THE INFLUENCE OF REWARD VALUE ON MEMORY AND DECISION MAKING

by

Christopher R. Madan

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Psychology
University of Alberta

Abstract

In our everyday lives we often make decisions based on our prior experiences, whether it be choosing to park without putting money in the meter or deciding what to buy as a gift for a loved one. Inevitably, our decisions in the present are informed by our memories of experiences past. In this dissertation I report the results from a series of studies examining how reward value influences memory, and how these reward-memory effects can in turn bias decision making such that people are generally more risk seeking for relative gains than relative losses. Specifically, these studies examined how previously learned reward values can subsequently influence memory for items, how more extreme reward outcomes influence decisions from experience, and how memory biases can drive risk preference in decision making. Together these convergent lines of research represent a theoretical advance in our understanding of memory and decision making.

Preface

This thesis is an original work by Christopher R. Madan. The research project, of which this thesis is a part, received research ethics approval from the University of Alberta Research Ethics Board, Project Name "Organisation and Retrieval Timecourse of Human Memory", No. Pro00009760, January 8, 2010 and Project Name "Neurobiology of decision-making in gambling" No. Pro00014795, June 24, 2010.

Chapter 2 of this thesis has been published as C. R. Madan, E. Fujiwara, B. C. Gerson, and J. B. Caplan, "High reward makes items easier to remember, but harder to bind to a new temporal context," Frontiers in Integrative Neuroscience, vol. 6, 61. I was responsible for the data collection and analysis as well as the manuscript composition. B. C. Gerson was involved with concept formation and assisted with data collection. E. Fujiwara and J. B. Caplan were involved with concept formation and manuscript composition.

Chapter 3 of this thesis has been published as C. R. Madan and M. L. Spetch, "Is the enhancement of memory due to reward driven by value or salience?," Acta Psychologica, vol. 139, 343–349. I was responsible for the data collection and analysis as well as the manuscript composition. M. L. Spetch was involved with concept formation and manuscript composition.

Chapter 4 of this thesis has been published as E. A. Ludvig, C. R. Madan, and M. L. Spetch, "Extreme outcomes sway risky decisions from experience," Journal of Behavioral Decision Making, vol. 27, 146–156. I was responsible for the data collection and analysis as well as the manuscript composition. E. A. Ludvig was involved with concept formation, analysis, and manuscript composition. M. L. Spetch was involved with concept formation and manuscript composition.

Chapter 5 of this thesis has been published as C. R. Madan, E. A. Ludvig, and M. L. Spetch, "Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions," Psychonomic Bulletin & Review, vol. 21, 629–636. I was responsible for the data collection and analysis as well as the manuscript composition. E. A. Ludvig was involved with concept formation, analysis, and manuscript composition. M. L. Spetch was involved with concept formation and manuscript composition.

Acknowledgements

First and foremost, I would sincerely like to thank my supervisor, Dr. Marcia Spetch, to whom I am grateful for her guidance. I would also like to thank my Ph.D. supervisory committee, which consisted of Drs. Marcia Spetch, Christopher Sturdy, and Anthony Singhal.

The work described in this dissertation is only a portion of the work I have done at the University of Alberta, but here I would like to thank all of those who have contributed to my overall Ph.D. experience. I am indebted to the professors with whom I have worked with, for the years of guidance and collaborations, for allowing me to use their lab resources to conduct the research that interests me and present at numerous conferences, and for the many fruitful discussions: Drs. Marcia Spetch, Jeremy Caplan, Esther Fujiwara, Anthony Singhal, Tobias Sommer (Universitätsklinikum Hamburg-Eppendorf), Nikolai Malykhin, Craig Chapman, and Alinda Friedman. The members of these labs have also played an enriching role in my Ph.D. experience, both inside and outside of the lab. In particular, I would also like to thank the many collaborators I have had the opportunity and pleasure of working with with: Elliot Ludvig, Eric Legge, Jean-François Nankoo, Yvonne Chen, Michelle Chan, Andrea Shafer, Stan Hrybouski, Matthew Brown, and Jeff Pisklak. Last but not least, I would also like to thank my family for their endless support.

Apart from the people who I have worked with, the research funding I have received has also played a critical role in my ability to conduct research. Throughout my Ph.D. I have directly received funding support from the Canadian Institutes of Health Research (CIHR), the Natural Sciences and Engineering Research Council of Canada (NSERC), the Alberta Gambling Research Institute (AGRI), the German Academic Exchange Service (Deutscher Akademischer Austauschdienst; DAAD), and the University of Alberta.

Sincerely,

Christopher R. Madan

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The overarching goal of the present work is to investigate the influence of reward value on memory and decision making. To introduce this work, several topics must first be discussed: the purpose of memory, the extant literature on the influence of rewards on memory, and the relationship between memory and decision making. Building on this background, my dissertation work aims to advance our understanding of how rewards bias our memory, and how these memory biases can subsequently influence decision making processes.

Memory a is vibrant topic of research within the field of psychology. Memory research revolves around the investigation of the processes underlying the encoding and retrieval of information, and the factors that can influence these processes. Common manipulations include the number of to-be-remembered items, properties of these items (e.g., emotionality, distinctiveness), presentation time and inter-trial interval during study, and the length of the delay between study and test. Memory retrieval can also be tested through a variety of methods and research questions, such as recognition or recall, and tests of item-memory or of order-memory. While all of these manipulations and procedures allow us to further our understanding of the *mechanisms* that underly memory, they do not provide insight into the *function* of memory. In our current environment, the ability to remember previous experiences is critical to our daily lives, ranging from personally relevant details such as where one parked their car to the birthdays of loved ones, as well as to memories for well-known facts, motor skills, and even one's lexicon. But, what is the fundamental function of memory, what is its basic purpose?

1.1 What is memory for?

Although the majority of memory researchers do not discuss the functional role of memory, a handful of researchers have broached the topic (e.g., Bruce, 1985; Glenberg, 1997; Howe & Otgaar, 2013; Klein, in press; Nairne, 2005, 2010; Schacter et al., 2012; Shettleworth, 2009; Sherry & Schacter, 1987). The general consensus of those who take a functionalist approach, is that memory is an evolutionary adaptation to improve biological fitness and ultimately survival. As an example, if an organism finds the location of a food source, 'memory' of this location would be beneficial such that the organism can again return directly to the food source without having to search for the food source anew. An example within an experiment setting is the Morris water, where rats incrementally take more direct paths to the hidden platform on each successive trial (Morris, 1984).

In the work presented here, I did not use naturalistic rewards such as food or erotic stimuli but instead used monetary rewards. This was done not only for practical reasons, such as to avoid issues of food preference, but also for theoretical reasons, as monetary rewards are more readily quantified and controlled. Nonetheless, neuroimaging work has shown that monetary, erotic, and food rewards share a common brain network (Sescousse, Caldú, Segura, & Dreher, 2013).

One of the goals of the present work was to investigate the influence of rewards on memory. Given the biological importance of remembering information related to rewards, it is unsurprising that this question has been asked previously — and before it is possible to ask novel questions in this domain, it is crucial to first appreciate what has already been done.

1.2 Studying the effects of reward value on memory

Recently, a number of research groups have investigated the influence of reward value on memory. For example, Adcock, Thangavel, Whitfield-Gabrieli, Knutson, and Gabrieli (2006) presented participants with images of scenes, preceded by a monetary cue ("$5.00" or "$0.10"; see Figure 1.1a). Participants were instructed that they would later have their memory tested for these scenes, with successful recognition of the images rewarded with the presented amount of money. Recognition hit rates were higher for the high-value scenes than for the low-value scenes. Gruber and Otten (2010) used a similar procedure, with words instead of scenes. Importantly, this can be viewed as a manipulation of prioritization – where participants should attend more to certain items relative to others. This approach was explained more directly by Castel, Benjamin, Craik, and Watkins (2002) and Watkins and Bloom (1999), where the research question was targeted specifically at participants' ability to selectively prioritize the learning of the higher-value items. In these studies, participants were presented with the monetary cue simultaneously with the to-be-remembered item (see Figure 1.1b). To explain the mechanism underlying this enhancement of memory due to prioritization, Castel et al. (2002) suggested that the effect may be mediated by working memory, where participants hold the higher-value items in mind during encoding and then can recall them more easily. Alternatively, or additionally, Castel et al. suggested that participants may devote more attention to the higher-value items and rehearse them more, strategies that have been shown to enhance memory. In later work this procedure was described as "value-directed remembering" (Castel, 2008). Importantly, these prioritization strategies could also have been used in the studies conducted by Adcock et al. (2006) and Gruber and Otten (2010).

While other reward prioritization manipulations have also been used, these other procedures do not differ substantially from the two described approaches. For instance, Shigemune et al. (2010) presented the monetary cue before blocks of to-be-remembered items, as illustrated in Figure 1.1c. Apart from cueing the monetary rewards for blocks of items rather than before each individual item, this procedure is quite similar to that used by Adcock et al. (2006) and Gruber and Otten (2010). In another procedure, Wittmann et al. (2005) rewarded participants for successful recognition, with reward value defined by semantic properties of the item – for example correct recognition of animate objects would net $0.50 in the successive recognition test, while inanimate objects were unrewarded. Note that this procedure was relatively similar to the procedure used by Castel et al. (2002), where reward values were presented simultaneous with the item itself, although in this case reward value was intrinsic to the item.

Importantly, in all of these procedures, rewards can be used to prioritize the learning of certain sets of items over others, without the reward value necessarily becoming a property of the item itself — higher value items need not

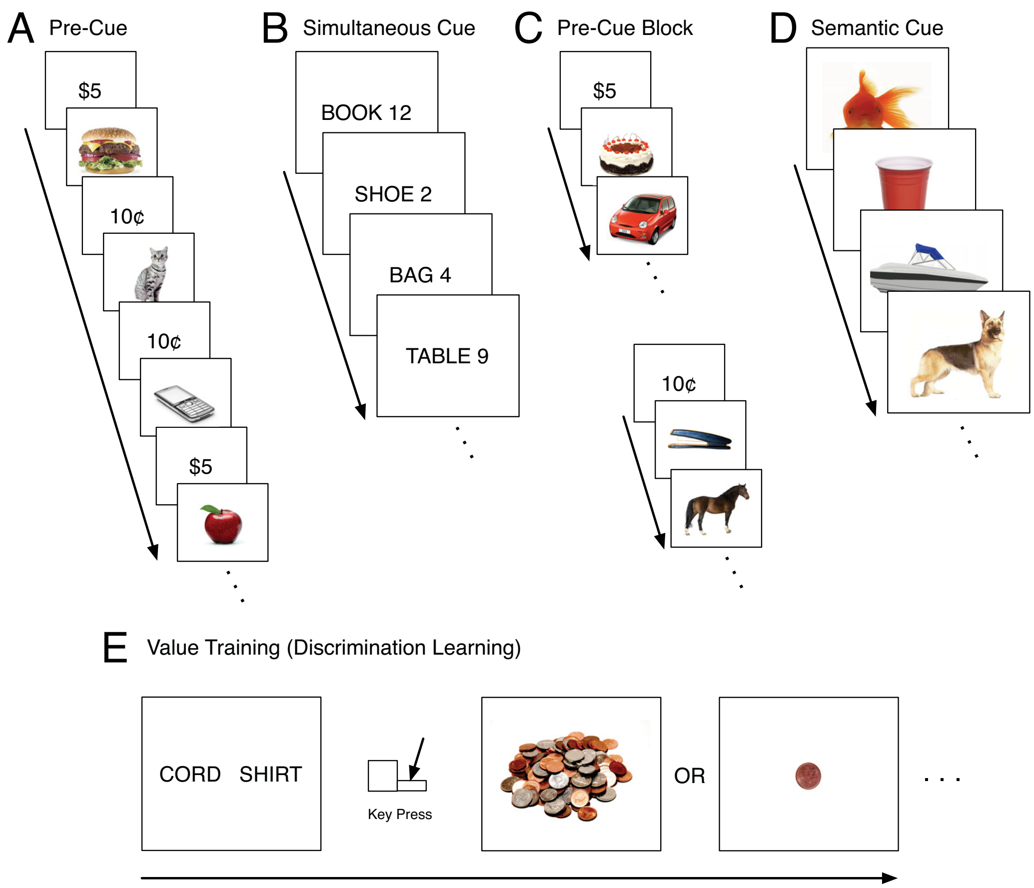

Figure 1.1: Illustration of the reward value procedures used in previous studies. (A) Presenting the monetary cue just prior to the to-be-remembered item (e.g., Adcock et al., 2006; Gruber & Otten, 2010). (B) Presenting the monetary cue simultaenously with the to-be-remembered item (e.g., Castel et al., 2002; Watkins & Bloom, 1999). (C) Presenting the monetary cue before a block of items (e.g., Shigemune et al., 2010). (D) Rewarding items based on their semantic properties (e.g., animate vs. inanimate Wittmann et al., 2005). (E) Having reward values be learned incrementally through trial-and-error (e.g., Chapters 2 and 3: Madan, Fujiwara, et al., 2012; Madan & Spetch, 2012a).

be 'valuable' *per se*, they just need to be explicitly and intentionally attended to or rehearsed to a greater degree.

While many studies have used this prioritization approach to investigating the effects of reward value on memory, there is another approach: Have participants incrementally learn to prefer certain items over others based on the reward outcomes they lead to, as illustrated in Figure 1.1e (e.g., Bayley, Frascino, & Squire, 2005; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006;

Raymond & O'Brien, 2009). This procedure is somewhat analogous to the operant conditioning procedure often used in animal studies. This value-learning procedure could then be followed by a memory test, to investigate if items associated with higher reward values became more memorable, while rewards were learned gradually and incrementally. This is the general approach used in Chapters 2 and 3.

In Chapter 2, a value-learning task was used to train participants that certain words are of higher value than others. Memory was then tested in subsequent memory tasks. Experiment 1 tested for effects of reward value on both implicit and explicit memory (lexical decision and free recall, respectively), despite the memory tasks themselves being unrewarded. Explicit memory was further tested in Experiment 2, using a study-test free recall task to test for effects of reward value on contextual discrimination.

In Chapter 3, a procedure similar to Experiment 1 of Chapter 2 was used. However, instead of using only two levels of reward (high and low), this procedure was extended to multiple levels of reward to determine if effects of reward value on memory are better explained by the reward's value or the reward's salience.

1.3 From memories of the past to decisions in the present

Given the basic result that reward values can influence memory, an open question is how this memory bias would then influence decision making. It is well known that memory effects can translate into decision making biases (e.g., Fredrickson, 2000; Hilbert, 2012; March, 1994). The clearest example of this phenomena is the availability heuristic (Tversky & Kahneman, 1973; also see Schwarz et al., 1991). In its simplest form this heuristic states that information that is more available or more easily retrieved from memory will be weighted more heavily in the decision making process. One typical example of the availability heuristic is as follows: "Consider the letter *R*. In the English language, is *R* more likely to appear in the first position of a word or the third position?"

(Tversky & Kahneman, 1973). Participants are more likely to respond that R occurs more often in the first position, despite this not being the case for all of the letters used in the experiment.

As a more general source of evidence for the availability heuristic, North, Hargreaves, and McKendrick (1997, 1999) conducted a study in a supermarket, where either French or German music was played in the section of the supermarket where the wines were shelved. In these shelves, an equivalent number of French and German wines were positioned — with many relevant factors controlled for (e.g., price, sweetness/dryness, shelf position). The researchers found that people were more likely to buy the bottles of French wine when the French music was playing, and much more likely to buy the German wine when the German music was playing, even though the shoppers were unaware that the ambient music had any influence on their wine selection.

If we continue to think of the availability heuristic more generally, any manipulation that makes certain items more memorable than others should also influence decision making, such that the more memorable items are also weighted more heavily in the decision-making process. One of the most robust and well-known results in the memory literature is the serial position curve (Ebbinghaus, 1885/1913; Hasher, 1973; Murdock, 1962). Briefly, if a list of items is presented to participants and they are subsequently asked to recall all of the items from the list that they can, the participants are more likely to recall the first and and last items in the list, relative to the intermediate items. These two effects are respectively referred to as the primacy and recency effects, as illustrated in Figure 1.2.

This generalized view of the availability heuristic, and of memory effects influencing decision making, is additionally supported by a number of lines of evidence in applied situations. For example, when hearing a series of arguments and counter-arguments, such as in jury persuasion, many studies have shown that, given all arguments are of equal strength, the first and last arguments made will have the greatest influence (Cromwell, 1950; Hovland, Campbell, & Brock, 1957; Lawson, 1968; Lund, 1925; McGuire, 1957; Miller & Campbell, 1959; Rosnow, 1966; Stone, 1969; Weld & Roff, 1938). Similarly,

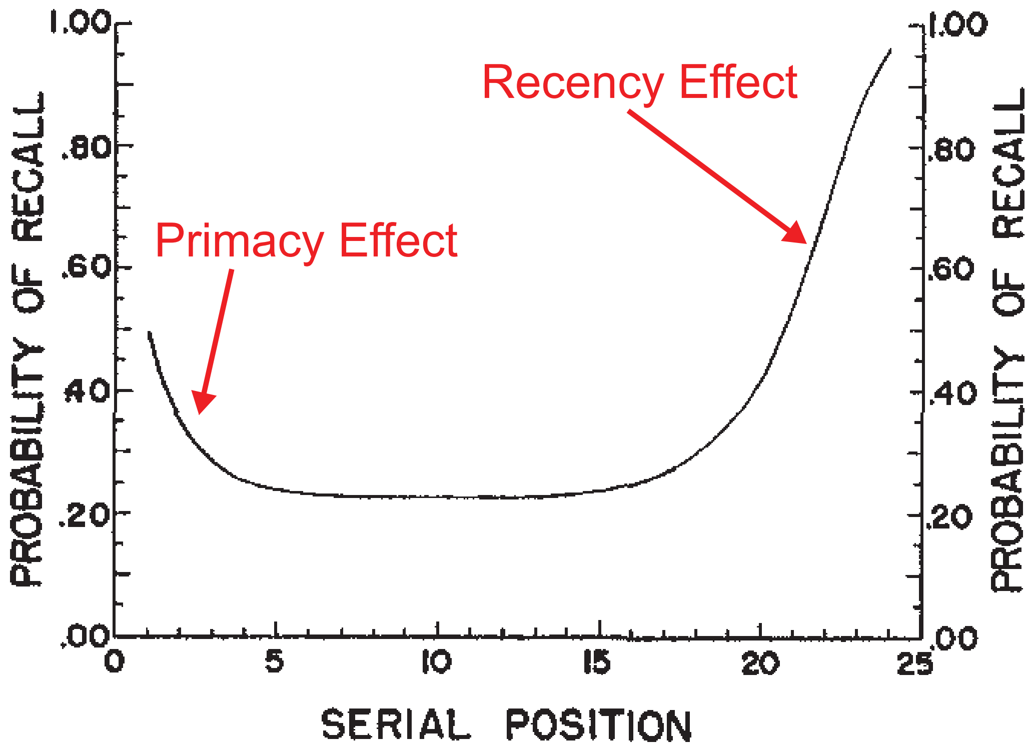

Figure 1.2: Example serial position curve with primacy and recency effects marked. Adapted from Murdock (1962, Figure 3).

when meeting someone for the first time, your opinion of them will be more strongly influenced by your first impressions than subsequent interactions, i.e., a primacy effect (Asch, 1946; Anderson, 1965; Highhouse & Gallo, 1997; Luchins, 1957; Shteingart, Neiman, & Loewenstein, 2013; Shaheen, 2010).

The applications of the availability heuristic extend beyond persuasion and first impressions though — these applications have also appeared in literature on experience-based decision making, but under different names. Consider a participant who is asked to make repeated decisions between an option that leads to a guaranteed outcome and a risky option that may lead to two possible outcomes, one better and one worse than the guaranteed alternative. If the participant chooses the risky option a few times and it happens to lead to the worse outcome, the participant may become risk averse, known as the "hot-stove effect" (Denrell & March, 2001). This effect could be described as being derived from a combination of the primacy effect and the availability heuristic.

The recency effect can also appear in a decision-making task, as a component of the "peak-ends effect". Briefly, the peak-ends effect (Fredrickson, 2000; Kahneman, Fredrickson, Schreiber, & Redelmeier, 1993; Mitchell, Thompson, Peterson, & Cronk, 1997; Redelmeier & Kahneman, 1996; Stone, Schwartz, Broderick, & Shiffman, 2005) suggests that the most salient experience — the 'peak,' and the most recent experience — the 'end,' will be weighted most heavily in the decision making process. Demonstrating this effect, Kahneman et al. (1993) asked participants to put one hand in painfully cold water for two trials. In one trial the hand was in cold water for 60 seconds. In the other trial, the hand was in cold water for 60 seconds (as before), but then the temperature was gradually increased over an additional 30 seconds. After experiencing both trials, participants were asked which trial they would rather repeat; most participants preferred to repeat the longer trial. Even though the long trial was more painful overall, the end of it was less painful, suggesting that a recency effect is a critical feature of this pain-preference result. Further research by Redelmeier and Kahneman (1996) suggests that this memory bias in decision making also generalizes to more applied situations, such as medical procedures.

The influence of the peak level of discomfort (or pleasure; see Mitchell et al., 1997), also suggests that the most salient experiences in memory can also influence decision making, converging with the findings reported in Chapter 3 (Madan & Spetch, 2012a). Given these many instances whereby memory biases can influence decision making, I sought to test for reward value effects on decision making, and how effects of reward value on memory can subsequently influence decision making.

1.4 Risky decision making

In Chapters 4 and 5, I focused on risky decision-making, where participants had to make decisions between an option that leads to a guaranteed, fixed outcome versus a risky option that probabilistically leads to one of two outcomes.

This procedure is also referred to as decision making under uncertainty. In studying risky decision making, it is important to first delineate the two general types of decision-making tasks that are of particular relevance: Decisions from description and decisions from experience. Decisions from description involve known odds. Within a casino setting, this would include games such as roulette and dice-based games such as craps. On the other hand, risky decisions can also be made from gradually acquired experience, where the odds are not readily available, such as in slot-machine games.

Demonstrating the importance of differentiating how information is acquired (i.e., description vs. experience), Ludvig and Spetch (2011) had participants make choices between fixed and risky options in a 2×2 design of reward valence (gains vs. losses) and decision type (description vs. experience). See Figure 1.3 for stimuli similar to that used in Ludvig and Spetch (2011). As an example of decisions from description, consider the following scenario: Given the choice between a guaranteed win of $100 or a 50/50 chance at $200 (or nothing), which would you prefer? While both options clearly have the same expected value, most people choose the guaranteed win. When the same question is re-framed in terms of losses, a guaranteed loss of $100 or a 50/50 chance at losing $200, most people take the risky option (see Figure 1.4). These general tendencies are indicative of an underweighting of more extreme values, e.g., a component of prospect theory (note that prospect theory also suggests an asymmetry in this function, where the function is steeper for losses, leading to loss aversion; Kahneman & Tversky, 1979), compare Figures 1.5a with 1.5b. In other words, in decisions from description, participants were more risk averse for gains than losses. In contrast, participants were more risk-seeking for gains than losses in decisions from experience, as shown in Figure 1.4. This contrast in risk preferences based on decision type is one example of an effect known as the description-experience gap (Hertwig & Erev, 2009).

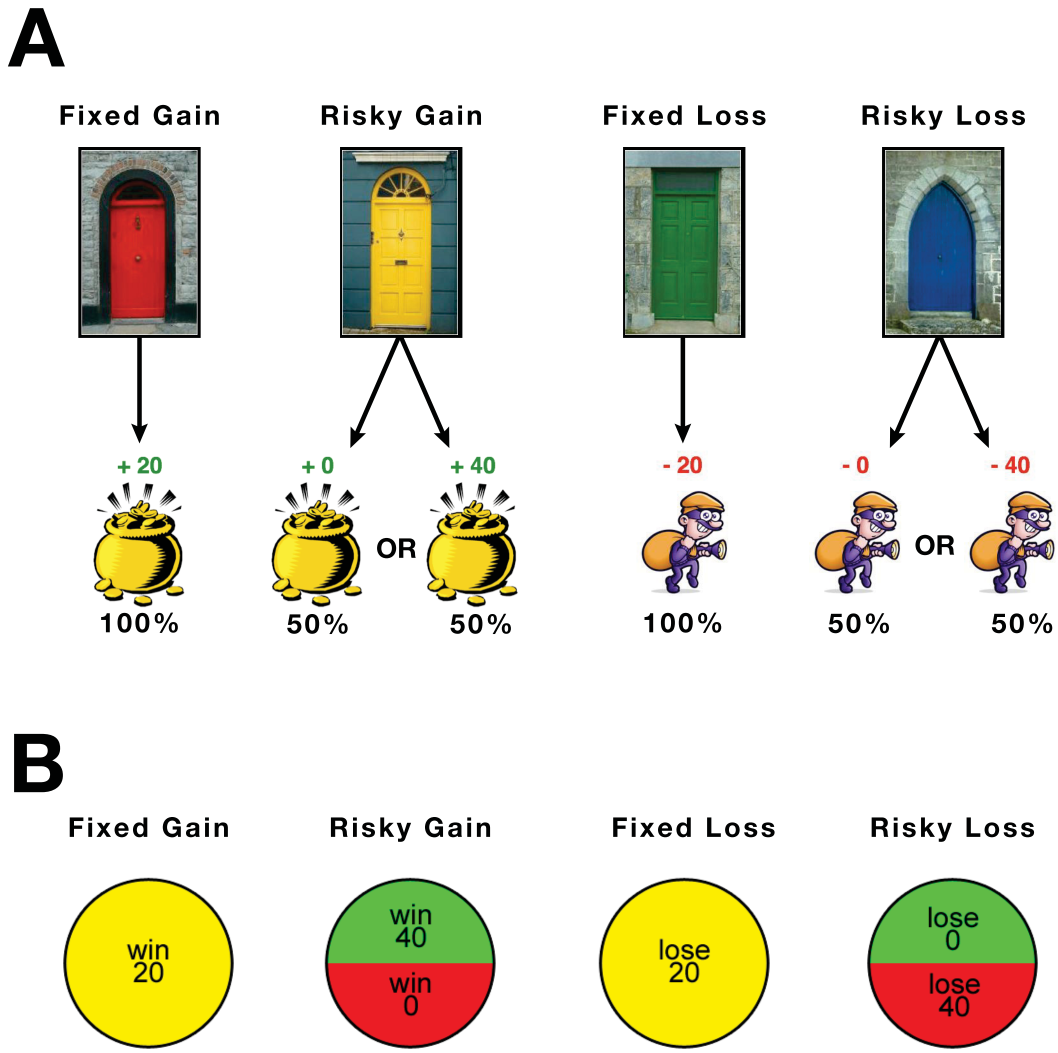

Figure 1.3: Illustration of the stimuli used in tasks investigating decisions from (a) experience and (b) description.

Relative to decisions from description, research into decisions from experience is still in its infancy. To explain the differences between these two types of decisions, Ludvig and Spetch (2011) suggest that extreme values may be overweighted in experience-based decisions, e.g., compare Figures 1.5a and 1.5b with 1.5c, possibly due to a memory-based process. This rationale is further supported by the results presented in Chapter 3 (Madan & Spetch, 2012a). In Chapters 4 and 5 we tested this hypothesis.

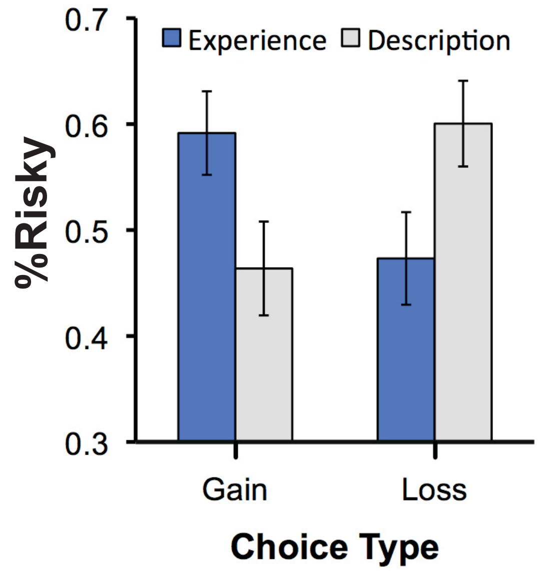

Figure 1.4: Risk preference for gains and losses in decisions from description and experience. Figure adapted from Ludvig and Spetch (2011, Figure 2c).

In Chapter 4, a series of experiments are used to investigate how reward values influence risky decision-making. Specifically, it is hypothesized that extreme values will be overweighted in the decision-making process. A number of supporting hypotheses are also tested, such as the possibility of 'zero' outcomes being underweighted, and the specificity of what determines what counts as an extreme value (i.e., the absolute value or the range of values experienced in the task).

In Chapter 5, the hypothesis that memory biases (such as those observed in Chapter 3) may drive overweighting of extreme values (as observed in Chapter 4) is tested directly. This is done by combining methods from earlier chapters, such that the choice task used in Chapter 4 is now followed by a memory task, where participants were asked to respond with the first outcome that came to mind for each of the door images, as well as were asked to make frequency judgments for all possible door × outcome pairings.

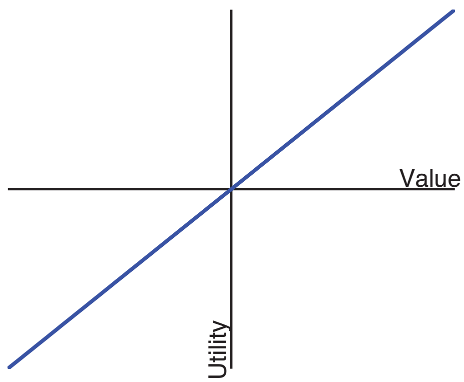

(a) Linear

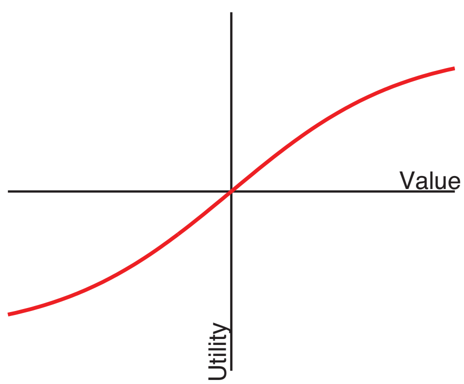

(b) Underweighted extreme values

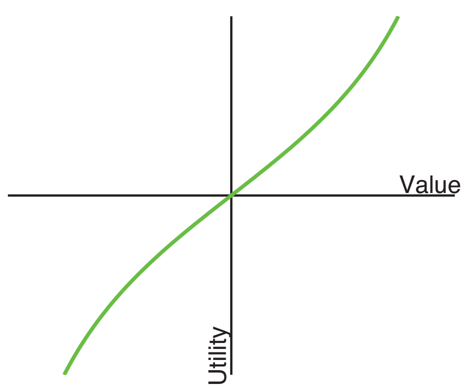

(c) Overweighted extreme values

Figure 1.5: Illustration of utility functions. (a) Linear. (b) Underweighting of extreme values. (c) Overweighting of extreme values.

1.5 Summary

Taken together, the work presented in my dissertation investigates the influence of previously learned reward values on memory, the link between memory and decision making, and how risky decision making may be influenced by these memory biases. Human memory is not a perfect representation of experiences past. By gaining a better understanding of the systematic biases present in memory and how these biases can then affect our ability to make decisions, we can also gain a better understanding of human behaviour as a whole, including the instances when these processes go awry.

1.6 References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517.

Anderson, N. H. (1965). Primacy effects in personality impression formation using a generalized order effect paradigm. *Journal of Personality and Social Psychology*, *2*, 1–9.

Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258–290.

Bayley, P. J., Frascino, J. C., & Squire, L. R. (2005). Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature*, *436*, 550–553.

Bruce, D. (1985). The how and why of ecological memory. *Journal of Experimental Psychology: General*, *114*, 78.

Castel, A. D. (2008). The adaptive and strategic use of memory by older adults: Evaluative processing and value-directed remembering. In A. S. Benjamin & B. H. Ross (Eds.), *The Psychology of Learning and Motivation, Volume 48* (pp. 225–270). London: Academic Press.

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*, 1078–1085.

Cromwell, H. (1950). The relative effect on audience attitude of the first versus the second argumentative speech of a series. *Communications Monographs*, *17*, 105–122.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523–538.

Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology.* New York: Teachers College, Columbia University.

Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion*, *14*, 577–606.

Glenberg, A. M. (1997). What is memory for. *Behavioral and Brain Sciences*, *20*, 1–55.

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

Hasher, L. (1973). Position effects in free recall. *American Journal of Psychology*, *86*, 389–397.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

Highhouse, S., & Gallo, A. (1997). Order effects in personnel decision making. *Human Performance*, *10*, 31–46.

Hilbert, M. (2012). Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making. *Psychological Bulletin*, *138*, 211–237.

Hovland, C. I., Campbell, E. H., & Brock, T. (1957). The effects of "commitment" on opinion change following communication. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 23–32). New Haven: Yale University Press.

Howe, M. L., & Otgaar, H. (2013). Proximate mechanisms and the development of adaptive memory. *Current Directions in Psychological Science*, *22*(1), 16–22.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*, 401–405.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Klein, S. B. (in press). The temporal orientation of memory: It's time for a change of direction. *Journal of Applied Research in Memory and Cognition*.

Lawson, R. G. (1968). Order of presentation as a factor in jury persuasion. *Kentucky Law Journal*, *56*, 523–555.

Luchins, A. S. (1957). Primacy-recency in impression formation. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 33–61). New Haven: Yale University Press.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

Lund, F. H. (1925). The psychology of belief: A study of its emotional, and volitional determinants. *Journal of Abnormal and Social Psychology*, *20*, 63–81 and 183–191.

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.

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

March, J. G. (1994). *Primer on decision making: How decisions happen*. New York: The Free Press.

McGuire, W. J. (1957). Order of presentation as a factor in "conditioning" persuasiveness. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 98–114). New Haven: Yale University Press.

Miller, N., & Campbell, D. T. (1959). Recency and primacy in persuasion as a function of the timing of speeches and measurements. *Journal of Abnormal and Social Psychology*, *59*, 1–9.

Mitchell, T. R., Thompson, L., Peterson, E., & Cronk, R. (1997). Temporal adjustments in the evaluation of events: The "rosy view". *Journal of Experimental Social Psychology*, *33*, 421–448.

Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods*, *11*, 47–60.

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482–488.

Nairne, J. S. (2005). The functionalist agenda in memory research. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer.* Washington, DC: American Psychological Association.

Nairne, J. S. (2010). Adaptive memory: Evolutionary constraints on remembering. *Psychology of Learning and Motivation*, *53*, 1–32.

North, A. C., Hargreaves, D. J., & McKendrick, J. (1997). In-store music affects product choice. *Nature*, *390*, 132.

North, A. C., Hargreaves, D. J., & McKendrick, J. (1999). The influence of in-store music on wine selections. *Journal of Applied psychology*, *84*, 271–276.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042–1045.

Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation. *Psychological Science*, *20*, 981–988.

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*, 3–8.

Rosnow, R. L. (1966). Whatever happened to the "law of primacy"? *Journal of Communication*, *16*, 10–31.

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron*, *76*(4), 677–694.

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, *61*, 195–202.

Sescousse, G., Caldú, X., Segura, B., & Dreher, J.-C. (2013). Processing of primary and secondary rewards: A quantitative meta-analysis and review of human functional neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*, 681–696.

Shaheen, J. (2010). The recency and primacy effects in the talent acquisition process. *Journal of Corporate Recruiting Leadership*, *5*, 13–15.

Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological review*, *94*(4), 439.

Shettleworth, S. J. (2009). *Cognition, evolution and behavior* (2nd ed.). New York: Oxford University Press.

Shigemune, Y., Abe, N., Suzuki, M., Ueno, A., Mori, E., Tashiro, M., ... Fujii, T. (2010). Effects of emotion and reward motivation on neural correlates of episodic memory encoding: A PET study. *Neuroscience Research*, *67*, 72–79.

Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology: General*, *142*, 476–488.

Stone, A. A., Schwartz, J. E., Broderick, J. E., & Shiffman, S. S. (2005). Variability of momentary pain predicts recall of weekly pain: A consequence of the peak (or salience) memory heuristic. *Personality and Social Psychology Bulletin*, *31*, 1340–1346.

Stone, V. A. (1969). A Primacy Effect in DecisionMaking by Jurors. *Journal of Communication*, *19*, 239–247.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*, 207–232.

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

Weld, H. P., & Roff, M. (1938). A study in the formation of opinion based upon legal evidence. *American Journal of Psychology*, *51*, 609–628.

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*, 459–467.

Chapter 2

High reward makes items easier to remember, but harder to bind to a new temporal context

A version of this work was previously published as: **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. doi:10.3389/fnint.2012.00061.

2.1 Abstract

Learning through reward is central to adaptive behaviour. Indeed, items are remembered better if they are experienced while participants expect a reward, and people can deliberately prioritize memory for high- over low-valued items. Do memory advantages for high-valued items only emerge after deliberate prioritization in encoding? Or, do reward-based memory enhancements also apply to unrewarded memory tests and to implicit memory? First, we tested for a high-value memory advantage in unrewarded implicit and explicit tests (Exp 1). Participants first learned high or low reward values of 36 words, followed by unrewarded lexical decision and free recall tests. High-value words were judged faster in lexical decision, and more often recalled in free recall. These two memory advantages for high-value words were negatively correlated suggesting at least two mechanisms by which reward value can influence later item-memorability. The ease at which the values were originally acquired explained the negative correlation: People who learned values earlier showed reward effects in implicit memory while people who learned values later showed reward effects in explicit memory. We then asked whether a high-value advantage would persist if trained items were linked to a new context (Exps 2a,b). Following the same value training as in Exp 1, participants learned lists composed of previously trained words mixed with new words, each followed by free recall. Thus, participants had to retrieve words only from the most recent list, irrespective of their values. High- and low-value words were recalled equally, but low-value words were recalled earlier than high-value words and high-value words were more often intruded (proactive interference). Thus, the high-value advantage holds for implicit and explicit memory, but comes with a side effect: High-value items are more difficult to relearn in a new context. Similar to emotional arousal, reward value can both enhance and impair memory.

2.2 Introduction

When faced with items of differing reward values, an individual has the possibility of prioritizing their efforts to learn as much as possible about the higher-valued items, likely at the expense of knowledge about the lower-value items. If people took advantage of this, they could maximize their accumulation of reward. In seeking reward, it may not only be beneficial to remember the values of items, but also related information such as the precise context in which the item was found, which we refer to as the reward-maximization hypothesis. Alternatively, reward value may be emotionally arousing; thus, effects of reward value on memory may resemble those found with emotional arousal. Emotionally arousing items are generally remembered better, but memory for related contextual information is often impaired (Burke, Heuer, & Reisberg, 1992; Christianson, 1992; Easterbrook, 1959; Madan, Caplan, Lau, & Fujiwara, 2012; Mather & Sutherland, 2011). Such impairment may be caused by diverting attention toward the arousing stimulus itself, and away from its context. If reward value functions like emotional arousal, then higher reward value should result in enhanced performance on some tests of memory (e.g., memory for the experienced items alone), but not others (e.g., judging whether an item was presented in a specific context), which we refer to here as the value-interference hypothesis. Whether higher reward value universally results in better item-memory across different types of memory tests (explicit and implicit), and whether reward value results in better memory for context is unknown. Finding a benefit for high-value items in rewarded memory tests tells us that participants are capable of prioritizing high-value items, but leaves open the question of whether participants favor high-value items when the procedure does not dictate that they should do so. Thus, our first objective was to test whether a memory advantage for words that were previously trained to have a high (versus a low) reward value persists in later *unrewarded* implicit and explicit memory tests (Experiment 1), to test for the generality of reward-value enhancements. Our second objective was to test whether an item-memory advantage for high-value words generalizes if the trained words have to be studied and memorized in a new context (Experiments 2a and 2b).

