DRAFT. PLEASE DO NOT CITE OR CIRCULATE WITHOUT PERMISSION OF THE AUTHORS

The Influence of Loss Aversion on Explore-Exploit Decisions

Alycia Chin, David Hagmann, and George Loewenstein

Author Note

Alycia Chin,*† Office of Research, Consumer Financial Protection Bureau, 1700 G Street NW, Washington, DC 20552, alycia.chin@cfpb.gov; David Hagmann, Carnegie Mellon University, Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, hagmann@cmu.edu; George Loewenstein, Carnegie Mellon University, Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Ave., Baker Hall 319D, Pittsburgh, PA 15213, (412) 268-8787, gl20@andrew.cmu.edu.

*Corresponding author

† This article is the result of the author's independent research and does not necessarily reflect the views of the Consumer Financial Protection Bureau or the United States.

Acknowledgements

We thank Vikram Jambulapati, Amanda Markey, and Eric M. VanEpps for comments given at various stages of this project.

Abstract

In “explore-exploit” situations, decision makers must choose between exploring unknown options and exploiting known options. We study how explore-exploit decisions vary under the influence of loss aversion, predicting that (1) people will be less likely to explore if doing so can lead to losses and that (2) people will be less likely to exploit when doing so would lead to repeated losses. To examine these predictions, we used a novel multiple round computer task in which participants explored one of two one-dimensional environments that were equivalent in terms of potential payoffs. In one decision environment, participants encountered only gains in all rounds of play. In the other, they were given an up-front payment that was offset in subsequent rounds, where participants experienced gains or losses. Across multiple studies, we find evidence for both of our predictions. Additionally, we demonstrate that loss aversion can be adaptive, leading participants in low payoff environments to higher total rewards.

Keywords: explore-exploit, loss aversion, decisions from experience

Introduction

The term "explore-exploit tradeoff" refers to a class of situations in which decision makers must choose between known and unknown options. At any given point, choosing one path precludes the others. For example, fire fighters choose between rescuing known survivors or searching for additional victims, consumers choose between an apartment they have already visited or a higher-quality option (Zwick, et al., 2003), and researchers divide their time between existing projects and new ones that might be more fruitful (see Cohen, McClure & Yu, 2007 for additional examples). Explore-exploit decisions are difficult because both types of choices have advantages and disadvantages. Exploring the unknown is risky because payoffs are uncertain, but revealed payoffs could be large and, once discovered, can be exploited. In contrast, exploiting a known option is safe because it takes advantage of existing information, but may lead a decision maker to miss a preferable, unexplored opportunity. Explore-exploit decisions have been studied extensively within multiple disciplines, with research concentrating on both optimal decision strategies (e.g., Gittins, 1979; Lee, et al., 2011) and actual decision making (see Cohen, McClure & Yu, 2007; Wilson, et al., 2014; and Mehlhorn, et al., 2015 for a review).

In this paper, we study how explore-exploit decisions are affected by loss aversion. Loss aversion refers to the observation that people are more sensitive to losses than gains of the same magnitude (Kahneman & Tversky, 1979). Loss aversion has been demonstrated in many areas of consumer behavior and is used to explain a range of effects (Rick, 2011). For example, investors who check their portfolio frequently are more likely to hold bonds than stocks, because bonds are less likely to show losses (Benartzi & Thaler, 1995; Thaler, et al., 1997). Additionally, homeowners who face nominal losses on their upcoming real estate transactions are reluctant to lower sale prices past the original purchase price, thereby reducing the probability of selling (Genesove & Mayer, 2001).

Given extensive literature on both the explore-exploit tradeoff and loss aversion, it is not surprising that some researchers have examined the potential effects of losses on exploration decisions. For example, research shows that decision makers explore more extensively when payoffs are negative, as compared to situations with positive payoffs of equal magnitude (Lejarraga, Hertwig, & Gonzales, 2014). Notably, however, this result comes from research using the "sampling paradigm," a common method for studying explore-exploit decisions within behavioral decision research (see Hertwig, et al., 2004; Lejarraga, Hertwig, & Gonzales, 2014; Rakow & Newell, 2010). In this paradigm, decision makers are typically presented with two options and can sample repeatedly from each one to gauge potential payoffs before making one "real" choice for pay. Exploration is measured by the number of draws that participants take before making their final choice. We suspect that the design of this paradigm, which does not require participants to internalize payoffs, may reduce the role of losses in decision making.

Indeed, a contrasting result comes from research in which participants repeatedly searched a grid (Teodorescu & Erev, 2013). For this research, exploration was defined as occurring any time a participant chose a new location while exploitation was defined as choosing a previously explored location. The grid search task was designed so that exploration could have positive or negative payoffs, depending on experimental condition, but exploitation always paid zero. Using this setup, the authors found that participants in an environment of infrequent negative returns explored less than those in an environment of infrequent positive rewards. In other words, when faced with the prospect of losses, participants tended to exploit options that yielded a fixed, zero payoff.

