

Adaptive Magnitude and Valence Biases in a Dynamic Memory Task

George Kachergis, Gabriel Recchia, and Richard M. Shiffrin

{gkacherg, grecchia, shiffrin}@indiana.edu

Department of Psychological & Brain Sciences, 1101 E. 10th St.

Bloomington, IN 47405 USA

Abstract

Given that human memory is fallible, it is likely adaptive for people to preferentially encode, retain, and retrieve important items better than insignificant ones. Using a dynamic decision-making paradigm with a response deadline, we find that humans demonstrate a bias to better remember 1) items with positive rather than negative value, and 2) items with high-magnitude values. Performance was greater when participants were shown all item-value pairs simultaneously, and were thus able to selectively attend to high-magnitude values. The same magnitude bias is observed for sequentially studied positive items, but not for negative items. Decision trajectories show participants sometimes change their minds during the course of a trial, choosing an item after first moving toward the other. Changes of heart occurred more often for trials with negative items. These findings suggest that memory is sensitive to value, and that real-time game paradigms can be used to reveal dynamic memory processes.

Keywords: adaptive memory, recognition, attention, loss aversion, magnitude bias

Introduction

Availability of specific items in memory is sensitive to frequency and recency, a principle codified in the rational approach to memory as the importance of likely need (Anderson & Milson, 1989; Anderson & Schooler, 1991; Adelman, Brown, & Quesada 2006). This property of memory is argued to be adaptive: items that are accessed frequently or recently tend to have greater relevance and utility, and thus having them highly available in memory confers an evolutionary advantage. If memory resources are preferentially allocated to high-utility items, one might expect memory to exhibit a magnitude bias: When participants must associate items with values, items associated with particularly positive or negative values (high magnitudes) should be more available in memory than those associated with less extreme values (low magnitudes). Given the considerable consequences of selecting a stimulus item with an extreme reward or an extreme penalty, such items are particularly important to remember in comparison to items with middling values. Atkinson and Shiffrin (1968) lays out a theory of control processes that emphasizes selective control processes in short term memory, including selective rehearsal strategies that could be used to explain magnitude effects.

Although merely finding a magnitude bias might seem unremarkable, there are several reasons to verify that such an effect exists. First, it is of intrinsic interest whether or not high-magnitude items are more readily retrieved; to our knowledge, no such magnitude bias has been explicitly demonstrated. Second, an underlying assumption of decision-making models is that the value of an item does not

influence its probability of correct retrieval. For example, in studies of experience-based decision making, in which participants must learn and maintain items associated with particular payoffs (or a distribution of possible payoffs) in memory, preferences are explained in terms of a value function that is applied only after retrieval. Although recent work has begun to explore memory's role in the decision process (Rakow & Newell, 2010; Hau, Pleskac & Hertwig, 2009; Lejarraga, 2010; Rakow & Rahim, 2009), the role of memory is limited to frequency, recency, and the number of sampled events that can be held in memory; the possibility that high-magnitude items may hold a privileged place in memory has not been investigated. Similarly, in instance-based learning theory (Gonzalez, Lerch, & Lebiere, 2003; Lejarraja, Dutt, & Gonzalez, 2010), models of repeated binary choice tasks include the probability of retrieval of prior items from memory, but retrieval probabilities are only influenced by item frequency and recency; the value of the item affects eventual choice, but not retrieval. Demonstration of a magnitude bias in memory would require significant modification of such models. At test, we utilize a novel decision-making game task that can elucidate the dynamics of the decision process, potentially yielding insight into the time-course for retrieval of magnitude information.

Third, we investigate whether memory is stronger for negative or positive items, which may help our understanding of memory for emotionally-valent events, in general. Negative events appear to command more attention overall (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Rozin & Royzman, 2001) and are remembered more accurately than positive events in some contexts, such as an emotionally charged public event (Kensinger & Schacter, 2006). Rozin & Royzman (2001) presents various theoretical accounts of the adaptive value of focusing on negative events, most of which apply primarily to dire threats; "in the extreme, negative events are more threatening than are positive events beneficial" (p. 314). A primary finding of the negative bias literature is that negative entities command more attention than positive ones, as they are more salient. Thus, our memory for extreme negative events may be at least partially explained by the fact that they command more attention. When we restrict selective attention, memory for positive events may dominate.