Rewarded memory tests in numerous studies have shown that people are able to prioritize learning of high-value over low-value items, both words and images (Adcock et al., 2006; Bjork & Woodward, Jr., 1973; Castel et al., 2002; Castel, Farb, & Craik, 2007; Castel, Balota, & McCabe, 2009; Eysenck & Eysenck, 1982; Gruber & Otten, 2010; Harley, 1965; Kuhl, Shah, DuBrow, & Wagner, 2010; Loftus & Wickens, 1970; Soderstrom & McCabe, 2011; Shohamy & Adcock, 2010; Tarpy & Glucksberg, 1966; Watkins & Bloom, 1999; Weiner, 1966; Weiner & Walker, 1966; Wolosin, Zeithamova, & Preston, 2012). For example, Castel et al. (2002) showed participants words along with numerical reward values ranging from 1 to 12. Participants were instructed to remember the words with the highest values as best as possible, to maximize the total value of their recalled words. High-value words were recalled more than low-value words. This suggests people were able to flexibly adjust the allocation of cognitive resources during learning to favor items with higher value over those with lower value, and thus maximize earned reward. Assuming a limited resource model, the authors also suggested that if a particular item is allocated more resources, it will be remembered better, but at the expense of the other studied items.

Prioritization effects are not limited to recall; Adcock et al. (2006) demonstrated an enhancement of memory due to reward value using a different explicit-memory test: recognition. They presented participants with a high- or low-value reward cue ("$5.00" or "$0.10") followed by a scene image. Participants were asked to remember the scenes (presented during reward anticipation) and were told that they would earn the respective reward amount if they successfully recognized the images in a memory task the following day. In the recognition test, participants earned the respective reward for recognition hits, and were penalized for false alarms. Hit rates were higher for high- than low-value items. Again, this result demonstrates people's ability to explicitly prioritize items associated with a higher-value reward over those with a lower-value reward, both during encoding and retrieval.

Such enhancements of memory due to reward value have been found with tests of explicit memory. However, reward value could influence implicit memory in equally powerful ways. That is, reward value might modulate behaviour even when the participant is not deliberately trying to retrieve item values. This would extend the prioritization findings beyond a deliberate encoding/retrieval strategy, and would suggest that in addition, participants may have a cognitive bias toward high-value items. Although it has never been tested directly, some findings are consistent with the hypothesis that higher reward-values lead to better implicit memory: Rewards that are presented subliminally can influence behaviour (reviewed in Custers & Aarts, 2010). For example, participants respond more quickly ($\sim$ 20 ms) in simple monetary incentive tasks when the trial is preceded by a high-value reward cue, than if it is preceded by a low-value reward cue (e.g., the participant is presented with the reward cue, and told to press a button once a target appears; Abler, Walter, & Erk, 2005; Sescousse, Redouté, & Dreher, 2010; Staudinger, Erk, & Walter, 2011). Furthermore, Pessiglione et al. (2007) presented participants with coin images of either 1-pound or 1-pence and asked them to squeeze a handgrip to earn the corresponding monetary reward. Coin images were presented either subliminally (for 17 or 50 ms) or supraliminally (100 ms). Participants squeezed the grip harder on the higher-value trials, even when the coin image was not consciously perceived. Hence, consciously and unconsciously processed reward cues can have analogous effects. Subliminally presented higher-value rewards also recruited more attention than lower-value rewards (pupil dilation: Bijleveld, Custers, & Aarts, 2009) and increased accuracy in arithmetic (Bijleveld, Custers, & Aarts, 2010). Though none of these studies have directly shown that reward value can enhance implicit memory, they provide at least indirect support for the hypothesis that high-value items could enhance implicit memory.

We also wanted to clearly separate the value-learning phase, which should be rewarded by necessity, from the later memory phase, which should be unrewarded. Our reasoning was as follows: To interpret the prioritization effects, one must consider that participants were instructed to prioritize. The positive prioritization results, therefore, tell us that participants are capable of prioritization. We ask here whether participants have a bias toward better memory for higher-value stimuli in an unrewarded memory test, even when there is no immediate need to favor the encoding of high-value stimuli. By clearly separating the value-learning phase from the memory study phase (Experiments 2a and 2b) and test phase (all experiments here), we can test whether people possess a learning bias universally favouring high over low-reward value items or reward-value might interfere with new learning.

Raymond and O'Brien (2009) conducted an experiment along these lines, testing for the non-deliberate effects of reward value on memory (see also Wittmann et al., 2005; Wittmann, Dolan, & Düzel, 2011), but it is difficult to determine whether their results were driven by implicit or explicit memory retrieval. In their value-learning task, stimulus values were learned with repeated experience, and the effects of the learned values on memory were later tested with an unrewarded, modified attentional blink (AB) task. Participants were first presented with pairs of faces and asked to choose one. Faces within-pair differed in their probability of reward (0.20 or 0.80; reward value across pairs was positive, negative or neutral). Unlike a conventional AB task, Raymond and O'Brien (2009) asked participants not simply to respond when they saw the target image, but instead to indicate whether the target image was an old face from the prior value-learning task, or a new face (i.e., old/new recognition). If a target image were to overcome the AB, it may also be better retrieved in explicit recognition-memory. Higher-value faces were indeed more often recognized as old than lower-value faces, even though, critically, performance in this task was unrewarded. Raymond and O'Brien (2009) concluded that more attentional resources are recruited for stimuli that previously acquired a higher value. Their results also demonstrate a prioritization from a value-learning task where target items are encoded incidentally. However, we suggest that the following interpretations are possible: (a) High-value faces were primed more during value-learning, leading to enhanced implicit memory for higher-value faces during the AB task. Greater priming for the higher-value faces may have led to increases in subjective experiences of familiarity in the recognition memory test in the AB task. (b) Old/new recognition is a test of explicit memory. Participants may have recognized the high-value faces in the AB task due to episodic recollection (i.e., explicit memory). (c) Recognition in the AB task may have resulted from a combination of implicit and explicit memory. Thus, while Raymond and O'Brien's results provide evidence of a reward-based enhancement of recognition-memory, it is unclear whether this was an enhancement of implicit or explicit memory or a mixture both.

In the current study, we first asked if previously learned reward values also enhance item accessibility in an implicit test of memory: lexical decision (Experiment 1). Participants were first presented with words in a two-alternative choice value-learning task, in which they learned, by trial-and-error with feedback, that half of the words led to a high-value reward and half of the words led to a low-value reward (also used by Madan & Spetch, 2012a). This value-learning task is notably similar to previous reward-learning procedures used by Estes and others (e.g., Allen & Estes, 1972; Estes, 1962, 1966, 1972; Humphreys, Allen, & Estes, 1968; Medin, 1972a, 1972b; Pubols, 1960) as well several more recent reward-learning studies (e.g., Bayley et al., 2005; Frank, Seeberger, & O'Reilly, 2004; Frank, O'Reilly, & Curran, 2006; Gradin et al., 2011; Johnsrude, Owen, Zhao, & White, 1999; Johnsrude, Owen, White, Zhao, & Bohbot, 2000; Pessiglione et al., 2006; Valentin & O'Doherty, 2009; Voon et al., 2010). Participants were then presented with an unrewarded lexical decision task, in which words from the value-learning task were shown again. Finally, in an unrewarded test, participants were asked to freely recall all the words from the session (value-learning phase and lexical decision). We predicted that explicit memory (free recall) would be enhanced by reward value. We further predicted that implicit memory would be enhanced due to reward value, as measured in the lexical decision task, if reward value enhances memory retrieval even when participants do not deliberately prioritize the retrieval of high-value items over low-value items. If memory is enhanced in both memory tests, we will then ask whether the two effects could have the same underlying cause or not. This will be done by correlating the high-value advantage in lexical decision with the high-value advantage in free recall across participants. If the correlation is large and positive, this would suggest that memory, both implicit and explicit, can be enhanced by reward value through a singular mechanism that globally enhances memory performance. However, implicit and explicit memory functions are supported by separable memory systems, both in behavior (e.g., Gopie, Craik, & Hasher, 2011; May, Hasher, & Foong, 2005) and in the brain (e.g., Rugg et al., 1998; Schott et al., 2005, 2006). If we instead find that performance in the two memory tasks is uncorrelated or even produce a negative correlation, this would suggest that enhancements of memory due to value are driven by separable modulations of different kinds of memory by reward value.

In a second pair of experiments, we asked if the enhancement of explicit memory due to reward value would persist if items with previously learned reward values were re-studied in a new context. Participants in Experiments 2a and 2b were first given the same value-learning task as in Experiment 1. Following this, participants were asked to study several lists composed of previously learned high- and low-value words, as well as new items, in an unrewarded free recall task. In this free recall task, participants had to disregard their memory for items from the value-learning task and instead, confine their memory retrieval to only the most recently studied list (a specific, temporally defined context). Experiments 2a and 2b were identical except that a faster presentation rate was used in Experiment 2b to test whether the results of Experiment 2a could be due to time-consuming processes applied during study, such as deliberate encoding of reward value. Because the list length was short (nine words per list), we expected that total probability of recall might not be a sensitive enough measure; we therefore additionally examined output order and intrusion rates to test whether high- or low-value items were remembered better.

According to the reward-maximization hypothesis, participants devote more resources to learning higher-value items than lower-value items. This should generalize to learning in a new context (determining whether an item was presented within a specific context), which leads to the prediction that free recall will be enhanced for high-value words in Experiments 2a and 2b. According to the value-interference hypothesis, cognitive resources may be diverted to high-value items, and this is at the expense of attention to other related information, including the list context. Thus, the value-interference hypothesis leads to the prediction that free recall will be worse for high-value items, and that high-value items will be intruded more than low-value items (due to failures of list discrimination).

2.3 Experiment 1

2.3.1 Methods

Participants

A total of 99 introductory psychology students at the University of Alberta participated for partial fulfillment of course credit. Five participants were excluded due to machine error. All participants had learned English before the age of six 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

Words were selected from the MRC Psycholinguistic database (Wilson, 1988). Imageability and word frequency were all held at mid-levels and all words had six to seven letters and exactly two syllables. We additionally used the Affective Norms for English Words (ANEW; Bradley & Lang, 1999) to exclude words with moderately arousing, positive or negative emotional connotations[1] (e.g., 'assault', 'hatred', 'heaven') which could interfere with learning reward values (e.g., participants may find it difficult to learn that 'hatred' is a high-value word, or that 'heaven' is a low-value word). Two words were removed manually as they were deemed by the authors to be emotional in nature, but were not included in ANEW (e.g., 'terror', 'regret'). A total of 21 words were

[1]Our criteria regarding the ANEW were to exclude words with an arousal rating greater than 5.5 (scored on a scale from 1 [not arousing] to 9 [highly arousing), and a valence rating (also on a scale from 1-9) of either (a) less than 4 [negative], or (b) greater than 7 [positive]. Note that we chose to keep two words that did meet the exclusion criteria: 'dancer' and 'rescue'.

	Concreteness	Imageability	Frequency	Length	Syllables
Mean	439	467	22	6.46	2
SD	99	80	12	0.50	0
Min	243	248	7	6	2
Max	580	578	52	7	2

Table 2.1: Word pool statistics, as obtained from the MRC Psycholinguistic database (Wilson, 1988).

excluded this way, and the final word pool consisted of 160 words (Table 2.1 reports word pool properties).

For the lexical decision phase, 160 pronounceable non–words were generated with the LINGUA non–word generator (Westbury, Hollis, & Shaoul, 2007), using a pre-compiled word frequency dictionary (Shaoul & Westbury, 2006). To match the length of the non–words to the words, we generated 87 six-letter and 73 seven-letter non–words.

Procedure

Prior to the experiment, participants were informed that the experiment was a 'word choice task,' and that they would receive a payment proportional to the total points earned in the value-learning task of the experiment, in addition to their partial course credit.

The experiment consisted of a sequence of four sequential tasks: value-learning, lexical decision, free recall, and a value judgement task. Participants were not provided with details of the subsequent task until the current task was completed.

Value learning. Participants were shown two words on the computer screen simultaneously. Words were selected at random from our word pool of 160 words. Participants were to choose one of the two words in each choice set by pressing the 'Z' or '/' key of a computer keyboard to choose the word presented on the left or right side of the computer screen, respectively.

For each participant, 36 words were randomly selected from the word pool, and each word was randomly assigned to one of two reward values: 1 or 10 points (low- or high-value, respectively). Trial choices were pseudorandomly generated, with each word used one time per choice set, but each set always consisted of one high- and one low-value word. This constraint was not revealed to the participant. After each choice, the participant saw the reward in the centre of the screen for 2000 ms; if they chose a high-value word, an image of a pile of coins was presented; if they chose a low-value word, an image of a penny was presented. The participant's current point balance was continually presented at the bottom of the screen throughout the duration of the value-learning task. There was no time limit on the choices and participants were given a 1000-ms delay before the next choice.

Training consisted of 18 choice sets per block for 13 blocks. At the end of the session, participants were paid $1.00 for every 500 points earned during the value-learning task, rounded up to the nearest 25-cent amount. Participants earned between $3.25 and $5.00 in this task.

Lexical decision. An additional 18 words, selected at random from the same pool as the trained words, were included as new words. Participants were asked to judge the lexical status of 108 items: 36 trained words, 18 new words and 54 non–words. Each item was presented for up to 10,000 ms, and the participant pressed either 'Z' on the computer keyboard to indicate that the item was a proper English word, or '/' to indicate that the item was a not a word. A fixation cross ('+') was presented for 1000 ms to separate each decision prompt.

The 108 items were preceded by 8 practice items (four words/four non–words) to attenuate a possible recency effect over the last words from the preceding value-learning task.

Free recall. In a final free recall task, participants were given 5 minutes to recall all of the words they could remember from the study, in any order. Participants were asked to type out their responses, terminated with the `ENTER` key. After each response, a blank screen was presented for 500 ms. Repeated recalls of the same words were ignored.

Value judgement. To measure participants' explicit memory of the reward values for each item, we included a value judgement task following free recall. At the end of the experiment, participants were presented with each of the words previously shown in the value-learning task, one at a time, and asked to judge how many points each word had been worth in the initial value-learning task. Participants were told to press the 'Z' key if they thought the word was worth 1 point, or '/' for 10 points.

Data Analysis

Effects were considered significant based on an alpha level of 0.05. For response time analyses, only correct responses were analyzed. Response time analyses were conducted on the within-subject mean accuracy and median response time for each condition.

For lexical decision, only responses made between 200 ms and the individual participant's mean plus three standard deviations were included in the analysis (0.61% trials excluded).

For free recall, responses were computer-scored and spelling errors were not corrected.

2.3.2 Results and Discussion

Accuracy in the value-learning task was measured as the proportion of trials on which the participant chose the high-value word. This measure began at chance, as the participant could not know which was the high-value word. In the last block of the value-learning task, accuracy was significantly greater than chance and near ceiling [$M \pm 95\%\ CI = 0.94 \pm 0.02$ correct; $t(93) = 37.34$, $p < .001$] (Figure 2.1a).

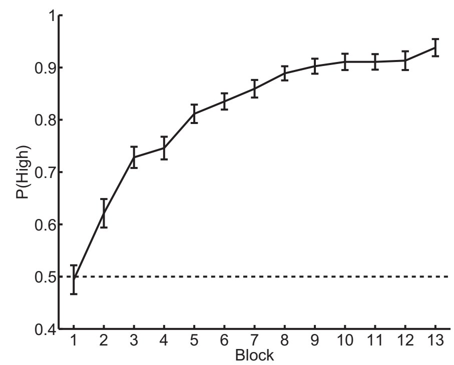

(a) Experiment 1

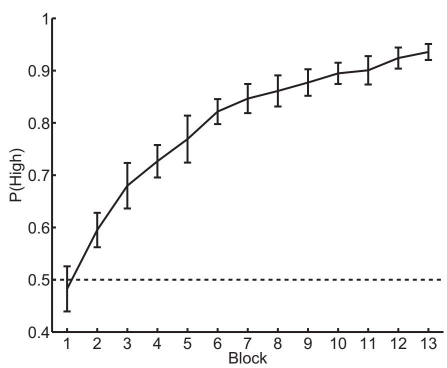

(b) Experiment 2a

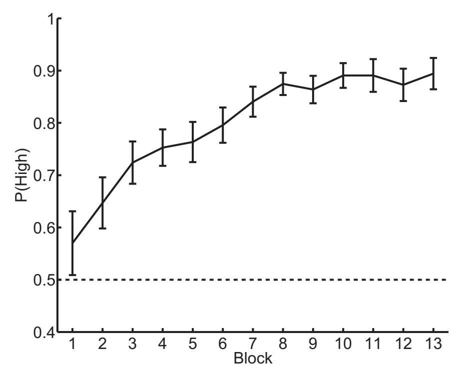

(c) Experiment 2b

Figure 2.1: Value-learning task results in experiments 1 (A), 2a (B), and 2b (C). Performance is shown as the probability of choosing the high-value word over the low-value word in each of the learning blocks 1 to 13. Chance probability of choosing the high-value word is indicated by the dashed line. Error bars are 95% confidence intervals around the mean, corrected for inter-individual differences (Loftus and Masson, 1994).

Lexical decision was significantly more accurate for the previously rewarded old words than for the untrained, new words [$t(93) = 6.94$, $p < .001$; old words: 0.99 ± 0.03 correct; new words: 0.95 ± 0.01 correct]. Participants also identified the old words significantly faster than the new words [$t(93) = 12.77$, $p < .001$, M(new) $= 708 \pm 12$ ms; M(old) $= 591 \pm 8$ ms]. There was no difference between accuracy for high- and low-value words [$t(93) = 0.00$, $p > .1$; high-value: $.99 \pm 0.04$ correct; low-value: 0.99 ± 0.04 correct]. Importantly, high-value words were identified significantly faster than low-value words [$t(93) = 2.42$, $p < .05$; M(high) $= 584 \pm 9$ ms; M(low) $= 599 \pm 8$ ms] (Figure 2.2a). Thus, trained words were primed, and high-value words were primed more than low-value words, a novel finding that suggests that reward value can influence implicit memory.

Probability of free recall (Figure 2.2b) was greater for high-value words than low-value words [$t(93) = 4.40$, $p < .001$; M(high) $= .50 \pm .04$; M(low) $= .40 \pm .03$]. 'New' words (from the lexical decision task) were also recalled, but far less often than the previously rewarded words [$t(93) = 23.80$, $p < .001$; M(new) $= .06 \pm .01$; M(old) $= .45 \pm .03$]. Thus, value also influenced explicit memory retrieval, replicating prior findings.

We next asked if the memory effects of value depended on the level of performance achieved during value training. However, the asymptotic accuracy in the value-learning task (averaged over the last four trials) did not correlate significantly with the value effects on either memory test [both r's $< .15$, p's $> .1$].

In the value judgement task, participants rated the value of the previously rewarded words much better than chance [M $= 0.87 \pm 0.03$ correct; $t(93) = 25.65$, $p < .001$]. The accuracy of judgements was similar for high-value words [M(high) $= 0.88 \pm 0.03$ correct] and low-value words [M(low) $= 0.87 \pm 0.03$ correct; $t(93) = 1.16$, $p > .1$]. That is, participants had substantial, though not perfect, explicit memory for the value of both high and low-value words. In the value-learning task, because all responses were a choice between a high- and a low-value item, those responses cannot be used to determine whether high and low values were learned to the same level. In the value-judgment task, items

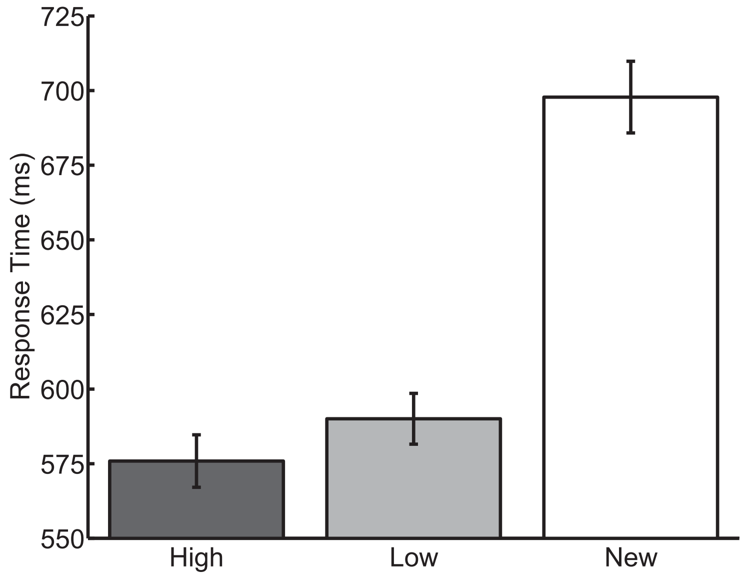

(a) Lexical Decision

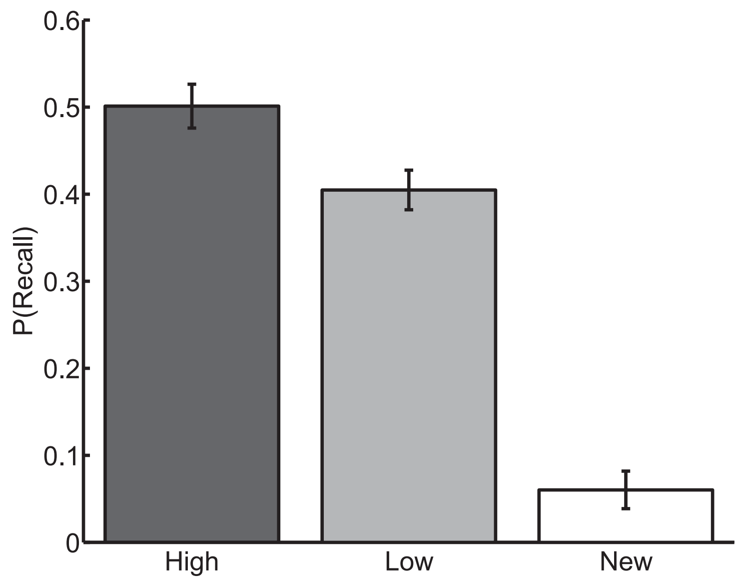

(b) Free Recall

Figure 2.2: Performance in the memory tasks in Experiment 1. (a) Response times from the lexical decision task. (b) Proportion of total words recalled from the free recall task. 'High' and 'Low' represent the high- and low-value words, respectively. 'New' represents words first used in the lexical decision task, that were not present in the value-learning task. Error bars are 95% confidence intervals, corrected for inter-individual differences (Loftus & Masson, 1994).

were judged individually; thus, the near-equivalence of value-judgements of high- and low-value items suggests that participants learned the values of high- and low-value words equally well. This rules out the possibility that participants simply remembered the high-value words better because they performed the value-learning task better for high- than low-value items. It could further be argued that the value judgements for both types of items could have been based on memory for high-value items alone: A participant then would decide to judge a high-value item as 'high' based on their memory for that item's value, but to judge all items for which they had no such memory as a 'low' item. That is, value-judgements would be made on a single value dimension. If only the strength of memory for high-value items was used to make judgements along this dimension, high-value words could be correctly classified as high (those with sufficient high-value item memory strength), low-value words could be correctly classified as low (those with insufficient high-value memory strength), and high-value words could be incorrectly classified as low (those with insufficient high-value memory strength). However, low-value words could not be incorrectly classified as high-value words this way. As reported, we did observe such errors in 13.2% of the low-value words. Note also that the probability of judging a low item as high was quite close to the probability of judging a high item as low (12.5%). Thus, regardless whether participants are basing their choices on a singular value dimension, they are doing so with the same accuracy for low as for high items. This suggests that the quality of memory (i.e., variance in memory strength for both word types along a value dimension) is equivalent for both types of words.

One plausible explanation of our results is that, instead of value, our effects on memory are due to choice behaviour: the more often a participant chose an item during value-learning, the more they remembered that item in the later memory tests (see Weber & Johnson, 2006). Since choice frequency and value are highly confounded (i.e., the task *requires* choosing high over low value items), a combined correlation spanning all items would not be possible either. As an alternative, we calculated choice frequency as the mean number of times a participant chose a high-value item, across all 13 blocks of the value-learning task, minus the mean number of times they chose a low-value item: DiffCF= $mean[choice\ frequency(H)] - mean[choice\ frequency(L)]$. DiffCF thus measures a participant's bias toward choosing high- over low-value words. DiffCF is, of course, expected to be highly correlated with accuracy, since participants are indeed asked to choose high items and to avoid low items. Confirming this, the correlation between participants' accuracy in the value training task and the DiffCF measures was highly significant [$\rho(93) = .48$, $p < .001$]). To rule out that choice frequency significantly co-varied with our effects of value on implicit and explicit memory, we then correlated DiffCF with: (a) the effect of value on lexical decision performance (the normalized difference in response times due to reward value: DiffLD= $[RT(low) - RT(high)] \div .5[RT(low) + RT(high)]$; (b) the effect of value on free recall performance, DiffFR = proportion of recalled high-value words, divided by the total number of words recalled. DiffCF correlated with neither the effect of reward value on implicit memory, nor the effect of reward value on explicit memory [lexical decision: $\rho(93) = .075$, $p > .1$; free recall: $\rho(93) = -.074$, $p > .1$]. Thus, the bias to choose high over low-value words in the value-learning task did not account for the effects of reward value on implicit- or explicit-memory. This is consistent with Madan and Spetch (2012a), who also ruled out choice frequency as a possible determining factor of subsequent memory with a similar training procedure.

Next, we asked whether the effects of reward value on our two memory tests were related, explaining common variance across participants, or unrelated, explaining different subject variability. Participants who demonstrated greater value-based facilitation in lexical decision had *less* value-based facilitation in free recall [$\rho(93) = -.20$, $p < .05$] (Figure 2.3). The fact that a positive correlation was not observed suggests that there are at least two partly dissociable mechanisms by which value can influence memory.

Because lexical decision always preceded final free recall, we were concerned that the negative correlation between the two value-based facilitation effects on memory could be due to the influence of lexical decision on free recall. If a word had a long response time in lexical decision, perhaps that would

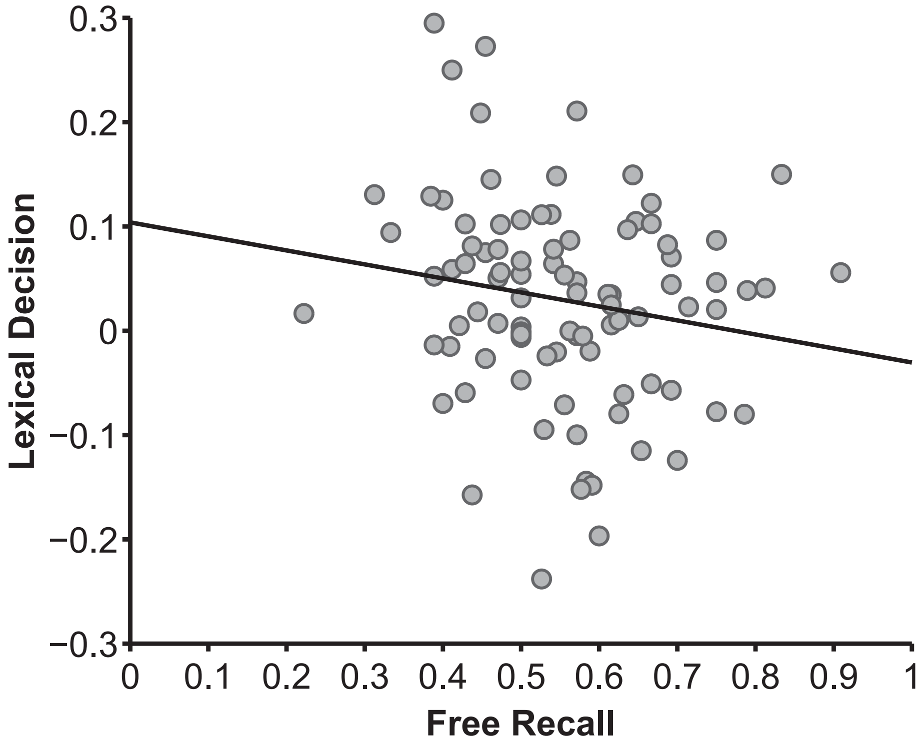

Figure 2.3: Correlation between lexical decision and free recall tasks in Experiment 1 [$\rho(93) = -.20$, $p < .05$]. The lexical decision measure was the facilitation of high-value words compared to low-value words (difference in response time) divided by the participants' average response time. The free recall measure was the proportion of recalled words that were high-value, divided by the total number of words recalled from the value-learning task. Each dot represents an individual participant.

correspond to increased encoding of the word; a poor lexical decision response might, then, turn into an increased probability of free recall. We tested for this kind of effect with within-subjects analyses. We compared lexical decision response times for words that were or were not free recalled, separately for high- and low-value words. Lexical decision response times were not significantly different between later recalled and later not recalled words [high-value: $t(93) = 1.13$, $p > .1$, Cohen's $d = .07$; low-value: $t(93) = 0.60$, $p > .1$, $d = .04$]. Thus, the effect of an item's value on explicit memory is not explainable by its effect on implicit memory, or vice versa, and this rules out explanations due to the fixed task order, i.e., the possibility that the negative correlation between value-based facilitation effects is merely due to further encoding during lexical decision. Instead, we found no relationship between lexical decision time for an item and its later recall probability, in line with our previous interpretation of the between-subjects correlations: the enhancements in the two tasks were driven by different mechanisms.

We had not expected the negative correlation between the effects of value on implicit and explicit memory. In an attempt to derive an explanation post-hoc, we took a closer look at our data. Perhaps the observed negative correlation between implicit and explicit memory had been driven by differences in participants' learning strategies in the value-learning task (even though, as reported above, the asymptotic accuracy in the value-learning task did not correlate significantly with the value effects on the two memory tests). We speculated that participants who learned values earlier may be the ones who showed greater effects of value on implicit memory, because they would have had a larger number of trials on which they knew the correct values. In contrast, participants who took longer to learn presumably found the value learning task more challenging early on; for these participants, value may have been used more as a deliberate retrieval cue in later explicit memory. For this purpose, we measured how long it took for participants to reach an 80% accuracy criterion in the value training task (i.e., trials-to-criterion, TTC: choosing the high-value item on 80% of all trials within a block). We then correlated TTC with reward effects on the lexical decision task (DiffLD) and with reward effects on free recall (DiffFR). Note that 6 participants never reached the 80% accuracy criterion; for these participants, the TTC was set to 14, i.e., one greater than the actual number of trial blocks presented in the value-learning task. This correction to the TTC measure served to denote that these participants required more learning trials to reach 80% accuracy. In line with our reasoning, we found that participants who reached the learning criterion *earlier* were exhibited stronger implicit memory effects due to reward value (i.e., greater priming in the lexical decision task, DiffLD) [$\rho(93) = -.22$, $p < .05$]. This may provide evidence that participants who had learned items values earlier (fewer trials-to-criterion) had more trials on which to accumulate value learning, which then enhanced implicit memory for high value items. Complementing this result, we found that participants who took *longer* to reach the learning criterion exhibited stronger explicit memory effects due to reward value (i.e., greater difference in recall probabilities in the free recall task, DiffFR) [$\rho(93) = .24$, $p < .05$]. This is consistent with the idea that participants for whom value learning was initially more challenging may have used value more as an explicit memory cue in free recall. Further, when controlling for trials-to-criterion in a partial correlation analysis, the negative correlation between the effects of value on lexical decision and free recall was no longer negative, and was not significant [$\rho_p(93) = .040$, $p > .1$]. Although our specific interpretation is post-hoc and should be considered with caution, the results of this analysis at least suggest that the way people learned the values initially mediated the mutually exclusive effects of value on implicit and explicit memory.

Summary

Experiment 1 revealed that high-value words were subsequently remembered better than low-value words in *both* implicit and explicit unrewarded memory tests. The effect of value on memory in these two memory tasks was slightly negatively correlated, suggesting the presence of at least two mechanisms mediating the memory enhancement by reward value, rather than a global enhancement of memory due reward value. Different initial value learning strategies may have contributed to this negative correlation.

The enhancement of implicit memory by value (i.e., an accessibility bias for high-value items) is novel in itself. Although the influence of reward value on response time in our lexical decision task was relatively small (~ 15 ms), this is consistent with studies that presented a reward cue in monetary incentive tasks and found that reward value facilitated response time by ~ 20 ms (Abler et al., 2005; Sescousse et al., 2010; Staudinger et al., 2011). Furthermore, nearly all prior studies demonstrating the reward-based enhancement of memory used procedures that led to the deliberate prioritization of encoding due to reward value. Here we used a procedure where participants incrementally learned values and found the enhancement of both implicit and explicit memory due to reward value. Because the memory tests were unrewarded, and no prioritization instructions were given, this suggests that not only can participants prioritize when asked to, but they exhibit a bias to learn high-value words better than low-value words. Such a bias may serve them well in naturalistic situations, however, if values change, this bias may become maladaptive.

2.4 Experiments 2a and 2b

We next asked whether the high-value item advantage observed following training in Experiment 1 would extend to a new learning situation involving the reward-value trained items. Having established that the value-learning procedure in Experiment 1 can enhance explicit-memory due to reward value, we used the same procedure to test for effects of reward value on new learning involving value-trained words in a different context. Following training as in Experiment 1 and a distractor task, we had participants learn word lists consisting of trained words and untrained words. In a study/test procedure, participants viewed each new list, followed by delayed free recall.

As in Experiment 1, our dependent measure in the free recall task was the proportion of words recalled of each word type (high-value, low-value or new). However, proportion of recalls is a rather coarse measure of memory, as it collapses across all responses given by a participant on a list. Apart from being a test of item retrievability, free recall is also a test of associations between items and a specific list-context. In other words, words output earlier in free recall represent the items that are easier to retrieve and also have the strongest associations with the current target-list context (e.g., Bjork & Whitten, 1974; Brown, Neath, & Chater, 2007; Crowder, 1976; Howard & Kahana, 1999; Raaijmakers & Shiffrin, 1981). Likewise, late in the recall sequence, responses are more likely to be guesses. Thus, in addition to recall accuracy, we tested if any word type (high-value, low-value, new) was recalled significantly earlier or later than any other word type.

We considered two hypotheses: Our reward-maximization hypothesis led to the prediction that participants will recall more high-value than low-value words, due to prioritized study of the words that had the high values previously, similar to previous studies finding an enhancement of memory due to reward value when rewards are earned for successful memory performance (e.g., Adcock et al., 2006; Castel et al., 2002). This hypothesis is also suggested by investigations of the effects of emotional arousal on memory, such as Hadley and MacKay's (2006) priority-binding hypothesis which proposes an enhancement of contextual binding due to arousal (also see Siddiqui & Unsworth, 2011). Alternatively, our value-interference hypothesis led to the opposite prediction: Words with a previously acquired high reward-value will be harder to learn and remember in a new context than words with a low reward-value if higher values direct attention toward the high-value items themselves, but away from other pertinent information. This hypothesis also suggests that for the high-value words, participants will find it difficult to constrain their memory retrieval processes to just the list-context of the most recently studied list and will instead erroneously recall more high-value words than low-value words. This hypothesis is based on studies finding an impairment of new associative memories between cues that had previously been predictive of emotionally arousing information (Mather & Knight, 2008; Nashiro, Sakaki, Huffman, & Mather, 2012; Novak & Mather, 2009; Sakaki, Niki, & Mather, 2011). For example, Mather and Knight (2008) found that participants had more difficulty learning new associations between sounds/faces and nearby presented digits (and other contextual information), if the sounds/faces had initially been paired with negative images, an effect that did not occur when they had been paired initially with neutral images. This suggests, emotional arousal may have interfered with participants' ability to learn subsequent associations. Furthermore, Novak and Mather (2009) had participants learn screen locations for neutral and negative images. When locations for individual pictures remained the same over several study-test cycles, participants made more location memory errors for emotional than neutral images in later cycles. Thus, an initial incorrect association between an emotional picture and a location may have led to more interference with learning the correct association than an initial incorrect association for a neutral picture. Moreover, when the locations for individual pictures were switched after three cycles, participants were worse at updating their memory with the new locations for negative images as opposed to neutral images. Together, these findings imply that emotional items are more affected by proactive interference from previous experience with the items, which may present as impaired learning of new associations with such items.

We conducted two variants of this experiment; Experiment 2b had a faster list presentation rate, to further test whether possible effects of previous reward value on new list learning are driven by a time-consuming strategy applied during study (i.e., value-based prioritization of encoding or retrieval).

2.4.1 Methods

Participants

A total of 72 introductory psychology students at the University of Alberta participated for partial fulfillment of course credit. All participants had normal or corrected-to-normal vision, learned English before the age of six, and were comfortable typing. Participants gave written informed consent prior to the study, which was approved by a University of Alberta Research Ethics Board. Participants never participated in more than one of Experiment 1, 2a, and 2b. Experiment 2a had 40 participants; Experiment 2b had 32 participants.

Materials

The same materials as used in the training phase of Experiment 1 were used in both Experiments 2a and 2b.

Six maze puzzles were generated for the distractor task using an online maze generator (`http://www.hereandabove.com/maze/mazeorig.form.html`). Mazes were made using the generator's default settings.

Procedure

The experiment consisted of three tasks performed in a fixed sequence: value-learning, maze distractor, and study/test free recall of six nine-word lists. Participants were not provided with details of the subsequent task until the current task was completed.

Value learning. The procedure was the same as in Experiment 1.

Maze distractor. To reduce the very high level of proactive interference from the value training phase on the free-recall phase, we included a non-verbal distractor task following value-training. Participants were given 5 min. to complete pencil-and-paper mazes. When participants finished one maze, they were provided with another maze. This procedure was repeated until the 5 min. had elapsed, at which point the maze was removed and the participant advanced to the study/test free recall task. On average, participants completed 2–3 mazes within the 5 min.

Study/test free recall. Participants were told to study each list of words and that their memory for the list would be tested, but that they would not earn any reward in this phase. Participants first studied one practice list of 9 words from the word pool that were excluded from analyses, and 6 experimental lists of 9 words each: 3 high-value words from the value-learning task, 3 low-value words from the value-learning task, and 3 new words (random order of presentation in each list).

Each word was presented for 1800 ms or 800 ms (Experiment 2a and 2b, respectively), after which the screen was cleared for 200 ms. After being presented with all 9 words, participants were given a distractor task that consisted of four equations in the form of $\mathsf{A} + \mathsf{B} + \mathsf{C} = \_\_\_$, where A, B, and C were randomly selected digits between 2 and 8. Each equation remained in the centre of the screen for 5000 ms. The participant was asked to type the correct answer during this fixed interval, after which the screen was cleared for 200 ms.