A final line of research relevant to loss aversion and explore-exploit decisions deals with animal foraging behavior. Foraging is modeled using the concept of an “energy budget,” which represents whether an animal’s current level of consumption is high enough to sustain survival (Bateson, 2002; McNamara & Houston, 1992). Specifically, an animal with a positive energy budget is receiving enough food to survive; one with a negative energy budget is not. Experiments show that certain animals make different choices depending on their energy budget. For instance, yellow-eyed juncos are more likely to choose a risky option when low temperatures cause them to be on a negative energy budget, and safe options when feeding at higher temperatures (Bateson, 2002). To the extent that repeatedly receiving a negative payoff approximates a negative energy budget, this research suggests that people faced with losses will find risky options more attractive.

Based on previous work on loss aversion and explore-exploit decisions, we expected that people making these decisions would demonstrate loss aversion, finding negative returns particularly painful. This reasoning led to two behavioral predictions. First, we predicted that participants would be deterred from exploring if they believed that doing so could lead to losses, consistent with Teodorescu and Erev (2013). Second, we predicted that participants receiving negative payoffs would be more likely to move to a different location to avoid repeated losses, that is, less likely to exploit negative returns. To our knowledge, this prediction has not been directly examined in previous research, but is consistent with experiments showing that animals avoid options that preclude survival (Bateson, 2002). In essence, our predictions reflect the view that exploration is driven by hope of positive returns and constrained by fear of negative returns.

Experimental Task: The Grain Game

To examine the effects of loss aversion on explore-exploit decisions, we developed an exploration task that would allow us to induce loss aversion via a framing effect. In our task, "The Grain Game," participants played the role of a farmer who repeatedly chose where to plant crops (see Figure 1 for a screenshot). In all of our studies, participants played the game for 70 rounds. Each round, they chose one of three locations to plant a crop: the same location as the previous turn (a decision to exploit), one space to the left, or one space to the right. Participants could revisit locations, so movement itself was not equivalent to exploration. Instead, we defined "exploration" as movement into a previously-untouched space and "retreat" as returning to a previously explored location. As shown in Figure 1, the most recent payoffs from up to 15 previously explored locations were displayed on the screen. Unexplored locations were indicated with a question mark icon.

We assigned participants to one of two conditions that differed in the framing of payoffs. Specifically, participants in the gain-only condition were in an environment in which payoffs ranged from 0 to 25 points per turn. Those in the gain-loss condition were in a matched environment with payoffs ranging from -10 to 15 points per turn, but, to equate objective payoffs between the two conditions, also received an extra 700 points at the outset (10 points for each of the 70 rounds). To gauge the influence of loss aversion, we defined a threshold at 0 in the gain-loss condition and 10 in the gain-only condition. We analyzed whether behavior was significantly different below that threshold across the two conditions.

All participants were told of the range of payoffs present in their environment and were asked to report this range prior to starting play in order to ensure that they understood the game. Once participants began the game, payoffs in each location were determined by a fixed value, which differed across locations, and noise of at most two points in either direction.[1] At the end of the experiment, participants received payment based on accumulated points (4 points = $0.01).[2]

Study 1a and 1b

Our first two studies were primarily intended to test the prediction that the presence of losses would cause participants to exhibit more conservative, exploitative behavior, thereby replicating findings from previous research (Teodorescu & Erev, 2013) in a different situation. We created two landscapes in which all participants started with positive payoffs that quickly decreased if they explored. In order to study the potential welfare effects of exploration, the two landscapes varied in terms of whether exploring past these low payoffs would be harmful (Study 1a) or beneficial (Study 1b). For both studies, all participants began on the left edge of the landscape and did not observe the landscape prior to playing.

Study 1a

Method

Participants. We recruited 140 U.S. residents online using Mechanical Turk. The average age was 31.4 years and 57.9% were male.

Stimuli. We created a version of the Grain Game (Figure 1) in which exploration was not worthwhile (Figure 2). Specifically, as shown, the highest payoff was available after only eight

[1] The probability distribution was as follows. For an expected value of 10, there was a .07 chance of receiving 8 or 12, a .15 chance of receiving 9 or 11, and a .56 chance of receiving 10. For values near the extremes of the possible range, probabilities were collapsed – for instance, the chance of receiving 25 was .78 (= .56 + .15 + .07), the chance of receiving 24 was .15 and the chance of receiving 23 was .07.

[2] For exploratory purposes, we also asked participants to respond to questions on their thought process while playing the game, the purpose of the game, six statements about their interest in exploring, and in Study 2, whether they copied the graph shown at the beginning of the study. For this paper, we concentrate on participant behavior rather than these responses.

moves. Participants in the gain-loss condition received negative payoffs starting at position 15, after which the highest expected payoff in was one (eleven in the gain-only condition).

Results

Exploration. Our primary research question was whether exploration varied between the two conditions. To study this question, we created a hazard plot containing the proportion of participants in each of the two conditions who explored to a given location (Figure 3). All participants started at the left-most location (location 0), so overall exploration was captured by the extent of their movement to the right. The graph shows that participants in the gain-loss condition were much less likely to explore beyond position 15, where they first started to experience losses. In particular, the average participant in the gain-only condition explored 32.5 (*SD* = 19.7) positions, but the average participant in the gain-loss condition explored 21.3 (*SD* = 16.0; $t(138) = 3.7$, $p < .001$). The largest difference between the two conditions came at position XX, which was reached by XX% of those in the gain-only and XX% in the gain-loss conditions.