Finally, we go beyond investigation of a simple magnitude bias, and also ask how the context—the values of other items on a list—affects memory of particular items.

For example, a valence bias that exists when half of the studied items are negative and half are positive may be different when most items are negative, or when most are positive. Such an effect would be reminiscent of a Von Restorff distinctiveness effect.

With so many issues at stake, it is unlikely that they will all be resolved with a single study. However, the lack of foundational work establishing the effects of magnitude and valence (reward or punishment) in memory motivated us to take a first step by examining memory for item-value pairs with respect to their valence and magnitude (i.e., extremeness vs. mildness of the reward/penalty) in three valence contexts. To examine differences between selective attention and selective rehearsal, some sets were simultaneously and some were studied in sequential fashion. In keeping with the previous discussion, we hypothesized stronger memory for high-magnitude vs. low-magnitude items, and stronger memory for high-magnitude negative items than high-magnitude positive items. To differentially impact attention and memory, item-value pairs were either studied simultaneously or sequentially, for the same total amount of time per pair. Selective attention may play a larger role during simultaneous study, in which participants may choose to study some pairs more than others, rather than during sequential study, when each pair is shown for the same amount of time. Effects in sequential study are more likely to be due to memory processes such as selective rehearsal. We believe our findings will be informative for models of memory and decision-making, and will suggest future manipulations. Moreover, we believe our novel decision-making task could be of great utility in understanding the dynamics of memory, learning, and categorization, in general.

Experiment

Traditional episodic memory experiments present a series of individual items with no indication that any item is more important than any other. In contrast, our experiment varies both the arrangement of items at study (simultaneous or sequential) and the value distribution of the studied items. Two value distributions employed were evenly spaced, but shifted to be mostly negative or mostly positive: will participants focus more on the oddball positive (or negative, resp.) stimuli in these conditions? The other two value distributions were symmetric about zero, but the magnitude of the most extreme values differed: will participants remember extreme-valued items twice the value of the middling items (e.g., 20 vs. 10) as well as when they are ten times the value (e.g., 100 vs. 10)? Overall, will participants better remember positive or negative items?

Item-value pairs were either studied simultaneously or sequentially, for the same total amount of time per pair. Selective attention may play a larger role during simultaneous study, in which participants see the value distribution and can choose to study some pairs more than others (and may selectively encode salient item features), rather than during sequential study, when each pair is shown

for the same amount of time. Effects in sequential study are more likely to be due to memory control processes such as selective rehearsals. Will sequential study nonetheless yield magnitude effects and a valence bias?

Subjects

68 undergraduates at Indiana University participated to receive course credit.

Stimuli

Each stimulus was a colored circle (50 pixels in diameter) inscribed with a differently colored polygon. The stimuli for each participant were chosen uniformly at random from 810 such stimuli (10 outer colors \times 9 inner colors \times 9 polygons). For each participant, 16 items were randomly assigned to each of the 16 blocks as targets and foils. Figure 1 shows an example set of studied item-value pairs. Thus, although the similarity of stimuli in any given set was arbitrary, and may have at times been advantageous (e.g., grouping similarly-valued objects) or disadvantageous (e.g., foils may be similar to targets), each participant saw a randomly-selected set, and saw no stimulus more than once.

In each condition, eight objects were studied: two each of four unique values. Four value distributions of objects were used: two were the same distribution, but shifted mostly negative (-50, -30, -10, 10) or mostly positive (-10, 10, 30, 50). The remaining two distributions were symmetric about zero, but varied in magnitude: small (-20, -10, 10, 20) or large (-100, -10, 10, 100).

During simultaneous study, participants were shown all eight target stimuli with their values on one display (e.g., Figure 1) for 40 seconds. During sequential study, each of the 8 item-value pairs was individually studied for five seconds, in a randomized order.



