-3765.

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, *14*(10), 464-472.

Soderstrom, N. C., & McCabe, D. P. (2011). The interplay between value and relatedness as bases for metacognitive monitoring and control: evidence for agenda-based monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(5), 1236-1242.

Tarpy, R. M., & Glucksberg, S. (1966). Effects of incentive and incentive-cue position on short-term retention. *Psychonomic Science*, *5*(8), 313-314.

Ward, G., & Tan, L. (2004). The effect of the length of to-be-remembered lists and intervening lists on free recall: a reexamination using overt rehearsal. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(6), 1196-1210.

Watkins, M. J., & Bloom, L. C. (1999). *Selectivity in memory: an exploration of willful control over the remembering process.* (Unpublished manuscript)

Weiner, B. (1966). Motivation and memory. *Psychological Monographs: General and Applied*, *80*(18), 1-22.

Weiner, B., & Walker, E. L. (1966). Motivational factors in short-term retention. *Journal of Experimental Psychology*, *71*(2), 190-193.

Westbury, C. F., Hollis, G., & Shaoul, C. (2007). LINGUA: the language-independent neighbourhood generator of the University of Alberta. *Mental Lexicon*, *2*(2), 273-286.

Wilson, M. D. (1988). The MRC psycholinguistic database: Machine readable dictionary, version 2. *Behavior Research Methods, Instruments, & Computers*, *20*, 6-11.

Winstanley, C. A., Cocker, P. J., & Rogers, R. D. (2011). Dopamine modulates reward expectancy during performance of a slot machine task in rats: evidence for a 'near-miss' effect. *Neuropsychopharmacology*, *36*(5), 913-925.

Wittmann, B. C., Dolan, R. J., & Düzel, E. (2011). Behavioral specifications of reward-associated long-term memory enhancement in humans. *Learning & Memory*, *18*, 296-300.

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*, *45*(3), 459-467.

Wolosin, S. M., Zeithamova, D., & Preston, A. R. (2012). Reward modulation of hippocampal subfield activation during successful associative encoding and retrieval. *Journal of Cognitive Neuroscience*, *24*(7), 1532-1547.

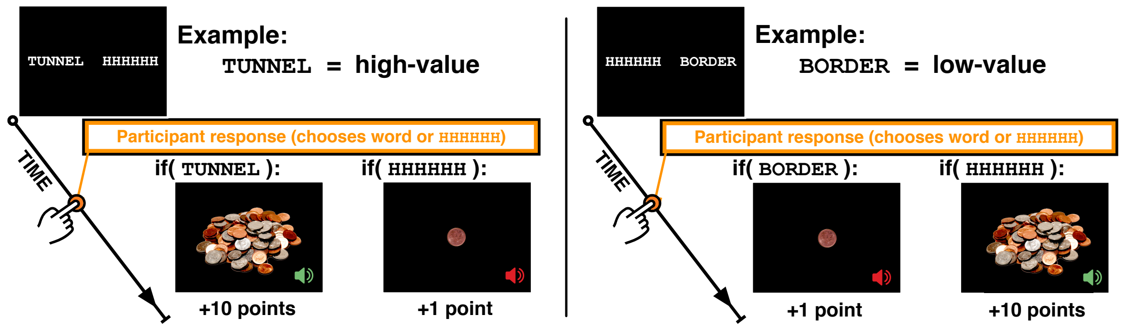

Figure 1. Illustration of the value-learning procedure used in Experiment 1. The green and red speaker icons represent the auditory feedback. The green speaker corresponds to the sound of many coins being dispensed (associated with gaining 10 points); the red speaker corresponds to the sound of a single coin hitting a table (associated with gaining 1 point).