In addition to examining overall exploration, we also analyzed decisions made each round. To do so, we ran a multinomial logit regression taking the following form:

$$\Pr\left(Y_{\{i,t\}} = j\right) = \beta_0 + \beta_1 \text{gainloss} + \beta_2 \text{adjusted payoff}_{\{t-1\}} + \beta_3 \text{below threshold}_{\{t-1\}} + \beta_4\left(\text{gainloss} \times \text{below threshold}_{\{t-1\}}\right)$$

Specifically, individual *i* at time *t* takes action $j \in \{exploit, explore, retreat\}$. *Gain-loss* is an indicator variable set to one for those randomly assigned to the gain-loss condition and zero for those in the gain-only condition. *Adjusted payoff* represents the payoff in the previous period, scaled to the gain-only range of [0, 25] by adding 10 to every payoff in the gain-loss condition. *Below threshold* is an indicator variable representing whether the payoff in the prior round was below the threshold of 10. Our predictions are that (1) the estimated coefficient on β_1 will be negative for exploration, representing a reluctance to explore in the gain-loss condition, and (2) β_4 will be positive, suggesting that participants will be unlikely to exploit negative payoffs.

The results of the multinomial logit are shown in Table 1. As predicted, participants in the gain-loss condition were less likely to explore than those in the gain-loss condition [add an interpretation of coefficient]. Additionally, the coefficient on adjusted payoff indicates that receiving a higher payoff is less likely to lead to exploration, regardless of condition. Contrary to our expectations, once accounting for prior payoffs, there was no difference in behavior for those receiving payoffs below the threshold, and no interaction between low payoffs and condition.

Payoffs. Conditional on playing in the same way, expected rewards in the study were the same for all participants. However, the underlying landscape of this study was designed to punish exploration. Consistent with this setup, we found a negative correlation between participants' payoffs and how far they explored (r = -.71; p < .001). Participants in the gain-only condition, who explored more, received significantly fewer points (M = 816.3, SD = 163.4) than those in the gain-loss condition (M = 881.8, SD = 172.5; t(138) = -2.30, p = .02).

Study 1b

This study followed the same design as Study 1a with a different landscape.

Method

Participants. We recruited 118 participants via Amazon Mechanical Turk. They were 30.7 (SD = 8.7) years old on average and 56% were male.

Stimuli. This study used an environment in which exploration was rewarded (Figure 4). As shown, participants in the gain-loss condition received negative payoffs from position 9 to 27. All participants who persisted past these low payoffs received the maximum possible payoff at position 39.

Results

Exploration. Following our methods from Study 1a, we first present a hazard plot of exploration. As shown in Figure 5, participants in the gain-loss condition explored less than those in the gain-only condition. Specifically, the average distance explored was 19.9 (*SD* = 13.7) for those in the gain-loss condition and 32.5 (*SD* = 19.7) for those in the gain-only condition ($t(116) = 2.34$, $p = .02$). The largest difference between the two conditions came at position XX, which was reached by XX% of participants in the gain-only condition and XX% of those in the gain-loss condition.

In addition to this overall pattern, we analyzed decisions to explore, exploit, and retreat using a multinomial logit regression as in Study 1a. The results show that participants in the gain-loss condition were more likely to exploit than those in the gain-only condition, as indicated by the negative estimates on the predicted probability of exploring and retreating (Table 1). The coefficients on *adjusted payoff* suggest that obtaining a high payoff was associated with a lower likelihood of moving to a different location through exploring or retreating. Additionally, receiving a payoff below the threshold was associated with a higher probability of exploring. Interestingly, while the estimated probability of exploring did not differ by condition, those in the gain-loss condition were more likely to retreat after receiving a low payoff than those in the gain-only condition. In other words, obtaining losses led participants to retreat to their small initial gains rather than encouraging them to explore further.

Payoffs. As in Study 1a, the reward landscape in the study was objectively the same for participants in both conditions. However, the landscape in this study was designed to reward exploration. Consistent with this design, there was a positive correlation between participants' payoffs and how far they explored ($r = .23$; $p = .01$). Participants in the gain-only condition, who explored more on average, also received significantly more points ($M = 810.0$, $SD = 185.7$) than those in the gain-loss condition ($M = 740.9$, $SD = 187.9$; $t(116) = 2.01$, $p = .05$).

Discussion of Studies 1a and 1b

We predicted that loss aversion would affect exploration in two ways. First, we expected that people would be less likely to explore when doing so could yield negative payoffs. In both Study 1a and 1b, we found support for this idea, as those in the gain-loss condition were less likely to explore and more likely to exploit. Second, we hypothesized that people receiving negative returns would be more likely to move to a new location, as losses would be experienced as especially aversive. We found that lower payoffs were associated with increased exploration in both studies (and in Study 1b, with increased "retreat" behavior); however, we observed inconsistent results regarding negative payoffs. Specifically, negative payoffs (i.e., those below the threshold) were associated with increased retreat behavior only in Study 1b. Upon reviewing these results, we suspected that these inconsistencies may have been caused by the landscapes that we designed. Specifically, because all participants had positive payoffs at the start of the game, everyone knew they could retreat to avoid negative returns. Additionally, because participants tended to avoid losses, we had less data available to examine potential differential effects of negative payoffs. To address these concerns and investigate our hypotheses in a larger number of situations, in Study 2 we created a larger set of landscapes for participants to explore and randomly assigned some participants to start in locations with negative returns.