After the distractor, participants were given 1 min. to recall as many of the words from the list that they could (i.e., free recall). Participants were asked to type out their responses. After each response, a blank screen was presented for 500 ms. Participants were allowed to pause prior to the presentation of the next list. This procedure (list encoding, math distractor, list free recall) was repeated for all 6 lists.

Data Analysis

Effects were considered significant based on an alpha level of 0.05. ANOVAs are reported with Greenhouse-Geisser correction for non-sphericity where appropriate and post-hoc pairwise comparisons are Bonferroni-corrected.

Because the value-learning task consisted of 13 blocks, which means participants had 13 presentations of each stimulus, we expected there to be a large amount of proactive interference. As participants could not advance to the next list until the full minute expired, we expected later responses to include a high level of guesses. However, we were also concerned that some participants may not have understood that they were to confine their responses to the very last list presented. Therefore, to identify such non-compliant participants, we screened out participants who had extremely low accuracy *early* in the output sequences.

We calculated the average proportion of correct recalls within the first four responses to ensure that participants included in the analysis attempted to recall items from the most recent list (i.e., that they followed instructions). We found that most participants responded with three correct recalls in their first four responses [Experiment 2a: $M = 3.46$; Experiment 2b: $M = 3.32$]. However, five participants produced an average of one or fewer correct recalls in their first four responses [Experiment 2a: two participants; Experiment 2b: three participants] and were excluded from further analyses. Excluding these participants, the number of correct recalls in the first four responses did not substantially change the mean correct recalls within the first four recalls of the entire samples [Experiment 2a: $M = 3.50$; Experiment 2b: $M = 3.35$]. Similarly, total number of correct recalls did not change much [Experiment 2a, total sample: $M = 9.35$; Experiment 2a, excluding 2 participants: $M = 9.45$; Experiment 2b, total sample: $M = 9.39$; Experiment 2b, excluding 3 participants: $M = 9.91$]. Thus, exclusion of these five participants did not substantially change the observed recall patterns. For the remaining participants, analyses were carried out after removing extra-experimental intrusions and within-list repetitions.

2.4.2 Results and Discussion

Value learning

The value training data resembled the data in Experiment 1 (Figures 2.1b and 2.1c). Performance again began near chance, and improved across blocks; in the last block (block 13), accuracy was significantly greater than chance [Experiment 2a: $t(37) = 31.16$, $p < .001$, $M = 0.94 \pm 0.03$ correct; Experiment 2b: $t(30) = 12.83$, $p < .001$, $M = 0.89 \pm 0.06$ correct].

Study/test free recall

Proportion of words recalled. In each of Experiments 2a and 2b, we conducted repeated-measures ANOVAs on Word Type (high-value, low-value, new) on the proportion of words recalled. Proportion recalled was defined as the average number of correct words recalled of each word type across lists, divided by 3 (the number of words of each type in each list). The main effects of Word Type were not significant in either experiment [Experiment 2a: $F(2, 67) = 1.08$, $p > .1$, $\eta_p^2 = .03$, $M(\text{high}) = .56 \pm .04$, $M(\text{low}) = .57 \pm .04$, $M(\text{new}) = .53 \pm .05$; Experiment 2b: $F(2, 52) = 2.74$, $p > .1$, $\eta_p^2 = .08$,

M(high) = .52 ± .05, M(low) = .51 ± .06, M(new) = .44 ± .06] (Figure 2.4). The lack of a difference in recall rates of high-value and low-value words suggests that, by this measure, effects of previously learned value on memory had been neutralized.

Output order. To analyze output order for each Word Type (high-value, low-value, new) directly, we borrowed the logic of the Wilcoxon–Mann–Whitney rank-sum test on the output positions of each word type to derive a measure of differences in median output position for each Word Type (as suggested by Hubert & Levin, 1978). For each list, for each pair-wise Word Type comparison, the U-statistic was Z-transformed and then averaged across lists to obtain a measure for each participant. Because these values were already Z-scores, they were then compared with a t-test against zero for each comparison between Word Types. Participants with no recalls of a given Word Type in two or more lists were excluded from this analyses as they did not contribute additional information to this analysis (leaving $N = 35$ and $N = 29$ in Experiments 2a and 2b, respectively). In both experiments, low-value words had significantly earlier median output positions than high-value words [Experiment 2a: $mean(Z_U) = 0.14$, $t(34) = 2.40$, $p < .05$, $d = .46$; Experiment 2b: $mean(Z_U) = 0.23$, $t(28) = 2.35$, $p < .05$, $d = .42$] (Figure 2.5), suggesting that low-value words were easier to recall. High- and low-value words did not differ significantly in median output position relative to new words [all p's $> .1$].

Intrusions. In Experiment 1, we found that high-value items were more retrievable in a final free recall test. Therefore, participants might be more likely to guess high- than low-value words in free recall of the 9-word lists in Experiments 2a and 2b. If list discrimination were enhanced for high-value words, following from the reward-maximization hypothesis, participants should make fewer intrusions of high-value words than for low-value words. However, if participants had a more difficult time determining whether high-value items belonged to the current list, following from the value-interference hypothesis, then we should instead find more intrusion responses for high-value words than

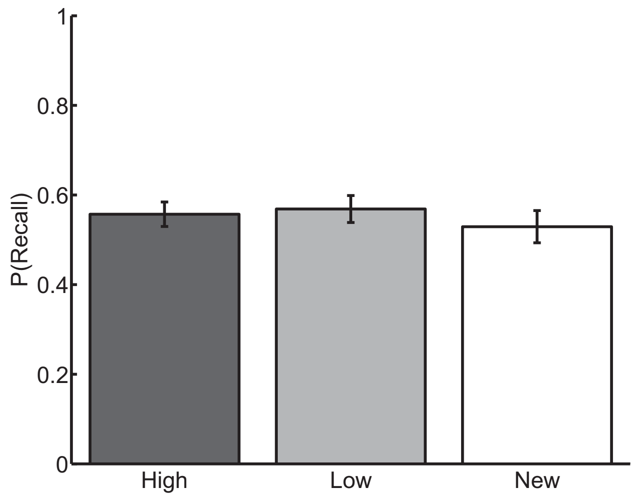

(a) Experiment 2a

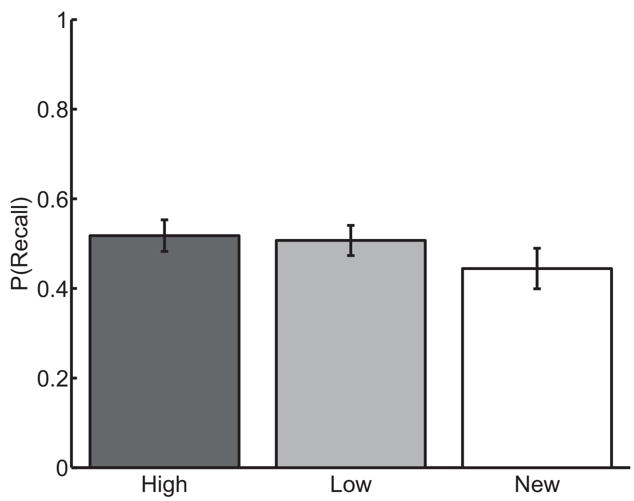

(b) Experiment 2b

Figure 2.4: Correct recall rates for Experiments 2a (a) and 2b (b). ‘High’ and ‘Low’ represent high- and low-value words, respectively, from the value-learning task. ‘New’ represents words that were not first present in the value-learning task, but only in study-test free recall. Error bars are 95% confidence intervals, corrected for inter-individual differences (Loftus & Masson, 1994).

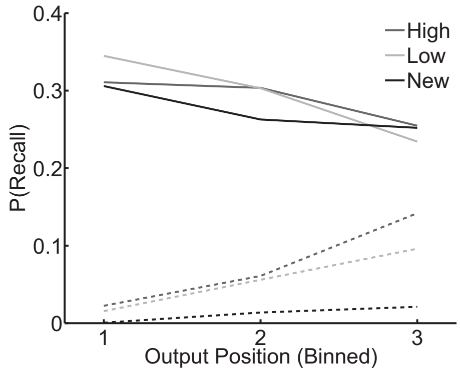

(a) Experiment 2a

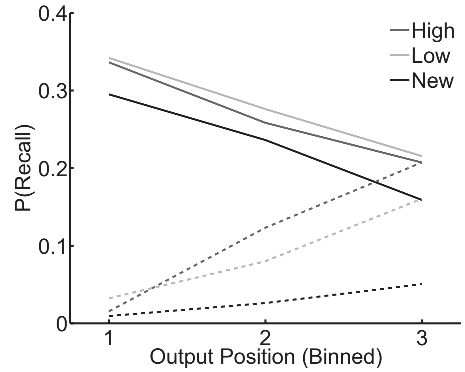

(b) Experiment 2b

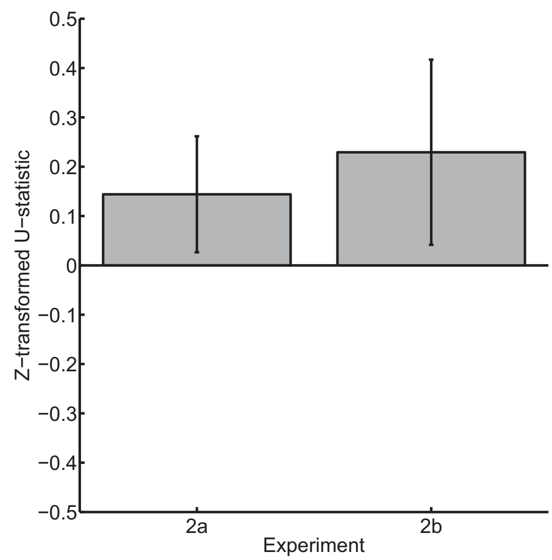

(c) Wilcoxon–Mann–Whitney rank-sum test

Figure 2.5: Output positions in the free recall task of Experiments 2a and 2b. Probability of recall of each word type for vincentized output position bins in the free recall task of Experiments 2a (A) and 2b (B). 'High' and 'Low' represent high- and low-value words, respectively, from the value-learning task. 'New' represents words first that were not present in the value-learning task. Solid lines and markers represent correct responses; dashed lines with hollow markers represent intrusion responses. Error bars were omitted for visual clarity. (C) Plots of the Z-transformed U-statistics comparing median output positions for high- vs. low-value words in both Experiments 2a and 2b. Larger values represent later output positions. Error bars are 95% confidence intervals.

for low-value words. As words were not re-used from one list to the next, intrusions were defined as words that were not on the target (most recently studied) list. Intrusions could come from the training or else from prior free-recall study lists.

Repeated-measures ANOVAs were conducted on intrusion rates. The measure was the proportion of all responses (excluding extra-experimental intrusions and repetitions) on a given list that were intrusions of each word type, averaged across lists. Participants with fewer than three intrusions in total were excluded from only this analysis as they provided an insufficient number of data points (leaving $N = 24$ and $N = 22$ included participants in Experiments 2a and 2b, respectively). The main effect of Word Type was significant in both experiments [Experiment 2a: $F(2, 34) = 22.11$, $p < .001$, $\eta_p^2 = .51$, $M(\text{high}) = .11 \pm .03$, $M(\text{low}) = .070 \pm .024$, $M(\text{new}) = .018 \pm .010$; Experiment 2b: $F(2, 36) = 23.14$, $p < .001$, $\eta_p^2 = .50$, $M(\text{high}) = .13 \pm .04$, $M(\text{low}) = .081 \pm .017$, $M(\text{new}) = .029 \pm .018$] (Figures 2.6a and 2.6b). High-value words were more likely to intrude than both low-value words [Experiment 2a: $t(23) = 2.55$, $p < .05$; Experiment 2b: $t(21) = 3.12$, $p < .01$] and new words [Experiment 2a: $t(23) = 6.73$, $p < .001$; Experiment 2b: $t(21) = 6.00$, $p < .001$]. Low-value words were also intruded more than new words [Experiment 2a: $t(23) = 4.64$, $p < .001$; Experiment 2b: $t(21) = 4.84$, $p < .001$]. This result also supports the value-interference hypothesis, which suggested that high-value words are harder to place uniquely within the target list (i.e., contextual binding). Moreover, the small advantage of high-value words over low-value words following training in Experiment 1 (ratio of ∼5:4) evolved into a much larger ratio (∼3:2) in the intrusion rates of Experiments 2a and 2b. If guessing were purely based on better retrievability caused by high value and measured by final free recall in Experiment 1, we would have expected the same ratio for intrusion rates, as the words would inherit the same distribution from the final free recall data. The fact that the ratio is exaggerated for intrusions here suggests that this measure is influenced by more than just item retrievability; we suggest that high-value words were not only sampled more often as candidate responses, but were also screened less well, and thus, were more likely to be recalled in error.

Summary

In both Experiment 2a and 2b, previously trained words were correctly recalled at equal rates, regardless of reward value; thus, the advantages we saw for high-value words in Experiment 1 did not carry forward to a situation in which participants had to relearn subsets of trained words and link them to a specific, new list context. Moreover, high-value words were output *later* in the recall sequence than low-value items, and were more likely to be retrieved erroneously (reflecting proactive interference). This suggests that they were more weakly linked to the current-list context, and were more difficult to accurately screen based on recent-list membership. This pattern of findings held both for a slower presentation rate (2 s/word in Experiment 2a) and for a faster presentation rate (1 s/word in Experiment 2b), and the magnitudes of the output-order and intrusion-rate effects were similar between experiments (Figures 2.5 and 2.6). This suggests that the effects of value unlikely result from a deliberate, effortful and time-consuming process during study (e.g., participants deliberately diverting attention toward the low-value words). It is more plausible that the difference between performance on high- and low-value words was due to persisting effects of reward value on memory from the value-learning task. This would make current-list membership more confusable for high-value items, and screening candidate responses more difficult for high- than for low-value items.

2.5 General Discussion

In two studies, we investigated the influence of previously trained reward value on unrewarded tests of memory. In Experiment 1, implicit memory (facilitated access in lexical decision) was enhanced by reward value, in addition to enhanced explicit memory due to reward value (probability of recall in final free recall). These two memory enhancements were negatively correlated across participants, suggesting the presence of at least two mechanisms whereby re-

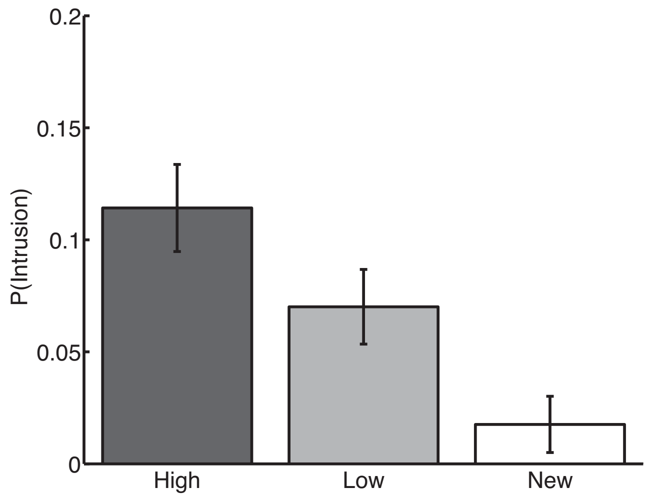

(a) Experiment 2a

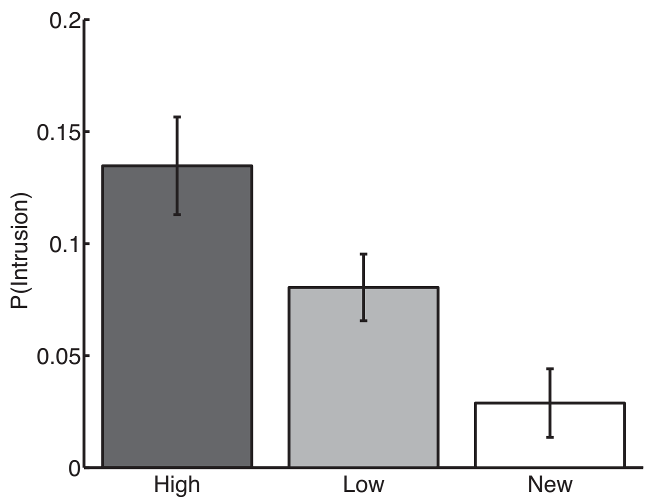

(b) Experiment 2b

Figure 2.6: Intrusion rates during free recall in Experiments 2a (a) and 2b (b). 'High' and 'Low' represent high- and low-value words, respectively, from the value-learning task. 'New' represents words first that were not present in the value-learning task. Error bars are 95% confidence intervals, corrected for inter-individual differences (Loftus & Masson, 1994).

ward value can influence memory. In Experiments 2a and 2b, we found that previously learned reward values can cause problems for contextual binding, when trained items needed to be tied to a new, specific context (namely, belonging to the most recent list). Low-value items were produced earlier in recall than high-value items, and high-value items intruded more often, suggesting that they were not effectively screened as belonging to the wrong list. The interactions between reward value and memory are thus multifaceted, with implicit and explicit memory being enhanced due to reward value through different mechanisms (Experiment 1), and reward value leading to impaired memory for contextual information (Experiments 2a and 2b).

Influence of reward value on implicit and explicit memory

In Experiment 1, we observed an enhancement of both explicit *and* implicit memory due to reward value. Since value enhanced both our memory measures, one may have expected that value-learning globally enhanced all kinds of learning of high-value items. For instance, enhanced memory could have been driven by value solely through the recruitment of additional attentional resources during the value-learning task: If participants paid more attention to the high-value items than the low-value items during this first phase, high-value items may then be more primed in lexical decision and more retrievable in free recall. Both enhancements would then have originated from a single, global, value-learning mechanism resulting in a high, positive correlation between the two measures. However, implicit and explicit memory are supported by distinct neural pathways (e.g., Rugg et al., 1998; Schott et al., 2005, 2006). Thus, it is also plausible that value may *separately* enhance implicit and explicit memory, and such enhancements would be uncorrelated or negatively correlated (also see Gopie et al., 2011; May et al., 2005). Our results favored the latter hypothesis: Implicit and explicit reward-based enhancements were negatively correlated across participants. In other words, participants who demonstrated greater reward value facilitation in lexical decision had *less* reward facilitation in free recall. Prior research also supports the notion of different value-based learning strategies leading to differential engagement of implicit- and explicit-memory (Wimmer & Shohamy, 2011, see also Bayley, Frascino, & Squire, 2005). This apparent trade-off between memory systems due to learning strategy is also supported by research on the effects of stress on memory, where deliberative (i.e., goal-directed) and procedural (i.e., habit-based) learning strategies can similarly be learned through two distinct memory systems (Schwabe & Wolf, 2011; Schwabe, Höffken, Tegenthoff, & Wolf, 2011; Schwabe, Joëls, Roozendaal, Wolf, & Oitzl, 2011; Schwabe, Dickinson, & Wolf, 2011).

One possible source of this negative correlation may be behaviour during the value-learning task, as suggested by the correlation analyses involving the trials-to-criterion measure. Trials-to-criterion explained the negative correlation between the effects of reward value on implicit- and explicit-memory. We suggest that the effects of value on implicit memory benefited from participants having more experience with the knowledge of the values of items. That is, the earlier someone learned to prefer the high- over the low-rewarding item during the training task, the more exposures to correct pairings of their own choice and high rewards they would have had. Such increased exposure and procedural practice of correct response-high reward pairings could then have selectively promoted the formation of an implicit memory bias. In contrast, free recall is a self-cued memory task; thus, value would be expected to influence free recall insofar as a participant includes value as part of their retrieval cue. Participants who initially found the value-learning task more challenging may have been oriented more toward value during the free recall test, thus producing a positive relationship between trials-to-criterion and the effect of value free-recall, opposite to what was observed with the effect of value on lexical decision. This indirect evidence of two distinct value-learning mechanisms may be related to similar dissociations in probabilistic value learning strategies reported by others (Allen & Estes, 1972; Estes, 1972; Humphreys et al., 1968; Medin, 1972a).

Although lexical decision and free recall test implicit and explicit memory, respectively, the two tests also differ in several other ways, so alternative interpretations of the cause of the dissociation must be considered. First, the dependent measure in lexical decision was response time, a measure of access speed; in free recall, the dependent measure was probability of recall, a measure that is sensitive to sampling probability and recovery processes, as well as memory cueing processes (e.g. Raaijmakers & Shiffrin, 1981). Our dissociation could therefore reflect differential influences of reward value on access speed versus sampling, recovery or cueing processes. Second, participants are presented with a copy-cue to judge in lexical decision, but in free recall, participants must apply their own retrieval cues to generate responses. Our dissociation could thus reflect distinct influences of reward value on judgement processes versus item-retrieval processes (cf. Humphreys, Bain, & Pike, 1989). Regardless of which of these accounts is correct, our findings extend the boundary conditions of reward-value enhancement of memory effects, and suggest that the effect of reward value on memory is non-unitary.

Influence of previously learned reward values on contextual binding

In Experiments 2a and 2b, what started as an advantage for high-value words (evident in Experiment 1) became a disadvantage when participants had to overcome proactive interference from the value-training phase and learn new sets of words that included both trained and untrained words. High- and low-value words were recalled at equivalent rates overall, but low-value words were produced earlier in output. High-value words were intruded more (and even more than expected based on the final free recall rates of Experiment 1). These findings suggest less effective contextual binding for high- than for low-value words. This contradicts our reward-maximization hypothesis, and suggests that there are limits to the degree to which participants are biased to modulate their learning to maximize cumulative reward; one such limit is in relearning high-valued items in new, specific contexts.

If the additional resources devoted to high-value items included processing items within their context (i.e., the most recent list), then one would also expect participants to be able to rule out words that were recalled from previous contexts (i.e., the value-learning task or previous lists in the free recall task), which is inconsistent with the elevated intrusion rate for high-value words in Experiments 2a and 2b. Thus, our findings are more consistent with our value-interference hypothesis, which posits that reward value impairs contextual binding. These results are also in line with findings obtained with manipulations of emotional arousal, where memory for the arousing items is enhanced, but the learning of new associations involving such items is impaired (Mather & Knight, 2008; Nashiro et al., 2012; Novak & Mather, 2009; Sakaki et al., 2011).

Although positive, as well as negative emotional items can be remembered better than emotionally neutral items (e.g., Dewhurst & Parry, 2000; Siddiqui & Unsworth, 2011), it would be reasonable to argue that the influence of reward value on memory may be more similar to the influence of positive— not negative— emotion on memory. While many studies have found that emotion can enhance memory for items and often impairs memory for associations, the majority of these findings used negatively valenced emotional stimuli (Fredrickson, 1998). Whereas negative emotions lead to attentional narrowing (e.g., the weapon focus effect; Loftus, Loftus, & Messo, 1987), positive emotion can to lead to a broadening of attention (Fredrickson, 1998). When participants are asked to learn associations containing emotionally positive, negative, or neutral items, participants are often better able to learn pairs with positive items than pairs with negative items (Okada et al., 2011; Pierce & Kensinger, 2011; Zimmerman & Kelley, 2010), suggesting that positive emotion can enhance participants' ability to form associations between items. (Note that sometimes an association-memory impairment has been observed even with positive stimuli, e.g., Mather & Knight, 2008.) If this interpretation is correct, and reward value functions similar to positive emotionality, then one would expect reward-value-based facilitation of free recall in Experiments 2a and 2b, inconsistent with our results. We recently showed previously reported arousal-based enhancements in association-memory could instead be attributed to enhanced memory for the target items, and that this item-memory effect can mask an underlying impairment of association-memory (Madan, Caplan, et al., 2012). Thus, it is similarly possible that prior findings regarding the effects of positive emotion on associative learning may be composed of conflicting effects. Finally, false memories can be viewed as failures of contextual discrimination. Emotion, both induced in the participant, and emotionality of items, can increase rates of false memories. This has been found for both negative and positive emotions (Corson & Verrier, 2007; Dehon, Laroi, & Van der Linden, 2010; Storbeck & Clore, 2005), and appears similar to the list-discrimination problems we found for high-value items here.

Implications for previous findings of reward-value enhancements of memory

Reconsidering Raymond and O'Brien (2009) we detailed in the Introduction, our results suggest that their findings may have resulted from a summation of two distinct enhancement effects, one acting on implicit and acting on explicit memory. Regarding studies that have found that participants can prioritize their memory processes based on specific item-values presented alongside stimuli (Adcock et al., 2006; Bjork & Woodward, Jr., 1973; Castel et al., 2002; Eysenck & Eysenck, 1982; Gruber & Otten, 2010; Harley, 1965; Kuhl et al., 2010; Soderstrom & McCabe, 2011; Tarpy & Glucksberg, 1966; Watkins & Bloom, 1999; Weiner & Walker, 1966), advantages in recall and recognition for high-value items resemble the enhancement effect we found in the final free recall measure of Experiment 1. However, in all these studies, values were presented with items, but participants were never asked to link those items to a new context. Our findings in study/test free recall in Experiments 2a and 2b raise the possibility that if participants have to learn new lists composed of previously prioritized items, their memory might be compromised by the kind of value-based interference effect found here. In particular, given that the intrusion pattern was the largest effect we observed in Experiments 2a and 2b, we would predict that items previously linked to higher values would be intruded more— that is, participants would continue to produce them as responses even when inappropriate. In turn, since prioritization procedures directly ask participants to favor high-value items, whereas our procedure did not, it is quite possible that the list discrimination procedure we found for high-value words could be overcome again if participants were asked to prioritize high-value words in later list learning.

2.6 Conclusion

Reward value can enhance memory for higher-valued items by increasing access speed and probability of retrieval. These dual enhancement effects of value on implicit and explicit memory measures may, in turn, be the results of dual value-learning styles. These enhancement effects come with a side effect of a poorer ability for participants to bind high-value items uniquely to a specific context, suggesting that items with high reward value can have a deleterious effect on subsequent memory tasks.

2.7 Acknowledgements

Supported by the Natural Sciences and Engineering Research Council of Canada, the Canadian Institutes of Health Research, and the Alberta Ingenuity Fund.

2.8 References

Abler, B., Walter, H., & Erk, S. (2005). Neural correlates of frustration. *Neuroreport*, *16*, 669–672.

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517.

Allen, G. A., & Estes, K. (1972). Acquisition of correct choices and value judgements in binary choice learning with differential rewards. *Psychonomic Science*, *27*, 68–72.

Bayley, P. J., Frascino, J. C., & Squire, L. R. (2005). Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature*, *436*, 550–553.

Bijleveld, E., Custers, R., & Aarts, H. (2009). The unconcious eye opener: Pupil dilation reveals strategic recruitment of resources upon presentation of subliminal reward cues. *Psychological Science*, *20*, 1313–1315.

Bijleveld, E., Custers, R., & Aarts, H. (2010). Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. *Cognition*, *115*, 330–335.

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, *6*, 173–189.

Bjork, R. A., & Woodward, Jr., A. E. (1973). Directed forgetting of individual words in free recall. *Journal of Experimental Psychology*, *99*, 22–27.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings* (Technical report No. C-1). Gainesville, FL: The Center for Research in Psychophysiology, University of Florida.

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*, 539–576.

Burke, A., Heuer, F., & Reisberg, D. (1992). Remembering emotional events. *Memory & Cognition*, *20*, 277–290.

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*, 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*, 1078–1085.

Castel, A. D., Farb, N., & 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*, 689–700.

Christianson, S.-A. (1992). Emotional-stress and eyewitness memory: A critical-review. *Psychological Bulletin*, *112*, 284–309.

Corson, Y., & Verrier, N. (2007). Emotions and false memories: Valence or arousal? *Psychological Science*, *18*, 208–211.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Custers, R., & Aarts, H. (2010). The unconcious will: How the pursuit of goals operates outside conscious awareness. *Science*, *329*, 47–50.

Dehon, H., Laroi, F., & Van der Linden, M. (2010). Affective valence influences participant's susceptibility to false memories and illusory recollection. *Emotion*, *10*, 627–639.

Dewhurst, S. A., & Parry, L. A. (2000). Emotionality, distinctiveness, and recollective experience. *European Journal of Cognitive Psychology*, *12*, 541–551.

Easterbrook, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review*, *66*, 183–201.

Estes, W. K. (1962). Theoretical treatments of differential reward in multiple-choice learning and two-person interactions. In J. Criswell, H. Solomon, & P. Suppes (Eds.), *Mathematical methods in small group processes.* Stanford, CA: Standord University Press.

Estes, W. K. (1966). Transfer of verbal discriminations based on differential reward magnitudes. *Journal of Experimental Psychology*, *72*, 276–283.

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

Eysenck, M. W., & Eysenck, M. C. (1982). Effects of incentive on cued recall. *The Quarterly Journal of Experimental Psychology Section A*, *34*, 489–498.

Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns : Midazolam enhances implicit inference in humans. *Psychological Science*, *17*, 700–707.

Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, *306*, 1940–1943.

Fredrickson, B. L. (1998). What good are positive emotions? *Review of General Psychology*, *2*, 300–319.

Gopie, N., Craik, F. I. M., & Hasher, L. (2011). A double dissociation of implicit and explicit memory in younger and older adults. *Psychological Science*, *22*, 634–640.

Gradin, V. B., Kumar, P., Waiter, G., Ahearn, T., Stickle, C., Milders, M., ... Steele, J. D. (2011). Expected value and prediction error abnormalities in depression and schizophrenia. *Brain*, *134*, 1751–1764.

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

Hadley, C. B., & MacKay, D. G. (2006). Does emotion help or hinder immediate memory? arousal versus priority-binding mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 79–88.

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

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 923-941.

Hubert, L. J., & Levin, J. R. (1978). Evaluating priority effects in free recall. *British Journal of Mathematical and Statistical Psychology*, *31*, 11–18.

Humphreys, M. S., Allen, G. A., & Estes, W. K. (1968). Learning of two-choice, differential reward problems with informational constraints on payoff combinations. *Journal of Mathematical Psychology*, *5*, 260–280.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*, *96*, 208–233.

Johnsrude, I. S., Owen, A. M., White, N. M., Zhao, W. V., & Bohbot, V. (2000). Impaired preference conditioning after anterior temporal lobe resection in humans. *Journal of Neuroscience*, *20*, 2649–2656.

Johnsrude, I. S., Owen, A. M., Zhao, W. V., & White, N. M. (1999). Conditioned preference in humans: A novel experimental approach. *Learning and Motivation*, *30*, 250–264.

Kuhl, B. A., Shah, A. T., DuBrow, S., & Wagner, A. D. (2010). Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *Nature Neuroscience*, *13*, 501–508.

Loftus, E. F., Loftus, G. R., & Messo, J. (1987). Some facts about 'weapon focus'. *Law and Human Behavior*, *11*, 55–62.

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

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

Madan, C. R., Caplan, J. B., Lau, C. S. M., & Fujiwara, E. (2012). Emotional arousal does not enhance association-memory. *Journal of Memory and Language*, *66*, 695–716.

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

Mather, M., & Knight, M. (2008). The emotional harbinger effect: Poor context memory for cues that previously predicted something arousing. *Emotion*, *8*, 850–860.

Mather, M., & Sutherland, M. R. (2011). Arousal-biased competition in perception and memory. *Perspectives on Psychological Science*, *6*, 114–133.

May, C. P., Hasher, L., & Foong, N. (2005). Implicit memory, age, and time of day: Paradoxical priming effect. *Psychological Science*, *16*, 96–100.

Medin, D. L. (1972a). Partial information and choice behaviour in differential reward magnitude learning. *Psychonomic Science*, *27*, 73–76.

Medin, D. L. (1972b). Role of reinforcement in discrimination learning set in monkeys. *Psychological Bulletin*, *77*, 305–318.

Nashiro, K., Sakaki, M., Huffman, D., & Mather, M. (2012). Both younger and older adults have difficulty updating emotional memories. *Journal of Gerontology: Psychological Sciences*, in press.

Novak, D. L., & Mather, M. (2009). The tenacious nature of memory binding for arousing negative items. *Memory & Cognition*, *37*, 945–952.

Okada, G., Okamoto, Y., Kunisato, Y., Aoyama, S., Nishiyama, Y., Yoshimura, S., ... Yamawaki, S. (2011). The effect of negative and positive emotionality on associative memory: An fMRI study. *PLoS ONE*, *6*, e24862.

Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., & Frith, C. D. (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science*, *316*, 904–906.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042–1045.

Pierce, B. H., & Kensinger, E. A. (2011). Effects of emotion on associative recognition: Valence and retention interval matter. *Emotion*, *11*, 139–144.

Pubols, B. H. (1960). Incentive magnitude, learning, and performance in animals. *Psychological Bulletin*, *57*, 89–115.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93–134.

Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation. *Psychological Science*, *20*, 981–988.

Rugg, M. D., Mark, R. E., Walla, P., Schloerscheidt, A. M., Birch, C. S., & Allan, K. (1998). Dissociation of the neural correlates of implicit and explicit memory. *Nature*, *392*, 595–598.

Sakaki, M., Niki, K., & Mather, M. (2011). Updating existing emotional memories involves the frontopolar/orbito-frontal cortex in ways that acquiring new emotional memories does not. *Journal of Cognitive Neuroscience*, *23*, 3498–3514.

Schott, B. H., Henson, R. N., Richardson-Klavehn, A., Becker, C., Thoma, V., Heinze, H.-J., & Düzel, E. (2005). Redefining implicit and explicit memory: The functional neuroanatomy of priming, remembering, and control of retrieval. *Proceedings of the National Academy of Sciences*, *102*, 1257–1262.

Schott, B. H., Richardson-Klavehn, A., Henson, R. N. A., Becker, C., Heinze, H.-J., & Düzel, E. (2006). Neuroanatomical dissociation of encoding processes related to priming and explicit memory. *Journal of Neuroscience*, *26*, 792–800.

Schwabe, L., Dickinson, A., & Wolf, O. T. (2011). Stress, habits, and drug addiction: A psychoneuroendocrinological perspective. *Experimental and Clinical Psychopharmacology*, *19*, 53–63.

Schwabe, L., Höffken, O., Tegenthoff, M., & Wolf, O. T. (2011). Preventing the stress-induced shift from goal-directed to habit action with a

β-adrenergic antagonist. *Journal of Neuroscience*, *31*, 17317–17325.

Schwabe, L., Joëls, M., Roozendaal, B., Wolf, O. T., & Oitzl, M. S. (2011). Stress effects on memory: An update and integration. *Neuroscience and Biobehavioural Reviews*, in press.

Schwabe, L., & Wolf, O. T. (2011). Stress-induced modulation of instrumental behavior: From goal-directed to habitual control of action. *Behavioural Brain Research*, *219*, 321–328.

Sescousse, G., Redouté, J., & Dreher, J.-C. (2010). The architecture of reward value coding in the human orbitofrontal cortex. *Journal of Neuroscience*, *30*, 13095–13104.

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

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

Siddiqui, A. P., & Unsworth, N. (2011). Investigating the role of emotion during the search process in free recall. *Memory & Cognition*, *39*, 1387–1400.

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*, 1236–1242.

Staudinger, M. R., Erk, S., & Walter, H. (2011). Dorsolateral prefrontal cortex modulates striatal reward encoding during reappraisal of reward anticipation. *Cerebral Cortex*, *21*, 2578–3588.

Storbeck, J., & Clore, G. L. (2005). With sadness comes accuracy; with happiness, false memory: Mood and the false memory effect. *Psychological Science*, *16*, 785–791.

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

Valentin, V. V., & O'Doherty, J. P. (2009). Overlapping prediction errors in dorsal striatum during instrumental learning with juice and money reward in the human brain. *Journal of Neurophysiology*, *102*, 3384–3391.

Voon, V., Pessiglione, M., Brezing, C., Gallea, C., Fernandez, H. H., Dolan, R. J., & Hallett, M. (2010). Mechanisms underlying dopamine-mediated reward bias in compulsive behaviors. *Neuron*, *65*, 135–142.

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

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memory. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

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

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

Westbury, C., Hollis, G., & Shaoul, C. (2007). LINGUA: The Language-Independent Neighbourhood Generator of the University of Alberta. *The Mental Lexicon*, *2*, 273–286.

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

Wimmer, G. E., & Shohamy, D. (2011). The striatum and beyond: Hippocampal contributions to decision making. In M. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Attention & Performance XXII.* Oxford, UK: Oxford University Press.

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*, 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*, 1532–1547.

Zimmerman, C. A., & Kelley, C. M. (2010). "I'll remember this!" effects of emotionality on memory predictions versus memory performance. *Journal of Memory and Language*, *62*, 240–253.

Chapter 3

Is the enhancement of memory due to reward driven by value or salience?

A version of this work was previously published as: **Madan, C. R.**, & Spetch, M. L. (2012). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica, 139*, 343–349. doi:10.1016/j.actpsy.2011.12.010.

3.1 Abstract

Past research using two levels of reward has shown that the higher-value items are remembered better than lower-value items and this enhancement is assumed to be driven by an effect of reward value. In the present study, multiple levels of reward were used to test the influence of reward salience on memory. Using a value-learning procedure, words were associated with reward values, and then memory for these words was later tested with free recall. Critically, multiple reward levels were used, allowing us to test two specific hypotheses whereby rewards can influence memory: (a) higher value items are remembered better than lower value items (reward value hypothesis), and (b) highest and lowest value items are remembered best and intermediate-value items are remembered worst (following a U-shaped relationship between value and memory; reward salience hypothesis). In two experiments we observed a U-shaped relationship between reward value and memory, supporting the notion that memory is enhanced due to reward salience, and not purely through reward value.

3.2 Introduction

In day-to-day life, people often remember stimuli associated with rewarding experiences better than those associated with less rewarding experiences. For example, you are likely to remember a person you enjoyed talking to at a party better than someone you found less interesting. Thus, experiences associated with higher reward values are likely to be remembered better than those with lower reward values and, given the choice, one would choose to repeat a more rewarding experience over a less rewarding experience. Thus, in the party example, you would be more likely to remember the person you enjoyed talking to, and at a future party you would likely choose to talk to that person over the less interesting one.

The enhancement of memory due to reward suggested by this anecdotal example has been the subject of a recent flurry of studies (e.g., Adcock et al., 2006; Castel et al., 2002; Gruber & Otten, 2010; Shigemune et al., 2010; Soderstrom & McCabe, 2011; Wittmann et al., 2005) This research has confirmed the effects of reward value on memory with two levels of reward (high-value versus low-value rewards, or reward versus no reward), and evidence suggests that reward enhances memory via dopaminergic modulation (reviewed in Shohamy & Adcock, 2010).