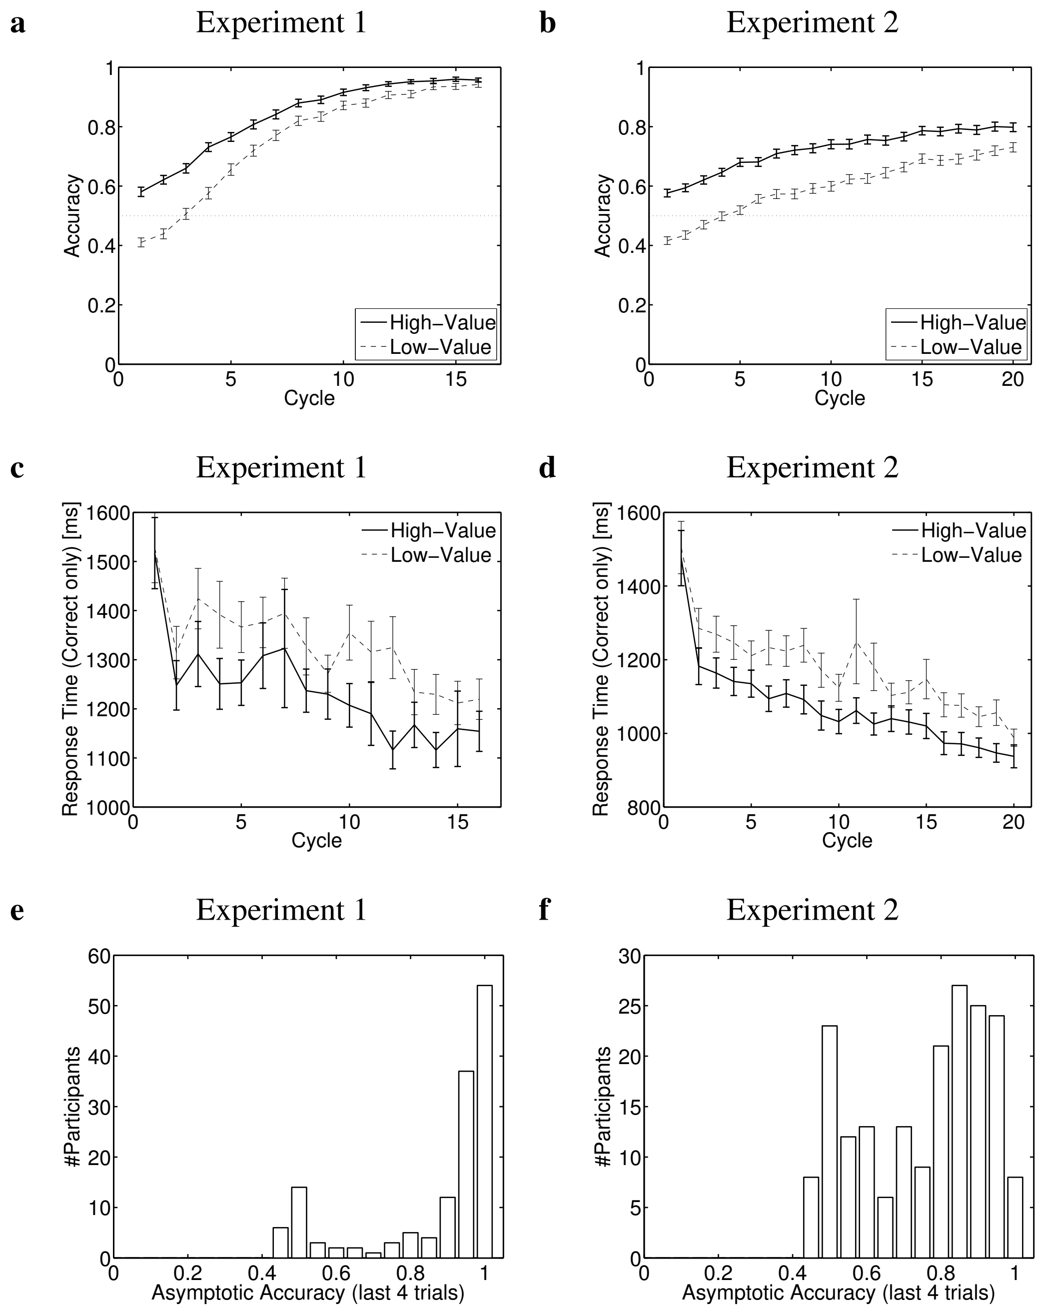

Figure 2. Learning curves for value learning for Experiment 1 (a,c) and Experiment 2 (b,d), plotting accuracy of responses (choose high or not-choose low; a,b) and response times for correct responses only (c,d). The dotted horizontal line in panels a and b denotes chance performance. The bottom panels plot the distribution of asymptotic accuracy values (mean accuracy over the last four cycles) in the value-learning phase of Experiment 1 (e) and Experiment 2 (f).