Study 2

Method

Participants. We recruited 1,068 U.S. residents online using Mechanical Turk. The average age was 31 years (*SD* = XX) and 61% were male. We recruited a large sample because we anticipated an increase in variability of behavior given the large variety of landscapes that we intended to test.

Stimuli. To investigate the effect of loss aversion on explore-exploit decisions in multiple environments, we created a landscape containing 700 locations. We randomly assigned participants to different starting positions within this landscape, subject to the constraint that they were at least 70 positions removed from either edge. This assignment ensured that participants were unconstrained in their ability to move in either direction.[3] In order to fix participant expectations, and unlike in the previous two experiments in which participants were not given information about the overall landscape of payoffs, all participants viewed a graph of potential payoffs before starting the game (Figure 6) and were told they would start in a random position on this graph. Although participants were initially shown a subset of payoffs, they could pan to the right or left using prominent arrows on the screen in order to see the entire set of payoffs.

Results

[3] Due to programming error, we failed to record the first and last decisions made by participants. Analyses are based on the remaining 68 choices.

Exploration. Our prediction was that participants would be less likely to explore when receiving gains and more likely to explore when incurring losses. To examine this pattern, we first rescaled payoffs to the gain-only range of [0, 25] by adding 10 to every gain-loss payoff. Next, we calculated the probability of exploring for each condition x payoff combination (Figure 7). The negative slope in this figure demonstrates that across both conditions, participants were less likely to explore when they had just received higher payoffs. Comparing the two conditions against each other shows that those in the gain-loss condition were more likely to explore than those in the gain-only condition when payoffs were below the threshold, and less likely to explore when payoffs were above it.

This pattern becomes clearer in Figure 8, which plots the difference in the probability of exploration between the two conditions for each payoff. Specifically, values above zero on the y-axis indicate that participants in the gain-loss condition were more likely to explore, and values below zero indicate that participants in the gain-only condition were more likely to explore. In all situations below the threshold (i.e., when participants in the gain-loss condition experienced negative returns), gain-loss participants were more likely to explore than those in the gain-only condition. However, gain-loss participants were less likely to explore above the threshold. There is a clear discontinuity at the point, indicated by a vertical line in the figure, separating gains from losses.

As in previous studies, we also ran a multinomial logit on each decision. Consistent with studies 1 and 2, those in the gain-loss condition were less likely to explore (Table 1). As indicated by the coefficients on *adjusted payoff*, participants in both conditions were less likely to move to a different location when receiving higher payoffs. Additionally, as hypothesized, the interaction between *gain-loss condition* and *below threshold* shows that the tendency to move to a new location after receiving a low payoff is different between the two conditions. Specifically, participants in the gain-loss condition who are receiving payoffs below the threshold (i.e., losses) are both more likely to explore and to retreat than those in the gain-only condition who are receiving low, but positive, payoffs.

Individual behavior. [Consider adding a section on the proportion of players that show behavior consistent with loss aversion – lots of questions at SJDM about individual differences].

Payoffs. There was no difference in payoffs between the two conditions (p = .41), reflecting the fact that different starting locations rewarded exploration to a different extent.

General Discussion

In this paper, we examined the effect of loss aversion on explore-exploit decisions using participants' behavior in a novel computer task. Our main contribution is to show that loss aversion can affect explore-exploit decisions in systematic ways. In all of our studies, we found that the presence of losses in an environment led people to explore less. In our last study, which tested a variety of environments, we also found that receiving losses (versus matched positive payoffs) led to less exploitation. As such, our findings are consistent with prior literature on loss aversion which states that negative returns are particularly painful (Kahneman & Tversky, 1979) and empirical research showing that environments with losses have less exploration (Teodorescu & Erev, 2013).

In addition to this main contribution, our research makes methodological advances by expanding beyond a simple two option framework (as suggested by Rakow & Newell, 2010) using a large environment (with 70 options) and "retreat" behavior. Such changes suggest new theoretical directions. In particular, much of the literature has focused on interpreting explore-exploit decisions in the context of decisions under risk (Hertwig et al., 2004; Lejarraga, Hertwig, & Gonzales, 2014; Rakow & Newell, 2010; although see Mehlhorn et al., 2015 for alternatives), with exploring usually considered risky and exploiting considered safe. We suspect that most scholars would also consider retreating to be a risk averse behavior since it builds upon a known alternative, although to our knowledge this has not been discussed in the literature. Given that position, it is particularly interesting to see that exploring and retreating (i.e., risky and safe behaviors) can both be triggered by the same situation – receiving negative payoffs (Study 2). Future research may consider additional methodological variation in order to better understand the relationship between the explore-exploit tradeoff and decisions under risk.

A natural question arising from our work is whether loss aversion is adaptive for explore-exploit decisions. Our first two studies served as existence proofs that showed loss aversion can be helpful or harmful depending on the degree to which the environment rewards exploration. Study 1a of this paper, which used an environment in which exploration was harmful, contributes to an expanding body of findings which document situations in which people with different types of biases make better decisions. Kahneman and Lovallo (1993), for example, showed that loss aversion and risk aversion can be beneficial to managers who are subject to overconfidence whereas Shiv et al. (2005) showed that people with damage to emotion-related brain areas made better decisions in an environment in which risk-taking was beneficial. There is also a large literature showing that people who are overconfident and over-optimistic tend to be happier, healthier and to exert more social influence (e.g., Taylor & Brown, 1988; Kunda, 1990). While we cannot say whether loss aversion is generally helpful, our results indicate that it should not automatically be considered a mistake.