Figure 1. An example of 8 studied objects and their values. During simultaneous study, participants were shown such a display for 40 seconds.

Procedure

Participants were instructed that they would be playing a game in which their goal would be to acquire points by

avoiding poisons and collecting foods (i.e., negatively- and positively-valued objects, respectively). Participants were told that after studying, they would choose between the studied objects and new objects, which are always worth zero points. After the 40-second study period (simultaneous or sequential), participants were told that pairs of objects would fall from the top of the screen, and that they would need to choose the more valuable object to catch by moving horizontally with the arrow keys. Moreover, they were told that if they are positioned below an object and ready to catch it, pressing the up arrow would shoot it down (see Figure 2). They were reminded that unstudied objects are neutral (worth 0 points), and thus should be chosen if the available studied item is negative.

On each test trial, two items—one studied, and one novel—appear, horizontally separated by a distance selected uniformly at random from [70, 255] pixels, vertically separated from the participant's agent by a distance selected uniformly at random from [270, 400] pixels, moving downward at a constant rate.

The frame rate of the experiment is 88 frames per second (12 ms/frame) on 15" CRT monitors with a resolution of 800x600 pixels. Objects drop at a rate of 1 pixel per frame, and 'bullets' move at a rapid rate of 8 pixels per frame. The participant-controlled agent, starting equidistant between the objects with 0 velocity, accelerates .6 pixels/frame when the left or right key is depressed: thus, the moving agent has inertia, and cannot immediately reverse direction. However, there is linear 'friction': if no key is pressed, the agent will lose ten percent of its speed on every frame. A trial ends when an object hits a bullet or the agent, at which time the value of the chosen object is shown for 1300 ms (e.g., Fig. 2). If the object passes below the horizontal plane of the agent, the participant is loses 30 points and is told to always pick an object. Participants' running score, tallied across all conditions, is shown throughout testing in the upper left corner of the screen.

Subjects participated in each of the eight unique study conditions twice, for a total of 16 blocks, each with eight trials. Condition order was counterbalanced across subjects.




Figure 2. Screenshot of a trial, in which the participant moved the agent (bottom) to the left and 'shot' a studied, poisonous item. Feedback appeared only after the decision was made. Shooting the unstudied alternative would have yielded '+0' points (as on every trial).

Results

Six participants were excluded from analysis because their overall performance was not significantly higher than chance (.535 for 128 trials), where accuracy is both a function of choosing old positive items, as well as the new item when a negative item is shown. The remaining 62 participants were first analyzed in terms of their probability of choosing the correct (i.e., non-negative) item. Note that this corresponds to hits when the studied item on a trial is positive, and correct rejections when the studied item is negative. An analysis of variance on the study arrangement (simultaneous or sequential), the valence of the old item (positive or negative), the absolute value (i.e., magnitude) of the old item (100, 50, 30, 20, or 10) nested by value distribution shows significant main effects of study arrangement ($F(1,61) = 5.53, p < .05$), valence ($F(1,61) = 15.73, p < .001$), and a main effect of absolute value ($F(4,61) = 2.31, p = .06$). A few interactions fell just short of significance: study arrangement by valence ($F(1,61) = 2.46, p = .12$), study distribution by valence ($F(3,183) = 1.58, p = .19$), and study arrangement by study distribution by magnitude ($F(2,122) = 1.62, p = .20$). All other interactions had F-values less than one.

Accuracy after simultaneous study was superior to accuracy after sequential study ($M = .68$ and $M = .64$, respectively), suggesting that participants benefit from being able to allocate attention to items of their choosing during study. Correct rejection of negative items was significantly worse than hit rate of positively valued items ($M = .62$ and $M = .71$, respectively), showing a bias for positive items that was not dependent on selective attention during simultaneous study. Finally, memory was better for large magnitude items than for items with middling value. In Figure 3, ceiling performance for negative values is 0, and for positive values is 1. Thus, monotonically increasing slopes indicate a magnitude bias. For simultaneously studied items, a magnitude bias is seen for both negative and positive items in every distribution, but only for the positive sequentially studied items.