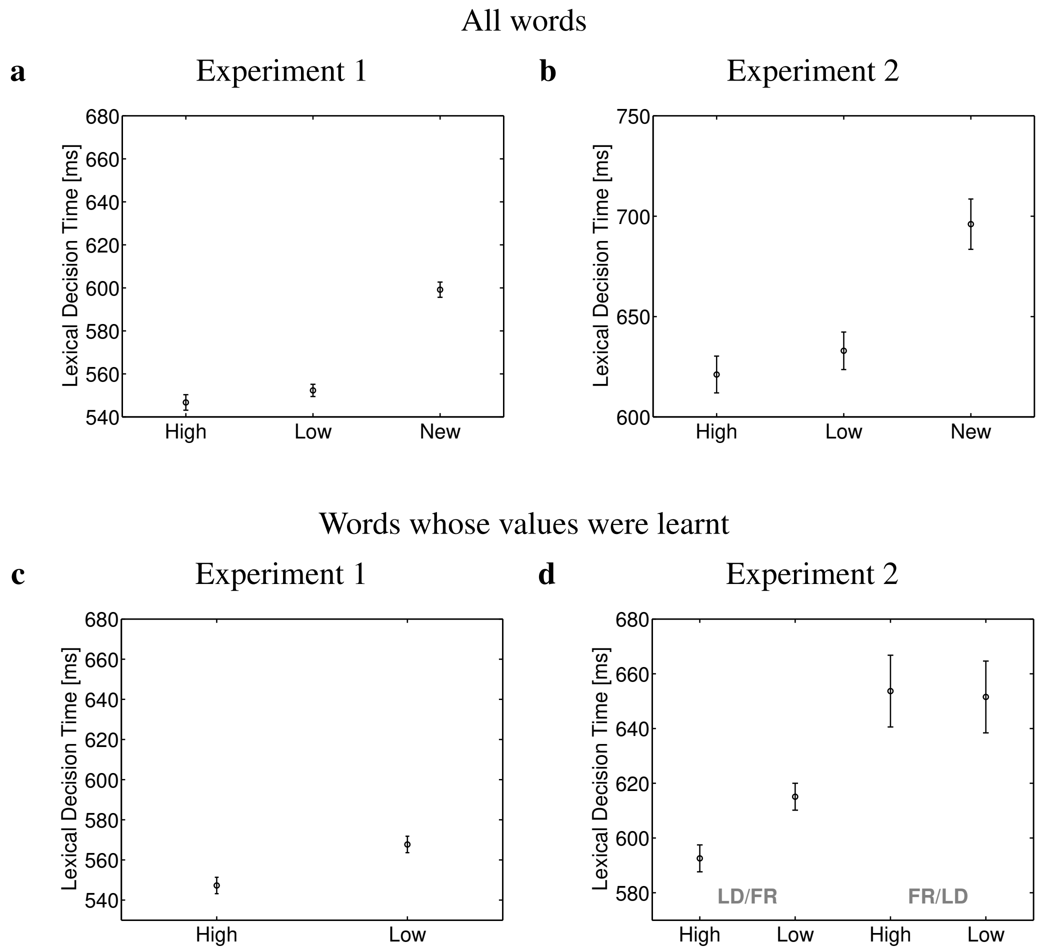

Figure 3. Lexical decision times for Experiment 1 (a) and Experiment 2 (b). Lexical decision times contingent on items having been learnt (correct decisions on last four cycles of value-learning) are plotted for Experiment 1 (c) and Experiment 2 (d), broken down by task-order. Group LD/FR had lexical decision before free recall, and vice versa for group FR/LD. Error bars plot standard error of the mean, corrected for between-subjects variability Loftus and Masson (1994).

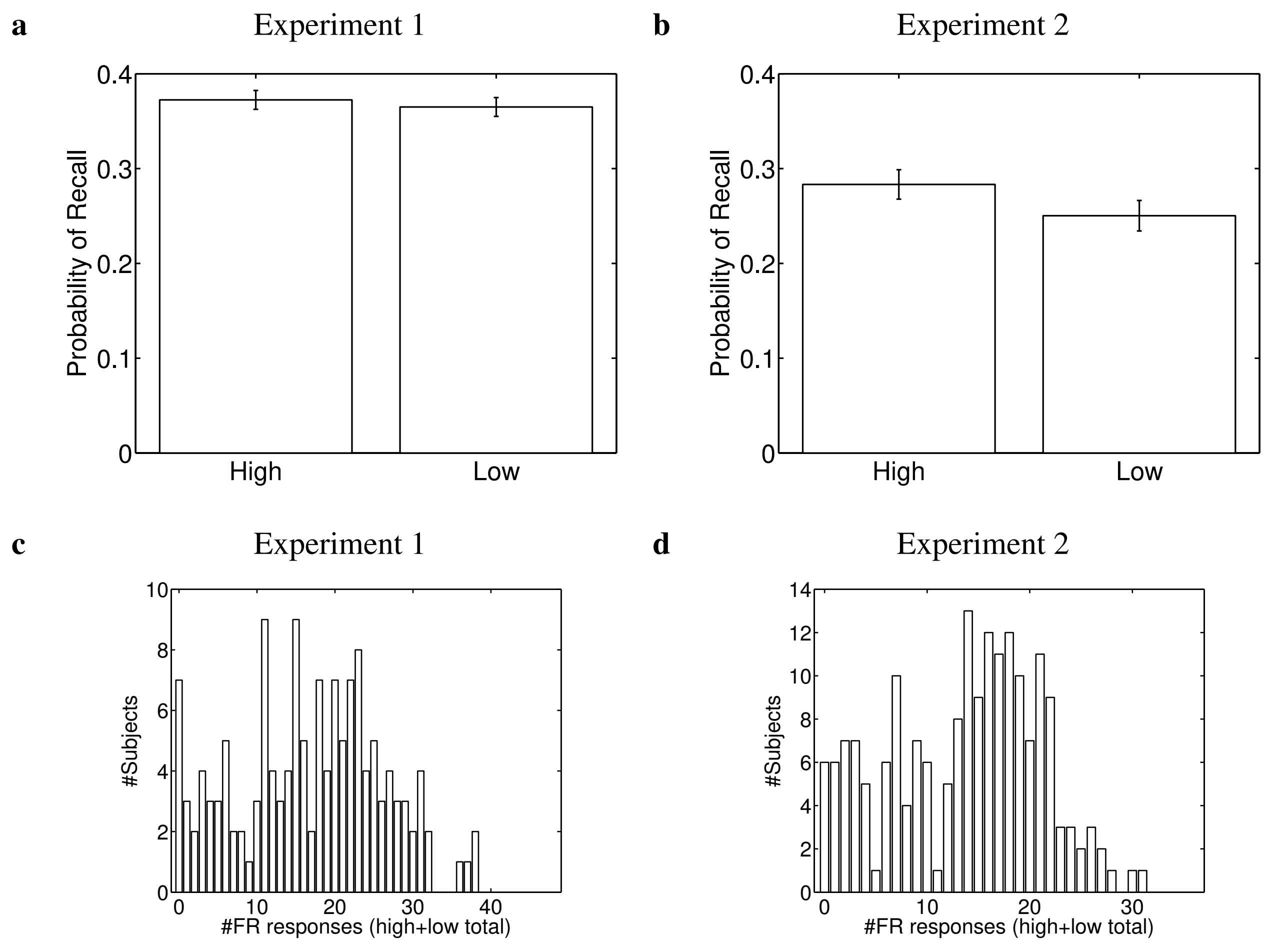

Figure 4. Probability of recall of high- and low-value words, for words that were learnt during value-learning, for Experiment 1 (a) and Experiment 2 (b). Error bars plot standard error of the mean, corrected for between-subjects variability (Loftus and Masson, 1994). Distributions of total number of high and low items recalled in free recall, for Experiment 1 (c), out of 48 total, and Experiment 2 (d), out of 36 possible.