Within the literature on explore-exploit tradeoffs, future work may want to consider additional psychological factors that could influence decisions. For instance, people may be influenced by social context, which has been shown to have an impact on children's curiosity and exploratory play (Engel, 2011; Henderson, 1984; Henderson, Charlesworth, & Gamradt, 1982). It also seems likely that individual attributes such as satisficer-maximizer tendencies (Schwartz, et al., 2002) will predict how satisfied people are with their current payoffs, and their associated willingness to explore. Such approaches could add to our understanding of decision making in these situations.

In the end, our results suggest that changing one's perspective can make a difference. Viewing a given outcome as a loss – relative to a low, positive payoff – may make you more likely to change, whether through exploring or retreating. But, viewing the world through rose colored glasses, with no losses in sight, may make you more likely to explore overall. If shortages can activate human ingenuity (Johnson, 2011), it is reasonable to wonder whether other seemingly unfavorable conditions, like the experience of losses, could, in the right circumstances, have favorable consequences for consumer outcomes.

References

Bateson, M. (2002). Recent advances in our understanding of risk-sensitive foraging preferences. *Proceedings of the Nutrition Society, 61*, 509-516. doi:10.1079/PNS2002181

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics, 110*(1), 73–92.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions: Biological Sciences, 362*(1481), 933-942.

Dam, G., & Körding, K. (2009). Exploration and exploitation during sequential search. *Cognitive science*, 33, 530-41. doi: 10.1111/j.1551-6709.2009.01021.x

Engel, S. (2011). Children's need to know: Curiosity in schools. *Harvard Educational Review, 81*(4), 625-45.

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*, 575-97. doi: 10.1002/bdm.602

Genesove, D., & Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics, 116*(4): 1233-1260.

Gittins, J. C. (1979). Bandit processes and dynamic allocation indices, *Journal of the Royal Statistical Society, Series B 9Methodological)*, 148-177.

Henderson, B. B. (1984). Social support and exploration. *Child Development, 55*(4), 1246-1251.

Henderson, B. B., Charlesworth, W. R., & Gamradt, J. (1982). Children's exploratory behavior in a novel field setting. *Ethology and Sociobiology, 3*, 93-9.

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*(8), 534-539.

Johnson, S. (2010). *Where good ideas come from: the natural history of innovation*. New York: Riverhead Books.

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management science*, *39*(1), 17-31.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263-291.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480-498.

Lee, M. D., Zhang, S., Munro, M., & Steyvers, M. (2011). Psychological models of human and optimal performance in bandit problems. *Cognitive Systems Research*, *12*, 164-174.

Lejarraga, T., Hertwig, R., & Gonzales, C. (2012). How choice ecology influences search in decisions from experience. *Cognition, 124(3)*, 334-342.

McNamara, J. M., & Houston, A. I. (1992). Risk-sensitive foraging: a review of the theory. *Bulletin of mathematical biology*, *54*(2), 355-378.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, 2, 3, 191-215.

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

Rick, S. (2011). Losses, gains, and brains: Neuroeconomics can help to answer open questions about loss aversion. *Journal of Consumer Psychology*, 21: 453-463.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: happiness is a matter of choice. *Journal of personality and social psychology*, *83*(5), 1178.

Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., & Damasio, A. R. (2005). Investment behavior and the negative side of emotion. *Psychological Science, 16*(6), 435-39. doi: 10.1111/j.0956-7976.2005.01553.x

Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social-psychological perspective on mental health. *Psychological Bulletin, 103,* 193-210.

Tedorescu, K., & Erev, I. (2013). On the decision to explore new alternatives: The coexistence of under- and over-exploration. *Journal of Behavioral Decision Making.* doi: 10.1002/bdm.1785

Thaler, R. H., Tversky, A., Kahneman, D., and Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics, 112*(2), 647-661.

Walker, P. M., Kolling, A., & Lewis, M. (2011). Human exploration patterns in unknown, time-sensitive environments. Paper contributed to IEEE International Conference on Systems, Man, and Cybernetics, Anchorage, AK. Retrieved from http://andreaskolling.academia.edu/AndreasKolling/Papers/1166000/Human_Exploration_Patterns_in_Unknown_Time-sensitive_Environments.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore-exploit dilemma. *Journal of Experimental Psychology: General, 143*(6), 2074-2081.

Zwick, R., Rapoport, A., Chung Lo, A. K., & Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, *22*(4), 503-19.

Table 1. Multinomial logit regression results predicting participant decisions.