A recent study by Madan, Fujiwara, et al. (2012) trained items to have reward values and later tested memory for these items in two unrewarded memory tasks: lexical decision (implicit memory) and free recall (explicit memory). They found enhanced memory for high-value items in *both* memory tasks, as well as a negative correlation between each task. Reward values were trained through a value-learning procedure (see Pessiglione et al., 2006), in which participants are presented with two items at a time and asked to choose one. At the beginning of training, performance is at chance, as item value is not yet known; however, through feedback, participants learn to choose higher-value items over lower-value items. This value-learning procedure is essentially a reinforcement learning procedure (Ludvig, Bellemare, & Pearson, 2011) and the effects of value on memory can be viewed in terms of prediction error

(Rescorla & Wagner, 1972; Sutton & Barto, 1998). Specifically, during the learning phase, the higher the reward value, the bigger the prediction error between what is expected for that item and the reward that occurs. Higher prediction errors trigger more activation of the dopamine system (Lisman & Grace, 2005; Shohamy & Adcock, 2010) and should make the higher-valued items more memorable. In the present study we closely follow the design of Madan, Fujiwara, et al. (2012) with one critical change: the use of multiple reward levels. While Madan, Fujiwara, et al. trained words to be either high- or low- value, we also include one (Experiment 1) or several (Experiment 2) intermediate values.

Prior studies investigating the enhancement of memory due to reward often explicitly instructed participants of an item's value when the item was being intentionally studied. For example, Adcock et al. (2006) presented a reward value cue (e.g., $5) to participants just prior to presenting the to-be-remembered item. Participants would only earn the rewarded amount for successfully remembering the item during the subsequent memory test. Since these studies only provided rewards for correctly remembering an item, and the reward given was based on an item's value, participants should deliberately prioritize their memory for the high-value items at the cost of the low-value items. Castel et al. (2002) tested this directly by presenting participants with twelve items, one-at-a-time, along with a unique point value ranging from 1 to 12. Points would only be earned for the successful recall of the associated item. Here the researchers also found better memory for the higher-value items relative to the low-value items. Studies using rewarded memory tests suggest that participants deliberately prioritize their memory for the higher-value items and suggest that memory performance should increase monotonically with reward value. In our procedure, participants are not told that they will need to recall the items and therefore our task, unlike those used in these prior studies, does not necessitate that participants deliberately prioritize their memory. Nevertheless, it is still possible that participants may attend more to the higher-value items at the cost of the lower-value items. Thus, one hypothesis is that reward value will monotonically predict memory performance, which we term the reward value hypothesis.

A second hypothesis regarding the relationship between reward value and memory is suggested by recent evidence of neural representations for reward salience, in addition to that for reward value (Cooper & Knutson, 2008; Jensen et al., 2007; Litt, Plassmann, Shiv, & Rangel, 2011; Zink, Pagnoni, Martin-Skurski, Chappelow, & Berns, 2004). This neural evidence suggests that both the most positive and the most negative values are more salient than intermediate values. Although our study includes only positive values, it is possible that the extremes of the range of values experienced are more salient than those closer to the center of the range. Thus a second hypothesis is that memory will show a U-shaped relationship to reward value, with both the highest- and the lowest-valued items being remembered better than the intermediate valued items, which we term the reward salience hypothesis. We should note that this hypothesis could also be consistent with a prediction error interpretation, because the results of Jensen et al. (2007) suggest that some brain regions may in fact represent prediction error based on reward salience.

Our experiment was designed to test the role of both reward value and reward salience on memory, using multiple reward values and an unrewarded memory task. Through the inclusion of multiple reward values, as well as the use of an unrewarded memory task, we will be able to determine if memory is enhanced due to reward value or reward salience. If there is a U-shaped (quadratic) relationship between reward value and memory as suggested by the reward salience hypothesis, then people may remember both the highest and the lowest-valued items better than the intermediate-valued items.

The inclusion of multiple reward values also allows an examination of the relationships between choice and memory. The reward value hypothesis predicts a monotonic relationship between choice and memory because higher valued items should be chosen more and remembered better. On the other hand, the reward salience hypothesis predicts that choice frequency and memory for items will not be monotonically related. Instead, the lowest valued item will be chosen the least but will be remembered better than intermediate valued items that were chosen more frequently.

Thus, our experiments were designed to test whether the reward salience hypothesis fares better than the reward value hypothesis in predicting the relationship between reward value and memory and between choice and memory. With our multiple level value-learning procedure, the reward salience hypothesis predicts that people will learn to choose items in direct relation to their value (i.e., choice will be monotonically related to reward value) but that the lowest and the highest valued items will be remembered best (memory will show a U-shaped relationship with reward value).

3.3 Experiment 1

3.3.1 Methods

Participants

71 introductory psychology students (40 female) at the University of Alberta participated for partial fulfillment of course credit. All participants were required to have learned English before the age of six and were required to be comfortable typing. Participants gave written informed consent prior to beginning the study, which was approved by a University of Alberta Research Ethics Board.

Materials

All of the words and non-words used in this study have been previously used in Madan, Fujiwara, et al. (2012). Words were selected from the MRC Psycholinguistic database (Wilson, 1988). Imageability and word frequency were all held at mid-levels and all words had six to seven letters and exactly two syllables. Words were additionally controlled to be of neutral emotional valence and low arousal using the Affective Norms for English Words (ANEW; Bradley & Lang, 1999). 21 words were removed manually due to possible confounding effects (e.g., 'reward', 'defeat', 'profit') or because they were deemed by the authors to be emotional in nature but were not included in ANEW (e.g., 'terror', 'regret'). The final word pool consisted of 160 words.

160 non-words were generated with the LINGUA non-word generator

(Westbury et al., 2007). Non-words were generated using a Markov chaining length of three. Half of the non-words were generated to be six letters in length, with the remaining half being seven letters in length, in order to match the length of the non-words to the words.

Procedure

Prior to the experiment, participants were informed that the experiment was a 'word choice task,' and that they would receive an honorarium proportional to the total points earned in the value-learning phase of the experiment, in addition to their partial course credit.

The experiment consisted of four sequential tasks: value-learning, lexical decision, free recall, and a value judgement task. Participants were not provided with details of the subsequent task until the current task was completed.

Value learning. Participants were shown two words on the computer screen simultaneously. Participants were instructed to choose one of the two words in each choice set. Participants pressed the 'Z' key of a computer keyboard to choose the word presented on the left side of the computer screen; to choose the word on the right side of the screen they were instructed to press the '/' key.

For each participant, 36 words were randomly selected from our pool of 160 words and randomly assigned to one of three reward levels: 2, 7, or 12 points. Thus, assignment of words to reward values varied across participants. Sets were pseudorandomly generated each round to never pair two words of the same reward level and to present pairings of all possible reward levels an equal number of times. This constraint was not revealed to the participant. Each choice by the participant immediately resulted in earning the respective reward. When a choice was made by the participant, an image of a pile of coins was shown to the participant in the centre of the computer screen for 2000 ms, where the number of coins in the image was directly proportional to the number of points earned (e.g., if the participant earned 12 points, the image displayed a pile of 12 coins). The participant's current point balance was continually presented at the bottom of the computer screen throughout the duration of the value-learning task. There was no time limit on how quickly the choices had to be made and participants were given a 1000 ms delay before the next choice was presented.

Training consisted of 18 choice sets per round for 13 rounds. At the end of the session, participants were paid \$1.00 for every 400 points earned during the value-learning task, rounded up to the nearest 25 cent amount. All participants earned between \$3.25 and \$4.50.

Lexical decision. 12 additional words, selected at random from the same pool as the previously rewarded words, were included as 'new' words. Participants were then asked to judge the lexical status of 96 items: 36 previously rewarded words, 12 new words and 48 non-words. Each item was presented for up to 10,000 ms, and the participant pressed either 'Z' on the computer keyboard to indicate that the item was a word, or '/' to indicate that the item was a non-word. A fixation cross ('+') was presented for 1000 ms to separate each decision prompt.

The 96 items were preceded by an additional 8 items (four words/four non-words). These 8 items were presented prior to the 96 items to attenuate a possible recency effect over the last words from the preceding value-learning task. These four words were not presented in the value-learning task and performance on these 8 items was not included in the analyses.

Free recall. In a surprise free recall task, participants were given 5 minutes to recall all of the words they could remember from the task. Participants were given a maximum of 45 s for each response. After each response, a blank screen was presented for 500 ms. Misspellings or variants of the correct word were scored as incorrect responses. Repetitions of correct responses were ignored.

Value judgement. To obtain an objective measure of a participant's explicitly learned reward value information, we included a value judgement task. At the end of the experiment, participants were presented with each of the words previously shown in the reward training phase, one at a time, and asked to judge how many points each word was worth from the initial reward training phase. Participants were reminded of the three possible reward levels and asked to type the number of points they thought the presented word was worth.

Data Analysis

All analyses are reported with Greenhouse-Geisser correction for non-sphericity where appropriate. Effects were considered significant based on an alpha level of 0.05 and post-hoc pairwise comparisons were always Bonferroni-corrected. Non-significant 'trend' effects ($p < .1$) are also reported.

For response time analyses, only correct responses were analyzed. In the lexical decision phase, only responses made between 200 ms and the individual participant's mean plus three standard deviations were included in the analysis. As response time distributions were not normally distributed, response times were log-transformed to accommodate parametric statistics.

As our hypotheses concern the relationship between reward level and memory, we tested for significant linear and quadratic effects of reward level. The reward value hypothesis would suggest a significant linear component. On the other hand, the reward salience hypothesis would suggest a significant quadratic component, corresponding to a U-shaped relationship.

3.3.2 Results

Value-learning

Accuracy in the value-learning task was measured as how often the participant chose the higher-value word. Performance in the training task began at chance, as the participant could not know which was the higher-value word. In the last round of the training task, accuracy was significantly greater than chance [$M = 77.9\%$ correct; $t(70) = 15.20$, $p < .001$]. In choice sets, the difference in values between the items within the set (V_{diff}) could be either 5 points (choice between a 2-point item and a 7-point item, or a 7-point item and a 12-point item) or 10 points (choice between a 2-point item and a 12-point item). At both value differences, participants were more likely than chance to choose the higher-value item [$V_{diff}(5) : M = 74.4\%$, $t(70) = 12.00$, $p < .001$; $V_{diff}(10) : M = 85.0\%$, $t(70) = 16.72$, $p < .001$] (see Figure 3.1a).

To directly examine if higher-value items were chosen more, we also tested for a relationship between choice frequency (the number of times an item was chosen) and reward level. A one-factor repeated-measures ANOVA, revealed a main effect of reward level on choice [$F(2, 140) = 204.05$, $p < .001$]. Pairwise comparisons found all differences to be significant, such that 2-point words < 7-point words < 12-point words [all p's $< .001$] (Figure 3.2a).

To keep analyses comparable between our value-learning and memory measures, we also tested for significant linear and quadratic components in the relationship between choice and reward level. We found a significant linear component of reward level [$F(1, 70) = 321.37$, $p < .001$]. The quadratic component was not significant [$F(1, 70) = 1.51$].

Lexical decision

Lexical decision accuracy was near ceiling, as expected, and was not significantly different between value conditions [$F(2, 140) = 0.11$]. Accuracy was higher for the previously rewarded 'old' words than for the untrained, 'new' words ['old' words: 98.2% correct; 'new' words: 91.3% correct; $t(70) = 5.11$, $p < .001$].

Participants also identified the 'old' words significantly faster than the 'new' words [$t(70) = 5.75$, $p < .001$]. No significant differences were found between lexical decision response times across the three reward levels [$F(2, 140) = 1.97$], although we found a trend quadratic effect of reward level [$F(1, 70) = 3.05$, $p < .10$]. The linear component was not significant [$F(1, 70) = 0.93$].

Free recall

As expected, participants recalled fewer of the untrained, new words than the previously rewarded words (Figure 3.2b). For the previously rewarded words we conducted a one-factor repeated-measures ANOVA (reward level) and found a main effect [$F(2, 140) = 11.87$, $p < .001$]. Pairwise comparisons

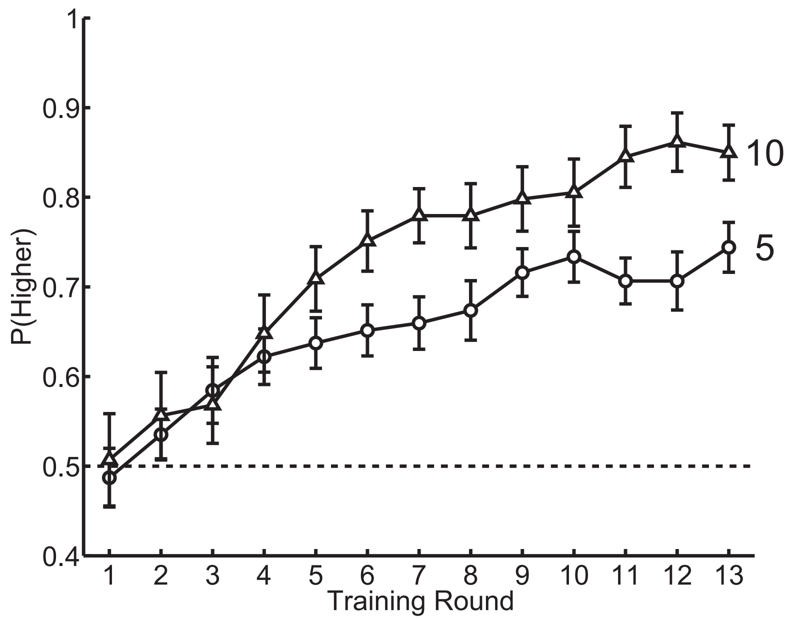

(a) Experiment 1

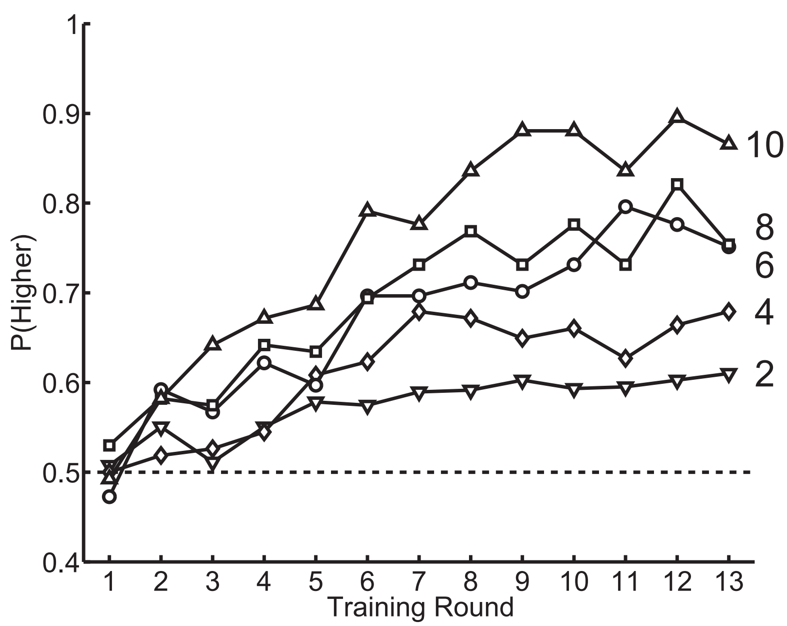

(b) Experiment 2

Figure 3.1: Percentage of trials in which the higher-value item was chosen in each round of the value-learning task, separated based on the difference in reward value between the two items, for (a) Experiment 1 and (b) Experiment 2. Error bars are 95% confidence intervals, corrected for inter-individual differences. Due to the number of reward levels, error bars were omitted from panel (b) for visual clarity.

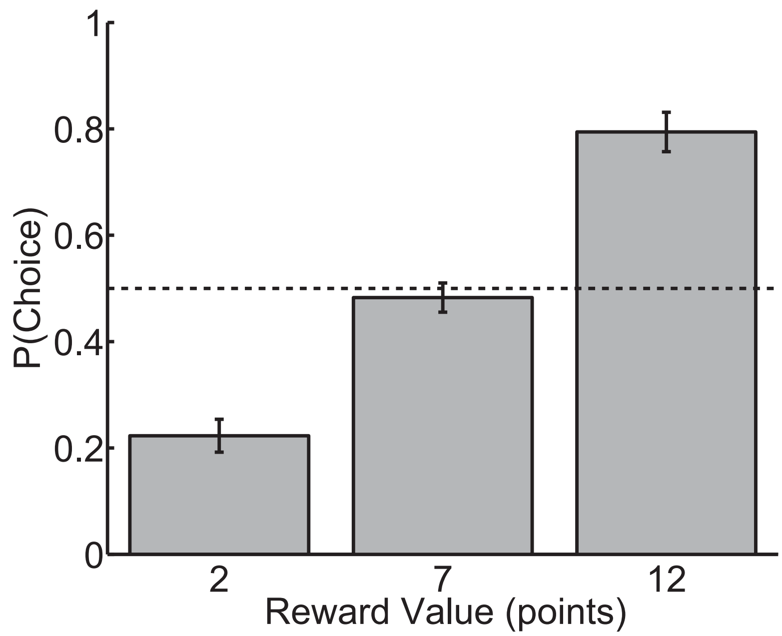

(a) Choice frequency

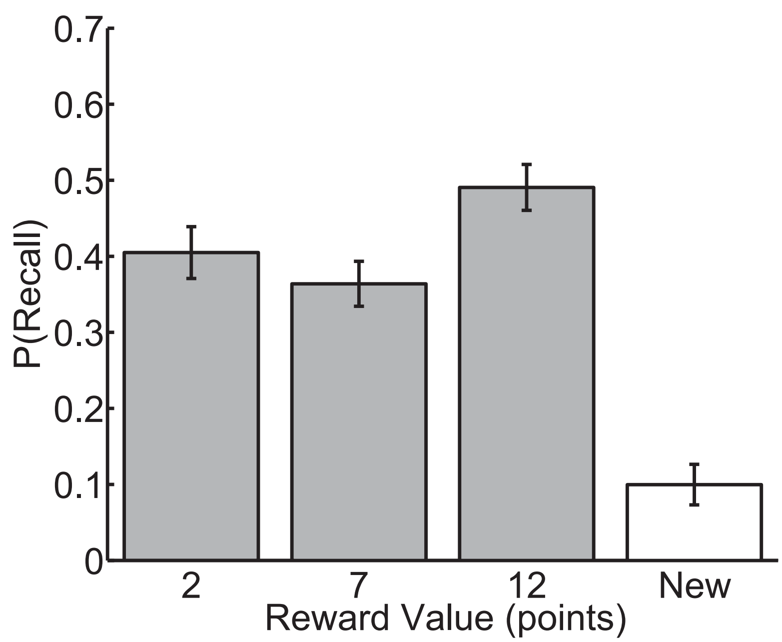

(b) Recall rates

Figure 3.2: Performance as a function of reward value for (a) choice frequency in the last three rounds of value-learning and (b) recall rates in free recall, in Experiment 1. Error bars are 95% confidence intervals, corrected for inter-individual differences.

found that significantly more 12-point words were remembered than 2-point and 7-point words [both p's $< .01$]. However, there was no significant difference between recall performance for 2-point and 7-point words [$t(70) = 1.44$]. We also found a significant linear effect of reward level [$F(1, 70) = 9.80$, $p < .01$], as well as a significant quadratic effect [$F(1, 70) = 14.21$, $p < .001$].

Value judgement

Participants correctly identified the value of the previously rewarded words at levels better than chance (33.3%) [$M = 61.4\%$ correct; $t(70) = 15.03$, $p < .001$]. To illustrate the responses in this task, we plotted the proportion of value judgement responses for each reward level, separated based on the actual reward level of the item (Figure 3.3a).

To further examine the relation between memory and value, we compared accuracy on the value judgement task for words that were recalled in free recall relative to words that were not recalled. Value was found to be judged better for words that were recalled than for words that were not recalled [$M_{recalled} =$ 64.4% correct; $M_{not-recalled} = 57.8\%$ correct; $t(70) = 3.00$, $p < .01$]. This supports the notion that value is learned as a property of the words themselves, and memory for value is less accurate when the items are also not remembered as well.

3.3.3 Discussion

Our finding of a significant quadratic component in the effects of reward level on free recall performance is consistent with the reward salience hypothesis. However, with only three reward levels, we are unable to statistically differentiate between linear and quadratic effects and conclusively state that one effect is more prominent than the other. Experiment 2 addressed this issue by using six different reward levels (2, 4, 6, 8, 10, and 12 points). The greater number of separable reward levels provided better resolution of reward as a continuous measure and allowed better comparison of the relative fits of both linear and quadratic models.

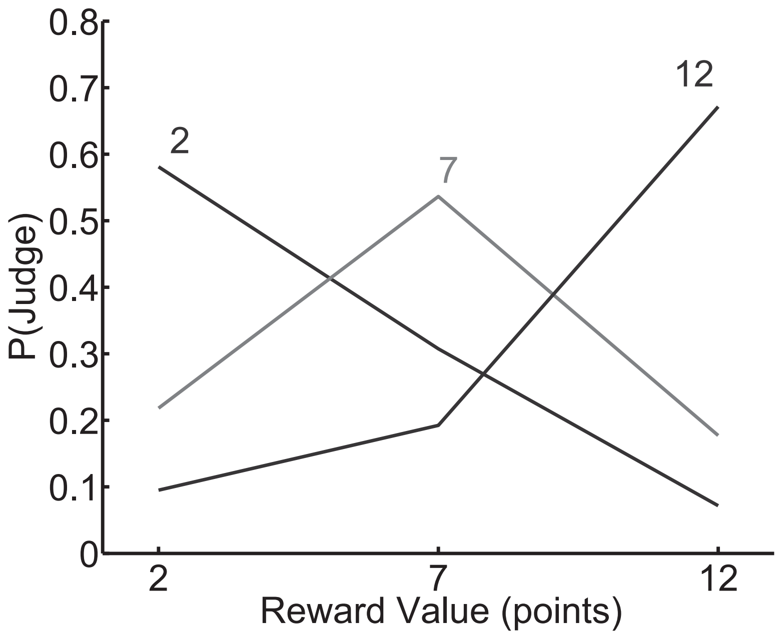

(a) Experiment 1

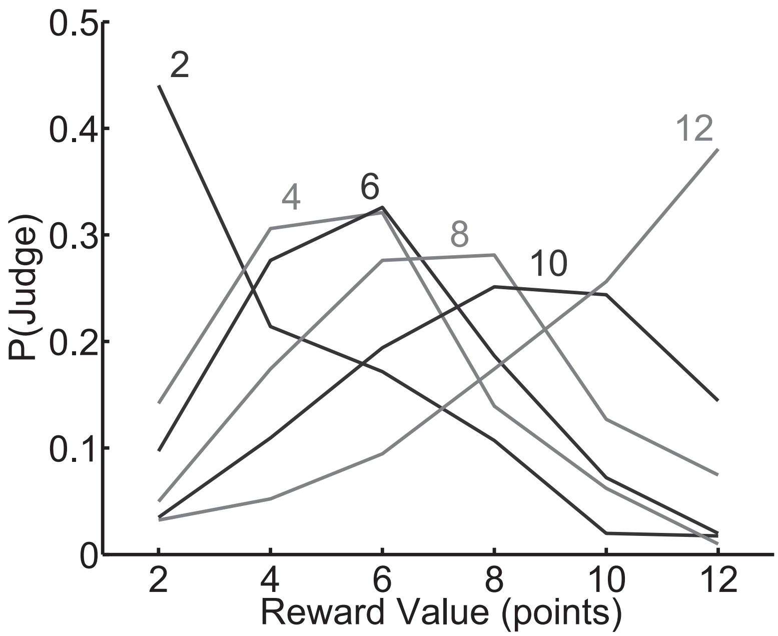

(b) Experiment 2

Figure 3.3: Proportion of value judgement responses for each reward level, separated based on the actual reward level of the item, for (a) Experiment 1 and (b) Experiment 2. Error bars were omitted from for visual clarity.

3.4 Experiment 2

3.4.1 Methods

Participants

67 introductory psychology students (48 female) at the University of Alberta participated for partial fulfillment of course credit. Restrictions for participating and informed consent were identical to Experiment 1. None of the participants from Experiment 1 participated in Experiment 2.

Materials

The same word and non-word pools were used as in Experiment 1.

Procedure

Instead of the three different reward levels used in Experiment 1, here we used six reward levels: 2, 4, 6, 8, 10, and 12 points. To keep the total number of previously rewarded words at 36, only six words were assigned to each reward level. All participants earned between $4.00 and $5.00 in the value-learning task. Additionally, only 6 new words were used in lexical decision task, rather than 12, to equate the number of new words with the words in any single reward level. The rest of the procedure remained the same as in Experiment 1.

3.4.2 Results

Value learning

In the last round of the training task, accuracy was significantly greater than chance [M = 67.9% correct; $t(66) = 9.46$, $p < .001$]. In choice sets, the difference in values for each item (V_{diff}) could be 2, 4, 6, 8, or 10 points. At all value differences, participants were more likely than chance to choose the higher-value item [$V_{diff}(2)$: $M = 61.0\%$, $t(66) = 5.32$, $p < .001$; $V_{diff}(4)$: $M = 67.9\%$, $t(66) = 5.79$, $p < .001$; $V_{diff}(6)$: $M = 75.1\%$, $t(66) = 7.86$, $p < .001$; $V_{diff}(8)$: $M = 75.4\%$, $t(66) = 7.08$, $p < .001$; $V_{diff}(10)$: $M = 86.6\%$, $t(66) = 8.71$, $p < .001$] (see Figure 3.1b).

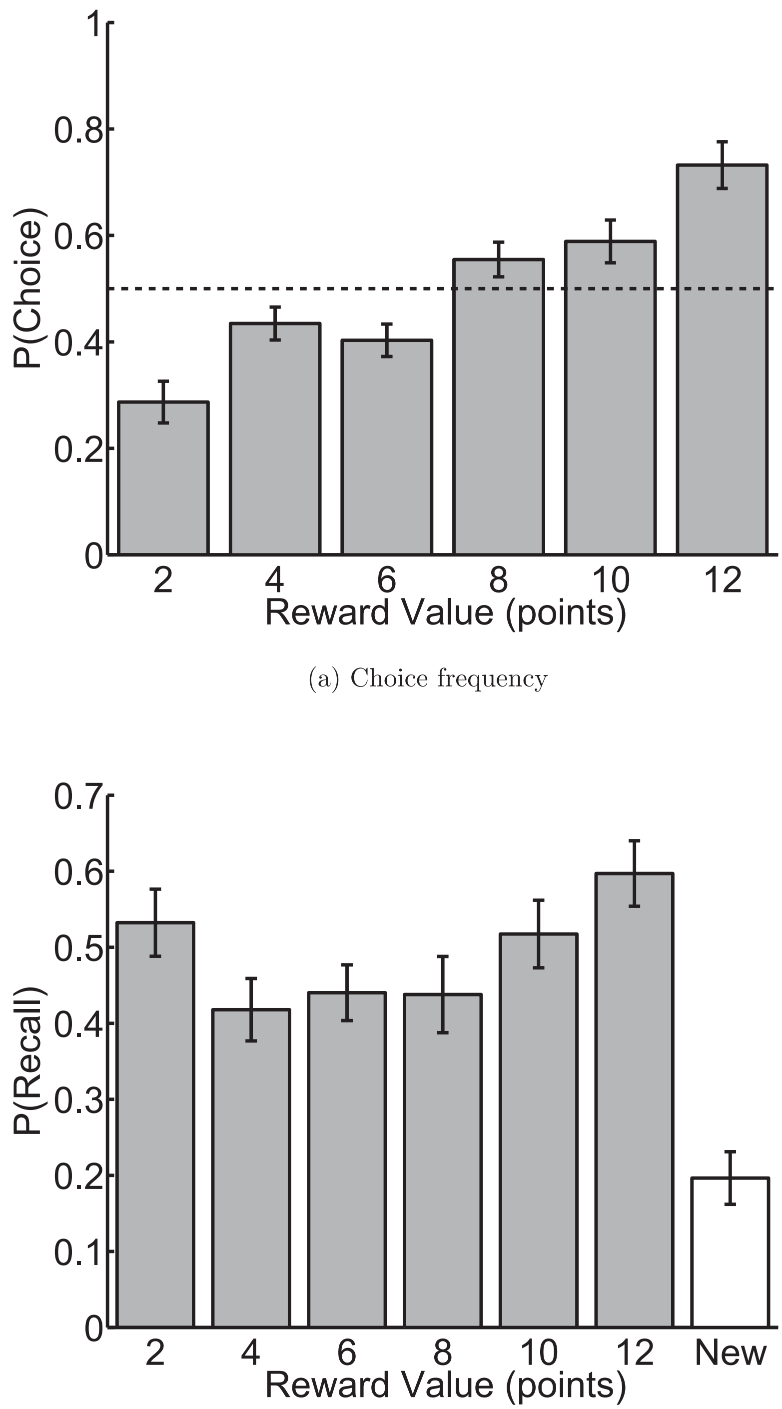

(a) Choice frequency

(b) Recall rates

Figure 3.4: Performance as a function of reward value for (a) choice frequency in the last three rounds of value-learning and (b) recall rates in free recall, in Experiment 2. Error bars are 95% confidence intervals, corrected for inter-individual differences.

A one-factor repeated-measures ANOVA revealed a main effect of reward level on choice [$F(4, 264) = 59.39$, $p < .001$]. Pairwise comparisons found all differences to be significant [all p's $< .001$], except for the difference between 4-point and 6-point words (Figure 3.4a). Similar to Experiment 1, we found a significant linear component of reward level [$F(1, 66) = 158.62$, $p < .001$]. The quadratic component was again not significant [$F(1, 66) = 2.48$].

Lexical decision

Lexical decision accuracy was near ceiling, as expected, and did not differ significantly between value conditions [$F(4, 269) = 0.56$]. Accuracy was higher for the previously rewarded 'old' words than for the untrained, 'new' words ['old' words: 98.8% correct; 'new' words: 94.0% correct; $t(66) = 4.21$, $p < .001$].

Participants also identified the 'old' words significantly faster than the 'new' words [$t(66) = 6.73$, $p < .001$]. No significant differences were found between lexical decision response times across the six reward levels [$F(4, 250) = 1.37$]. We found neither a significant linear effect of reward level [$F(1, 66) = 0.55$] nor a quadratic effect [$F(1, 66) = 0.75$] (see Figure 3.5).

Free recall

Recall was again higher for the previously rewarded words than for the untrained, new words (Figure 3.4b). For the previously rewarded words, we again conducted a one-factor repeated-measures ANOVA (reward level) and found a main effect [$F(5, 305) = 8.42$, $p < .001$]. Pairwise comparisons revealed no significant differences between 2-, 10-, and 12-point words [$p > .5$]. Importantly, more 2-point words were recalled than 4-point words [$p < .05$], and more 12-point words were recalled than 4-, 6-, or 8-point words [all p's $< .001$]. We also found a significant linear effect of reward level [$F(1, 66) = 12.27$, $p < .01$], as well as a significant quadratic effect [$F(1, 66) = 27.37$, $p < .001$].

To better test our hypotheses, we fit constant, reward value (linear), reward salience (quadratic), and reward value + salience (linear and quadratic) models to our free recall data. The constant model assumed that reward had

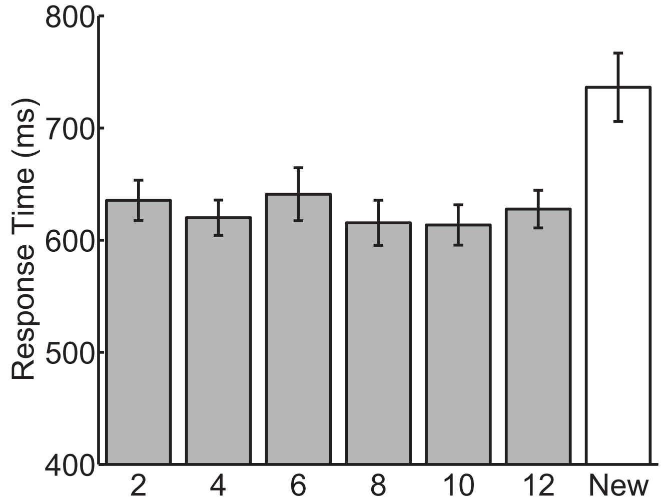

Figure 3.5: Response times from the lexical decision task in Experiment 2. Error bars are 95% confidence intervals, corrected for inter-individual differences.

no effect on memory performance and thus that recall should be equivalent regardless of reward level and only contained one parameter. The reward value model assumed that recall performance is monotonically related to reward value (e.g., higher value items are remembered better than lower-value items). This model contained two parameters: a term that varied monotonically with reward value and a constant. The reward salience model assumed that recall performance follows a U-shaped (quadratic) function with relation to reward value. In other words, high-salience items (those at the extremes of the range of values experienced) are remembered better than low-salience items (those in the middle of the range). This model contained two free parameters: a term that varied quadratically with reward value and a constant. This function was shifted along the x-axis such that the minima of the function occurred at the middle of the range of values experienced (i.e., 7 points). The reward value + salience model was based on the reward salience model, but additionally incorporated a parameter that shifted the function to best overlap with the data (i.e., the reward value where function's minima occurred). The reward value + salience model therefore allows reward value, in addition to reward

Table 3.1: Fitness measures for the model fits to the free recall data of Experiment 2.

Model	*df*	ΔBIC	R^2	$RMSD$
Constant	1	4.25	.00	.064
Reward Value	2	5.28	.22	.056
Reward Salience	2	2.41	.70	.035
Reward Value + Salience	3	0.00	.93	.017

salience, to be a predictor of item-memory.

To compare the different models, we calculated the BIC (Bayesian Information Criterion) as a measure of model fitness that takes into account the number of free parameters (see Table 3.1). By convention, if the difference between two model fits is less than two, neither of the models' fit to the data is significantly better — thus we report all scores as ΔBIC relative to the best-fitting model (Burnham & Anderson, 2002). We also report R^2 and $RMSD$ (root mean squared deviation) as additional measures of model fitness. Note that a best fitting model would be characterized by low ΔBIC, low $RMSD$, and high R^2 values. As evident in Figure 3.6, the reward value + salience model is the best-fitting model, a conclusion that was supported by all three of our quantitative model fitting measures ($R^2 = .93$). More specifically, it appears that this model was primarily characterized by reward salience ($R^2 = .70$), along with a much smaller contribution from reward value ($R^2 = .22$).

Value judgement

Participants correctly identified the value of the previously rewarded words at levels better than chance (16.7%) [$M = 33.0\%$ correct; $t(66) = 11.84$, $p < .001$]. As in Experiment 1, we plotted the proportion of value judgement responses for each reward level, separated based on the actual reward level of the item (Figure 3.3b). Here it is evident that even when participants were incorrect at judging the value, they often responded with an adjacent

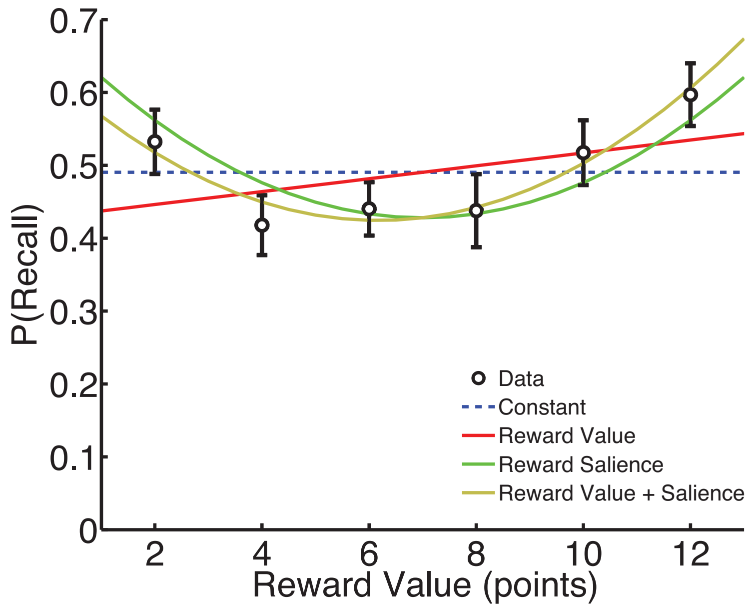

Figure 3.6: Model fits to the free recall data from Experiment 2 (see Table 1 for model fitness measures). Error bars are 95% confidence intervals, corrected for inter-individual differences.

value (e.g., judging an 6-point item as being worth either 4- or 8-points). We also observe a large degree of overlap in value judgement responses for items worth 4-, 6-, and 8-points, converging with the clustering observed in memory performance (Figure 3.4b).

As in Experiment 1, we compared accuracy on the value judgement task for words that were recalled in free recall relative to words that were not recalled. Value was found to be judged better for words that were recalled than for words that were not recalled [$M_{recalled} = 37.0\%$ correct; $M_{not-recalled} = 27.5\%$ correct; $t(66) = 5.28$, $p < .001$]. Here we again find evidence that our procedure may cause value to be learned as an attribute of the word itself.

Ruling out simple decision heuristics

One possible explanation for a U-shaped function is that participants relied on simple decision heuristics of choosing the highest-valued item and avoiding the lowest-valued item. However, two sources of evidence argue against the possibility that participants relied primarily on these simple heuristics. First, the value judgment results presented in Figure 3.3b show that even for the intermediate items, the value was discriminated from some of the other intermediate items and not just from the two extreme items (i.e., participants rarely judged a 4-point item as having a value of 10 points or vice versa). Second, for the 2-point difference decisions, reliance on a simple choose/avoid heuristic would result in highly accurate choices for 2-point versus 4-point and for 10-point versus 12-point choice sets, but much lower accuracy for the intermediate choice sets, such as 6-point versus 8-point. To test this empirically, we conducted a one-factor repeated-measures ANOVA on the compared choice sets for the last four rounds of training (i.e., when participants had largely learned the values). Our dependent measure was the proportion of choices for which participants correctly chose the higher value item when the difference in value of the two items was only 2-points. We did not find a main effect, suggesting that there were no significant differences in accuracy between choice sets [$F(3, 229) = 2.42$] (Figure 3.7). Thus, even though the words associated with extreme values were remembered better, this enhancement did not appear to be caused by a simple decision heuristic based on only the extreme values.

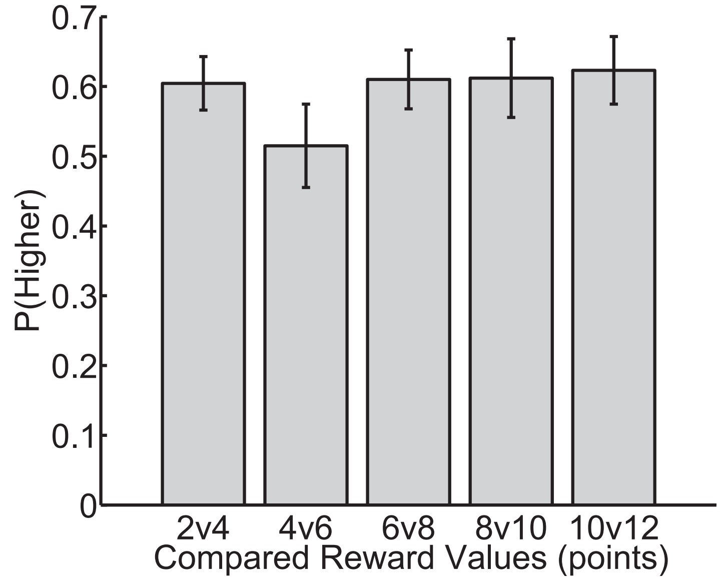

Figure 3.7: Accuracy across the last four rounds of the value-learning task in Experiment 2, for choice sets where the difference in value was only 2-points. Error bars are 95% confidence intervals, corrected for inter-individual differences.