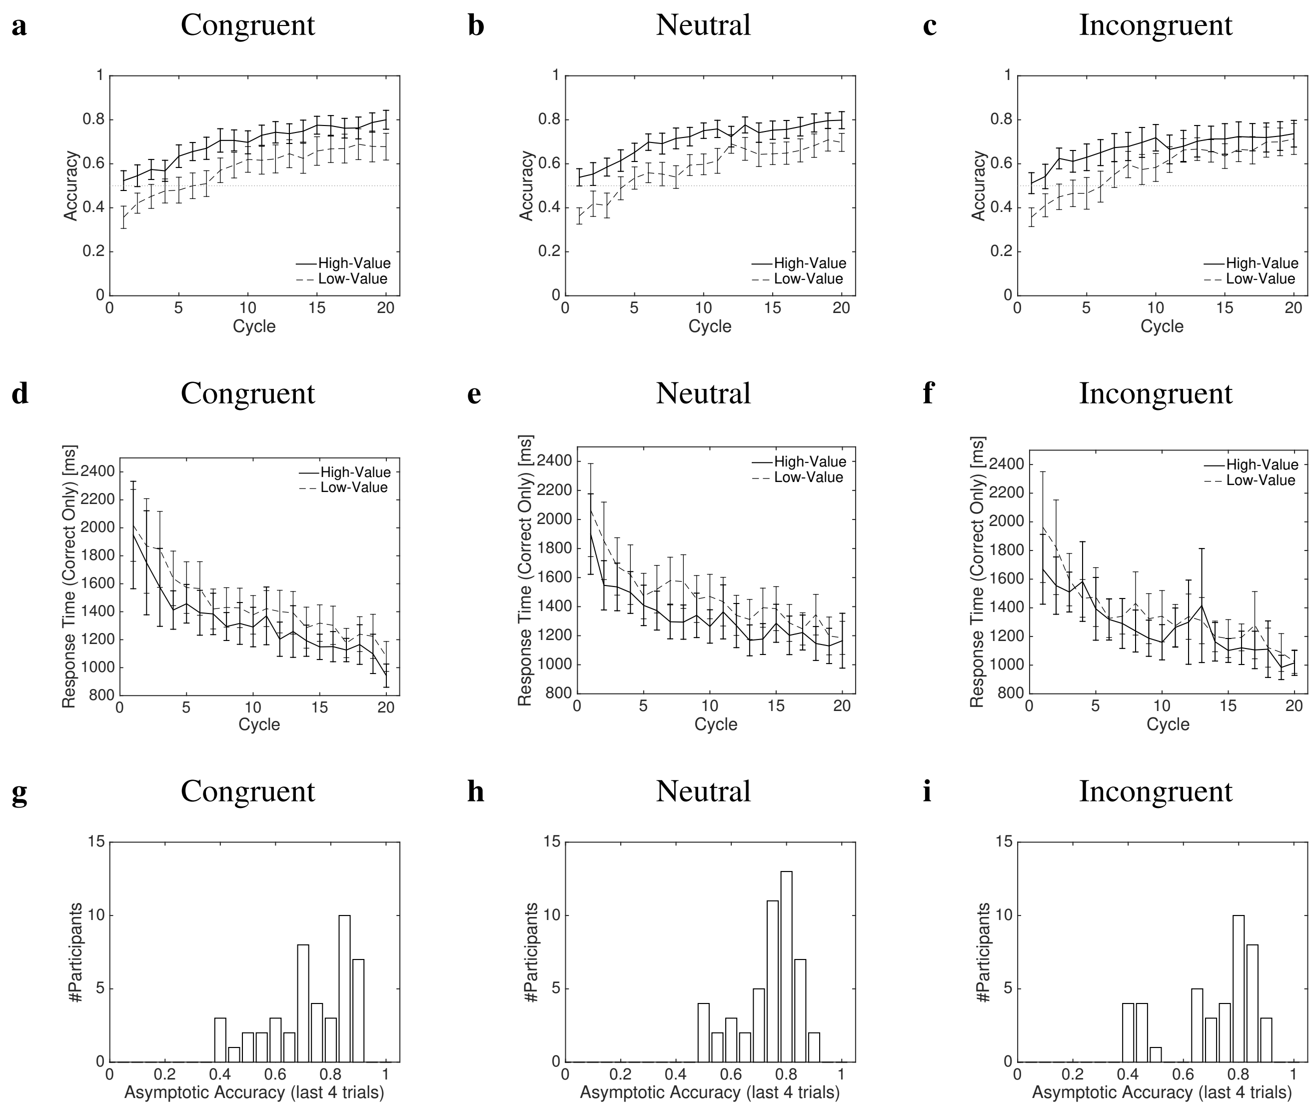

Figure 5. Learning curves for value learning for Experiment 3 (a–f), plotting accuracy of responses (choose high or not-choose low: a, b, c) and response time for correct responses only (d, e, f). The dotted horizontal line in panels a, b and c denotes chance performance. The bottom panels plot the distribution of asymptotic accuracy values (mean accuracy over the last four cycles) in the value-learning phase of Experiment 3 (g, h, i).