	Study 1a		Study 1b		Study 2	
	Pr(Explore)	Pr(Retreat)	Pr(Explore)	Pr(Retreat)	Pr(Explore)	Pr(Retreat)
Gain-Loss Condition	-0.585*** (0.055)	-0.123* (0.060)	-0.625*** (0.063)	-0.813*** (0.073)	-0.167*** (0.022)	0.099*** (0.023)
Adjusted Payoff [t-1]	-0.149*** (0.011)	0.018 (0.013)	-0.043*** (0.007)	-0.068*** (0.008)	-0.110*** (0.002)	-0.040*** (0.002)
Below threshold [t-1]	-0.088 (0.106)	0.198 (0.134)	0.542*** (0.109)	-0.101 (0.126)	-0.090** (0.033)	-0.082* (0.038)
Gain-Loss x Below threshold [t-1]	-0.227 (0.117)	-0.103 (0.148)	-0.157 (0.118)	0.690*** (0.131)	0.621*** (0.040)	0.308*** (0.046)
Intercept	2.022*** (0.148)	-0.795*** (0.180)	0.622*** (0.112)	0.606*** (0.128)	1.426*** (0.040)	0.049 (0.043)
R2	0.042		.033		.041	
Log Likelihood	-9733.948		-8370.258		-74,864.810	
N (obs)	9520		8024		72,420	
N (participants)	140		118		1068	

Note. Participant actions were classified as exploring a new field, retreating to a previously explored field, or exploiting the current option. *Adjusted payoff* represents the payoff in the previous period, scaled to the gain-only range of [0, 25] by adding 10 to every payoff in the gain-loss condition. *Below threshold* is an indicator variable representing whether the payoff in the prior round was below the threshold of 10.

Figure 1. Screenshot of The Grain Game, gain-only condition.

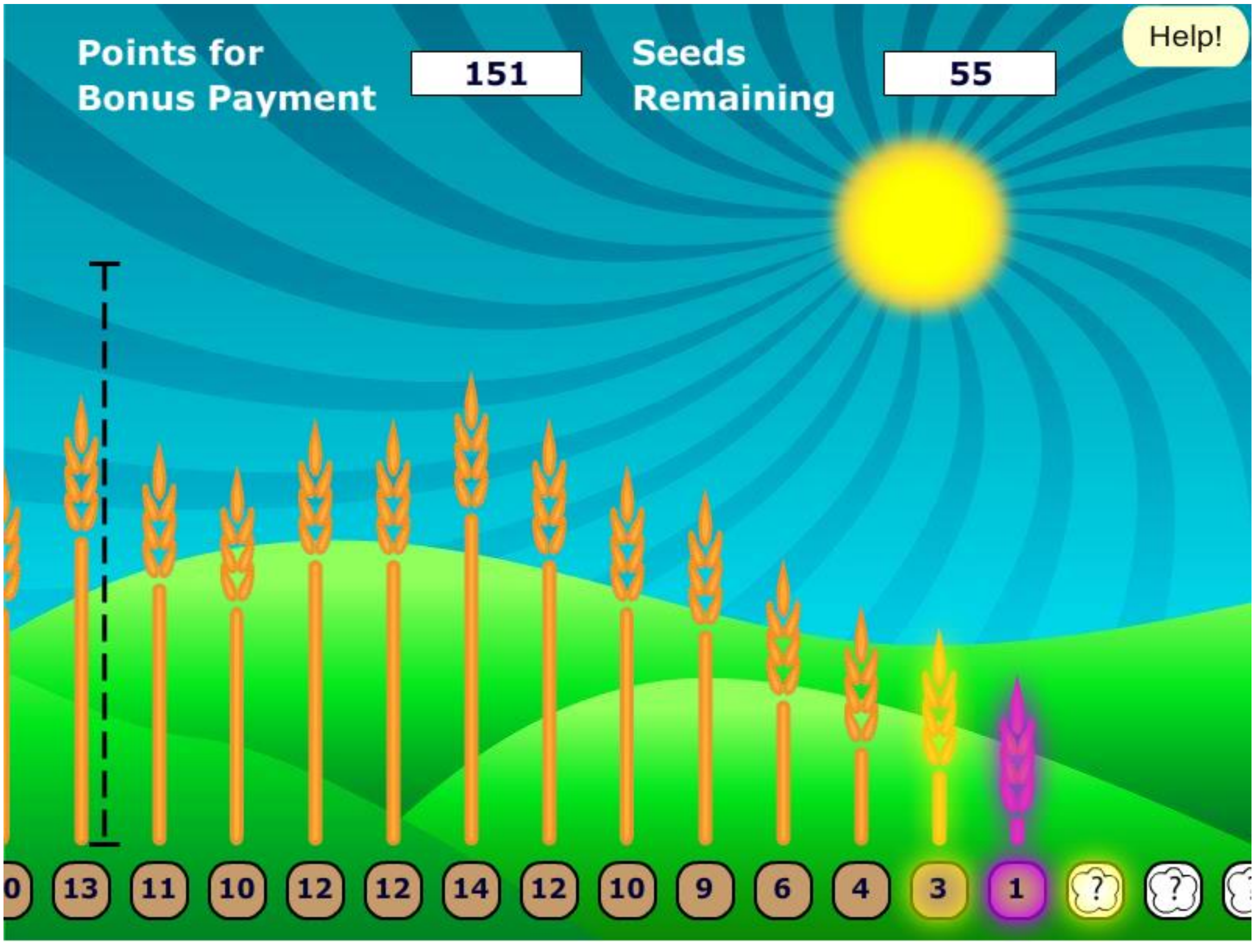

Note: The upper left shows accumulated points and the upper right shows the number of turns remaining. Unexplored positions are designated by question marks. The most recent payoffs received from each explored position are given in the boxes below the grain. The participant's current position is highlighted in purple; on the next turn, they can exploit that position, retreat to the left, or explore to the right. The Help button displays instructions for playing the game.