3.5 General Discussion

In the present study we provide clear evidence that memory for items does not correspond monotonically with increases in reward value. Instead, we find that the most and least rewarding items are remembered best and that memory follows a U-shaped function (see Figure 3.6), suggesting an enhancement of memory driven by reward salience. Thus, we provide evidence that free recall performance is not driven solely by reward value or the number of times an item was chosen (choice frequency), as these accounts would both correspond to a monotonic (linear) relationship between reward value and memory.

Studies have shown that positive and negative values are learned through different neural substrates (e.g., Yacubian et al., 2006). In the present study, however, we found an enhancement of memory for extreme values that did not necessitate the presence of both positive and negative values, as all of our values were positive. That is, the lowest value items were remembered better than intermediate-level items, even though our lowest value items were near zero in absolute value. Our results therefore suggest that reward salience is relative to the *range of values experienced*, and is not necessarily driven by the use of positive and negative values (e.g., Cooper & Knutson, 2008) or appetitive and aversive stimuli (e.g., Litt et al., 2011). Although the significant linear effect of reward value on free recall performance suggests that increases in value do enhance memory, this effect was much weaker than the quadratic component and explained less of the variance. Thus, memory in free recall was explained better by the reward salience of the stimuli than by their reward value alone.

Our evidence suggesting that reward salience is driven by the range of values experienced, may also be related to another psychological mechanism: the anchoring effect. The anchoring effect suggests that the ends of the stimulus continuum play an important role in judgements of absolute value (e.g., Eriksen & Hake, 1957) and that the anchoring effect is driven by memory for the end-points (Petrov & Anderson, 2005; Weber & Johnson, 2006). Anchoring effects could also be at play in our value judgement task (see Figure 3.3). While it can be argued that temporal effects (i.e., primacy and recency) in memory are the result of anchoring effects, they have not previously been considered within a value-learning procedure. Thus, it is possible that reward salience is intertwined with the anchoring effect.

Though prior studies have utilized fMRI in conjunction with reward-based memory paradigms, none of these studies utilized multiple levels of reward and are thus unable to test a reward saliency hypothesis. These studies have implicated the activation of several reward-related brain regions as predictors of which items will be later remembered, including the ventral tegmental area, striatum, substantia nigra, and orbitofrontal cortex (Adcock et al., 2006; Shigemune et al., 2010; Wittmann et al., 2005). Of these regions, the orbitofrontal cortex has been related to reward value (Jensen et al., 2007; Litt et al., 2011). However, activation in the striatum has been associated with either reward salience (Cooper & Knutson, 2008; Jensen et al., 2007; Zink et al., 2004) or both reward value and reward salience (Litt et al., 2011). Currently, none of the fMRI studies investigating reward salience mention the substantia nigra or the ventral tegmental area. Based on the results of studies investigating reward salience, the striatum appears to be the reward-related brain region most likely to be responsible for the memory results in the present study. Specifically, activations in the striatum correspond to a combination of reward value and reward salience (e.g., see Litt et al., 2011, Figure 3.4c), with a stronger contribution of reward salience, similar to the results of our model fit. Nonetheless, the neural underpinnings of our result are open to debate and will likely be the focus of future research.

Additionally, although we found an effect of reward on explicit memory in our free recall task, we did not observe an effect of reward on implicit memory in our lexical decision task. Although an enhancement of memory due to reward in lexical decision has been reported previously (Madan, Fujiwara, et al., 2012), the current study, which had fewer data points per reward level for each participant, may have been less sensitive to subtle enhancements of implicit priming due to reward. Nonetheless, we did observe a trend quadratic effect of reward on lexical decision response times in Experiment 1, suggesting that the enhancement of implicit memory due to reward seen in the previous study may also be driven in part by reward salience.

3.6 Conclusion

In two experiments we demonstrate that the enhancement of memory due to reward is driven not only by reward value, but also by reward salience. Most previous studies that suggested a monotonic influence of reward value on memory used only two levels of reward and thus were unable to capture the full relationship between reward and memory. Through the addition of intermediate reward levels we are able to determine that memory is enhanced for both the highest- and lowest-value items. This U-shaped (quadratic) relationship between value and memory occurred even though all reward values were positive and is best characterized as an enhancement of memory due to reward salience.

3.7 Acknowledgements

We would like to thank Elliot Ludvig and Eric Legge for insightful feedback in the design and analysis of this study, as well as Esther Fujiwara and Jeremy Caplan for feedback on an earlier draft of the manuscript. This research was partly funded by a Discovery grant from the National Science and Engineering Research Council of Canada held by MLS and by a scholarship from the Canadian Institutes of Health Research to CRM.

3.8 References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings* (Technical report No. C-1). Gainesville, FL: The Center for Research in Psychophysiology, University of Florida.

Burnham, K. E., & Anderson, D. R. (2002). *Model selection and multimodel inference* (2nd ed.). New York: Springer-Verlag.

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*, 1078–1085.

Cooper, J. C., & Knutson, B. (2008). Valence and salience contribute to nucleus accumbens activation. *NeuroImage*, *39*, 538–547.

Eriksen, C. W., & Hake, H. W. (1957). Anchor effects in absolute judgments. *Journal of Experimental Psychology*, *53*, 132–138.

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

Jensen, J., Smith, A. J., Willeit, M., Crawley, A. P., Mikulis, D. J., Vitcu, I., & Kapur, S. (2007). Separate brain regions code for salience vs. valence during reward prediction in humans. *Human Brain Mapping*, *28*, 294–302.

Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron*, *46*, 703–713.

Litt, A., Plassmann, H., Shiv, B., & Rangel, A. (2011). Dissociating valuation and saliency signals during decision-making. *Cerebral Cortex*, *21*, 95–102.

Ludvig, E. A., Bellemare, M. G., & Pearson, K. G. (2011). A primer on reinforcement learning in the brain: Psychological, computational, and neural perspectives. In E. Alonso & E. Mondragon (Eds.), *Computational neuroscience for advancing artificial intelligence: Models, methods and applications* (pp. 111–144). Hershey, PA: IGI Global.

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.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042–1045.

Petrov, A. A., & Anderson, J. R. (2005). The dynamics of scaling: A memory-based anchor model of category rating and absolute identification. *Psychological Review*, *112*, 383–416.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement.

In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning ii* (pp. 64–99). New York: Appleton-Century-Crofts.

Shigemune, Y., Abe, N., Suzuki, M., Ueno, A., Mori, E., Tashiro, M., … Fujii, T. (2010). Effects of emotion and reward motivation on neural correlates of episodic memory encoding: A PET study. *Neuroscience Research*, *67*, 72–79.

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Science*, *14*, 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*, 1236–1242.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press.

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memory. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

Westbury, C., Hollis, G., & Shaoul, C. (2007). LINGUA: The Language-Independent Neighbourhood Generator of the University of Alberta. *The Mental Lexicon*, *2*, 273–286.

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*, 459–467.

Yacubian, J., Gläscher, J., Schroeder, K., Sommer, T., Braus, D. F., & Büchel, C. (2006). Dissociable systems for gain and loss related value predictions and errors of prediction in the human brain. *Journal of Neuroscience*, *26*, 9530–9537.

Zink, C. F., Pagnoni, G., Martin-Skurski, M. E., Chappelow, J. C., & Berns, G. S. (2004). Human striatal response to monetary reward depend on saliency. *Neuron*, *42*, 509–517.

Chapter 4

Extreme outcomes sway risky decisions from experience

A version of this work was previously published as: Ludvig, E. A.*, **Madan, C. R.***, & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making, 27*, 146–156. doi:10.1002/bdm.1792. (* – Authors contributed equally.) This work has been reproduced with permission.

4.1 Abstract

Whether buying stocks or playing the slots, people making real-world risky decisions often rely on their experiences with the risks and rewards. These decisions, however, do not occur in isolation, but are embedded in a rich context of other decisions, outcomes, and experiences. In this paper, we systematically evaluate how the local context of other rewarding outcomes alters risk preferences. Through a series of 4 experiments on decisions from experience, we provide evidence for an extreme-outcome rule, whereby people overweight the most extreme outcomes (highest and lowest) in a given context. As a result, people should be more risk seeking for gains than losses, even with equally likely outcomes. Across the experiments, the decision context was varied so that the same outcomes served as either the high extreme, low extreme, or neither. As predicted, people were more risk seeking for relative gains, but only when the risky option potentially led to the high-extreme outcome. Similarly, people were more risk averse for relative losses, but only when the risky option potentially led to the low-extreme outcome. We conclude that, in risky decisions from experience, the biggest wins and the biggest losses seem to matter more than they should.

4.2 Introduction

Many behavioral economic studies on risky decisions present people with scenarios in which the outcomes and their probabilities are explicitly described (e.g., Kahneman & Tversky, 1979). For example, people might be explicitly asked whether they would prefer a guaranteed $20 or a 50/50 chance at $40. When faced with these risky decisions from description, people are usually risk averse for gains and risk seeking for losses—a pattern of risk preference known as the reflection effect. In life, however, people often make economic decisions based on their past experience with the consequences of those decisions. People frequent certain stores but not others, decide which products to buy, and risk whether or not to pay for the parking meter for a short dash into the store—all based in part on their own experiences. These decisions from experience can sometimes lead to markedly different behavior than those from description (e.g., Barron & Erev, 2003; Camilleri & Newell, 2011; Hertwig, Barron, Weber, & Erev, 2004; Hertwig & Erev, 2009; Ludvig & Spetch, 2011; Weber, Shafir, & Blais, 2004).

The most prominent difference between description and experience occurs when one outcome is relatively rare. People overweight rare outcomes in description whereas they underweight rare events in experience—often known as the description-experience (DE) gap. In the absence of rare events, there are usually no systematic differences between the described and experienced cases (e.g., Erev et al., 2010). In a recent study, however, we found exactly such a DE gap. Using 50/50 outcomes (no rare events), we found a clear reversal of the reflection effect in experience, but not in description (Ludvig & Spetch, 2011). People were more risk seeking for gains than losses in experience, but, conversely, were more risk seeking for losses than gains in description. In those experiments, unlike many studies that examine the DE gap, we used a within-subject design (but see Camilleri & Newell, 2009). This design intermingled decisions from experience and descriptions as well as those between gains and losses, suggesting that perhaps the decision context is crucial for determining the pattern of risky choice (Ludvig & Spetch, 2011). In this paper, we focus solely on decisions from experience and present a series of experiments that systematically evaluates how the decision context influences this experience-based risky choice.

In perception, context effects abound: A surface can appear different colors because of the surrounding colors (e.g., Lotto & Purves, 2000), and the apparent length of a line can depend on the direction of the arrowheads (Müller-Lyer, 1889). Similarly, in choice, the local context in which a decision is made can greatly influence that decision (Simonson, 1989; Simonson & Tversky, 1992). For example, when choosing between two items, people's preference between the two options can be altered by introducing a third, seemingly irrelevant, option. If the third option is similar, but slightly worse than one of the original two options, this change will often lead to an increase in preference for the similar option (Heath & Chatterjee, 1995; Huber, Payne, & Puto, 1982). This preference bump occurs even though nothing has changed about the original two options except for the surrounding context of other options. Risky decisions can also be influenced by how the decision is framed (Tversky & Kahneman, 1981) or the alternative options available (Erev, Glozman, & Hertwig, 2008; Stewart, Chater, Stott, & Reimers, 2003). Thus, as with psychophysical judgments (e.g., Helson, 1947; Thomas & Jones, 1962), risky decisions can be context-dependent and altered by the comparison set experienced.

In the present studies, we examined how the decision context in which repeated choices are made affects risky decisions from experience. Specifically, we were interested in whether the edges of the experienced distribution (the biggest wins and losses) might be overweighted in decisions from experience. A similar effect, termed the peak-end rule, is observed in retrospective judgments of affective experiences. The post-hoc valuation depends primarily on the point of maximum intensity (the peak) and on the end (Fredrickson, 2000). In a classic illustration of this rule, post-operative pain judgments of patients undergoing a colonoscopy were strongly correlated with the peak intensity of pain (as judged in real time) and with the pain intensity at the end, but not with the duration of the procedure (Kahneman et al., 1993; Redelmeier &

Kahneman, 1996; Stone et al., 2005). Delayed judgments relating to positive experiences, such as vacations, are similarly influenced by peak and end intensities (Mitchell et al., 1997).

Inspired by the peak portion of the peak-end rule, we hypothesized that the extreme outcomes in a decision context may be overweighted in decisions from experience, as they are in delayed judgments. Following this *extreme-outcome rule*, when a risky option occasionally leads to the best possible gain in a given context, that extreme gain should be overweighted in the valuation of that risky option. As a result, when pitted against another option of similar expected value, but without the possibility of an extreme gain, that option would be chosen more frequently. Similarly, when a risky option occasionally leads to the worst possible outcome in a context, people should overweight that extreme loss. When pitted against another option of similar expected value, but without the possibility of an extreme loss, that option should be chosen less frequently. Thus, the extreme-outcome rule predicts that people will become more risk seeking for gains than for losses in decisions from experience, but only when the risky choices include the most extreme outcomes in the decision context (i.e., the biggest gain or loss). This prediction about the effects of rewarded experience on subsequent decisions from experiences runs opposite to the usual reflection effect in decisions from description (e.g., Kahneman & Tversky, 1979).

Some evidence for this potential extreme-outcome rule comes from studies with non-human animals, which can only rely on experience for learning about outcomes. Many of these studies have also reported risk seeking for gains (e.g., Hayden, Heilbronner, Nair, & Platt, 2008; Heilbronner & Hayden, 2013; Kacelnik & Bateson, 1996; O'Neill & Schultz, 2010; McCoy & Platt, 2005), and some evidence suggests that this risk seeking for gains may be driven by extreme outcomes in a context. For example, the risk-seeking behavior of rhesus macaques was shown to be sensitive to the magnitude of the jackpot or largest reward on a 2-outcome choice. When the jackpot was reduced, risk seeking declined; when the jackpot was increased, risk seeking increased. An identical manipulation of the smaller, non-extreme reward had no influence on risk preference (Hayden et al., 2008).

In this paper, we present 4 experiments that test the extreme-outcome rule by systematically manipulating the decision context. Figure 4.1 illustrates the basic task (cf. Ludvig & Spetch, 2011). On most trials, people decided between two doors (Fig. 4.1A). Picking one door always led to a fixed outcome, and picking the other (risky) door led with a 50/50 chance to more or less than the fixed outcome. For example, in Experiment 1, the fixed gain door always led to +20, while the risky gain door led to a 50/50 chance of 0 or +40. Conversely, the fixed loss door always led to -20, while the risky loss door led to a 50/50 chance of -40 or 0. In this case, the extreme outcomes were +40 and -40. By the proposed extreme-outcome rule, these highest- and lowest-valued outcomes in the decision context would be overweighted, leading to more risk seeking for gains, yet more risk aversion for losses.

A partial-feedback procedure was used in which participants only saw the outcome for the chosen option, but not the foregone option (see Hertwig & Erev, 2009; Camilleri & Newell, 2011). Outcomes for the forgone option were not included out of concern that people might confuse which option led to which outcome. There were three types of trials. On decision trials (Fig. 4.1A), which always involved a choice between two loss doors or between two gain doors, the objective expected value of the fixed and risky door was equal. Interspersed catch (Fig. 4.1B) and single-choice (Fig. 4.1C) trials ensured that people indeed learned and experienced the correct contingencies. This partial-feedback procedure was designed to highlight the relationship between the option chosen and the outcome received.

The series of four experiments each examined a different facet of the extreme-outcome rule. Table 1 details the exact contingencies in each experiment. Experiment 1 examined decisions from experience with gains and losses intermingled in a single decision context, without concurrent decisions from description (cf. Ludvig & Spetch, 2011). Experiment 2 directly tested the alternative hypothesis that zero values were underweighted rather than extreme values being overweighted. Experiment 3 tested whether all larger magnitude options were overweighted or only extreme ones, by adding into the decision context additional doors that potentially led to more extreme outcomes. Experiment 4 evaluated a novel prediction of the extreme-outcome rule that only relative extremes matter, independent of whether they are absolute gains or losses. In all cases, based on the extreme-outcome rule, we predict more risk seeking for risky options that potentially led to the high extreme in the experiment and, conversely, less risk seeking for risky options that potentially led to the low extreme.

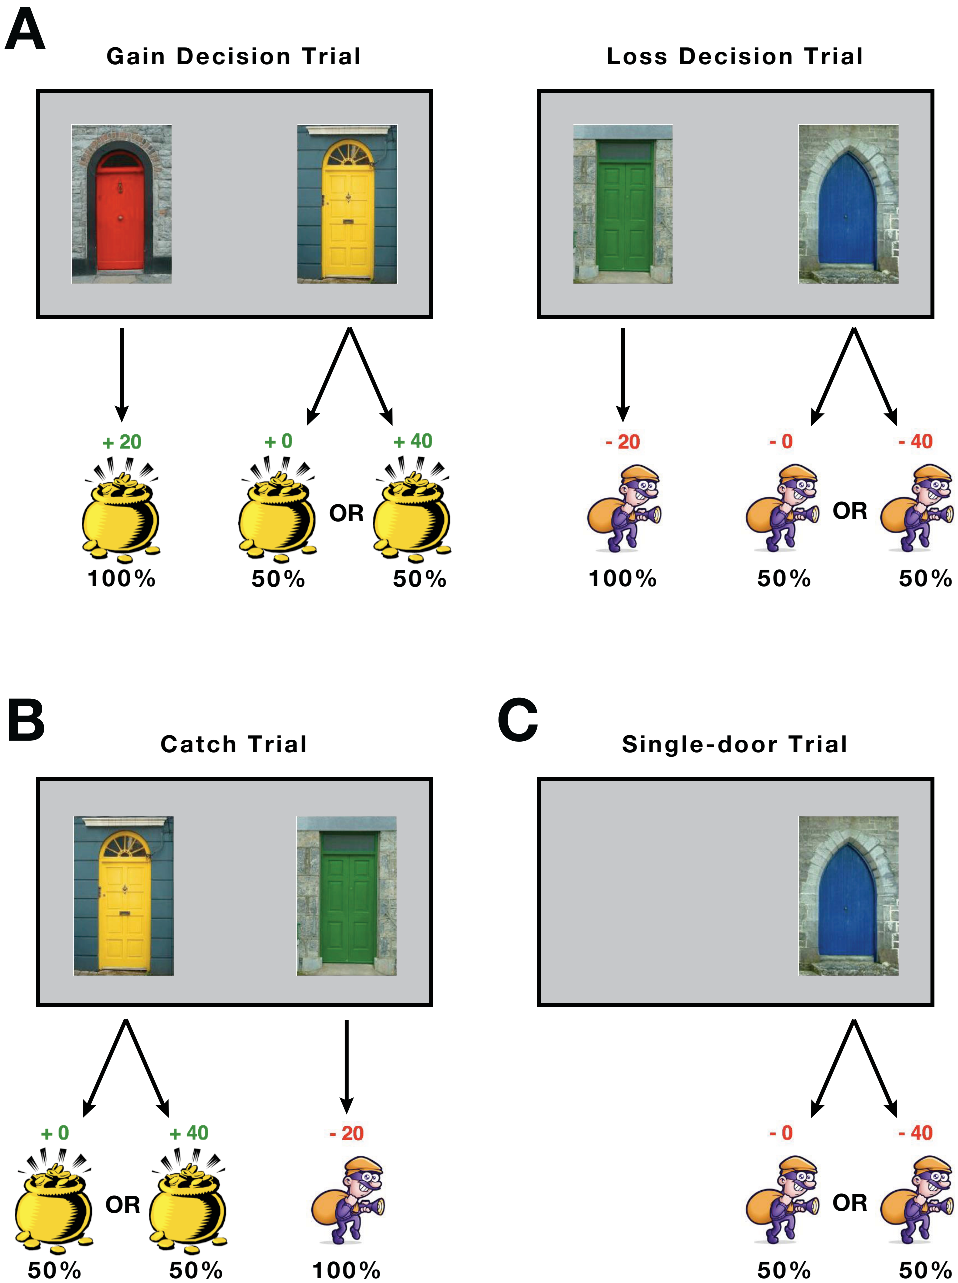

Figure 4.1: Schematic of the general method used. Specific values correspond to Experiment 1. A. Decision trials involved choices between 2 gain doors or 2 loss doors. One door always led to a gain (or loss) of a fixed number of points, and the other door led equiprobably to one of 2 possible outcomes. Choices were followed by feedback about the amount gained or lost. B. Catch trials involved choices between one gain door and one loss door and ensured that participants paid attention to their decisions. C. Single-door trials only presented one door and ensured that participants occasionally experienced the planned reward contingencies.

Table 4.1: Details of all experimental manipulations and summary of main results. Bold font indicates extreme outcome in an experiment. X = Extreme outcome. NX = No Extreme. HX = High Extreme. LX = Low Extreme. BX = Both Extremes.

Experiment	Manipulation	Decision Type	Fixed Outcome	Risky Outcomes (50/50)	Degree of Risk Seeking
1	Mixed Problems	Gain	+20	0 / **+40**	Gains > Losses
		Loss	-20	**-40** / 0	
2	No zeroes	Gain	+25	+5 / **+45**	Gains > Losses
		Loss	-25	**-45** / -5	
3	Magnitude	X-Gain	+40	0 / **+80**	X: Gains > Losses
		X-Loss	-40	**-80** / 0	
		NX-Gain	+20	0 / +40	NX: No difference
		NX-Loss	-20	-40 / 0	
4G	All Gains	HX	+60	+40 / **+80**	HX > BX > LX
		BX	+40	**0** / **+80**	
		LX	+20	**0** / +40	
4L	All Losses	HX	-20	-40 / **0**	HX > BX > LX
		BX	-40	**-80** / **0**	
		LX	-80	-80 / -40	

4.3 Experiment 1: Mixed Gains and Losses

Experiment 1 evaluated risk preferences in a decision context with intermixed gain and loss problems. There were four possible doors that each led to a different outcome (see Table 1): a fixed-gain door (+20), a risky-gain door (50/50 chance of either 0 or + 40), a fixed-loss door (-20), and a risky-loss door (50/50 chance of either -40 or 0). People repeatedly made choices between pairs of these doors. In this decision context, +40 and -40 were the two extremes. According to the extreme-outcome rule, people should overweight these two extremes in the decision process, leading to more risk seeking in the gain case and more risk aversion in the loss case. This experiment also evaluated more generally whether a reverse-reflection effect (greater risk seeking for gains than losses) would occur in decisions from experience without concurrent decisions from description (cf. Ludvig & Spetch, 2011).

4.3.1 Methods

Participants

29 introductory psychology students at the University of Alberta participated for course credit (15 females; $M_{age} = 18.9$ years, $SD = 1.7$). Each participant gave written informed consent. The study was approved by a university ethics board.

Procedure

Participants played a computer-based task in which they were told to try and earn as many points as possible. As illustrated in Figure 4.1, on most trials, participants were presented with pictures of two doors, and they indicated their choice by clicking on one of those doors. Choices were immediately followed by feedback for 1.2 s, which showed the number of points won or lost along with a cartoon graphic. Feedback was only given for the chosen door, as in a partial-feedback procedure (e.g., Hertwig & Erev, 2009; Camilleri & Newell, 2011). The total accumulated points were continuously displayed at the bottom of the screen. An inter-trial interval of 1 to 2 s separated each trial.

Sessions were each organized into 5 blocks of 48 trials, separated by a brief break. Each block included a mixture of three trial types. *Decision trials* involved choices between 2 gain doors or 2 loss doors (Figure 4.1A). For both gains and losses, the *fixed door* always led to the same outcome, and the *risky door* led equiprobably (i.e., with a 50/50 chance) to one of two outcomes: one smaller and one large than the fixed outcome. The objective expected value on these decision trials was always equal for the fixed and risky doors. Across the risky option in all 4 experiments, the experienced likelihood of receiving either outcome never deviated significantly from .5. On *single-door* trials, there was only one door presented, which had to be clicked on to continue (Figure 4.1C). These trials ensured that all doors were sometimes selected and that participants occasionally experienced all reward contingencies. *Catch trials* presented 2 doors with substantially different objective expected values—typically a choice between a gain door and a loss door (Figure 4.1B). These trials provided the opportunity to gain points over the session and ensured that participants paid attention to their decisions. Participants that chose the gain door on fewer than 60% of these catch trials, corresponding to behavior not distinguishable from chance responding at $p < .05$ with 80 catch trials, were excluded from further analyses. In Exp. 1, data from one participant was excluded for poor performance on catch trials.

As detailed in Table 1, there were 4 different-colored doors in Experiment 1: a fixed gain (100%: +20 points), risky gain (50%: 0 and 50%: +40), fixed loss (100%: -20), or risky loss (50%: 0 and 50%: -40). The 48 trials in each block were divided among the 3 trial types as follows: 24 decision trials between the 2 gains or the 2 losses (12 of each), 16 catch trials, and 8 single-door trials. Each trial was equally incentivized with people earning (or losing) points on all 240 trials in the experiment. There were no separate sampling and choice phases, mirroring a repeated-choice design, rather than a sampling design (e.g., Hertwig & Erev, 2009; Erev et al., 2010). Trial order was randomized, but the total number of trials (240) and their distribution was constant across all participants. The door color associated with the fixed or risky gain or loss was counterbalanced across participants.

The numbers of different trial types were chosen to ensure that each door appeared equally often on both sides of the screen to prevent any side biases. Both the decision trials and the single-door trials in each block were equally divided between gain and loss trials, and thus had a total expected value of 0. Only the catch trials provided an opportunity to earn a net gain of points. Participants were encouraged to maximize their number of points, and good performance on the catch trials provided independent evidence that participants were adequately incentivized by the points.

Data Analysis

Risk preference was calculated as the probability of choosing the risky door. For each experiment, risk preference was averaged over the final 3 blocks and compared using t-tests or ANOVAs as appropriate. The final 3 blocks were selected as the primary dependent measure because that afforded participants sufficient opportunity to learn the outcomes associated with each option, while providing a long enough sample to get a reliable measure of their risk preference over time. Linear-trend analyses across all blocks were conducted to look for learning effects. All tests were repeated measures. Given the a priori hypotheses, all t-tests were one-tailed. The Greenhouse-Geisser correction for non-sphericity was applied where appropriate. Inferential statistics were calculated using SPSS (IBM Inc.; Armonk, NY) and MATLAB (The Mathworks Inc.; Natick, MA).

4.3.2 Results and Discussion

The left bars in Figure 4.2A depict the average risk preference (proportion of risky choices) over the final 3 blocks, split by gains and losses. Participants were more risk seeking for gains than losses, displaying a reversal of the usual reflection effect [$t(27) = 1.82$, $p < .05$, $d = .51$]. In this experiment, +40 and -40 were the extreme outcomes in the decision context, and thus we predicted that these outcomes would be overweighted in the decision process. The observed difference in risk seeking for gains and losses was consistent with this prediction. Figure 4.2B plots these risk preferences by experimental block and shows an interaction between valence (gain or loss) and block [$F(4, 71) = 3.63$, $p < .05$, $\eta_p^2 = .12$]. There was an increase in risk aversion for losses across the experiment [linear effect of Block: $F(1, 27) = 8.76$, $p < .01$, $\eta_p^2 = .25$], but no change in risk preference for gains [linear effect of Block: $p > .1$, $\eta_p^2 = .009$].

Across the population, risk preference hovered around .5, raising the possibility that people have a tendency toward equal allocation of their responses across the two options. To evaluate this possibility, Figure 4.2C plots the risk preferences for each individual for both gains and losses. There is clearly a wide spread of risk preferences in the population (from almost never to almost always). Independent of this individual variability in overall risk preference, most people show the same pattern: more risk seeking for gains than losses [22 out of 28 participants = 79%; $p < .001$ by a binomial test].

These results extend our previous results (Ludvig & Spetch, 2011), which also showed greater risk seeking for gains than losses in decisions from experience, but that design also included interspersed decisions from description. Here, we found more risk seeking for gains than losses in experience-based choice, even in the absence of described choices.

4.4 Experiment 2: No Zeroes

Experiment 1 provided evidence for an extreme-outcome rule, whereby the largest and smallest outcomes in a decision context are overweighted in choice. In that experiment, however, the potential non-extreme outcomes for the risky options were always 0 (see also Ludvig & Spetch, 2011). Zero, however, is neither an absolute gain nor an absolute loss and may instead be treated as a special number (e.g., Shampanier, Mazur, & Ariely, 2007). Thus a possible alternative hypothesis to the overweighting of extreme outcomes is that, instead, the 0 outcomes may be underweighted. In addition, because the 0 outcome potentially followed the risky option in both the gain and the loss case in Experiment 1, it occurred twice as frequently as the extreme outcomes. As a result, perhaps this increased frequency of the 0 outcome led to a reduced

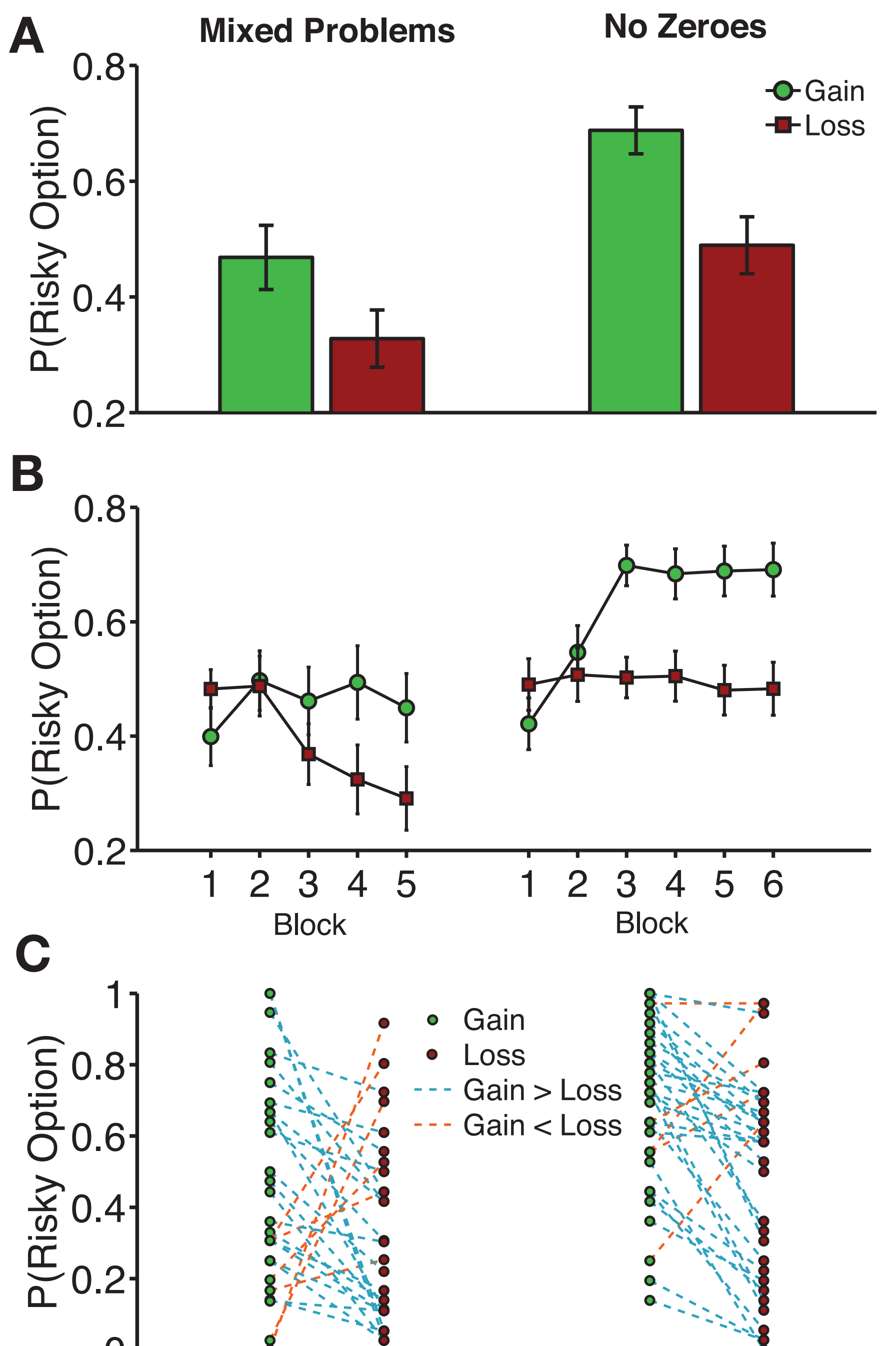

Figure 4.2: Results of Experiments 1 and 2. A. Mean risk preference ($\pm$ SEM) for gain and loss doors averaged over the last 3 blocks. B. Mean risk preference ($\pm$ SEM) as a function of block for all blocks. C. Mean risk preference for individuals over the last 3 blocks. Dashed lines connect data from the same individual. From left to right, the plots show results for mixed problems (Experiment 1; $N = 28$) and no zeroes (Experiment 2; $N = 34$). As predicted, in both cases, risk preferences are higher for gains than losses, consistent with the extreme-outcome rule.

weighting of those 0 outcomes relative to the extreme (+40 or -40) outcomes. This second alternative hypothesis is unlikely, however, as it goes against the wealth of existing evidence that rare events are underweighted in experience-based choice (e.g., Hertwig et al., 2004; Hertwig & Erev, 2009; Weber et al., 2004).

This experiment aimed to rule out both of these alternative hypotheses by shifting the absolute values of the different outcomes by 5 points. Thus, as indicated in Table 1, people were presented with 4 possible doors: a fixed-gain door (+25), a risky-gain door (50/50 chance of either +5 or + 45), a fixed-loss door (-25), and a risky-loss door (50/50 chance of either -45 or -5). People repeatedly made choices between pairs of these doors. The extreme outcomes are now +45 and -45. By the extreme-outcome rule, we expect more risk seeking for gains and more risk aversion for losses. The two alternative hypotheses, in contrast, would predict that this reversal of the reflection effect should disappear when the more frequent 0 outcome is removed from the mix.

4.4.1 Methods

Participants

34 students drawn from the same subject pool as Experiment 1 participated (15 females; $M_{age} = 19.1$ years, $SD = 2.0$).

Procedure

The procedure was identical to Experiment 1, except that the absolute values of the potential outcomes that followed the 4 doors (see Table 1) were all shifted by 5 points and participants were run for an additional (6th) block of 48 trials. The outcomes for the 4 doors were: a fixed gain (100%: +25), risky gain (50%: +5 and 50%: +45), fixed loss (100%: -25), or risky loss (50%: -5 and 50%: -45). No participants in this experiment were excluded for poor performance on the catch trials.

4.4.2 Results and Discussion

The right bars in Figure 4.2A plot the average risk preference on gain and loss trials over the final 3 blocks in Experiment 2. Participants were once again more risk seeking for gains than losses—a clear reversal of the reflection effect [$t(33) = 4.46$, $p < .001$, $d = .77$]. Figure 4.2B clearly shows how risk preference changed across the 6 blocks [main effect of Block: $F(3, 99) = 10.15$, $p < .001$, $\eta_p^2 = .24$; Block×Valence interaction: $F(3, 101) = 5.47$, $p < .001$, $\eta_p^2 = .14$]. Across the blocks, risk seeking for gains increased [linear effect of Block: $F(1, 37) = 16.93$, $p > .001$, $\eta_p^2 = .34$], but there was little change in risk preference for losses [linear effect of Block: $p > .1$, $\eta_p^2 = .002$]. Figure 4.2C plots the individual differences, where a significant majority of people are more risk seeking for gains than losses [29 out of 34 participants = 85%; $p < .001$ by a binomial test]. The higher level of risk seeking for gains than losses in the absence of zero values rules out the two alternative hypotheses. The pattern of risk preferences observed in Experiments 1 and 2 were not due to underweighting the zero outcomes. The results, however, offer further support for the extreme-outcome rule, as they are congruent with an overweighting of the two extremes in the decision context, which were +45 and -45 in this experiment.

4.5 Experiment 3: Outcome Magnitude

One clear prediction of the extreme-outcome rule is that a given outcome will only be overweighted when it is either the largest and smallest outcomes in a decision context. In Experiment 3, we directly tested this prediction by adding new doors with higher-magnitude outcomes into a decision context with the same gain and loss problems as in Experiment 1. In this case, as detailed in Table 1, participants encountered 8 possible doors. The four non-extreme (NX) doors led to the same outcomes as in Experiment 1: a guaranteed gain/loss of 20 or a 50/50 chance of a gain/loss of 40. The other 4 extreme (X) doors led to exactly double those outcomes: a guaranteed gain/loss of 40 or a 50/50 chance of a gain/loss of 80. The highest and lowest outcomes were +80 and

-80 respectively, and the extreme-outcome rule predicts more risk seeking for gains than losses for decisions involving those extreme outcomes. For the NX doors, even though the outcomes were identical to those in Experiment 1, the decision context has changed so that +40 and -40 are no longer extreme outcomes. As a result, the extreme-outcome rule predicts that the reverse reflection effect observed in Experiment 1 (see Fig. 4.2A; Ludvig & Spetch, 2011) should not be present for these NX doors.

4.5.1 Methods

Participants

39 students drawn from the same subject pool as Experiments 1-2 participated (16 females; $M_{age} = 18.9$ years, $SD = 2.6$).

Procedure

The procedure followed the same protocol as Experiments 1 and 2 with the following minor changes because there were now 8 total doors. The outcomes for the 4 non-extreme (NX) doors were the same as in Experiment 1: a fixed gain (100%: +20 points), risky gain (50%: 0 and 50%: +40), fixed loss (100%: -20), or risky loss (50%: 0 and 50%: -40); The outcomes for the extreme (X) doors were twice the magnitude: a fixed high gain (100%: +40), risky high gain (50%: 0 and 50%: +80), fixed high loss (100%: -40), or risky high loss (50%: 0 and 50%: -80). Trials were presented in 5 blocks of 72 trials that each contained 32 decision trials between fixed and risky gains or losses of the same magnitude level (8 of each), 24 catch trials, and 16 single-door trials. Three types of catch trials were used: choices between the extreme gain and loss doors, choices between the non-extreme gain and loss doors, and choices between 2 gain or 2 loss doors of different expected values (e.g., +20 vs. 0/+80). Only the choices between a gain door and a loss door were used to exclude participants. Two participants were excluded.

4.5.2 Results and Discussion

Figure 4.3 shows how, as predicted, there was an interaction between decision type (X/NX) and reward valence in the final 3 blocks [$F(1, 36) = 4.27$, $p < .05$, $\eta_p^2 = .11$; no main effects]. Risk seeking was greater for gains than for losses in X decisions [$t(36) = 1.84$, $p < .05$, $d = .40$], but not in NX decisions [$t(36) = 0.80$, $p > .1$, $d = .17$]. Note how the outcomes that followed the NX doors were identical to the ones that followed the doors in Experiment 1, yet the results were reversed in the new decision context (cf. Figure 4.2). Across the session, Figure 4.3B shows how, for the X doors, participants became increasingly risk averse for losses [linear effect of Block: $F(1, 36) = 5.05$, $p < .05$, $\eta_p^2 = .12$], but the visual trend toward an increase in risk-seeking for gains was not significant [$p > .1$]. For the NX doors, risk preference did not change across blocks for either gains or losses [both p's$> .1$]. The greater risk seeking for gains than losses for X decisions, but not NX decisions, provides strong support for the prediction that only the most extreme outcomes in the decision context would be overweighted.