In the asymmetric distributions, in which the oddball is worth either +10 or -10, there is a clear advantage of simultaneous study over sequential study: the small-magnitude oddballs get more attention when all other items are seen to be of the opposite valence. In contrast, there was little or no advantage of simultaneous study in asymmetric distributions for the middle and extreme values (which always were part of the majority). Thus, we found a distinctiveness effect after simultaneous, but not sequential study. Although selective attention yielded an overall performance advantage beyond sequential study, participants still showed magnitude and valence effects in the sequential conditions. One possible explanation is that participants selectively rehearsed the items with large magnitude or positive values, yielding higher fidelity memory traces for these items, and a stronger familiarity signal at test.




Figure 3. Probability of choosing an old item by value distribution, item value and study arrangement. For negative items, ceiling performance is 0—the old item should never be chosen. For positive items, ceiling is 1. Error bars: \pm SE.

To determine whether the magnitude sensitivity and valence bias in accuracy extend to the time domain, we first examine participants' median time to make a first movement. Time until first movement in our continuous time paradigm may be equivalent to reaction time in traditional experiments, although it is also bounded by the drop speed of objects. Time pressure was not extreme: mean trial duration was 1948 ms ($SD = 218$ ms). An ANOVA on median time to make a first movement by study arrangement, valence, absolute value, and correctness, nested in value distribution finds significant main effects of study arrangement ($F(1,61) = 9.97, p < .01$) and valence ($F(1,61) = 4.66, p < .05$). We also find significant interactions of valence and correctness ($F(1,61) = 7.84, p < .01$), distribution and valence ($F(3,183) = 3.39, p < .05$), and a marginal interaction of study arrangement and distribution ($F(3,183) = 2.16, p = .09$). Participants were slower to initiate movements on test trials of simultaneously studied items ($M = 1143$ ms) than sequentially studied items ($M = 1115$ ms), and they were much slower on trials with a negative item ($M = 1173$ ms) than trials with a positive item ($M = 1081$ ms). For correct decisions, first movements were faster for positive items ($M = 1043$ ms) than for negative items ($M = 1150$), which were nearly the same as incorrect choice initiation times on positive item trials ($M = 1135$ ms), and much faster than incorrect choice initiation times on negative item trials ($M = 1200$ ms). Thus, time to first movement shows an advantage for positive items, and especially for correctly remembered positive items. Does

this imply that for well-known items, the decision is made by the time of first movement?

To address this question, we looked at whether participants change their minds (i.e., crossed the mid-point between the two items) after their first movement on a trial. Indeed, on 9.1% of trials, participants changed their minds, reversed directions and crossed the midline—a result that could not be found in a traditional memory experiment. An ANOVA on crossover rate shows no significant main effects, although valence approached significance ($F(1,61) = 2.34, p = .13$). However, there were significant interactions of study arrangement, value distribution, and valence ($F(3,183) = 2.97, p < .05$); as well as study arrangement, value distribution, and correctness. Figure 4 shows that there are fewer changes of mind for positive items after simultaneous study, whereas after sequential study, crossover rates vary significantly within and between value distributions. Fewer changes of mind for positive objects after simultaneous study—which accuracy tells us participants prefer to focus on—complements the story told by time until first movement: participants are quickly identifying the well-known positive item and moving toward it, without going back.




Figure 4. Crossovers (changes of mind) as a function of value distribution, study arrangement, and item valence. After simultaneous study, positive items had fewer crossovers, reflecting participants' greater certainty about them after presumably preferentially studying them. After sequential study, there is no clear pattern.

Figure 5 shows that after simultaneous study, there are fewer crossovers for correct responses except for the mostly-negative asymmetric distribution. However, after sequential study, only the mostly-positive asymmetric distribution shows fewer crossovers for correct responses. The ambigousness of crossover rates after sequential rehearsal in both Figures 4 and 5 reflects the difficulty of sequential rehearsal when the item and value distribution is unknown.