Figure 2. Point gradient from Study 1a.

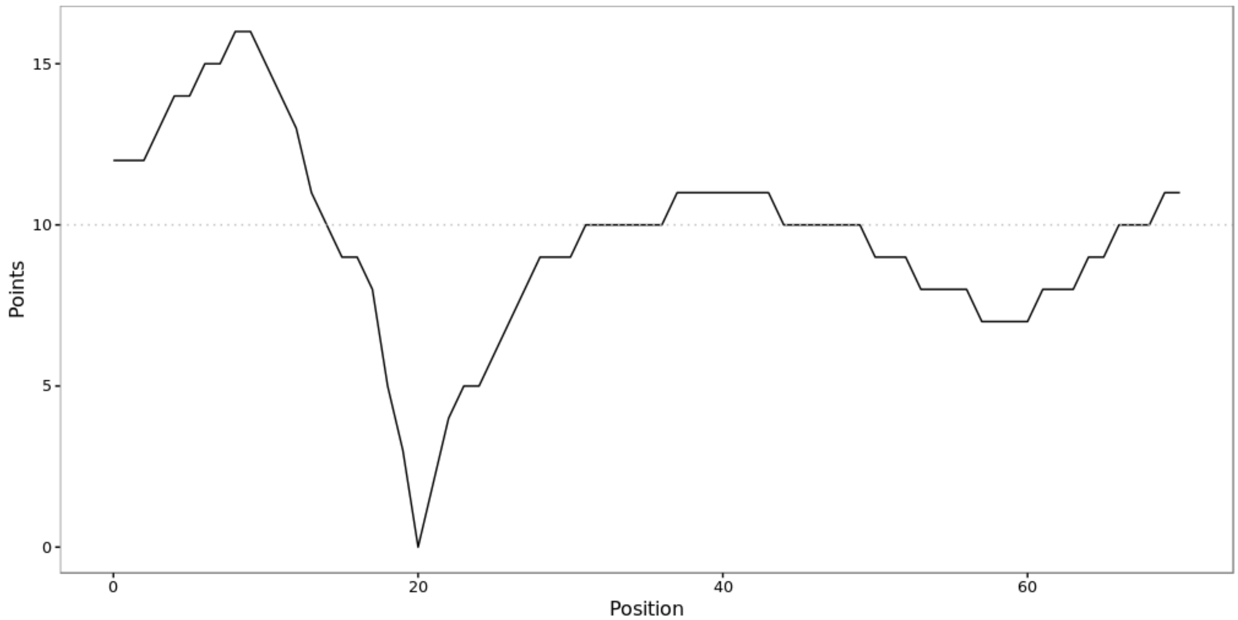

Note. The landscape from Study 1a is unfavorable to exploration. The highest payoff (16) is obtained at positions 8 and 9. Participants in the gain-loss condition first experience losses in expectation at position 15.

Figure 3. Hazard plot for furthest exploration position (Study 1a).

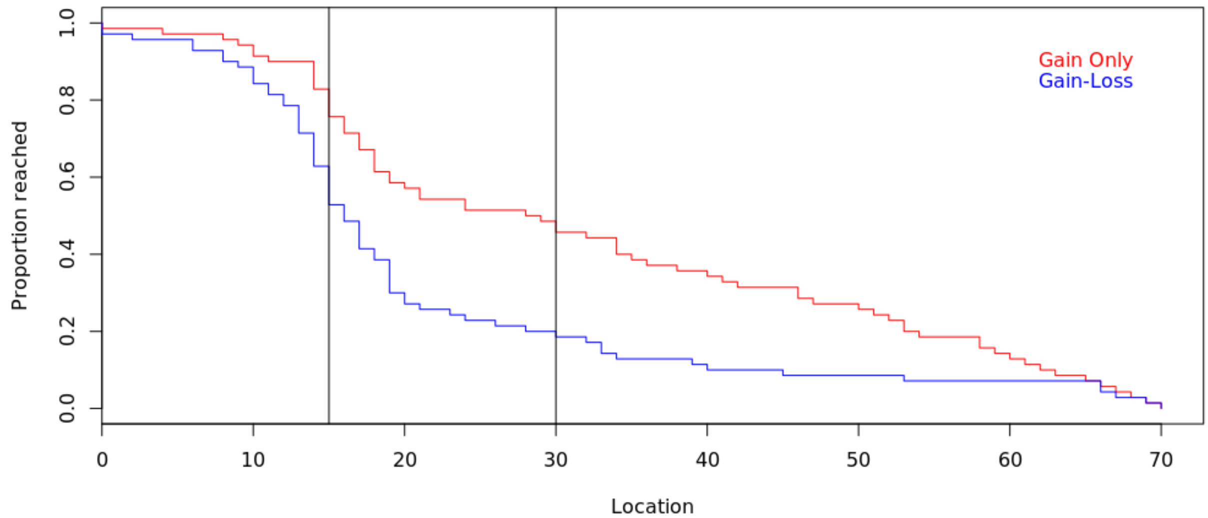

Note. Vertical lines at positions 15 and 30 shows the range in which those in the gain-loss condition have a negative expected return.