4.6 Experiment 4: Framing Effects

In the first 3 experiments, the extreme values were always gains and losses of the same magnitude (40 in Exp 1, 45 in Exp 2, and 80 in Exp 3). A further prediction of the extreme-outcome rule is that the largest and smallest values in a decision context should be overweighted in the decision process, independent of whether they are absolute gains or losses. For example, in a decision context with all gains, the lowest gain would be an extreme and should be overweighted.

To test this prediction, Experiment 4 split gains and losses across participants to test whether relative extremes are sufficient to elicit changes in risk sensitivity (see Table 1). In the All-Gain group (Exp 4G), the largest possible gain (+80) was the high-extreme (HX) outcome, and the smallest possible gain (0) was the low-extreme (LX) outcome. Conversely, in the All-Loss group (Exp 4L), the smallest possible loss (0) was the HX outcome, and the largest

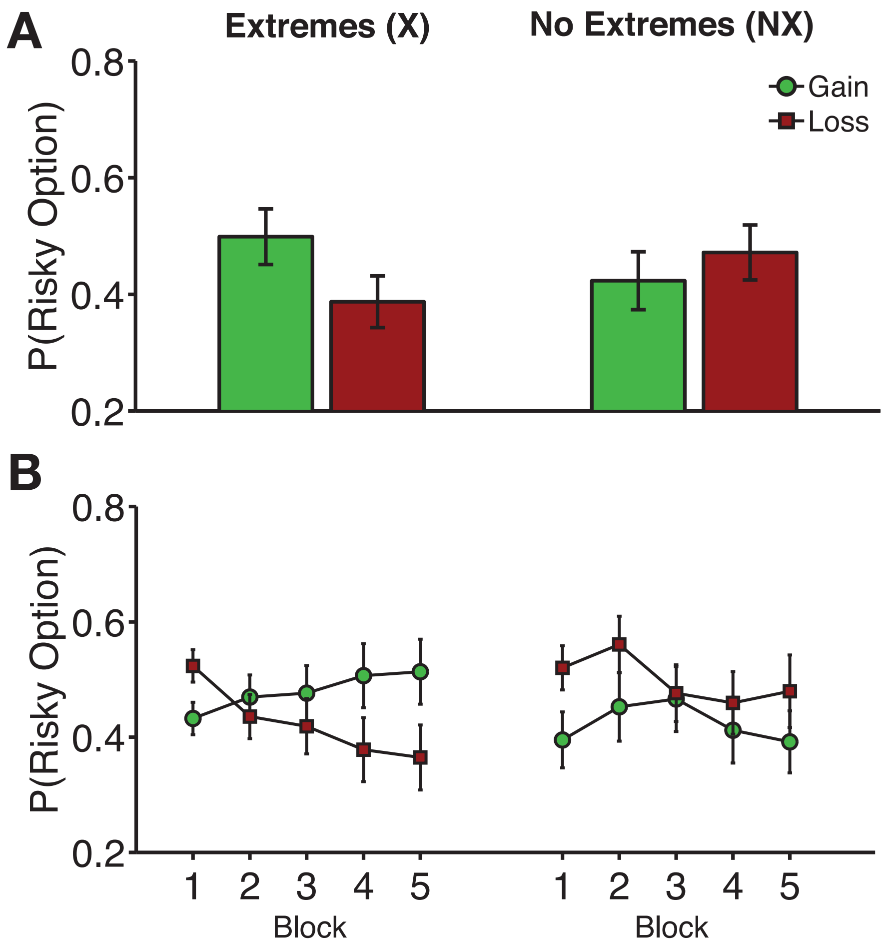

Figure 4.3: Results of Experiment 3. A. Mean risk preference (± SEM) for gain and loss doors averaged over the last 3 blocks. B. Mean risk preference (± SEM) as a function of block for all blocks. Decisions with extreme (X) outcomes are on the left, and decisions with no extreme (NX) outcomes are on the right ($N = 37$). As predicted by the extreme-outcome rule, risk preferences are higher for gains than losses in the X condition, but not in the NX condition.

possible loss (-80) was the LX outcome. Following the extreme-outcome rule, in both groups, there should be more risk seeking for the doors with potential HX outcomes and more risk aversion for the doors with potential LX outcomes. In addition, we attempted to evaluate the relative weightings of the high and low extremes in the decision process. To do so, we also included risky options that potentially led to both extremes (BX). Given the loss aversion that characterizes many decisions (e.g., Kahneman & Tversky, 1979; Tom, Fox, Trepel, & Poldrack, 2007; Tversky & Kahneman, 1992), we might expect that the LX outcome (the relative loss) would be more heavily weighted, leading to risk aversion in both the gain and loss groups. In experience-based choice, however, significant loss aversion is often not observed (e.g., Erev, Ert, & Yechiam, 2008; Yechiam & Hochman, 2013), suggesting that the two extremes might be more evenly weighted, leading to an intermediate level of risk preference.

4.6.1 Methods

Participants

A total of 79 students from the same subject pool participated (58 females; M_{age} = 19.5 years, SD = 2.3; N = 39 and 40 for the All-Gain and All-Loss groups, respectively).

Procedure

The basic procedure was the same as in Experiments 1-3 (see Figure 4.1) except that participants in the All-Gain group (Exp 4G) experienced only gain doors and participants in the All-Loss group (Exp 4L) experienced only loss doors. For both groups, there were 6 different-colored doors. As detailed in Table 1, for the All-Gain group, low-extreme (LX) decisions were between a fixed gain (100%: +20 points) or a risky gain (50%: 0 and 50%: +40), the HX decisions were between a fixed high gain (100%: +60) or a risky high gain (50%: +40 or 50%: +80), and the BX decisions were between a fixed intermediate gain (100%: +40), and a risky gain that included both extremes (50%: 0 and 50%: +80). For the All-Loss group, the doors led to: a LX fixed loss (100%: -60), LX risky loss (50%: -40 and 50%: -80), BX fixed loss (100%: -40), BX risky loss (50%: 0 and 50%: -80), HX fixed loss (100%: -20), and HX risky loss (50%: 0 and 50%: -40). Thus, for gains, the high and low extremes were +80 and 0, and for losses, the high and low extremes were 0 and -80. In both groups, we expected more risk seeking for HX decisions than for LX decisions, and we expected that risk seeking for the BX decisions would be closer to the LX decisions or fall between the other two decision types.

For both groups, trials were presented in 5 blocks of 60 trials that were divided among the 3 trial types: 36 decision trials between fixed and risky gains or losses of the same value level (12 of each), 12 catch trials, and 12 single-door trials. Two types of catch trials were used. Easy catch trials consisted of a choice between an HX door and an LX door (e.g., -20 vs. -40/-80). Subtle catch trials consisted of a choice between a BX door and either an HX or LX door (e.g., 40 vs. -20/-40). Subtle catch trials were included in the design to match the number of presentations of each door, but performance on these trials was not used to exclude participants. A total of 18 participants (10 in the All-Loss group and 8 in All-Gain group) were excluded for poor performance on the easy catch trials.

To ensure that both groups ended the experiment with a similar number of points, the All-Loss group started with approximately twice the number with which they would end the experiment (24,000 points). The All-Gain group started with zero points as in Experiments 1-3. The two groups did not differ in the number of points remaining/earned at the end of the experiment [$t(77) = 0.28$, $p > .1$].

4.6.2 Results and Discussion

Consistent with the extreme-outcome rule, Figure 4.4A shows that people were more risk seeking for HX doors than for LX doors in the final 3 blocks for both groups [$F(2, 108) = 16.68$, $p < .001$, $\eta_p^2 = .22$]. Across the session, there was an increase in risk aversion for the LX decisions [$F(1, 60) = 11.00$, $p < .01$, $\eta_p^2 = .16$], but there was no change in risk preference for HX and BX decisions [both p's¿.1, η_p^2's$< .05$]. Note how, in the All-Gain group, the LX decision was between a guaranteed +20 and a 50/50 chance at +40 (see

Table 1). This decision was identical to the gain decision in Experiment 1, yet, in this decision context, people were now a lot more risk averse (cf. Fig. 4.2; $t(59) = 2.40$, $p < .05$, $d = 0.63$). Similarly, in the All-Loss Group, the HX decision was between a guaranteed -20 and a 50/50 chance at -40. This decision was identical to the loss decision in Experiment 1, yet, in this decision context, people were now a lot more risk seeking (cf. Fig. 4.2; $t(55) = 3.88$, $p < .001$, $d = 1.05$). This comparison across groups and experiments provides strong evidence for the extreme-outcome rule.

For the BX decisions, risk preferences were intermediate to the HX and LX decisions in both All-Gain and All-Loss groups [linear effect of decision type: $F(1, 60) = 32.46$, $p < .001$, $\eta_p^2 = .35$], but were slightly risk averse and closer to the risk preference for the LX decision (particularly in the All-Gain group). This result suggests that the low extreme is weighted slightly more heavily than the high extreme.

4.7 General Discussion

Across all 4 experiments, when decisions from experience contained the most extreme outcomes in the decision context, people were more risk seeking for relative gains than for relative losses. This result is opposite to the usual reflection effect observed with decisions from description (Kahneman & Tversky, 1979), but accords with some recent results with decisions from experience (e.g., Ludvig & Spetch, 2011; Tsetsos, Chater, & Usher, 2012) as well as the risk seeking often observed in non-human animals (Hayden et al., 2008; O'Neill & Schultz, 2010). In accord with the extreme-outcome rule, this reversed reflection effect only occurred when the risky option potentially led to an extreme outcome. Thus, risky choice behavior was dependent not only on the potential outcomes of the current decision but also on the decision context in which it occurred (cf. Table 1).

This overall pattern can be clearly seen in Figure 4.5, which plots the difference between risk preferences for relative gains and losses in each experiment. In all 4 experiments, the risky option that led to the best possible

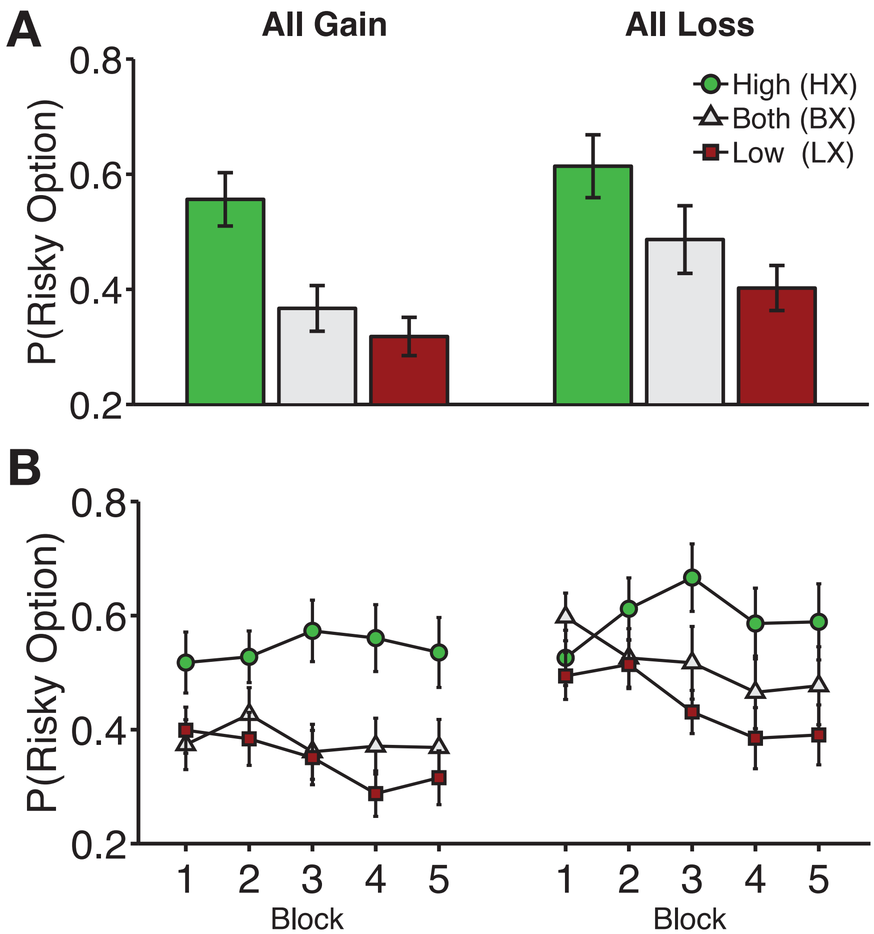

Figure 4.4: Results of Experiment 4. A. Mean risk preference (± SEM) for each door averaged over the last 3 blocks. B. Mean risk preference (± SEM) as a function of block for all blocks. From left to right, the plots show results for gain doors that led to high-extreme (HX), low-extreme (LX), or both extreme (BX) outcomes (Experiment 4G; $N = 32$), and loss doors that had HX, LX, or BX outcomes (Experiment 4L; $N = 29$). In both cases, risk preferences were highest for HX doors and lowest for LX doors, as predicted by the extreme-outcome rule.

outcome in that experiment produced more risk-seeking behavior, whereas the risky option that led to the worst possible outcome in that experiment led to more risk-averse behavior. This reversal held even when all the experienced outcomes were absolute gains or losses (Exp 4). The pattern was not apparent when the risky option never led to an extreme outcome (Exp 3). In short, people chased the potential big win, but avoided the potential big loss.

One key factor in generating this pattern of behaviour is the intermingling of multiple problems in the same decision context. As a result, the outcomes

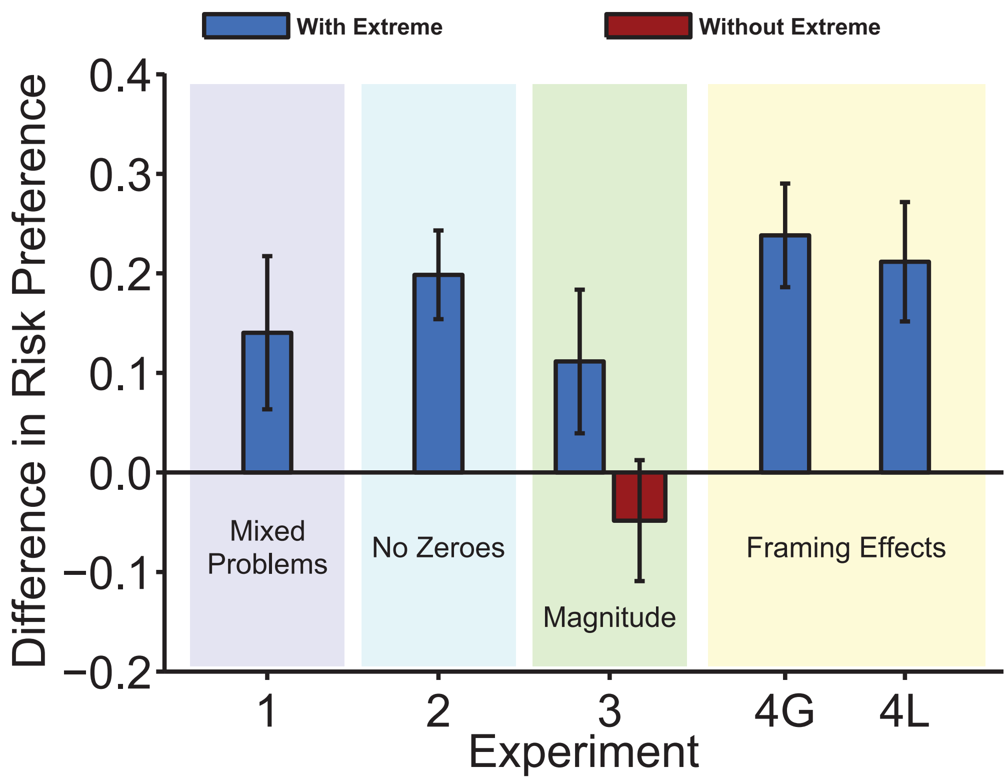

Figure 4.5: Mean difference (± SEM) in risk preference across the last 3 blocks in each experiment. For Experiments 1 to 3, the difference is calculated between the risk preferences for gains and losses. For Experiment 4, the difference is calculated between the risk preferences for high-extreme (HX) and low-extreme (LX) decisions. In all cases with extreme outcomes, the difference in risk preference was significantly above zero, indicating greater risk seeking for relative gains than relative losses.

that follow a risky option can be a high extreme, low extreme, or neither, pending the decision context, influencing risk preference (see Fig. 4.5). In other experiments on decisions from experience, participants occasionally encounter multiple problems sequentially (e.g., Erev et al., 2010), but they are not intermixed with one another. In those cases, even with 50/50 outcomes, there is no oversensitivity to the extreme outcomes, and the usual reflection effect is observed (e.g., Ert & Yechiam, 2010). In these other tasks with sequential presentation and 50/50 outcomes on the risky option, both outcomes become extremes—most similar to the both-extremes (BX) option in Exp. 4 here. Indeed, isolating out that BX option, we also find the usual reflection effect, with more risk seeking for losses than gains (compare the grey bars in Fig. 4.4).

An important question is thus what exactly constitutes the decision context. In all our experiments, there was clearly only one decision context. All problems were intermingled within each experimental block and throughout the whole experimental session. In contrast, the other studies that looked at 50/50 outcomes in decisions from experience, but found the usual reflection effect, presented problems sequentially within a session (e.g., Erev et al., 2010; Ert & Yechiam, 2010). This difference suggests that the relevant unit for the decision context is smaller than a full session. At the opposite end, we only presented one or two options at a time, yet the outcomes of other, non-presented, options influenced risk preference. Thus, the decision context is clearly larger than the options immediately available (cf. Heath & Chatterjee, 1995; Huber et al., 1982; Simonson, 1989; Simonson & Tversky, 1992). More generally, we think that the decision context is the comparison set of all the other options and outcomes that are considered when making a decision—similar to the decision environment in the decision-by-sampling framework (Stewart, Chater, & Brown, 2006). One possibility is that these other options and outcomes are linked to the context of the current decision through stimulus-stimulus associations as has been supposed in some models of animal learning (Miller & Matzel, 1988; Stout & Miller, 2007). In that case, options and outcomes that have previously co-occurred with or immediately preceded or followed either of the options under consideration would fall into the decision context.

Although opposite in direction, the context effects observed here for experience-based choices complement the previous evidence for framing effects in decisions from description (Tversky & Kahneman, 1981). The 2 groups in Experiment 4 highlight this relativity: zero outcomes served as both the high-extreme value for the All-Loss group, leading to more risk seeking, and the low-extreme value for the All-Gain group, leading to more risk aversion. The low and high extremes may not be weighted equally, however. In Experiment 4, the risky door for BX decisions, which led to both the best and the worst possible outcomes, produced moderate risk aversion. This finding suggests that the worst outcomes (relative losses) were weighted more heavily than the best outcomes (relative gains), reminiscent of loss aversion.

In experience-based choice, decisions must be made based on the memories of past outcomes. This dependence on the past suggests that memory biases may play a role in the overweighting of extreme outcomes. In other contexts, choice is indeed influenced by the biases inherent in human memory (see Weber & Johnson, 2006). There is a well-known bias in which highly salient and emotional events are overweighted in memory tasks (e.g., Brown & Kulik, 1977; Phelps & Sharot, 2008; Talarico & Rubin, 2003; Madan & Spetch, 2012b; Madan, Fujiwara, et al., 2012) and retrospective judgments, as in the peak-end rule (e.g., Fredrickson, 2000; Kahneman et al., 1993). Thus, one possibility is that the extreme outcomes are more likely to be retrieved at the time of the decision (cf. Johnson, Häubl, & Keinan, 2007; Stewart et al., 2006), perhaps as a simplifying heuristic to limit the number of outcomes considered. More frequent retrieval of the extreme outcome would thus lead to more risk seeking for relative gains and more risk aversion for relative losses.

An alternate possibility is that the extreme outcomes are more salient (and thus overweighted) at the time of their occurrence, biasing the encoding of the learned values for the risky options (see Tsetsos et al., 2012; Niv, Edlund, Dayan, & O'Doherty, 2012). High extremes would increase the learned values for risky gains, causing more risk seeking, whereas low extremes would decrease the learned values for risky losses, causing more risk aversion in line with what we observe. The current results do not allow us to disambiguate these potential interpretations, but suggest directions for future research.

The bulk of the literature on decisions from experience focuses on the key finding that rare events are underweighted in choice (e.g., Barron & Erev, 2003; Hertwig et al., 2004; Hertwig & Erev, 2009). Our experiments are complementary to that literature: there were no rare events, and all risky options led to 2 equiprobable outcomes. We found that common, extreme outcomes were overweighted in the decision process, leading to more risk seeking for relative gains than for relative losses (Figure 4.5; Ludvig & Spetch, 2011). The extreme-outcome rule does not explain the underweighting of rare events, but does make a clear prediction: If a rare event is also an extreme, then there should be less underweighting than a parallel situation where the rare event is non-extreme. For example, take the range of outcomes from Exp 1 (-40 to +40). The prediction is that there would be more underweighting of the rare event for an option that led to 95% +40 and 0 elsewise than an option that led to +40 only 5% of the time and 0 elsewise. That is, a rare, extreme outcome should be underweighted less than a rare, non-extreme outcome. Note that this prediction only holds in situations where multiple problems are intermingled, so that not all outcomes are extreme in the decision context. Alternatively, it is also possible that the extreme-outcome rule would be dominated or non-applicable in situations with rare events.

Our results are not likely to be caused by a sampling bias—a particular concern with protocols that use rare events (Hertwig & Erev, 2009; Fox & Hadar, 2006; Rakow, Demes, & Newell, 2008; Ungemach, Chater, & Stewart, 2009). All outcomes were experienced numerous times by participants, and the mean proportion of outcomes for the risky option never deviated significantly from .5 in any experiment. The risk seeking with relative gains (most notably in Exp 2 and Exp 4L) also rules out the "hot stove effect" (Denrell & March, 2001) that can occur from sequential sampling—whereby risk aversion emerges in experience-based choice due to the avoidance of ephemerally unlucky risky options (March, 1996; Niv et al., 2012). Furthermore, the inclusion of single-choice trials in each run assured that the planned contingencies were occasionally experienced by participants. Our results also cannot be attributed to a wealth effect (Thaler & Johnson, 1990), which would predict a consistent increase in risk seeking across the experiment, rather than the observed divergence in risk preferences between relative gains and losses.

Our task also adds some new methodological wrinkles to the study of decisions from experience (e.g., Erev et al., 2010; Camilleri & Newell, 2011; Hertwig et al., 2004; Hertwig & Erev, 2009; Ungemach et al., 2009; Weber et al., 2004). In some repeated-choice experiments, people repeatedly select (by clicking) from the same two options whose physical location is constant (for a given participant). This fixed location may introduce a "switch cost" in that it can be faster and easier for subjects to continue clicking in the same location rather than to move the mouse over to the other side. This potential switch cost may help induce some choice inertia or perseveration bias, which indeed has been observed (e.g., Erev et al., 2010).

The current task has several features that mitigate this potential bias. First, each trial was separated from the next by a short inter-trial interval (1-2 s) instead of immediately following the previous choice. Second, the cursor was re-centered after each trial, forcing an equivalent movement to the left or the right on each choice. Third, from trial to trial, the location of the different doors were randomly counterbalanced, appearing on either side half the time—a common feature in studies with animals (e.g., Vasconcelos & Urcuioli, 2008). Thus, any perseverative side bias would appear as random (50/50) choice and not a preference for either option. This last design feature also removes any potential confounds between choices based on aliasing stimulus identity and stimulus location. Finally, as part of our primary manipulation, several different problems were intermingled, thus the identity of the doors on a given trial could not be known in advance. Collectively, these design features produce a task where each choice requires active engagement with the available stimuli and their physical locations.

These features of the experimental design do indeed neutralize the perseverative bias. For example, in the problem with the 50/50 outcomes from the Technion Prediction Competition dataset (Problem 49; Erev et al., 2010), the perseveration rate was $87.0 \pm 2.4\%$ (calculated from the 20 subjects in the estimation set who encountered that problem). In contrast, in Exp 1 here, people only selected the same side on the next trial $49.7 \pm 3.9\%$ of the time, meaning they were equally likely to perseverate or alternate. This result is not that surprising given that the location and identity of the stimuli changed from trial to trial.

As a further countermeasure to the potential disengagement of subjects, we included catch trials that are designed to incentivize people to attend to their choices. These catch trials provide an explicit means of ensuring that participants are, in fact, paying attention to their choices. On catch trials, the expected value of the options differs significantly. For example, in Experiment 1, some catch trials gave participants a choice between a guaranteed gain of +20 or a guaranteed loss of -20. Independent of variations in risk sensitivity, participants should choose the option with the higher expected value on these trials. By excluding participants who perform poorly (below 60%) on catch trials, we were able to ensure that the remaining participants were sufficiently incentivized. Importantly, unlike most outlier removal, this exclusion was not based on our primary dependent measure (risk preference on decision trials), but rather on a secondary measure (risk preference on the catch trials). Thus, our main results cannot be due to participants who may have ignored the stimuli and responded randomly.

In conclusion, we found that, in decisions from experience, people chase the big win, but avoid the big loss. The results provide evidence for an extreme-outcome rule, whereby the highest and lowest outcomes in a decision context are overweighted in choice. This potent role of extreme values in decision making has important real-world implications. For example, when gambling, people often choose between a smaller loss (the bet) and a larger win (the jackpot). Our results suggest that the overweighting of the largest wins with experience might contribute to an increased tendency to gamble.

4.8 Acknowledgements

This research was funded by grants from the Alberta Gambling Research Institute (AGRI) and the National Science and Engineering Research Council of Canada (NSERC) held by MLS. CRM was supported by NSERC, AGRI, and the Canadian Institutes of Health Research. EAL was supported by NIH Grant #P30 AG024361. We thank the Explore-Exploit Group at Princeton for insightful discussions and Ashley Rodgers for help with data collection. Door images were extracted from "Irish Doors" on fineartamerica.com with permission from Joe Bonita.

4.9 References

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.

Brown, R., & Kulik, J. (1977). Flashbulb memories. *Cognition*, *5*, 73–99.

Camilleri, A. R., & Newell, B. R. (2009). Within-subject preference reversals in description- and experience-based choice. *Proceedings of the Annual Conference of the Cognitive Science Society*, *31*, 449–454.

Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*, 377–384.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523–538.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., . . . Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*, 575–597.

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*, 153–177.

Ert, E., & Yechiam, E. (2010). Consistent constructs in individuals risk taking in decisions from experience. *Acta Psychologica*, *134*, 225–232.

Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling theory + prospect theory: Reconsidering Hertwig, Barron, Weber, & Erev (2004). *Judgment & Decision Making*, *1*, 159–161.

Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion*, *14*, 577–606.

Hayden, B., Heilbronner, S. R., Nair, A. C., & Platt, M. L. (2008). Cognitive influences on risk-seeking by rhesus macaques. *Judgment & Decision Making*, *3*, 389–395.

Heath, T. B., & Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, *22*, 268–284.

Heilbronner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, *7*, 7.

Helson, H. (1947). Adaptation-level as frame of reference for prediction of psychophysical data. *The American Journal of Psychology*, *60*, 1–29.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90–98.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning Memory & Cognition*, *33*, 461–474.

Kacelnik, A., & Bateson, M. (1996). Risky theories: The effects of variance on foraging decisions. *American Zoologist*, *36*, 402–434.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*, 401–405.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Lotto, R. B., & Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy Sciences USA*, *97*, 12834–12839.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

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.

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

March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309–319.

McCoy, A. N., & Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, *8*, 1220–1227.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51–92). San Diego: Academic Press.

Mitchell, T. R., Thompson, L., Peterson, E., & Cronk, R. (1997). Temporal adjustments in the evaluation of events: The "rosy view". *Journal of Experimental Social Psychology*, *33*, 421–448.

Müller-Lyer, F. C. (1889). Optische urteilstäuschungen. *Archiv für Physiologie*, *Suppl*, 263–270.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*, 551–562.

O'Neill, M., & Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron*, *68*, 789–800.

Phelps, E. A., & Sharot, T. (2008). How (and why) emotion enhances the subjective sense of recollection. *Current Directions Psychological Science*, *17*, 147–152.

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior And Human Decision Processes*, *106*, 168–179.

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*, 3–8.

Shampanier, K., Mazur, N., & Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, *26*, 742–757.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, *16*, 158–174.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, *29*, 281–295.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology: General*, *132*, 23–46.

Stone, A. A., Schwartz, J. E., Broderick, J. E., & Shiffman, S. S. (2005). Variability of momentary pain predicts recall of weekly pain: A consequence of the peak (or salience) memory heuristic. *Personality and Social Psychology Bulletin*, *31*, 1340–1346.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*, 759–783.

Talarico, J. M., & Rubin, D. C. (2003). Confidence, not consistency, characterizes flashbulb memories. *Psychological Science*, *14*, 455–461.

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, *36*, 643–660.

Thomas, D. R., & Jones, C. (1962). Stimulus generalization as a function of the frame of reference. *Journal of Experimental Psychology*, *64*, 77–80.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision making under risk. *Science*, *315*, 515–518.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences USA*, *109*, 9659–9664.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473–479.

Vasconcelos, M., & Urcuioli, P. J. (2008). Deprivation level and choice in pigeons: A test of within-trial contrast. *Learning & Behavior*, *36*, 12–18.

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memory. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*, 497–518.

Chapter 5

Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions

A verison of this work was previously published as: **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. doi:10.3758/s13423-013-0542-9 (* – Authors contributed equally.) This work has been reproduced with permission.

5.1 Abstract

When making decisions based on past experiences, people must rely on their memories. Human memory has many well-known biases, including the tendency to better remember highly salient events. We propose an extreme-outcome rule, whereby this memory bias leads people to overweight the largest gains and largest losses, leading to more risk seeking for relative gains than relative losses. To test this rule, in two experiments, people repeatedly chose between fixed and risky options, where the risky option led equiprobably to more or less than the fixed option. As predicted, people were more risk seeking for relative gains than losses. In subsequent memory tests, people tended to recall the extreme outcome first and also judged the extreme outcome as having occurred more frequently. Across individuals, risk preferences in the risky-choice task correlated with these memory biases. This extreme-outcome rule presents a novel mechanism through which memory influences decision making.

5.2 Introduction

In decisions from experience, people must rely on their memories of past outcomes to evaluate the available options (Hertwig et al., 2004; Hertwig & Erev, 2009; Ludvig & Spetch, 2011; Weber et al., 2004). As a result, systematic biases in memory may affect experience-based decisions (Weber & Johnson, 2006). One well-known memory bias is the tendency to recall more salient experiences (Phelps & Sharot, 2008; Talarico & Rubin, 2003). This bias toward peak moments has a strong influence on affective judgments of past events (e.g., Fredrickson, 2000; Yu, Lagnado, & Chater, 2008). We propose that a memory bias for extreme outcomes also occurs in risky decisions from experience, making people more sensitive to the biggest gains and losses they encounter. Consequently, people should become more risk seeking for relative gains than losses—contrary to the risk preferences in decisions from description (Kahneman & Tversky, 1979), but congruent with recent results with decisions from experience (e.g., Ludvig, Madan, & Spetch, 2014; Tsetsos et al., 2012).

People tend to vividly remember highly emotional events, such as the Kennedy assassination, September 11th, or, more positively, the birth of a child (Brown & Kulik, 1977; Phelps & Sharot, 2008; Talarico & Rubin, 2003). Salient events are also recalled more readily and used to forecast future affective reactions (Morewedge, Gilbert, & Wilson, 2005). Similarly, people show better memory for the most and least rewarding events (e.g., Madan, Fujiwara, et al., 2012; Madan & Spetch, 2012b). This tendency to remember the best and worst of times also influences affective judgments of painful or pleasant episodes (e.g., Kahneman et al., 1993). These judgments depend primarily on the points of maximal and final intensity—an effect aptly summarized as the peak-end rule (Fredrickson, 2000). Such a bias also appears in judgments about past risky outcomes: in a simulated gambling task, people gave higher estimates of total payouts after sessions in which the payouts included high peak and end values compared to sessions with higher overall payouts but no extreme values (Yu et al., 2008).

We hypothesized that a memory bias for extreme values (both the highest and lowest) may cause these extremes to be overweighted in risky decisions from experience (see Ludvig et al., 2014). Following this *extreme-outcome rule*, when a risky option occasionally leads to the best possible gain in a context, that large gain would tend to be better remembered than other outcomes and overweighted in subsequent decisions; consequently, the risky option would be chosen more often. By the same logic, risky options that occasionally lead to the worst outcome in a context would be chosen less often. With repeated experience, people should thus become more risk seeking for relative gains than losses.

The pattern of risk preference predicted by the extreme-outcome rule is opposite to the reflection effect observed in decisions from description (e.g., Kahneman & Tversky, 1979), but accords with the greater risk seeking for relative gains than losses observed in recent experiments on repeated decisions from experience (e.g., Ludvig & Spetch, 2011; Ludvig et al., 2014; Tsetsos et al., 2012). For example, Tsetsos et al. (2012) had people choose between two reward distributions, which were learned about through a rapid visual stream of possible outcomes. When the task was to select a distribution, thereby highlighting the highest values, people were risk seeking, but when the task was to reject a distribution, thereby highlighting the lowest values, people were risk averse. Similarly, we found that people were more risk seeking for relative gains than losses in decisions from experience (Ludvig et al., 2014). Increased risk seeking for gains is also regularly observed in rhesus monkeys when they make rapid, repeated, small-stakes decisions (Heilbronner & Hayden, 2013).

The present experiments directly tested whether a memory bias for extreme outcomes drives these risky decisions. In both experiments, participants completed a choice task and a memory task. Figure 5.1 illustrates the choice task. People repeatedly chose between pairs of doors. One door always led to the same fixed outcome, whereas the other (risky) door led to more or less than the fixed outcome with a 50/50 chance. Importantly, gain and loss problems (Exp. 1) or high- and low-value problems (Exp. 2) were randomly intermingled in the task. This provided a context in which each risky option led to either an extreme or a non-extreme outcome with a 50/50 chance. The choice task was followed by two memory tests, which were the focus of the current paper. First, participants reported the first outcome that came to mind for each door. Second, participants judged the frequency that each door was followed by a particular outcome. Both memory tests revealed systematic biases that correlated with risky choice.

5.3 Experiment 1

5.3.1 Methods

Participants

114 introductory psychology students at the University of Alberta participated for course credit and a performance-based monetary bonus (80 females; $M_{age} =$ 19.6 years). The research was approved by a university ethics board.

Procedure

Choice Task. Figure 5.1 illustrates the task. On each trial, participants saw pictures of 1 or 2 doors on a computer screen and selected one by clicking on it. Choices were immediately followed by feedback in which the points won/lost along with a cartoon graphic were displayed for 1.2 s. Feedback was only given for the chosen door as in a partial-feedback procedure (Hertwig & Erev, 2009). Total accumulated points were continuously displayed on the screen. An interval of 1 to 2 s separated trials.

Sessions consisted of 5 blocks of 48 trials. Each block included a mixture of trial types: There were 24 *decision trials* that required a choice between either 2 gain doors or 2 loss doors (12 of each: Figure 5.1A). In both cases, the *fixed door* always led to the same outcome (+20 or -20), and the *risky door* led equiprobably to double the fixed outcome (+40 or -40) or nothing (0). There were 16 *catch trials* that required a choice between a gain door and a loss door (see Figure 5.1B). These trials ensured that participants were engaged in the task. Data from 7 participants who chose the gain door on fewer than 60% of these catch trials were excluded. On 8 *single-door trials*, there was only one

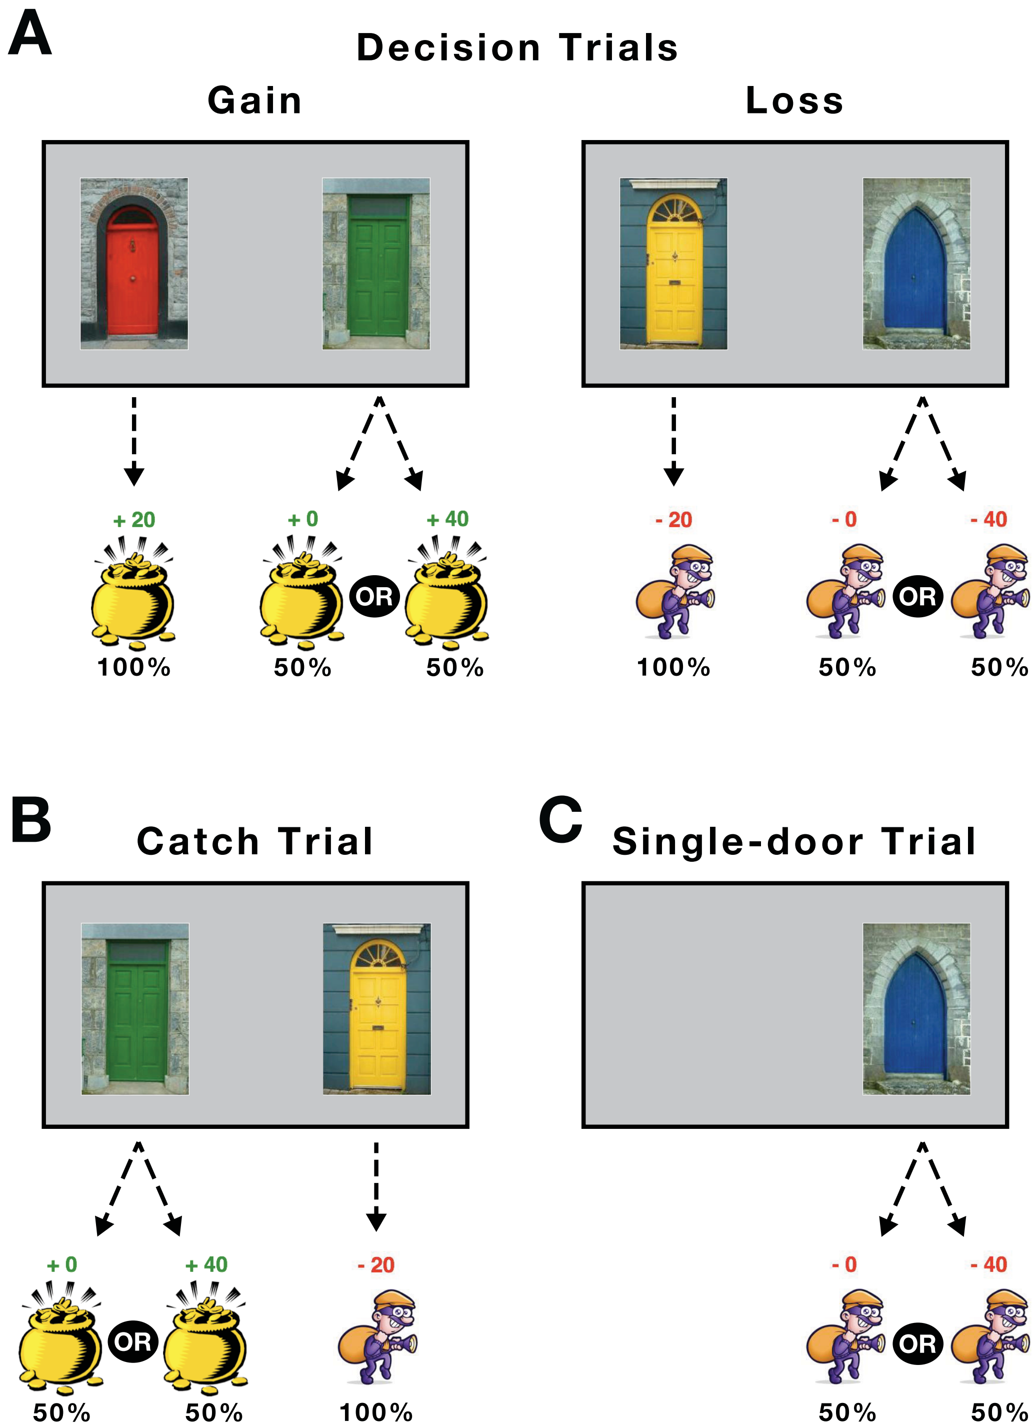

Figure 5.1: Choice task. A. *Decision trials* involved choices between two gain or two loss doors. One door always led to a gain (or loss) of a fixed number of points, and the other door led equiprobably to one of 2 possible outcomes. Choices were followed by feedback about the amount gained (or lost). B. *Catch trials* involved choices between one gain door and one loss door. C. *Single-door trials* only presented one door that had to be chosen.

door, which had to be selected to continue (2 of each; Figure 5.1C). These trials guaranteed that all reward contingencies were experienced, even if the doors were initially unlucky, thereby limiting any hot-stove effects (Denrell & March, 2001).