Figure 5. Crossovers as a function of value distribution, study arrangement, and correctness. After simultaneous study most conditions show fewer crossovers for correct responses, but this pattern is not seen after sequential study.

Finally, we analyze the decision trajectories themselves by finding the agent's integrated distance from the midline over time on each trial. For example, if a participant never moves, or if the participant wavers back and forth on either side of the midline the entire trial, the integral will be 0. However, if a participant decisively moves toward one item at the beginning of the trial, the integral (i.e., decisiveness) will be large. Thus, the integral is partially a function of both time until first movement (earlier movement yields a larger integral) and crossovers (more crossovers will yield a smaller integral). An ANOVA finds significant main effects of valence ($F(1,61) = 7.30, p < .01$) and correctness ($F(1,61) = 7.31, p < .01$), and a marginally significant interaction of valence and correctness ($F(1,61) = 2.78, p = .09$). Corroborating other measures, the integral is larger for correct items, and particularly for positive items (correct: positive $M = 67,016$, negative $M = 57,400$; incorrect: positive $M = 55,500$, negative $M = 54,680$).

Discussion

Episodic memory tasks typically ask participants to simply remember all studied memory items, with no explicit indication that any items are more important than others. However, given that human memory is far from perfect, it is not unreasonable to expect that our memory system has evolved to preferentially store, retain, and retrieve information that is deemed to be important. Using point values to indicate item importance, we set out to determine whether participants selectively attend to and better retrieve more important items. By including items with both positive and negative values, we also aimed to abstractly test the broad hypothesis that episodic memory is stronger for bad events than for good events (Rozin & Royzman, 2001;

Baumeister et al., 2001). Faced with items of varying magnitude and valence that will be tested against novel, neutral objects, one can imagine several reasonable strategies, including: 1) focus on the positive objects, and avoid the unknown, 2) focus on and avoid the negative objects, or 3) attempt to remember the extreme-valued items of both positive and negative valence, at the risk of confusing one for the other. Other strategies, mixtures of the above, and nonstrategic differences are also possible.

We have presented the first evidence for a value-based recognition memory bias. Participants showed significantly better memory for positive items, as well as for items that were of greater absolute value. These biases were present when participants were shown all item-value pairs at once, and were thus able to choose which to attend to, but were also somewhat evident when study time for each item-value pair was controlled by presenting them sequentially, for the same total amount of time per pair. Thus, we suggest that these valence and magnitude biases in memory are not solely based on selective attention, but can also result from encoding and retrieval processes. Although sequential study does not preclude the possibility that participants simply disregarded items they considered less important, or selectively rehearsed more important items, the unknown and sometimes asymmetric value distributions, randomized order of item presentation, and roughly equal performance for the same item values across conditions reduce the likelihood that these are the sole explanations. Using item value distributions asymmetric about zero (i.e., with a preponderance of negative or positive items, and only two low-magnitude oddballs) we found a list context effect: under simultaneous study, participants were as sensitive to the oddball items as to the highest-magnitude items. After sequential study, participants showed only the same sensitivity benefit for high-magnitude items, and no advantage for the oddballs. Given limited memory for confusable items, the bias that people exhibit for remembering high magnitude items is adaptive. The preference for positive items is more pervasive in our data.

We found further evidence of valence effects during retrieval in our examination of time until first movement: participants move sooner for positive items than negative ones. This decisiveness for positive items was also reflected in integrated distance. Finally, we found that participants change their minds—crossing the midline well after their initial movement—a significant proportion of the time. The rate of preference reversals was affected by study arrangement, value distribution, valence, and correctness of the eventual response. After simultaneous study, participants changed their minds fewer times for correct responses than incorrect responses, and fewer times for positive items than negative items. After sequential study, the rates of crossovers were inconsistent, and beg further teasing apart. Greater time pressure by increasing drop speed, larger differences in value magnitude may reveal more interpretable effects.