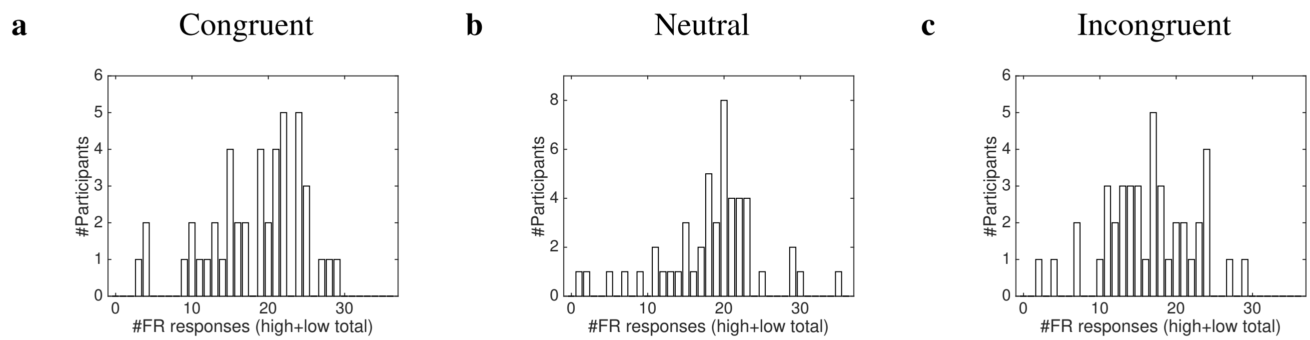

Figure 6. Distribution of total number of high and low items recalled in free recall, for Experiment 3, out of 36 possible.

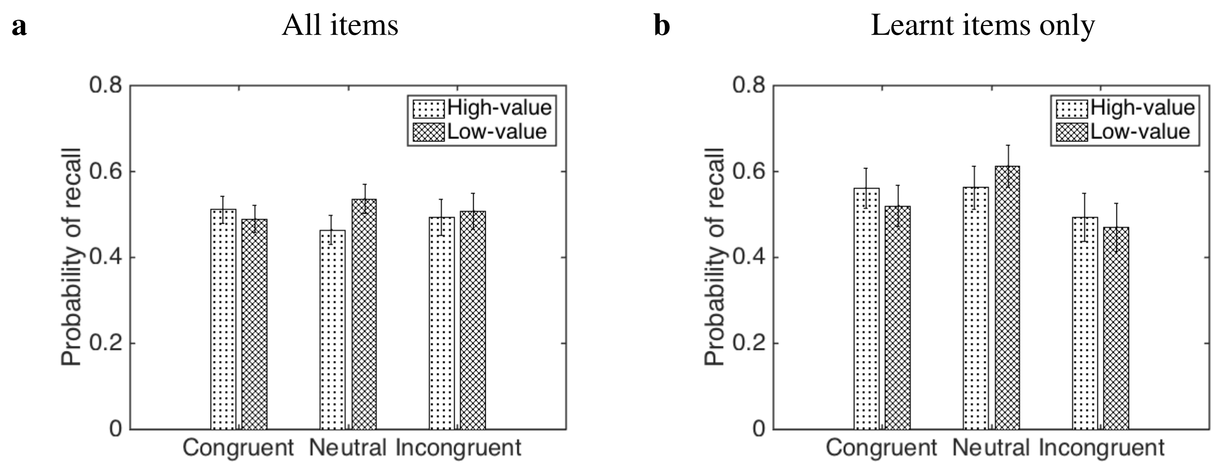

Figure 7. Probability of recall of high- and low-value words, (a) considering all words (b) considering words that were learnt during value-learning for Experiment 3. Error bars plot standard error of the mean, corrected for between-subjects variability (Loftus and Masson, 1994).