Participants won or lost points on all 240 trials and were paid $1 for every 400 points to a maximum of $5. Trial order was randomized within blocks. Each door appeared equally often on either side of the screen and in combination with the other doors. Door color was counterbalanced across participants.

Memory Tests. After the choice task, participants' memory for the outcomes associated with each door was tested in two ways. First, participants were shown the four doors in random order and asked to report for each the first outcome that came to mind. Second, participants were again shown the four doors in random order and asked to judge the frequency in percent of each of the possible outcomes (-40, -20, 0, +20, and +40). For each door, these outcomes were displayed simultaneously, and participants typed a number from 0 to 100 beside each outcome.

Data Analysis

Risk preference was operationalized as the probability of choosing the risky door over the final three blocks—after sufficient opportunity to learn the outcomes associated with each door. To assess the relationships between risky choice and memory, we used partial correlations (see Abdi, 2007) that controlled for the actual outcomes experienced. This approach allowed us to measure the relationship between risky choice and memory that occurred over and above any effect of the actual outcomes that participants experienced. Thus, the correlations we report cannot be attributed to differences in reward history. Given the *a priori* hypotheses, all tests were one-tailed, except where indicated.

5.3.2 Results

As predicted, Figure 5.2A shows how, on the choice task, participants were significantly more risk seeking for gains than losses in the final three blocks [$t(106) = 1.86$, $p < .05$, $d = 0.23$]. Across blocks (Figure 5.2B), the proportion of risky choices decreased for losses [linear effect of Block: $F(1, 106) = 39.49$, $p < .001$, $\eta_p^2 = .27$], but stayed constant for gains [linear effect of Block: $F(1, 106) < .001$, $p > .1$, $\eta_p^2 < .001$]. These risk preferences were modulated by recent outcomes. Fig. 5.2C shows the proportion of risky choices for gains and losses, split by whether the most recently experienced risky option of that valence yielded the good (+40 for gain or 0 for loss) or bad (0 for gain or -40 for loss) outcome. There was more risk seeking following the good outcome for both gains [$t(106) = 2.68$, $p < .01$, $d = 0.10$] and losses [$t(106) = 4.21$, $p < .001$, $d = 0.36$], resembling a win–stay/lose–shift pattern.

When asked after the task the first outcome to come to mind, people were significantly more likely to report the extreme outcome (+40 or -40) than the zero outcome for both gains [$\chi^2(1, N = 90) = 10.00$, $p < .01$] and losses [$\chi^2(1, N = 88) = 38.23$, $p < .001$]. Figure 5.3B plots risk preference in the

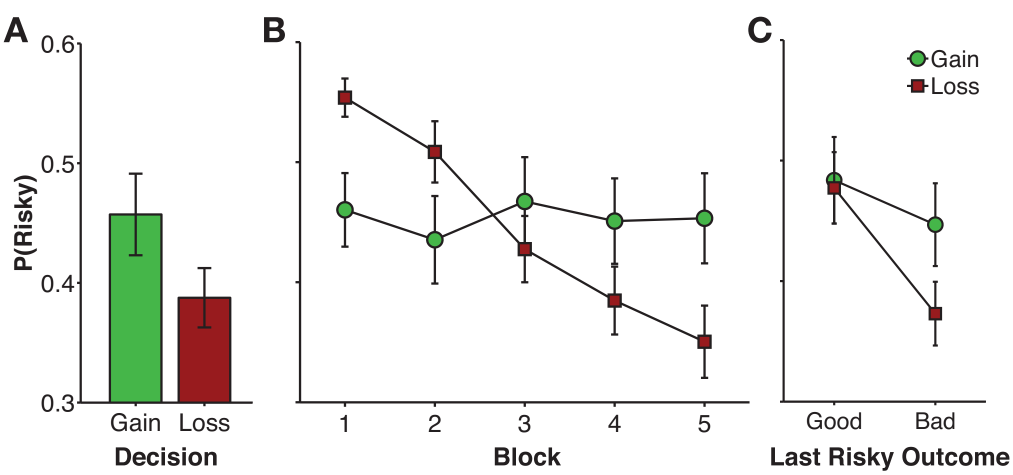

Figure 5.2: Choice results for Experiment 1. A. Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks. B. Mean risk preference (± SEM) for gain and loss doors for each block. C. Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks, separated by the most recent risky outcome experienced on that door.

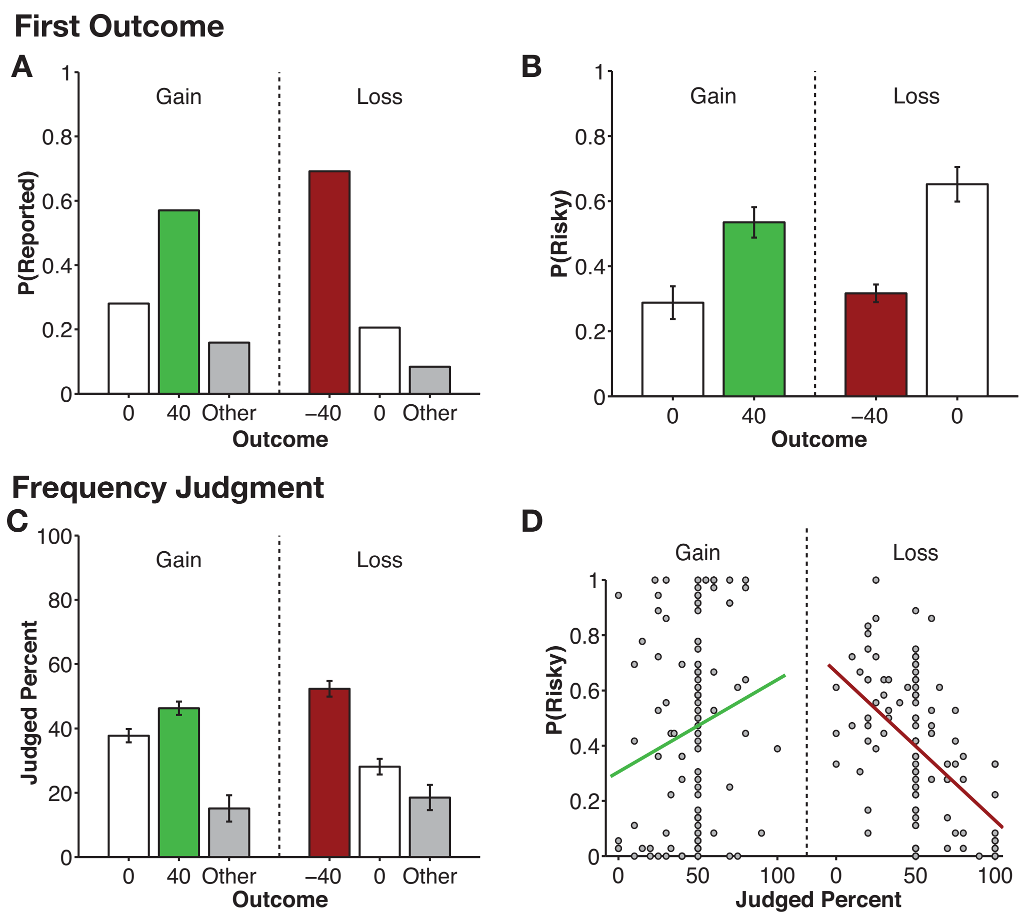

Figure 5.3: Memory results for Experiment 1. A. Proportions of participants that responded with +/- 40, 0, or neither in the first-outcome question. B. Mean risk preference (± SEM) for gains and losses, split based on what the participant reported to the first-outcome question. C. Mean judged percent (± SEM) for the +/- 40 and 0 outcomes from the frequency-judgment question. For simplicity, all other values were coded as "Other". D. Scatterplot of risk preference and frequency judgment responses for the gain and loss doors. Each dot represents an individual participant.

choice task based on participants' responses to this first-outcome question. For gains, people who reported +40 were more risk seeking than those who reported 0 [$r_p(87) = .31$, $p < .01$], and for losses, people who reported -40 were less risk seeking than those who reported 0 [$r_p(86) = -.44$, $p < .001$], even after controlling for any effect of outcomes received.

Frequency judgments showed a similar pattern (Figure 5.3C): people reported significantly higher percentages for the extreme outcome (+40 or -40) than the zero outcome for both gains [$t(106) = 2.70$, $p < .01$, $d = 0.38$] and losses [$t(106) = 6.78$, $p < .001$, $d = 1.05$]. Figure 5.3D plots risk preference in the choice task against frequency judgments for the extreme outcomes (+40 or -40). For gains, risk seeking increased with the judged frequencies of the +40 outcome [$r_p(105) = .16$, $p < .05$], whereas for losses, risk seeking decreased with the judged frequency of the -40 outcome [$r_p(105) = -.48$, $p < .001$], even after controlling for the outcomes received. We also tested for potential primacy and recency effects in the memory tests. For both gains and losses, neither the first outcome experienced nor the last outcome experienced correlated with the results on either memory test (all ps$>.1$, two-tailed).

Nearly half the participants reported the correct proportions of the outcomes (50/50). Nonetheless, these participants showed the same memory biases in the first-outcome question. Of those who correctly reported gains, 23% of subjects reported 0 and 53% reported +40 as the first outcome [$\chi^2(1, N = 45) = 9.80$, $p < .01$]. Of those who correctly reported losses, 73% of subjects reported -40 and 9% reported 0 as the first outcome [$\chi^2(1, N = 36) = 21.78$, $p < .001$; cf. Figure 5.3A]. These participants also showed the same pattern of greater risk seeking for gains (45%) than losses (38%) (cf. Fig. 5.2A). Thus, even participants who accurately reported the contingencies showed biases in memory accessibility and risky choice.

5.3.3 Discussion

In Exp. 1, the extreme outcomes were overweighted in memory and the relative weighting correlated with risky choice across individuals (see Fig. 5.3). These outcomes were the biggest gains and losses experienced, leaving open the question as to whether the absolute or relative extremes are important. Tsetsos et al. (2012), for example, found that both the high and low extremes could be overweighted, even with all gains (see also Ludvig et al., 2014). To test this possibility, Exp. 2 restricted all outcomes to the gain domain by shifting outcomes up by 40 points from Exp. 1. The high extreme was thus +80, and the low extreme was 0. If the relative extremes are critical, then we should see similar overweighting of these extremes in the memory tests and risky choice.

5.4 Experiment 2

5.4.1 Methods

Participants

72 participants were drawn from the same pool as Experiment 1 (47 females; Mage = 19.4 years).

Procedure

The procedure was almost identical to Experiment 1 except all outcomes were gains. On *high-value* decisions, the fixed door led to +60, and the risky door led equiprobably to +40 or +80. On *low-value* decisions, the fixed door led to +20, and the risky door led equiprobably to 0 or +40 (identical to the gain problem of Exp. 1). Thus, the extreme outcomes were +80 (best outcome) and 0 (worst possible). The number and distribution of decision, single-door, and catch trials in each run were the same as in Exp. 1. Catch trials were between a high-value door and a low-value door, with 4 participants excluded because they chose correctly <60% of the time. The session contained 6 blocks of 48 trials. Participants were paid $1 for every 3600 points to a maximum of $5. The memory tests were identical to Exp. 1, except that the frequency test displayed outcomes of 0, +20, +40, +60 and +80.

5.4.2 Results

As predicted, Fig. 5.4A shows how participants were more risk seeking in high-value decisions than in low-value decisions over the final blocks [$t(67) = 9.41$, $p < .001$, $d = 1.44$]. Over blocks, high-value decisions showed no significant change [linear effect of Block: $F(1, 67) = 0.60$, $p > .1$, $\eta_p^2 = .009$], whereas, risk preference decreased for low-value decisions [linear effect of Block: $F(1, 67) = 31.08$, $p < .001$, $\eta_p^2 = .32$] (Figure 5.4B). Figure 5.4C displays how risk seeking was greater following a recent good outcome on the risky option for both the high-value decisions [$t(67) = 2.78$, $p < .01$, $d = 0.13$] and the low-value decisions [$t(67) = 4.61$, $p < .001$, $d = 0.50$]. Independent of recent outcomes, however, risk seeking was greater in the high-value than the low-value decisions.

On the first-outcome memory tests, more people reported the extreme outcomes (80 or 0) than the non-extreme outcomes (40) for both the high-value [$\chi^2(1, N = 61) = 2.77$, $p < .05$] and the low-value risky doors [$\chi^2(1, N = 62) = 47.03$, $p < .001$] (Figure 5.5A). Figure 5.5B shows that for high-value decisions, people who reported 80 were more risk seeking in the choice task than those

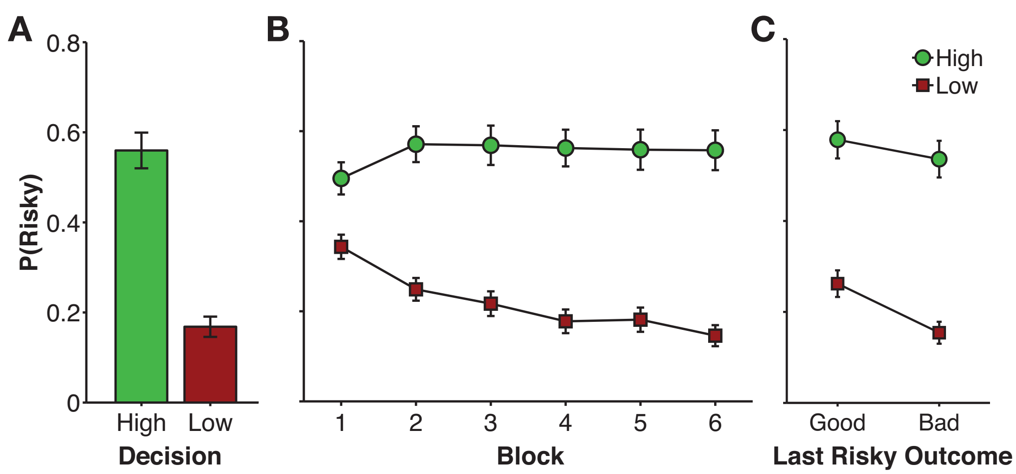

Figure 5.4: Choice results for Experiment 2. A. Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks. B. Mean risk preference (± SEM) for gain and loss doors for each block. C. Mean risk preference (± SEM) for gain and loss doors averaged over the last three blocks, separated by the most recent risky outcome experienced on that door.

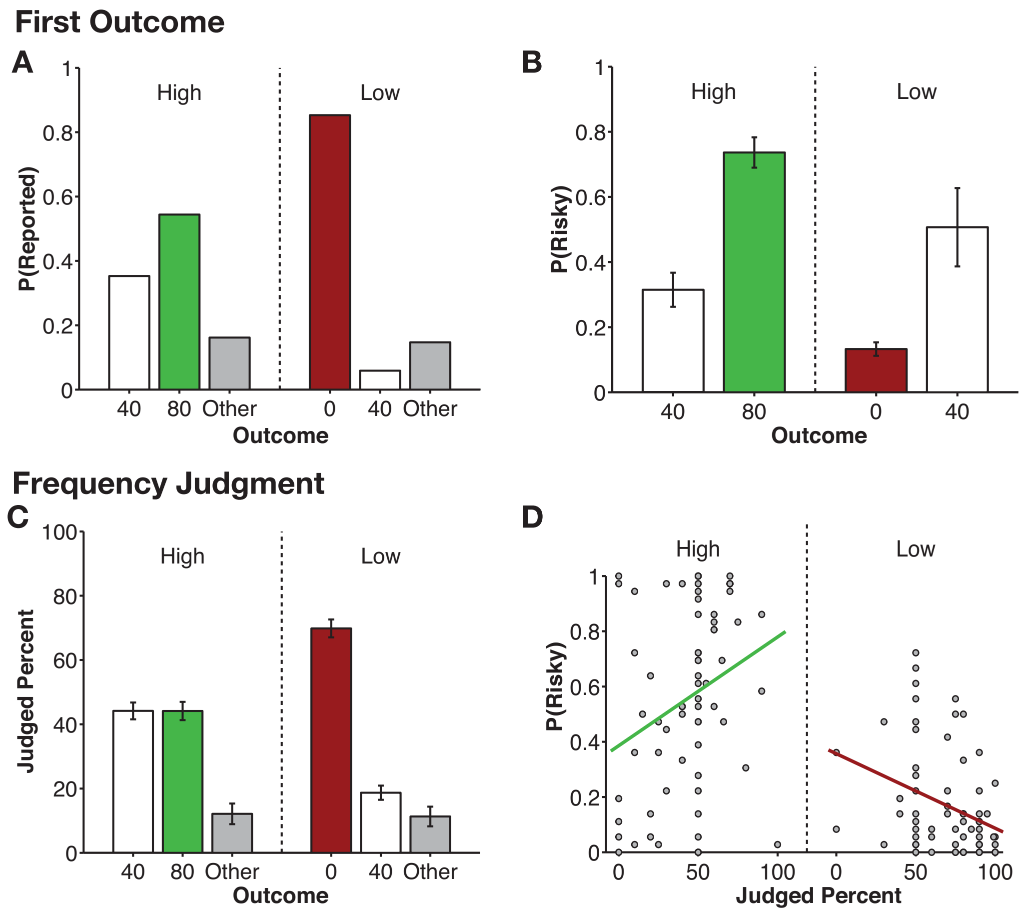

Figure 5.5: Memory results for Experiment 1. A. Proportions of participants that responded with 0, 40, 80, or 'other' in the first-outcome question. B. Mean risk preference (± SEM) for high- and low-value, split based on what the participant reported to the first-outcome question. C. Mean judged percent (± SEM) for the 0, 40, and 80 outcomes from the frequency-judgment question. For simplicity, all other values were coded as "Other". D. Scatterplot of risk preference and frequency judgment responses for the high- and low-value doors. Each dot represents an individual participant.

who reported 40 [$r_p(58) = .58$, $p < .001$]. For low-value decisions, people who reported 0 were less risk seeking than those who reported 40 [$r_p(59) = -.46$, $p < .001$].

For the frequency judgments, the extreme outcome was judged as more frequent than the non-extreme outcomes for the low-value decisions [$t(65) = 11.46$, $p < .001$, $d = 1.41$], but not for the high-value decisions [$t(65) = 0.09$, $p > .1$, $d = 0.01$] (Figure 5.5C). In both cases, however, Fig. 5.5D shows how the judged frequencies correlated with risky choices, even after controlling for the experienced outcomes. Higher judged frequencies for the negative extreme (0) correlated with less risk seeking in the low-value decisions [$r_p(63) = -.24$, $p < .05$], and higher judged frequencies for the positive extreme (+80) correlated with more risk seeking in the high-value decisions [$r_p(63) = .23$, $p < .05$].

5.5 General Discussion

These two experiments provide evidence that a memory bias for extreme outcomes influences risk preference in decisions from experience. As predicted by the extreme-outcome rule, people were more risk seeking for relative gains than losses—contrary to the usual results in decisions from description (Kahneman & Tversky, 1979), but in agreement with recent results in decisions from experience (Ludvig & Spetch, 2011; Ludvig et al., 2014; Tsetsos et al., 2012). Furthermore, people tended to recall the extreme outcomes more readily than the non-extreme outcomes and tended to overestimate their relative frequency. Across individuals, both memory biases correlated with risky choice: overweighting of the high extreme led to more risk seeking, whereas overweighting of the low extreme led to more risk aversion. These results support an extreme-outcome rule, whereby the biggest gains and losses are better remembered and shift risk preferences in decisions from experience.

Our results provide a new addition to the literature on decisions from experience, which has highlighted how rare events are underweighted in experience (Hertwig et al., 2004; Hertwig & Erev, 2009; Weber et al., 2004). In the current study, however, there were no rare events, and both risky options led to equiprobable outcomes. Our extreme-outcome rule does make a clear prediction for cases with rare events: when the rare event is also the extreme outcome in that context, any underweighting should be diminished. Moreover, unlike many studies on decisions from experience, our choice task intermingled a gain and a loss problem or a high-value and a low-value problem. This intermingling established a decision context in which the largest gain and the largest loss followed different risky options, thereby allowing an overweighting of extremes to bias risky choice (see also Ludvig et al., 2014).

The high and low extremes do not, however, carry equal weight. The overweighting of the low extreme in memory was more pronounced. Similarly, when faced with a risky option that led to the low extreme, people were significantly risk averse, but only risk neutral (or moderately risk seeking) when faced with a risky option that led to the high extreme. The risk neutrality is not immediately concordant with our extreme-outcome rule. If the extreme outcome were truly overweighted, then absolute risk seeking should be observed. One possible resolution is that worse outcomes are more heavily weighted than better outcomes, akin to loss aversion, whereby losses loom larger than gains (Kahneman & Tversky, 1979; Yechiam & Hochman, 2013). As a result, when the worse outcome is also the extreme (as with the relative losses), there is significant risk aversion, but when the better outcome is the extreme (as with the relative gains), the risk aversion is reduced (Exp. 1) or reversed (Exp. 2). In other situations when both (or neither) outcomes is extreme, then this negatively-biased weighting would produce risk aversion, as has been observed in other experiments on risky decisions from experience with equiprobable outcomes (Erev et al., 2010; Ert & Yechiam, 2010; Niv et al., 2012).

One possible reason that intermingling multiple problems results in a bias toward the extreme outcomes in both risky choice and memory is that the pairing implicitly introduces a choose/reject problem frame (Shafir, 1993; Tsetsos et al., 2012). Following this idea, in the gain and high-value cases, people focus on which option to choose, whereas, in the loss and low-value cases, people focus on which option to reject. In verbally described problems, people were influenced by the most positive attributes when selecting an option, but were focused on the most negative attributes when rejecting one (Shafir, 1993). In our experiments, people were similarly influenced by the positive extremes in the highest-value decisions, but the negative extremes in the lowest-value decisions (see also Tsetsos et al., 2012).

The observed bias toward remembering extreme values is consistent with other memory studies (Phelps & Sharot, 2008; Talarico & Rubin, 2003; Madan & Spetch, 2012b) and is a partial extension of the peak-end rule to risky choice (e.g., Fredrickson, 2000; Yu et al., 2008). This memory bias could readily be incorporated into recent theories that posit memory retrieval as a key component of decision making, such as the decision-by-sampling framework (Stewart et al., 2006), query theory (Johnson et al., 2007), or instance-based learning (Gonzalez & Dutt, 2011). In demonstrating that extreme outcomes bias both memory and risky choice in the same task, our study provides a novel link between the memory and decision-making literatures.

The two memory tests yielded similar, but not identical, results. In all instances, the first-outcome memory test showed a bias toward the extreme outcome and a significant correlation with the proportion of risky choices. The frequency judgments generally showed a bias toward the extreme outcome, except for the high-value decision in Exp. 2, and the correlations with risky choice, though statistically reliable, were less robust than the first-outcome results. Moreover, the same bias in the first-outcome tests (and risky choice pattern) appeared even for participants who reported the exact 50/50 outcomes in the frequency judgments test. Taken together, these results suggest that the memory effect seems to be more one of relative accessibility than one of explicit misjudgement (though this happens too). There is also the possibility that the biases observed in the memory tests and choice patterns were both produced by a common cause, such as the increased saliency of the extreme outcomes when they occurred (e.g., Niv et al., 2012; Tsetsos et al., 2012). We did not explicitly manipulate nor measure the saliency of the different outcomes, leaving this question open for future research. Finally, independent of these memory biases, there was still significant inter-individual variability in risk preference (Figs. 3D/4E), which likely reflects factors outside of experimental control, such as personality traits or socioeconomic status (e.g., Ginley, Whelan, Meyers, Relyea, & Pearlson, in press; Griskevicius, Tybur, Delton, & Robertson, 2011).

People are continuously confronted with risky decisions—be it shopping at the mall, picking medical treatments, or gambling at a casino. Our research suggests that when people base choices on past experience, their propensity for risk can be influenced by memory biases—in particular, a tendency to remember extreme outcomes. These results highlight the degree to which memory processes can inform and influence our decision making.

5.6 Acknowledgements

Research was funded by grants from the Alberta Gambling Research Institute (AGRI) and the National Science and Engineering Research Council (NSERC) of Canada held by MLS. CRM was supported by scholarships from AGRI and NSERC. EAL was supported by NIH Grant #P30 AG024361. We thank the Explore-Exploit Group at Princeton for insightful discussions and Ashley Rodgers for research help. Door images were extracted from “Irish Doors” on fineartamerica.com with permission from Joe Bonita.

5.7 References

Abdi, H. (2007). Part and partial correlations. In N. J. Salkind (Ed.), *Encyclopedia of measurement and statistics* (pp. 736–740). Thousand Oaks, CA: Sage.

Brown, R., & Kulik, J. (1977). Flashbulb memories. *Cognition*, *5*, 73–99.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523–538.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., ... Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47.

Ert, E., & Yechiam, E. (2010). Consistent constructs in individuals risk taking in decisions from experience. *Acta Psychologica*, *134*, 225–232.

Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion*, *14*, 577–606.

Ginley, M., Whelan, J., Meyers, A., Relyea, G., & Pearlson, G. (in press). Exploring a multidimensional approach to impulsivity in predicting college student gambling. *Journal of Gambling Studies*.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*, 523–551.

Griskevicius, V., Tybur, J., Delton, A., & Robertson, T. (2011). The influence of mortality and socioeconomic status on risk and delayed rewards: A life history theory approach. *Journal of Personality and Social Psychology*, *100*, 1015–1026.

Heilbronner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, *7*, 7.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning Memory & Cognition*, *33*, 461–474.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*, 401–405.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*, *24*, 146–156.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

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.

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

Morewedge, C., Gilbert, D., & Wilson, T. (2005). The least likely of times: How remembering the past biases forecasts of the future. *Psychological Science*, *16*, 626–630.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*, 551–562.

Phelps, E. A., & Sharot, T. (2008). How (and why) emotion enhances the subjective sense of recollection. *Current Directions Psychological Science*, *17*, 147–152.

Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory & Cognition*, *21*, 546–556.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Talarico, J. M., & Rubin, D. C. (2003). Confidence, not consistency, characterizes flashbulb memories. *Psychological Science*, *14*, 455–461.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences USA*, *109*, 9659–9664.

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memory. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*, 497–518.

Yu, E., Lagnado, D., & Chater, N. (2008). Retrospective evaluations of gambling wins: Evidence for a 'peak-end' rule. *Conference of the Cognitive Science Society*, *30*, 77–82.

Chapter 6

General Discussion

The studies reported in this dissertation examined the influence of reward value on memory and decision making, focusing on the incremental learning of reward values for each item through experience. This contrasts with prior research that explicitly instructed participants of an item's value to investigate reward effects on memory (i.e., 'value-directed remembering'; e.g., Adcock et al., 2006; Castel et al., 2002; Watkins & Bloom, 1999; Wittmann et al., 2005) and with studies of decisions made from description (e.g., Kahneman & Tversky, 1979; Weber et al., 2004).

To examine the influence of reward value on memory and decision making, I set out to address the following questions: (1) can previously learned reward values subsequently influence the memory for items in both implicit and explicit memory? (Chapter 2; Madan, Fujiwara, et al., 2012); (2) are reward value effects on memory better explained as reward value or reward salience? (Chapter 3; Madan & Spetch, 2012a); (3) how do more extreme reward outcomes influence decisions made from experience? (Chapter 4; Ludvig et al., 2014); and (4) are reward value effects on decision making driven by memory biases? (Chapter 5; Madan, Ludvig, & Spetch, 2014). To ask these questions, I used two general approaches: In Chapters 2 and 3, I used a memory-based approach where items were trained to be associated with reward values in a preceding value-learning task, which was then followed by memory tests. In Chapters 4 and 5, I used a risky choice task and measured risk preference and it's correspondence with extreme vs. non-extreme outcomes. In this final chapter, I provide a summary of the theoretical and methodological contributions of this work and outline some future directions.

6.1 Summary of Novel Findings

6.1.1 Influences of reward value on implicit and explicit memory

The results presented in Chapter 2 (published in Madan, Fujiwara, et al., 2012) demonstrate that previously learned reward values can persist and influence memory: Words previously associated with high-value rewards were remembered better than those that were previously associated with low-value rewards, even when the memory task itself was unrewarded. These results are important in that they demonstrate an influence of reward value on memory that cannot be explained as reward prioritization, as in many previous studies (e.g., Adcock et al., 2006; Castel et al., 2002; Watkins & Bloom, 1999; Wittmann et al., 2005). Importantly, participants in my study were not aware of the subsequent memory test, but were still biased in recalling and recognizing high-value items more readily than low-value items. However, this memory enhancement did not come without a cost, as high-value items exhibited worse temporal discriminability. These results demonstrated that gradually learned reward values can generalize across experiences and influence memory in different contexts.

6.1.2 Reward effects on memory: Value versus salience

Chapter 3 (published in Madan & Spetch, 2012a) further probed the influence of reward on memory to test the shape of this relationship. Specifically, using a procedure similar to that used in Chapter 2, I used multiple reward values, rather than just two reward values. Here a U-shaped relationship between reward value and memory was found, where the highest- and lowest-value words were remembered better than those that led to intermediate reward values. These results clearly showed that reward does not linearly influence memory, but instead modulates the relative salience of the memory.

Our results are also informed by our understanding of the neurobiology of reward, as prior studies have identified only a few brain regions that are activated following from a U-shaped function (Cooper & Knutson, 2008; Jensen et al., 2007; Litt et al., 2011; Zink et al., 2004). However, when this list is combined with reward-related regions that were identified as supporting memory in previous reward-memory fMRI studies (Adcock et al., 2006; Shigemune et al., 2010; Wittmann et al., 2005), only one brain region was present in both sets of studies: the striatum.

6.1.3 Influences of extreme values on experienced-based decision-making

The studies discussed in Chapter 4 (published in Ludvig et al., 2014) investigated the effects of contextual salience on experience-based decision making. Specifically, recent research found that people's risk preferences are strongly influenced by the format that information is conveyed: If the information is described verbally or learned through experience, risk preferences for gains and losses can reverse (e.g., Ludvig & Spetch, 2011). To elucidate the mechanisms underlying experience-based risky decision-making, where memory necessarily plays a role, we conducted a series of experiments where we systematically varied the range of the reward outcomes experienced. The results suggest that the most extreme outcomes experienced within the experiment (i.e., biggest gains and losses) weighed more heavily than they should have been, which we label as the 'extreme-outcome rule.' This result demonstrated that the most extreme experiences have a disproportionate influence on people's subsequent decisions.

6.1.4 Memory biases may drive risky choices

Chapter 5 (published in Madan et al., 2014) directly tested for relationships between memory biases and risk preference. The results of Chapter 3 indicated that the highest- and lowest-valued words were remembered best. In combination with the extreme-outcome rule proposed in Chapter 4, I thought it would be possible that this same memory bias may be driving risk preferences. To test this, the study presented in Chapter 5 involved a choice task, similar to those used in Chapter 4, followed by two memory tasks to test for accessibility of the experienced outcomes. Results indicated that participants remembered the extreme outcomes better than their non-extreme counterparts in both memory tasks. This memory bias was also found to correlate with risk preferences, such that stronger memory biases for gains correlated with greater risk seeking for gains, while stronger memory biases for losses were related to greater risk aversion for losses. Taken together, these results provide evidence that memory biases influence risk preference. This finding extends prior research showing that memory can influence decision making (e.g., the availability heuristic; Tversky & Kahneman, 1973) to also include memory biases due to reward values.

6.2 Future Directions

While the studies reported in my dissertation provide novel insights into the influences of reward value on memory and decision-making processes, they also raise a number of questions that require further investigation and are the focus of current research. Here I briefly describe some of these questions that need to be pursued to further improve our understanding of memory and decision making.

6.2.1 Risky decision-making in non-human animals

In our daily lives, we make countless decisions. Though we often make these decisions based on our prior experiences, i.e., experience-based decision-making, we often also acquire some information through descriptions of the odds. In contrast, described information cannot be conveyed to non-human animals and their decisions must be made solely based on their prior experiences. Thus, one avenue for future research is to conduct analogous experiments as those conducted in Chapters 4 and 5 (Ludvig et al., 2014; Madan et al., 2014) with non-human animals and test for the generalizability of the extreme-outcome rule. These results will allow us to determine if this rule is a uniquely human characteristic that relies on so-called 'higher-level cognition' or if it is a more basic process that is intrinsic to memory function.

6.2.2 Extreme outcomes and decisions from description

In the decision-making studies presented in this dissertation, only experience-based decision-making was examined. While description-based decisions have been studied more extensively than their experience-based counterpart, our results warrant further research to test the boundary conditions of the extreme-outcome rule. For instance, the results of experiments presented in Chapters 4 and 5 demonstrate that it does not matter if the experienced outcomes are gains or losses *per se*, but rather if the outcomes are relatively the biggest gains and losses within the experienced context. This question can be further applied to decisions from description to see if context will matter here as well, using a design similar to Ludvig and Spetch (2011).

6.2.3 Gambling tendencies and memory biases

As the risky decision-making task used in Chapters 4 and 5 is intended to study a subset of the processes involved in gambling, a next step would be to test for between-population differences between individuals who self-report being frequent gamblers versus those that do not consider themselves as such. While risk preferences for gains and losses, as well as differences between decisions from description and experience, are of interest, testing for differences in memory biases would be the focus of this study. In other words, within the same experimental design as Chapter 5 (Madan et al., 2014), do gamblers demonstrate greater memory biases than non-gamblers? An additional study can also be conducted with more quantitative measures of gambling behaviour, such as the Problem Gambling Severity Index (PGSI; Ferris & Wynne, 2001). Along similar lines, a correlational study of gambling-related personality traits (e.g., impulsivity, risk taking behaviour, reward sensitivity) with our task would also be a fruitful avenue of future research.

6.2.4 Neurobiology of risky decision-making

The studies presented in this dissertation provide insights into the neurobiology of risky decision-making that warrant futher investigation. Specifically, through the use of fMRI, one specific question would be: How are brain regions differentially involved in decisions from description versus experience? One way to view these two types of decisions is with respect to the "hot" and "cold" systems proposed by Metcalfe and Mischel (1999). Briefly, this two-system framework suggests that decisions are made by a combination of two systems, one based on cognitive, emotionally neutral, slow ("cold") processes and one based on emotional, impulsive processes ("hot"). Within this framework, decisions from description would map onto the cold system, based on rational, contemplative decision-making. In contrast, decisions from experience would be made through hot processes, made relatively more impulsively and hedonistcally.

By hypothesizing that decisions from description and experience can map onto the hold-cold systems framework, we can make predictions about the brain regions involved in each type of decision. In an fMRI study of logical reasoning, Goel and Dolan (2003) found a dissociation in the prefrontal cortex, where greater activation was found in the lateral and dorsolateral prefrontal cortex (LPFC and DLPFC, respectively) for cold than hot reasoning. The ventromedial prefrontal cortex (VMPFC) was activated to a greater degree for hot than cold reasoning. Dolcos and McCarthy (2006) found similar results in the prefrontal cortex using an emotional distraction task, but also reported greater lateral parietal cortex (LPC) engagement for cold processing, and greater activity in the amygdala and fusiform gyrus (FFG) for hot processing. A review by Dolcos, Iordan, and Dolcos (2011) came to similar conclusions, as illustrated in Figure 6.1.

Additionally, while several brain regions have been associated with reward-related processing (e.g., Litt et al., 2011), the striatum is the only region that has also been found to be important to reward-related memory effects. Thus, it is likely that the striatum plays a particularly important role in experience-based decisions. In contrast, numerical values are more prominent in decisions made from description. Thus, brain regions associated with numerical processing, such as the intraparietal suclus (IPS; e.g., Cantlon, Brannon, Carter, & Pelphrey, 2006), likely will be activated to a greater degree for decisions from description than decisions from experience.

As a further test of the degree of involvement of these regions in decision making, I plan to test for correlations across participants between risk preference and brain activation. Specifically, for decisions from experience, I predict that individuals that demonstrate greater risk seeking for gains than losses in decisions from experience, i.e., a stronger bias due to extreme outcomes,

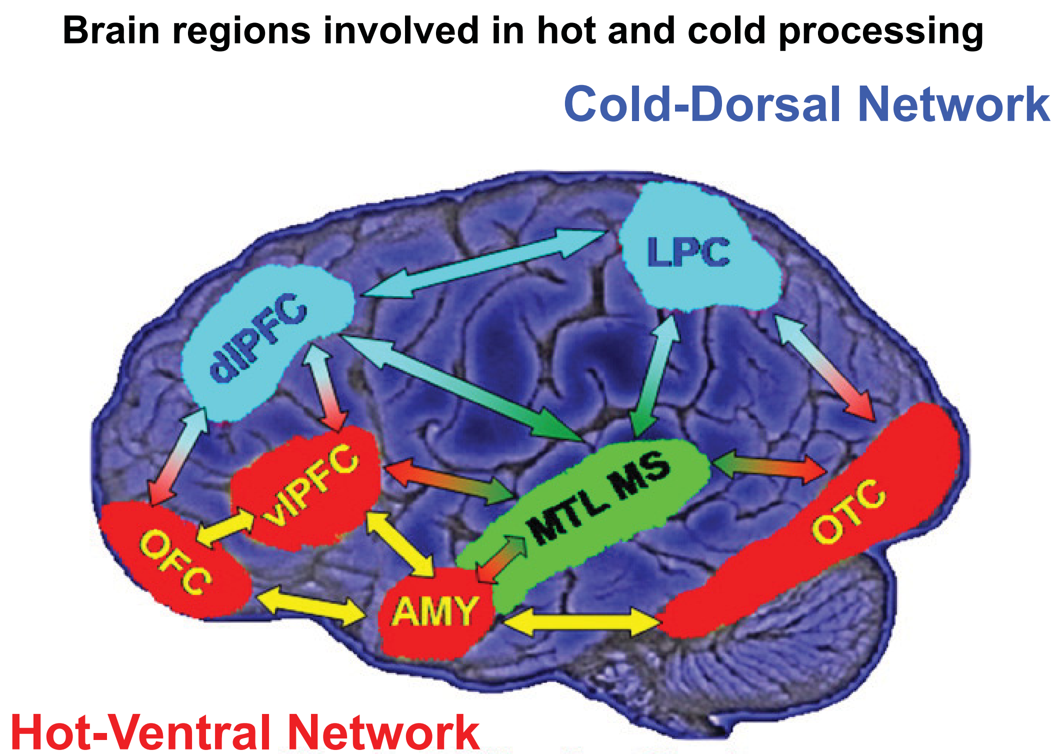

Figure 6.1: Brain regions involved in hot and cold processing. Brain regions activated to a greater degree for 'cold' processing include the dorsolateral prefrontal cortex (dlPFC) and lateral parietal cortex (LPC). Regions activated to a greater degree as part of 'hot' processing include the amygdala (AMY), ventrolateral prefrontal cortex (vlPFC), orbitofrontal cortex (OFC), and occipitotemporal cortex (OTC). The medial temporal lobe memory system (MTL MS) plays a significant role in both types of processing. Figure adapted from Figure 2 of Dolcos et al. (2011).

would also have relatively greater activation of the VMPFC, amygdala, FFG, and striatum, than participants that demonstrate a relatively weaker bias. In decisions from description, I predict individuals who show greater risk seeking for losses than gains (i.e., a reflection effect, as proposed by prospect theory; Kahneman & Tversky, 1979), would also exhibit greater activation of the LPFC, DLPFC, LPC, and IPS.