Consideration of changes of mind may be able to address an important debate in the memory literature. It has been suggested that decisions in both item recognition (Hintzman & Curran, 1994) and associative recognition (Clark, 1992) may involve two memory processes. Using a response-signal recognition paradigm, Hintzman & Curran (1994) found that participants faced with rejecting test items that were quite similar to studied targets (plural vs. singular nouns) showed biphasic false alarm curves, suggesting that an early familiarity process is supported by a slower recollection process that can better reject similar, unstudied items. However, reanalysis, further empirical data, and modeling have questioned the necessity of positing two processes (Rotello & Heit 1999; Ratcliff, Van Zandt, & McKoon, 1995). In an associative recognition paradigm, Dosher (1984) found that participants initially accepted pairs with semantically-related lures, and only later correctly rejected them. Our task is somewhat like a response-signal paradigm, but with a few advantages: 1) because they see the falling items, participants can visually judge how long they have to respond, alleviating anticipatory uncertainty, 2) the response deadline is hard, and times should thus be less variable, and 3) with only a single deadline on a trial, we can observe how the response trajectory changes over time, so it may be unnecessary to include several signal intervals. Although it is not yet clear whether our task should be treated as associative recognition, or as attentionally-modulated single item recognition, we believe our paradigm can provide unique insight into this debate.

The potential to discover changing effects during retrieval warrants the further use of dynamic decision-making tasks such as the paradigm used here. We have demonstrated that our point-based memory game can be utilized to find novel recognition memory effects, measurable both by traditional accuracy and response time variables, as well as new dependent variables such as crossover rate and integrated trajectory. We hope that other researchers will find our paradigm useful for studying not only memory, but also learning and categorization.

Acknowledgments

The authors would like to thank Gregory Cox, Christopher Donkin, and Caitlin Fausey for helpful comments. We would also like to thank Brennen Jenkins, Kyung Lee, Colleen Monn, Kim Mullen, Gladys Raintama and Autumn Salisbury for running participants. This research was funded by an AFOSR grant to Richard Shiffrin.

References

Adelman, J. S., Brown, G. D. A., & Quesada, J. F. (2006). Contextual diversity, not word frequency, determines word-naming and lexical decision times. *Psychological Science*, 17, 814–823.

Anderson, J.R., & Milson, R. (1989). Human memory: An adaptive perspective. *Psychological Review*, 96, 703–719.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.

Atkinson, R. C. & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K.W. Spence and J. T. Spence (Eds.), *The Psychology of Learning and Motivation: Advances in Research and Theory*, Vol. 2, 89–195). New York: Academic Press.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Volhs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5, 323–370.

Clark, S. E. (1992). Word frequency effects in associative and item recognition. *Memory & Cognition*, 20, 231–243.

Dosher, B. A. (1984). Discriminating pre-experimental (semantic) information from learned (episodic) associations: A speed-accuracy study. *Cog Psych.*, 16, 519–555.

Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27, 591–635.

Hau, R., Pleskac, T. J., & Herwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, 23, 48–68.

Hintzman, D. L., & Curran, T. (1994). Retrieval dynamics of recognition and frequency judgments: Evidence for separate processes of familiarity and recall. *Journal of Memory and Language*, 33, 1–18.

Kensinger, E. A. & Schacter, D. L. (2006). When the Red Sox shocked the Yankees: Comparing negative and positive memories. *Psychonomic Bulletin & Review*, 13(5), pp. 757–763.

Lejarraga, T. (2010). When experience is better than description: time delays and complexity. *Journal of Behavioral Decision Making*, 23, 100–116.

Rakow, T. & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, 23, 1–14.

Rakow, T., & Rahim, S. B. (2009). Developmental insights into experience-based decision making. *Journal of Behavioral Decision Making*, 23, 69–82.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1995). Process dissociation, single process theories, and recognition memory. *Journal of Experimental Psychology: General*, 124, 352–374.

Rotello, C. M. & Heit, E. (1999). Two-Process Models of Recognition Memory: Evidence for Recall-to-Reject? *Journal of Memory and Language*, 40, 432–453.

Rozin, P., & Royzman, E. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5, 296–320.