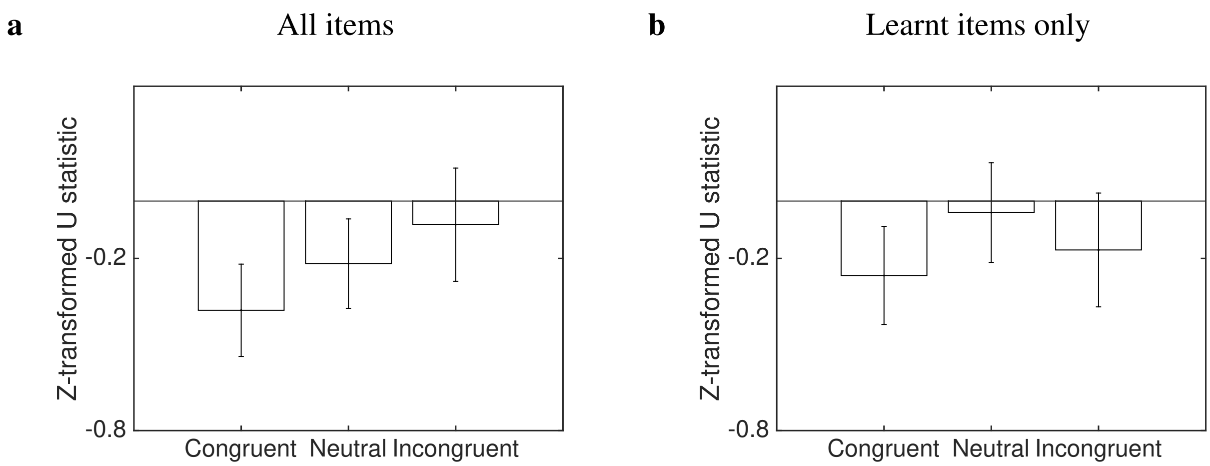

Figure 8. Output position for recall of high- and low-value words, (a) considering all words (b) considering words that were learnt during value-learning for Experiment 3. Error bars plot standard error of the mean.

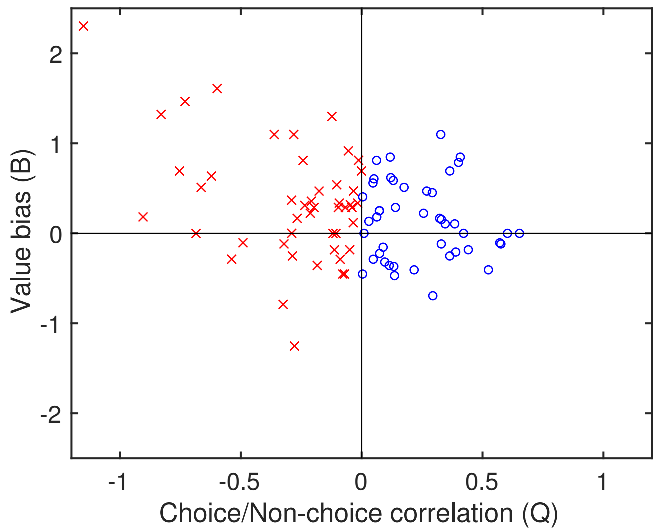

Figure 9. The value bias (B) of free recall of Madan et al. (2012) Experiment 1, plotted as a function of the Choice/Non-choice correlation (Q). Participants with $Q \leq 0$ are coded in red cross and those with $Q \geq 0$ are coded in blue circles.