6.2.5 Causal effects of memory on decision making

Probably the most informative future study of all would be to test directly the causality of memory biases on risky choice. While the experiments presented in Chapter 5 (Madan et al., 2014) provide strong evidence of a relationship between memory and decision making, the directionality of this effect is not tested. One approach to addressing this is would be to prime participants with one outcome just prior to the doors being presented, to experimentally improve the outcome's accessibility during the choice.

In this experiment, instead of presenting the outcomes of choices with the pot of gold and robber images used in Chapters 4 and 5 (see Figures 4.1 and 5.1), I would use unique images for each outcome, as illustrated in Figure 6.2a. By using unique images for each outcome, it is now possible to experimentally improve the accessibility of one outcome over another by cueing the participant with the unique image. To implement this manipulation in the experimental design, I would present a unique images just prior to the decision trials (see Figure 6.2b). Only unique images associated with a possible outcome would be presented (i.e., a unique image associated with a loss outcome would not be presented before a gain decision trial). If the manipulation is successful, the prime images will cause participants to incidentally retrieve the prime's associated outcome value from memory and thereby influence the subsequent choice.

To test this manipulation, I will compare risk preference for trials that are primed with either the good or bad outcome. Given the example shown in Figure 6.2b, I predict greater risk seeking on gain decision trials primed with the pear (good, +40 outcome) than the banana (bad, +0 outcome).

Trials primed with the strawberry (+20 outcome) would serve as a baseline. Similarly, for loss trials I predict greater risk seeking for trials primed with the lemon (good, −0 outcome) than the grapes (bad, −40 outcome); here the watermelon (−20 outcome) would serve as a baseline.

If no difference in risk preference is found based on the prime, this could suggest that the priming procedure was not effective. However, if follow-up studies also fail to produce a priming effect, this would suggest against a bias in memory retrieval as driving the risk preference effects observed against Chapters 4 and 5 and instead indicate that the overweighting of extreme outcomes is driven by another mechanism that co-varies with both risk preference and memory.

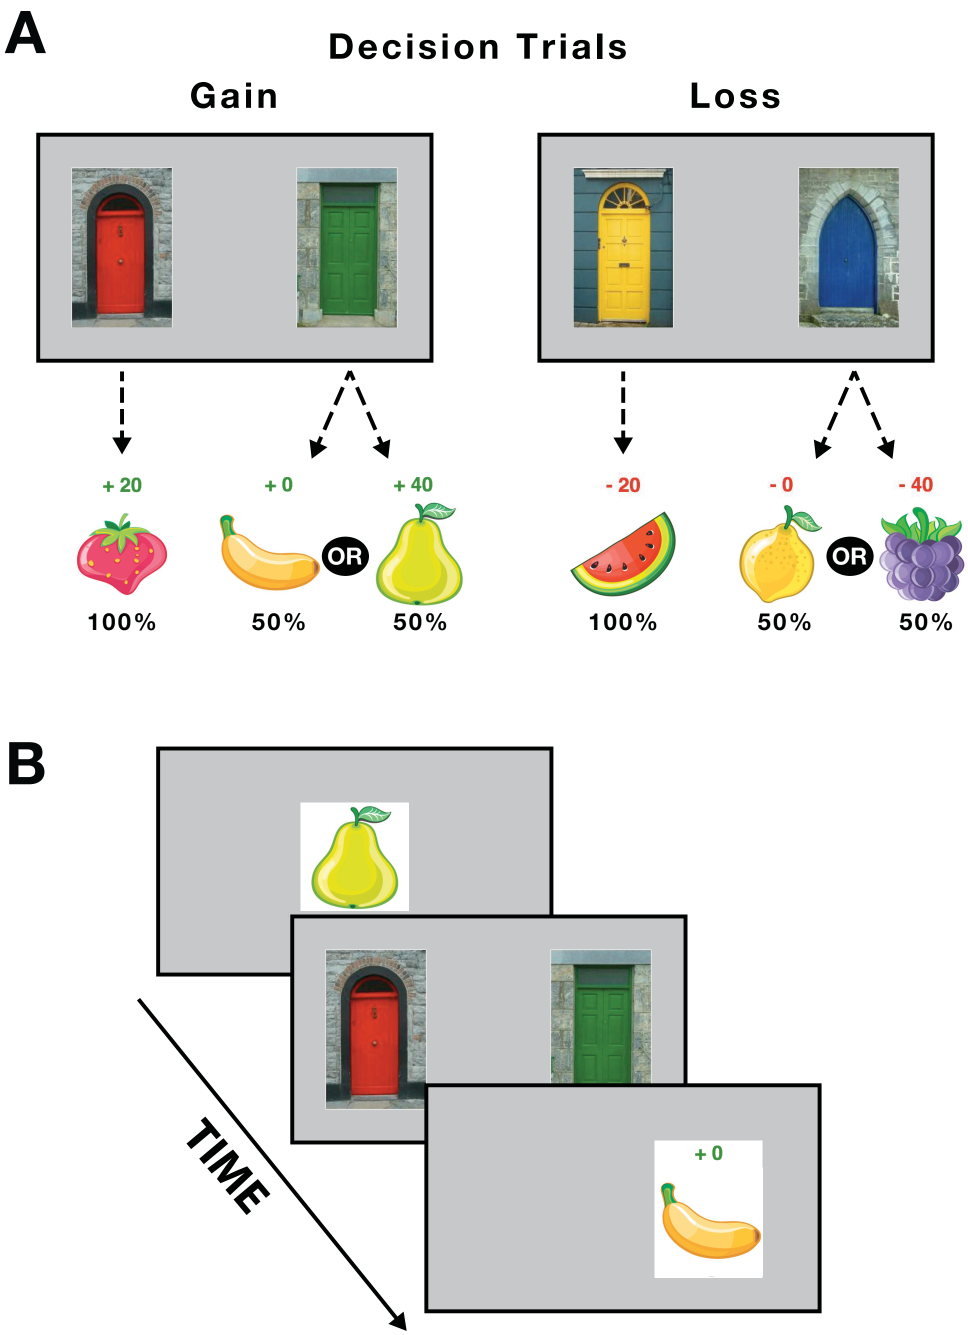

Figure 6.2: Schematic of the method for the priming study. (a) Instead of presenting the pot of gold and robber images, as in Figure 5.1, unique images would be presented for each outcome, such as fruit images. (b) Illustration of a single decision trial in the priming study. Participants will first see the prime image (a fruit image with no outcome value present), followed by the two door images. After the choice, participants will be presented with their outcome (both the fruit image and the outcome value).

6.3 Summary

The ability to remember previous experiences and make decisions based on these remembered experiences is an integral aspect of daily life. Chapters 2 and 3 demonstrate a novel approach for testing for effects of reward value on memory through a choice task. Chapters 4 and 5 provide evidence that the extreme values, relative to the range of values experienced, play an important role in decision making, which may be mediated by memory processes. Together these studies represent a theoretical advance in our understanding of influences of reward value on memory and decision making and lay the foundation for future research.

6.4 References

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517.

Cantlon, J. F., Brannon, E. M., Carter, E. J., & Pelphrey, K. A. (2006). Functional imaging of numerical processing in adults and 4-y-old children. *PLoS Biology*, *4*, 844–854.

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*, 1078–1085.

Cooper, J. C., & Knutson, B. (2008). Valence and salience contribute to nucleus accumbens activation. *NeuroImage*, *39*, 538–547.

Dolcos, F., Iordan, A. D., & Dolcos, S. (2011). Neural correlates of emotion–cognition interactions: A review of evidence from brain imaging investigations. *Journal of Cognitive Psychology*, *23*, 669–694.

Dolcos, F., & McCarthy, G. (2006). Brain systems mediating cognitive interference by emotional distraction. *Journal of Neuroscience*, *26*, 2072–2079.

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

Goel, V., & Dolan, R. J. (2003). Reciprocal neural response within lateral and ventral medial prefrontal cortex during hot and cold reasoning. *NeuroImage*, *20*, 2314–2321.

Jensen, J., Smith, A. J., Willeit, M., Crawley, A. P., Mikulis, D. J., Vitcu, I., & Kapur, S. (2007). Separate brain regions code for salience vs. valence during reward prediction in humans. *Human Brain Mapping*, *28*, 294–302.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Litt, A., Plassmann, H., Shiv, B., & Rangel, A. (2011). Dissociating valuation and saliency signals during decision-making. *Cerebral Cortex*, *21*, 95–102.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*, *24*, 146–156.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

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.

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

Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, *106*, 3–19.

Shigemune, Y., Abe, N., Suzuki, M., Ueno, A., Mori, E., Tashiro, M., . . . Fujii, T. (2010). Effects of emotion and reward motivation on neural correlates of episodic memory encoding: A PET study. *Neuroscience Research*, *67*, 72–79.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*, 207–232.

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

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

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*, 459–467.

Zink, C. F., Pagnoni, G., Martin-Skurski, M. E., Chappelow, J. C., & Berns, G. S. (2004). Human striatal response to monetary reward depend on saliency. *Neuron*, *42*, 509–517.

Works Cited

Abdi, H. (2007). Part and partial correlations. In N. J. Salkind (Ed.), *Encyclopedia of measurement and statistics* (pp. 736–740). Thousand Oaks, CA: Sage.

Abler, B., Walter, H., & Erk, S. (2005). Neural correlates of frustration. *Neuroreport*, *16*, 669–672.

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron*, *50*, 507–517.

Allen, G. A., & Estes, K. (1972). Acquisition of correct choices and value judgements in binary choice learning with differential rewards. *Psychonomic Science*, *27*, 68–72.

Anderson, A. K. (2005). Affective influences on the attentional dynamics supporting awareness. *Journal of Experimental Psychology: General*, *134*, 258–281.

Anderson, N. H. (1965). Primacy effects in personality impression formation using a generalized order effect paradigm. *Journal of Personality and Social Psychology*, *2*, 1–9.

Andre, T. (1975). The priority of new items effect: A strategy model. *Journal of General Psychology*, *92*, 187–204.

Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258–290.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233.

Battig, W. F., Allen, M., & Jensen, A. R. (1965). Priority of free recall of newly learned items. *Journal of Verbal Learning and Verbal Behavior*, *4*, 175–179.

Bayley, P. J., Frascino, J. C., & Squire, L. R. (2005). Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature*, *436*, 550–553.

Bijleveld, E., Custers, R., & Aarts, H. (2009). The unconcious eye opener: Pupil dilation reveals strategic recruitment of resources upon presentation of subliminal reward cues. *Psychological Science*, *20*, 1313–1315.

Bijleveld, E., Custers, R., & Aarts, H. (2010). Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. *Cog-*

nition, *115*, 330–335.

Bjork, R. A., & Whitten, W. B. (1974). Recency-sensitive retrieval processes in long-term free recall. *Cognitive Psychology*, *6*, 173–189.

Bjork, R. A., & Woodward, Jr., A. E. (1973). Directed forgetting of individual words in free recall. *Journal of Experimental Psychology*, *99*, 22–27.

Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings* (Technical report No. C-1). Gainesville, FL: The Center for Research in Psychophysiology, University of Florida.

Brainerd, C. J., Reyna, V. F., & Howe, M. L. (1990). Children's cognitive triage: Optimal retrieval or effortful processing? *Journal of Experimental Child Psychology*, *49*, 428–447.

Brainerd, C. J., Reyna, V. F., Howe, M. L., & Kevershan, J. (1990). The last shall be first: How memory strength affects children's retrieval. *Psychological Science*, *1*, 247–252.

Brainerd, C. J., Reyna, V. F., Howe, M. L., & Kevershan, J. (1991). Fuzzy-trace theory and cognitive triage in memory development. *Developmental Psychology*, *27*, 351–369.

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*, 539–576.

Brown, R., & Kulik, J. (1977). Flashbulb memories. *Cognition*, *5*, 73–99.

Bruce, D. (1985). The how and why of ecological memory. *Journal of Experimental Psychology: General*, *114*, 78.

Budson, A. E., Todman, R. W., Chong, H., Adams, E. H., Kensinger, E. A., Krangel, T. S., & Wright, C. I. (2006). False recognition of emotional word lists in aging and alzheimer disease. *Cognitive and Behavioral Neurology*, *19*, 71–78.

Burke, A., Heuer, F., & Reisberg, D. (1992). Remembering emotional events. *Memory & Cognition*, *20*, 277–290.

Burnham, K. E., & Anderson, D. R. (2002). *Model selection and multimodel inference* (2nd ed.). New York: Springer-Verlag.

Camilleri, A. R., & Newell, B. R. (2009). Within-subject preference reversals in description- and experience-based choice. *Proceedings of the Annual Conference of the Cognitive Science Society*, *31*, 449–454.

Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*, 377–384.

Cantlon, J. F., Brannon, E. M., Carter, E. J., & Pelphrey, K. A. (2006). Functional imaging of numerical processing in adults and 4-y-old children. *PLoS Biology*, *4*, 844–854.

Castel, A. D. (2008). The adaptive and strategic use of memory by older adults: Evaluative processing and value-directed remembering. In A. S. Benjamin & B. H. Ross (Eds.), *The Psychology of Learning and Motivation, Volume 48* (pp. 225–270). London: Academic Press.

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*, 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*, 1078–1085.

Castel, A. D., Farb, N., & 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*, 689–700.

Christianson, S.-A. (1992). Emotional-stress and eyewitness memory: A critical-review. *Psychological Bulletin*, *112*, 284–309.

Cooper, J. C., & Knutson, B. (2008). Valence and salience contribute to nucleus accumbens activation. *NeuroImage*, *39*, 538–547.

Corey, D. M., Dunlap, W. P., & Burke, M. J. (1998). Averaging correlations: Expected values and bias in combined Pearson rs and Fisher's z transformations. *Journal of General Psychology*, *125*, 245–261.

Corson, Y., & Verrier, N. (2007). Emotions and false memories: Valence or arousal? *Psychological Science*, *18*, 208–211.

Cromwell, H. (1950). The relative effect on audience attitude of the first versus the second argumentative speech of a series. *Communications Monographs*, *17*, 105–122.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Custers, R., & Aarts, H. (2010). The unconcious will: How the pursuit of goals operates outside conscious awareness. *Science*, *329*, 47–50.

Dehon, H., Laroi, F., & Van der Linden, M. (2010). Affective valence influences participant's susceptibility to false memories and illusory recollection. *Emotion*, *10*, 627–639.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*, 523–538.

Dewhurst, S. A., & Parry, L. A. (2000). Emotionality, distinctiveness, and recollective experience. *European Journal of Cognitive Psychology*, *12*, 541–551.

Dolcos, F., Iordan, A. D., & Dolcos, S. (2011). Neural correlates of emotion–cognition interactions: A review of evidence from brain imaging investigations. *Journal of Cognitive Psychology*, *23*, 669–694.

Dolcos, F., & McCarthy, G. (2006). Brain systems mediating cognitive interference by emotional distraction. *Journal of Neuroscience*, *26*, 2072–2079.

Easterbrook, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review*, *66*, 183–201.

Ebbinghaus, H. (1885/1913). *Memory: A contribution to experimental psychology.* New York: Teachers College, Columbia University.

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., . . . Lebiere, C. (2010). A choice prediction competition: Choices from experience and

from description. *Journal of Behavioral Decision Making*, *23*, 15–47.
Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*, 575–597.
Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*, 153–177.
Eriksen, C. W., & Hake, H. W. (1957). Anchor effects in absolute judgments. *Journal of Experimental Psychology*, *53*, 132–138.
Ert, E., & Yechiam, E. (2010). Consistent constructs in individuals risk taking in decisions from experience. *Acta Psychologica*, *134*, 225–232.
Estes, W. K. (1962). Theoretical treatments of differential reward in multiple-choice learning and two-person interactions. In J. Criswell, H. Solomon, & P. Suppes (Eds.), *Mathematical methods in small group processes.* Stanford, CA: Standord University Press.
Estes, W. K. (1966). Transfer of verbal discriminations based on differential reward magnitudes. *Journal of Experimental Psychology*, *72*, 276–283.
Estes, W. K. (1972). Reinforcement in human behavior: Reward and punishment influence human actions via informational and cybernetic processes. *American Scientist*, *60*, 723–729.
Eysenck, M. W., & Eysenck, M. C. (1982). Effects of incentive on cued recall. *The Quarterly Journal of Experimental Psychology Section A*, *34*, 489–498.
Ferris, J., & Wynne, H. (2001). *The Canadian problem gambling index* (Final report). Ottawa, ON: Canadian Centre on Substance Abuse.
Fisher, R. A. (1921). On the 'probable error' of a coefficient of correlation deduced from a small sample. *Metron*, *1*, 1–32.
Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling theory + prospect theory: Reconsidering Hertwig, Barron, Weber, & Erev (2004). *Judgment & Decision Making*, *1*, 159–161.
Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns : Midazolam enhances implicit inference in humans. *Psychological Science*, *17*, 700–707.
Frank, M. J., Seeberger, L. C., & O'Reilly, R. C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science*, *306*, 1940–1943.
Fredrickson, B. L. (1998). What good are positive emotions? *Review of General Psychology*, *2*, 300–319.
Fredrickson, B. L. (2000). Extracting meaning from past affective experiences: The importance of peaks, ends, and specific emotions. *Cognition & Emotion*, *14*, 577–606.
Friedman, P., & Gildemeister, J. E. (1980). A negative priority effect. *Perceptual and Motor Skills*, *51*, 423–430.
Geller, A. S., Schleifer, I. K., Sederberg, P. B., Jacobs, J., & Kahana, M. J. (2007). PyEPL: A Cross-Platform Experiment-Programming Library. *Behavior Research Methods, Instruments, & Computers*, *65*, 50–64.

Ginley, M., Whelan, J., Meyers, A., Relyea, G., & Pearlson, G. (in press). Exploring a multidimensional approach to impulsivity in predicting college student gambling. *Journal of Gambling Studies*.

Glenberg, A. M. (1997). What is memory for. *Behavioral and Brain Sciences*, *20*, 1–55.

Goel, V., & Dolan, R. J. (2003). Reciprocal neural response within lateral and ventral medial prefrontal cortex during hot and cold reasoning. *NeuroImage*, *20*, 2314–2321.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*, 523–551.

Gopie, N., Craik, F. I. M., & Hasher, L. (2011). A double dissociation of implicit and explicit memory in younger and older adults. *Psychological Science*, *22*, 634–640.

Gradin, V. B., Kumar, P., Waiter, G., Ahearn, T., Stickle, C., Milders, M., ... Steele, J. D. (2011). Expected value and prediction error abnormalities in depression and schizophrenia. *Brain*, *134*, 1751–1764.

Griskevicius, V., Tybur, J., Delton, A., & Robertson, T. (2011). The influence of mortality and socioeconomic status on risk and delayed rewards: A life history theory approach. *Journal of Personality and Social Psychology*, *100*, 1015–1026.

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

Hadley, C. B., & MacKay, D. G. (2006). Does emotion help or hinder immediate memory? arousal versus priority-binding mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 79–88.

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

Hasher, L. (1973). Position effects in free recall. *American Journal of Psychology*, *86*, 389–397.

Hayden, B., Heilbronner, S. R., Nair, A. C., & Platt, M. L. (2008). Cognitive influences on risk-seeking by rhesus macaques. *Judgment & Decision Making*, *3*, 389–395.

Heath, T. B., & Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, *22*, 268–284.

Heilbronner, S. R., & Hayden, B. Y. (2013). Contextual factors explain risk-seeking preferences in rhesus monkeys. *Frontiers in Neuroscience*, *7*, 7.

Helson, H. (1947). Adaptation-level as frame of reference for prediction of psychophysical data. *The American Journal of Psychology*, *60*, 1–29.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523.

Highhouse, S., & Gallo, A. (1997). Order effects in personnel decision making. *Human Performance*, *10*, 31–46.

Hilbert, M. (2012). Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making. *Psychological Bulletin*, *138*, 211–237.

Hovland, C. I., Campbell, E. H., & Brock, T. (1957). The effects of "commitment" on opinion change following communication. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 23–32). New Haven: Yale University Press.

Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*, 923-941.

Howard, M. W., & Kahana, M. J. (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, *46*, 85–98.

Howe, M. L., & Otgaar, H. (2013). Proximate mechanisms and the development of adaptive memory. *Current Directions in Psychological Science*, *22*(1), 16–22.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90–98.

Hubert, L. J., & Levin, J. R. (1978). Evaluating priority effects in free recall. *British Journal of Mathematical and Statistical Psychology*, *31*, 11–18.

Humphreys, M. S., Allen, G. A., & Estes, W. K. (1968). Learning of two-choice, differential reward problems with informational constraints on payoff combinations. *Journal of Mathematical Psychology*, *5*, 260–280.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review*, *96*, 208–233.

Jensen, J., Smith, A. J., Willeit, M., Crawley, A. P., Mikulis, D. J., Vitcu, I., & Kapur, S. (2007). Separate brain regions code for salience vs. valence during reward prediction in humans. *Human Brain Mapping*, *28*, 294–302.

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning Memory & Cognition*, *33*, 461–474.

Johnsrude, I. S., Owen, A. M., White, N. M., Zhao, W. V., & Bohbot, V. (2000). Impaired preference conditioning after anterior temporal lobe resection in humans. *Journal of Neuroscience*, *20*, 2649–2656.

Johnsrude, I. S., Owen, A. M., Zhao, W. V., & White, N. M. (1999). Conditioned preference in humans: A novel experimental approach. *Learning and Motivation*, *30*, 250–264.

Kacelnik, A., & Bateson, M. (1996). Risky theories: The effects of variance on foraging decisions. *American Zoologist*, *36*, 402–434.

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*, 401–405.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–292.

Kensinger, E. A., & Corkin, S. (2003). Memory enhancement for emotional words: are emotional words more vividly remembered than neutral words? *Memory & Cognition*, *31*, 1169–1180.

Klein, S. B. (in press). The temporal orientation of memory: It's time for a change of direction. *Journal of Applied Research in Memory and Cognition*.

Kuhl, B. A., Shah, A. T., DuBrow, S., & Wagner, A. D. (2010). Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *Nature Neuroscience*, *13*, 501–508.

Lawson, R. G. (1968). Order of presentation as a factor in jury persuasion. *Kentucky Law Journal*, *56*, 523–555.

Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron*, *46*, 703–713.

Litt, A., Plassmann, H., Shiv, B., & Rangel, A. (2011). Dissociating valuation and saliency signals during decision-making. *Cerebral Cortex*, *21*, 95–102.

Loftus, E. F., Loftus, G. R., & Messo, J. (1987). Some facts about 'weapon focus'. *Law and Human Behavior*, *11*, 55–62.

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

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

Lotto, R. B., & Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy Sciences USA*, *97*, 12834–12839.

Luchins, A. S. (1957). Primacy-recency in impression formation. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 33–61). New Haven: Yale University Press.

Ludvig, E. A., Bellemare, M. G., & Pearson, K. G. (2011). A primer on reinforcement learning in the brain: Psychological, computational, and neural perspectives. In E. Alonso & E. Mondragon (Eds.), *Computational neuroscience for advancing artificial intelligence: Models, methods and applications* (pp. 111–144). Hershey, PA: IGI Global.

Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*, *24*, 146–156.

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PLoS One*, *6*, e20262.

Lund, F. H. (1925). The psychology of belief: A study of its emotional, and volitional determinants. *Journal of Abnormal and Social Psychology*, *20*,

63–81 and 183–191.
Madan, C. R., Caplan, J. B., Lau, C. S. M., & Fujiwara, E. (2012). Emotional arousal does not enhance association-memory. *Journal of Memory and Language*, *66*, 695–716.
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.
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., & Spetch, M. L. (2012a). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica*, *139*, 343–349.
Madan, C. R., & Spetch, M. L. (2012b). Is the enhancement of memory due to reward driven by value or salience? *Acta Psychologica*, *139*, 343–349.
Maratos, E. J., Allan, K., & Rugg, M. D. (2000). Recognition memory for emotionally negative and neutral words: an erp study. *Neuropsychologia*, *38*, 1452–1465.
March, J. G. (1994). *Primer on decision making: How decisions happen.* New York: The Free Press.
March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*, 309–319.
Mather, M., & Knight, M. (2008). The emotional harbinger effect: Poor context memory for cues that previously predicted something arousing. *Emotion*, *8*, 850–860.
Mather, M., & Sutherland, M. R. (2011). Arousal-biased competition in perception and memory. *Perspectives on Psychological Science*, *6*, 114–133.
Mathewson, K. J., Arnell, K. M., & Mansfield, C. A. (2008). Capturing and holding attention: The impact of emotional words in rapid serial visual presentation. *Memory & Cognition*, *36*, 182–200.
May, C. P., Hasher, L., & Foong, N. (2005). Implicit memory, age, and time of day: Paradoxical priming effect. *Psychological Science*, *16*, 96–100.
McCoy, A. N., & Platt, M. L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. *Nature Neuroscience*, *8*, 1220–1227.
McGuire, W. J. (1957). Order of presentation as a factor in "conditioning" persuasiveness. In C. I. Hovlund et al. (Ed.), *The order of presentation in persuasion* (pp. 98–114). New Haven: Yale University Press.
Medin, D. L. (1972a). Partial information and choice behaviour in differential reward magnitude learning. *Psychonomic Science*, *27*, 73–76.
Medin, D. L. (1972b). Role of reinforcement in discrimination learning set in monkeys. *Psychological Bulletin*, *77*, 305–318.
Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, *106*, 3–19.
Miller, N., & Campbell, D. T. (1959). Recency and primacy in persuasion as a function of the timing of speeches and measurements. *Journal of*

Abnormal and Social Psychology, *59*, 1–9.

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51–92). San Diego: Academic Press.

Mitchell, T. R., Thompson, L., Peterson, E., & Cronk, R. (1997). Temporal adjustments in the evaluation of events: The "rosy view". *Journal of Experimental Social Psychology*, *33*, 421–448.

Morewedge, C., Gilbert, D., & Wilson, T. (2005). The least likely of times: How remembering the past biases forecasts of the future. *Psychological Science*, *16*, 626–630.

Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods*, *11*, 47–60.

Müller-Lyer, F. C. (1889). Optische urteilstäuschungen. *Archiv für Physiologie*, *Suppl*, 263–270.

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482–488.

Nairne, J. S. (2005). The functionalist agenda in memory research. In A. F. Healy (Ed.), *Experimental cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer.* Washington, DC: American Psychological Association.

Nairne, J. S. (2010). Adaptive memory: Evolutionary constraints on remembering. *Psychology of Learning and Motivation*, *53*, 1–32.

Nashiro, K., Sakaki, M., Huffman, D., & Mather, M. (2012). Both younger and older adults have difficulty updating emotional memories. *Journal of Gerontology: Psychological Sciences*, in press.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*, 551–562.

North, A. C., Hargreaves, D. J., & McKendrick, J. (1997). In-store music affects product choice. *Nature*, *390*, 132.

North, A. C., Hargreaves, D. J., & McKendrick, J. (1999). The influence of in-store music on wine selections. *Journal of Applied psychology*, *84*, 271–276.

Novak, D. L., & Mather, M. (2009). The tenacious nature of memory binding for arousing negative items. *Memory & Cognition*, *37*, 945–952.

Okada, G., Okamoto, Y., Kunisato, Y., Aoyama, S., Nishiyama, Y., Yoshimura, S., ... Yamawaki, S. (2011). The effect of negative and positive emotionality on associative memory: An fMRI study. *PLoS ONE*, *6*, e24862.

O'Neill, M., & Schultz, W. (2010). Coding of reward risk by orbitofrontal neurons is mostly distinct from coding of reward value. *Neuron*, *68*, 789–800.

Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., & Frith, C. D. (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science*, *316*, 904–906.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042–1045.

Pesta, B. J., Murphy, M. D., & Sanders, R. E. (2001). Are emotionally charged lures immune to false memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 328–338.

Petrov, A. A., & Anderson, J. R. (2005). The dynamics of scaling: A memory-based anchor model of category rating and absolute identification. *Psychological Review*, *112*, 383–416.

Phelps, E. A., & Sharot, T. (2008). How (and why) emotion enhances the subjective sense of recollection. *Current Directions Psychological Science*, *17*, 147–152.

Pierce, B. H., & Kensinger, E. A. (2011). Effects of emotion on associative recognition: Valence and retention interval matter. *Emotion*, *11*, 139–144.

Poldrack, R. A., & Packard, M. G. (2003). Competition among multiple memory systems: converging evidence from animal and human brain studies. *Neuropsychologica*, *41*, 245–251.

Postman, L., & Keppel, G. (1968). Conditions determining the priority of new items in free recall. *Journal of Verbal Learning and Verbal Behavior*, *7*, 260–263.

Pubols, B. H. (1960). Incentive magnitude, learning, and performance in animals. *Psychological Bulletin*, *57*, 89–115.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, *88*, 93–134.

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior And Human Decision Processes*, *106*, 168–179.

Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation. *Psychological Science*, *20*, 981–988.

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*, 3–8.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning ii* (pp. 64–99). New York: Appleton-Century-Crofts.

Rosnow, R. L. (1966). Whatever happened to the "law of primacy"? *Journal of Communication*, *16*, 10–31.

Rugg, M. D., Mark, R. E., Walla, P., Schloerscheidt, A. M., Birch, C. S., & Allan, K. (1998). Dissociation of the neural correlates of implicit and explicit memory. *Nature*, *392*, 595–598.

Rundus, D. (1974). Output order and rehearsal in multi-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, *13*, 656–663.

Sakaki, M., & Mather, M. (2012). How reward and emotional stimuli induce different reactions across the menstrual cycle. *Social and Personality Psychology Compass*, *6*, 1–17.

Sakaki, M., Niki, K., & Mather, M. (2011). Updating existing emotional memories involves the frontopolar/orbito-frontal cortex in ways that acquiring new emotional memories does not. *Journal of Cognitive Neuroscience*, *23*, 3498–3514.

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron*, *76*(4), 677–694.

Schott, B. H., Henson, R. N., Richardson-Klavehn, A., Becker, C., Thoma, V., Heinze, H.-J., & Düzel, E. (2005). Redefining implicit and explicit memory: The functional neuroanatomy of priming, remembering, and control of retrieval. *Proceedings of the National Academy of Sciences*, *102*, 1257–1262.

Schott, B. H., Richardson-Klavehn, A., Henson, R. N. A., Becker, C., Heinze, H.-J., & Düzel, E. (2006). Neuroanatomical dissociation of encoding processes related to priming and explicit memory. *Journal of Neuroscience*, *26*, 792–800.

Schwabe, L., Dickinson, A., & Wolf, O. T. (2011). Stress, habits, and drug addiction: A psychoneuroendocrinological perspective. *Experimental and Clinical Psychopharmacology*, *19*, 53–63.

Schwabe, L., Höffken, O., Tegenthoff, M., & Wolf, O. T. (2011). Preventing the stress-induced shift from goal-directed to habit action with a β-adrenergic antagonist. *Journal of Neuroscience*, *31*, 17317–17325.

Schwabe, L., Joëls, M., Roozendaal, B., Wolf, O. T., & Oitzl, M. S. (2011). Stress effects on memory: An update and integration. *Neuroscience and Biobehavioural Reviews*, in press.

Schwabe, L., & Wolf, O. T. (2011). Stress-induced modulation of instrumental behavior: From goal-directed to habitual control of action. *Behavioural Brain Research*, *219*, 321–328.

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, *61*, 195–202.

Sescousse, G., Caldú, X., Segura, B., & Dreher, J.-C. (2013). Processing of primary and secondary rewards: A quantitative meta-analysis and review of human functional neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*, 681–696.

Sescousse, G., Redouté, J., & Dreher, J.-C. (2010). The architecture of reward value coding in the human orbitofrontal cortex. *Journal of Neuroscience*, *30*, 13095–13104.

Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory & Cognition*, *21*, 546–556.

Shaheen, J. (2010). The recency and primacy effects in the talent acquisition process. *Journal of Corporate Recruiting Leadership*, *5*, 13–15.

Shampanier, K., Mazur, N., & Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, *26*, 742–757.

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

Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological review*, *94*(4), 439.

Shettleworth, S. J. (2009). *Cognition, evolution and behavior* (2nd ed.). New York: Oxford University Press.

Shigemune, Y., Abe, N., Suzuki, M., Ueno, A., Mori, E., Tashiro, M., . . . Fujii, T. (2010). Effects of emotion and reward motivation on neural correlates of episodic memory encoding: A PET study. *Neuroscience Research*, *67*, 72–79.

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

Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology: General*, *142*, 476–488.

Siddiqui, A. P., & Unsworth, N. (2011). Investigating the role of emotion during the search process in free recall. *Memory & Cognition*, *39*, 1387–1400.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, *16*, 158–174.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, *29*, 281–295.

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*, 1236–1242.

Staudinger, M. R., Erk, S., & Walter, H. (2011). Dorsolateral prefrontal cortex modulates striatal reward encoding during reappraisal of reward anticipation. *Cerebral Cortex*, *21*, 2578–3588.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology: General*, *132*, 23–46.

Stone, A. A., Schwartz, J. E., Broderick, J. E., & Shiffman, S. S. (2005). Variability of momentary pain predicts recall of weekly pain: A consequence of the peak (or salience) memory heuristic. *Personality and Social Psychology Bulletin*, *31*, 1340–1346.

Stone, V. A. (1969). A Primacy Effect in DecisionMaking by Jurors. *Journal of Communication*, *19*, 239–247.

Storbeck, J., & Clore, G. L. (2005). With sadness comes accuracy; with happiness, false memory: Mood and the false memory effect. *Psychological Science*, *16*, 785–791.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, *114*, 759–783.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press.

Talarico, J. M., & Rubin, D. C. (2003). Confidence, not consistency, characterizes flashbulb memories. *Psychological Science*, *14*, 455–461.

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

Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, *36*, 643–660.

Thomas, D. R., & Jones, C. (1962). Stimulus generalization as a function of the frame of reference. *Journal of Experimental Psychology*, *64*, 77–80.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision making under risk. *Science*, *315*, 515–518.

Tops, M., van der Pompe, G., Baas, D., Mulder, L. J. M., den Boer, J. A., Meijman, T. F., & Korf, J. (2003). Acute cortisol effects on immediate free recall and recognition of nouns depend on stimulus valence. *Psycophysiology*, *40*, 167–173.

Touryan, S. R., Marian, D. E., & Shimamura, A. P. (2007). Effect of negative emotional pictures on associative memory for peripheral information. *Memory*, *15*, 154–166.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences USA*, *109*, 9659–9664.

Tsukiura, T., & Cabeza, R. (2008). Orbitofrontal and hippocampal contributions to memory for facename associations: The rewarding power of a smile. *Neuropsychologica*, *46*, 2310–2319.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, *5*, 207–232.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)?

Psychological Science, *20*, 473–479.

Valentin, V. V., & O'Doherty, J. P. (2009). Overlapping prediction errors in dorsal striatum during instrumental learning with juice and money reward in the human brain. *Journal of Neurophysiology*, *102*, 3384–3391.

Vasconcelos, M., & Urcuioli, P. J. (2008). Deprivation level and choice in pigeons: A test of within-trial contrast. *Learning & Behavior*, *36*, 12–18.

Voon, V., Pessiglione, M., Brezing, C., Gallea, C., Fernandez, H. H., Dolan, R. J., & Hallett, M. (2010). Mechanisms underlying dopamine-mediated reward bias in compulsive behaviors. *Neuron*, *65*, 135–142.

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

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memory. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445.

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

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

Weld, H. P., & Roff, M. (1938). A study in the formation of opinion based upon legal evidence. *American Journal of Psychology*, *51*, 609–628.

Westbury, C., Hollis, G., & Shaoul, C. (2007). LINGUA: The Language-Independent Neighbourhood Generator of the University of Alberta. *The Mental Lexicon*, *2*, 273–286.

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

Wimmer, G. E., & Shohamy, D. (2011). The striatum and beyond: Hippocampal contributions to decision making. In M. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Attention & Performance XXII.* Oxford, UK: Oxford University Press.

Windmann, S., & Kutas, M. (2001). Electrophysiological correlates of emotion-induced recognition bias. *Journal of Cognitive Neuroscience*, *13*, 577–592.

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., Schiltz, K., Boehler, C. N., & Düzel, E. (2008). Mesolimbic interaction of emotional valence and reward improves memory formation. *Neuropsychologica*, *46*, 1000–1008.

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*, 459–467.

Wixted, J. T., Ghadisha, H., & Vera, R. (1997). Recall latency following pure- and mixed-strength lists: A direct test of the relative strength model of free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 523–538.

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*, 1532–1547.

Yacubian, J., Gläscher, J., Schroeder, K., Sommer, T., Braus, D. F., & Büchel, C. (2006). Dissociable systems for gain and loss related value predictions and errors of prediction in the human brain. *Journal of Neuroscience*, *26*, 9530–9537.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, *139*, 497–518.

Yu, E., Lagnado, D., & Chater, N. (2008). Retrospective evaluations of gambling wins: Evidence for a 'peak-end' rule. *Conference of the Cognitive Science Society*, *30*, 77–82.

Zimmerman, C. A., & Kelley, C. M. (2010). "I'll remember this!" effects of emotionality on memory predictions versus memory performance. *Journal of Memory and Language*, *62*, 240–253.

Zink, C. F., Pagnoni, G., Martin-Skurski, M. E., Chappelow, J. C., & Berns, G. S. (2004). Human striatal response to monetary reward depend on saliency. *Neuron*, *42*, 509–517.