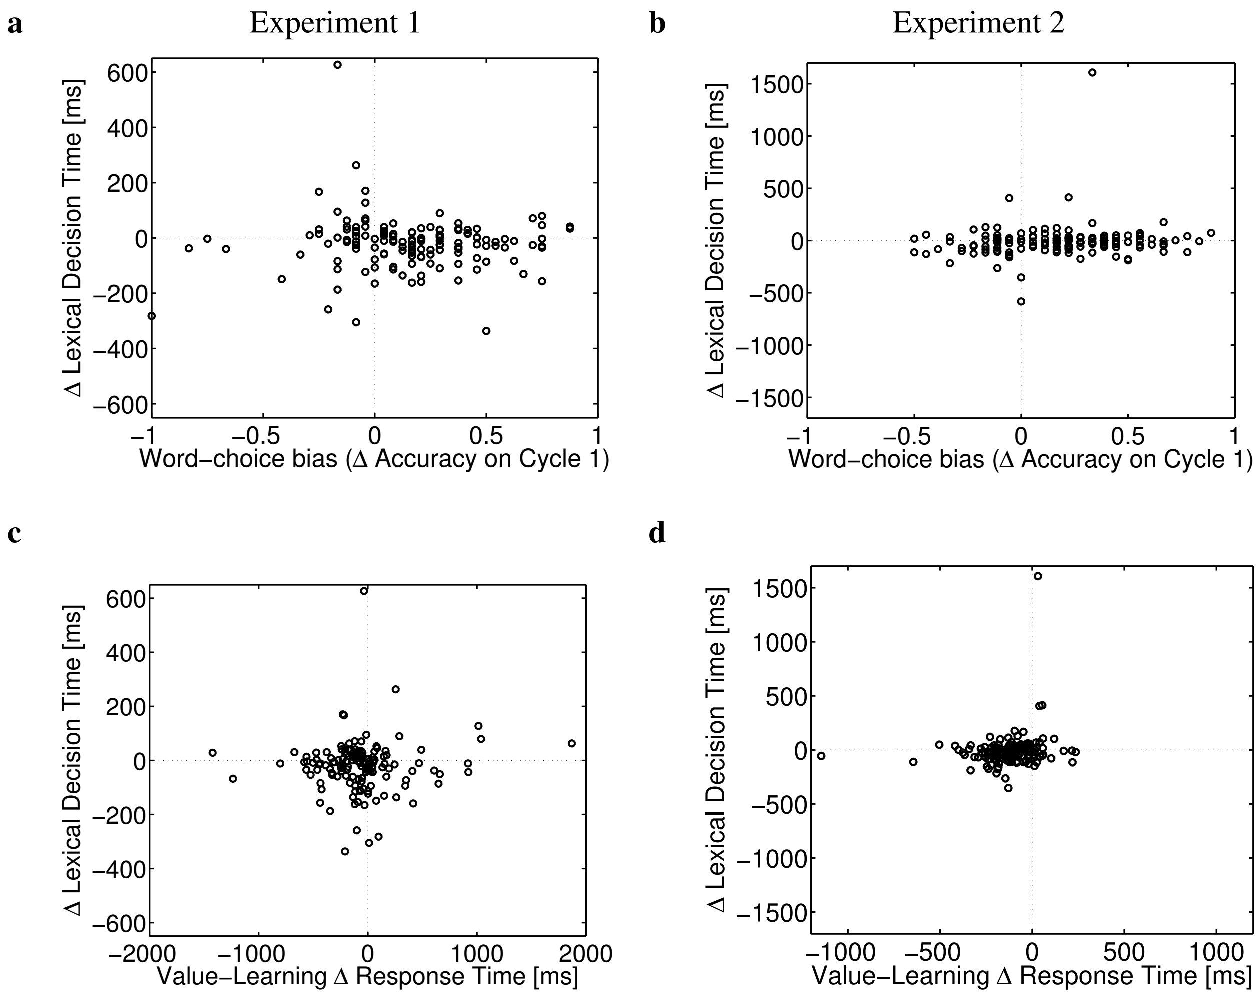

Figure 10. The value bias of lexical decision (difference in lexical decision time between high- and low-value words, restricted to words whose value had been learnt) is plotted as a function of the word-choice bias (a,b) and the difference in response time during value-learning (c,d), for Experiment 1 (a,c) and Experiment 2 (b,d). The measure of the word-choice bias is the difference in accuracy during the first cycle of value-learning, at which time no values are known, and all responses must be guesses. The difference in value-learning response time is computed for correct responses only, on the last cycle of value-learning. The scatter plots show that there are no grounds to suspect that either the word-choice bias or the difference in value-learning response time could have been singly responsible for the value-bias of lexical decision.

Supplementary Materials

1. **Additional analysis on free recall (Experiment 2):** Since we found a significant effect of value on free recall probabilities for the FR-LD group, we wondered how this effect may map on to individual words during value learning rather than proportions. Thus we tested two additional hypotheses: 1) if increased number of correct responses to a word during value learning also increases its chances of being recalled later and 2) if this happens differently for high and low-value words. To test it, we applied a paired t test to the number of accurate responses to an item during value-learning for later-free-recalled versus later-not-free-recalled words. This difference was significant for high-value words, $t(177) = 1.89,\ p = 0.06$, but was nowhere near significant for low-value words, $t(177) = 0.43,\ p = 0.67$, and neither was the difference of these differences between high- and low-value words ($p > 0.2$). For Group LD-FR, both were non-significant ($p > 0.5$). For Group FR-LD, the comparison was nearly significant for high-value words, $t(83) = 1.94,\ p = 0.056$, but not for low-value words, $t(83) = -0.026,\ p = 0.98$; and the difference between the effect of value on free recall, was, itself, not significantly different for high-value than low-value words, $t(83) = 1.50,\ p = 0.14$. This suggests that learnt value may parametrically increase probability of free recall of high-value words, but not increase (possibly even decrease) probability of recall for low-value words, although these effects are quite small in magnitude.

2. **Restriction on participation (Experiment 3):** We were concerned that participants who are possible problem gamblers might relate to reward differently and thus included the PGSI questionnaire to screen them out. Notably, with the introductory psychology participant pool, rates of pathological PGSI scores are quite low (only 4 participants out of 201 in Experiment 2 had a score greater or equal to 3, the lower bound for “moderate” gambling problems), too low to meaningfully analyze. In Experiment 3, we simply opened the study only to participants with non-zero PGSI scores in order to obtain a more uniform sample. The PGSI was administered in an online, mass-testing questionnaire taken as part

of the course requirement, prior to the experiment. The Problem Gambling Severity Index (PGSI) (Ferris & Wynne, 2001) is a 9-item questionnaire, developed as part of the Canadian Problem Gambling Index, measuring gambling involvement, problem gambling behaviour, and adverse consequences. It was developed to measure gambling severity and problem gambling in the general population. The PGSI has good internal consistency ($\alpha = .84$), re-test reliability ($\rho = .78$), and convergent validity with other measures of problem gambling (i.e., SOGS; Ferris & Wynne, 2001).

3. **Crossed-random mixed-effects analysis for free recall (Experiment 3):** We report mixed-effects models for free recall data from Experiment 3, conducted with the R-package lme4 (Bates, 2007). Recall is the dependent variable; value (high, low) and group (congruent, neutral and incongruent) are the fixed effects. We also include the random effects: subjects and words. The data family is binomial.

   - **Mixed Model 1:** The main model included adjustments to both intercepts and slope due to each random factor but it fails to converge.

     $m1 < -glmer(recall\ class + gID + (1|word) + (1|subID) + (1 + recall|word) + (1 + recall|subID), data = d, family = binomial)$

     Warning message:

     In $checkConv(attr(opt, "derivs"), opt\$par, ctrl = control\$checkConv,:$ Model failed to converge with $max|grad| = 0.0018826$ $(tol = 0.001, component1)$

   - **Mixed Model 2:** Thus we simplified m1 by removing the correlation parameter and by assuming homoskedasticity of subjects and words w.r.t. recall.

     $m2 < -glmer(recall\ class + gID + (1|word) + (1|subID) + (1|recall : word) + (1|recall : subID), data = d, family = binomial)$

     > print(m2,corr=FALSE)

     Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']

Family: binomial (logit)

$Formula : recall\ class + gID + (1|word) + (1|subID) + (1|recall : word) + (1|recall : subID)$

Data: d

AIC BIC logLik deviance df.resid

199.8126 251.7820 -91.9063 183.8126 4888

Random effects:

Groups Name Std.Dev.

recall:subID (Intercept) 0.000e+00

subID (Intercept) 8.197e-05

recall:word (Intercept) 8.931e+01

word (Intercept) 0.000e+00

Number of obs: 4896, groups: recall:subID, 272; subID, 136; recall:word, 108; word, 54

Fixed Effects:

(Intercept) classlow gID2 gID3

15.80472 0.02740 -0.10358 0.01372

We found that all the variances due to words and subjects were very small, specially the variance for the by subject adjustment to recall was almost zero, so was the random intercept for words and subjects.

- **Mixed Model 3:** Thus, we further simplified m2 by including only adjustment to the slope due to words.

  $x < -glmer(recall\ class + gID + (1 + recall|word), data = d, family = binomial)$

  > print(x,corr=FALSE)

  Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) ['glmerMod']

Family: binomial (logit)

Formula: recall class + gID + (1 + recall | word)

Data: d

AIC BIC logLik deviance df.resid

144.9205 190.3937 -65.4602 130.9205 4889

Random effects:

Groups Name Std.Dev. Corr

word (Intercept) 29.87

recallTRUE 59.77 -1.00

Number of obs: 4896, groups: word, 54

Fixed Effects:

(Intercept) classlow gID2 gID3

-0.010611 0.031408 -0.098991 0.009597

A follow up Chi-square test showed that model m3 is justified as it does not differ significantly from the original model m1.

anova(m1,m3)

Data: d

Models:

$m3 : recall\ class + gID + (1 + recall|word)$

$m1$:

$recall\ class + gID + (1|word) + (1|subID) + (1 + recall|word) + (1 + recall|subID)$

Df AIC BIC logLik deviance Chisq Chi Df Pr(>Chisq)

m3 7 144.92 190.39 -65.460 130.92

m1 12 154.92 232.87 -65.458 130.92 0.0037 5 1

Still we see no main effects or interaction for the fixed effects in the final model m3.

summary(m3)

Generalized linear mixed model fit by maximum likelihood (Laplace Approximation) [glmerMod]

Family: binomial (logit)

Formula: recall class + gID + (1 + recall | word) Data: d

AIC BIC logLik deviance df.resid

144.9 190.4 -65.5 130.9 4889

Scaled residuals:

Min 1Q Median 3Q Max

-0.011764 -0.010509 -0.009732 0.010650 0.012195

Random effects:

Groups Name Variance Std.Dev. Corr

word (Intercept) 892.3 29.87

recallTRUE 3572.0 59.77 -1.00

Number of obs: 4896, groups: word, 54

Fixed effects:

Estimate Std. Error z value Pr(>|z|)

(Intercept) -0.010611 1.883722 -0.006 0.996

classlow 0.031408 0.403237 0.078 0.938

gID2 -0.098991 0.503916 -0.196 0.844

gID3 0.009597 0.484626 0.020 0.984

Correlation of Fixed Effects:

(Intr) clsslw gID2

classlow -0.110

gID2 -0.128 0.000

gID3 -0.134 0.002 0.500

Overall, we found that this approach reiterated our previous results, and including Word as a random factor did not greatly change our sensitivity to the effects of interest. Thus, in order to maintain consistency with Madan et al. (2012), we did not include Word as a random factor in the main analyses of our experiments.

4. **Additional analysis on free recall (Experiment 3):** As in Experiment 2, we asked if the chances of a word being recalled was related to the number of correct responses it received during value learning. To test this, for each subject, we measured the number of correct responses to each words during value learning, separated by Value and later Recall status (i.e., recalled or not). Averages across words were taken for each subject for each Value and Recall status. Then, we conducted a 2×2×3 ANOVA on this measure with two within-subjects factors: Value (high and low) and Recall Status (recalled and not recalled) and one between-subjects factor Group (Congruent, Neutral and Incongruent). This revealed a main effects of Value, $F(1,133) = 81.71,\ p < 0.01$, as well as Recall Status, $F(1,133) = 4.95,\ p < 0.05$. However, the main effect of Group was not significant, $F(2,133) = 0.62,\ p = 0.54$, nor were any interactions (all $p > 0.2$). The main effect of Value was due to the average number of correct responses toward high-value words (13.8) being greater than that towards low-value words (11.6), reflecting more accurate responses to high- than low-value words during value-learning, due to the word-choice bias. The main effect of Recall Status was due to slightly more correct responses during value training for words that were later recalled (12.83) than for words that were later not recalled (12.56). Thus, the number of correct responses during value learning may have influenced probability of free-recall, but unlike in Experiment 2, not differentially based on value, and not differently for the three groups.