Figure 4. Point gradient from Study 1b.

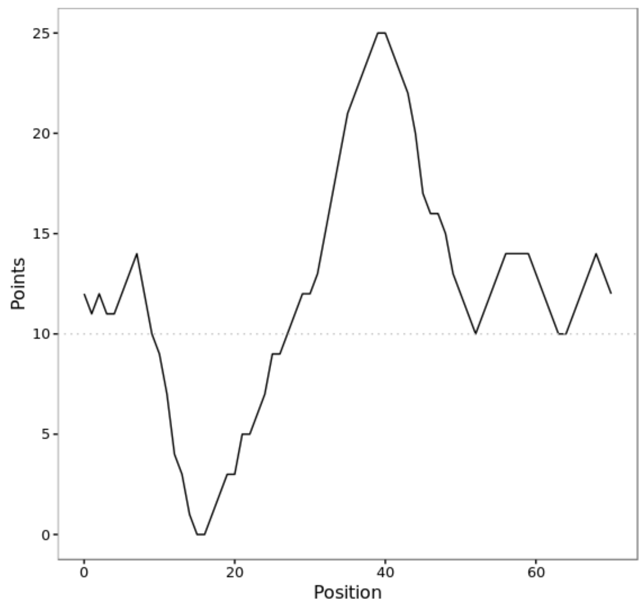

Note. The landscape from Study 1b is favorable to exploration. The highest payoff (25) is obtained at position 39. Participants in the gain-loss condition first experience losses in expectation at position 9.

Figure 5. Hazard plot of exploration from Study 1b.

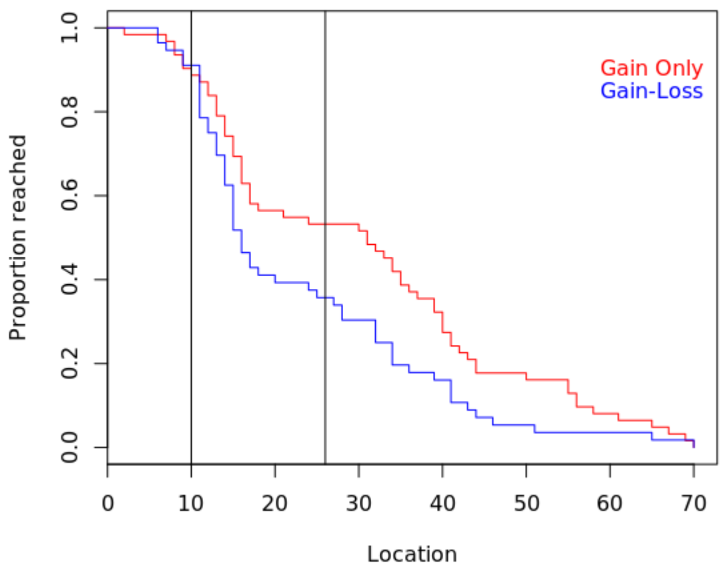

Note. This chart shows the proportion of participants in each condition who reached each position. Dotted lines at positions 10 and 26 show the range in which those in the gain-loss condition have a negative expected return.

Figure 6. Preview of landscape payoffs, Study 2.

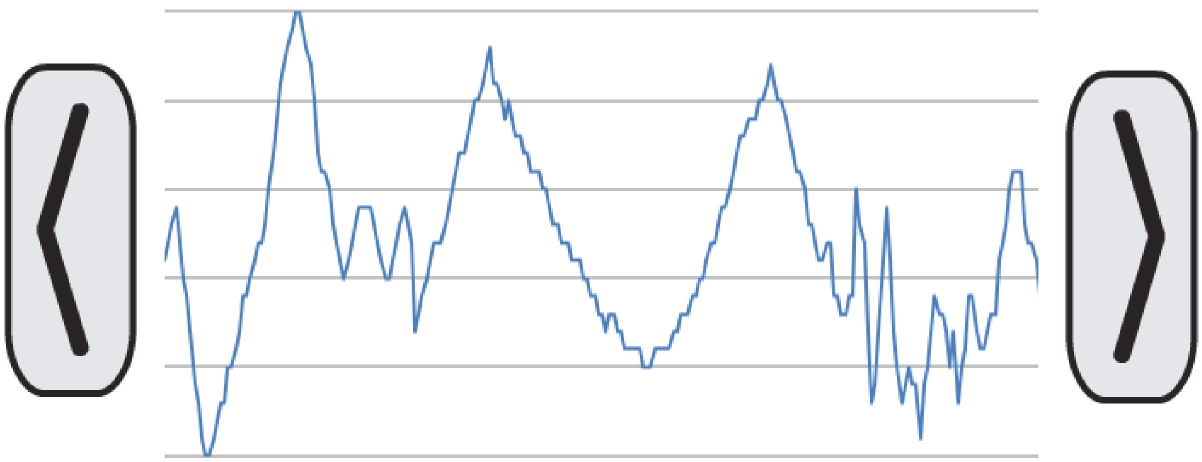

Note. All participants viewed this graphic prior to starting the study. The arrows on the left and right can be used to pan over the entire landscape.

Figure 7. Probability of exploration by condition, Study 2.

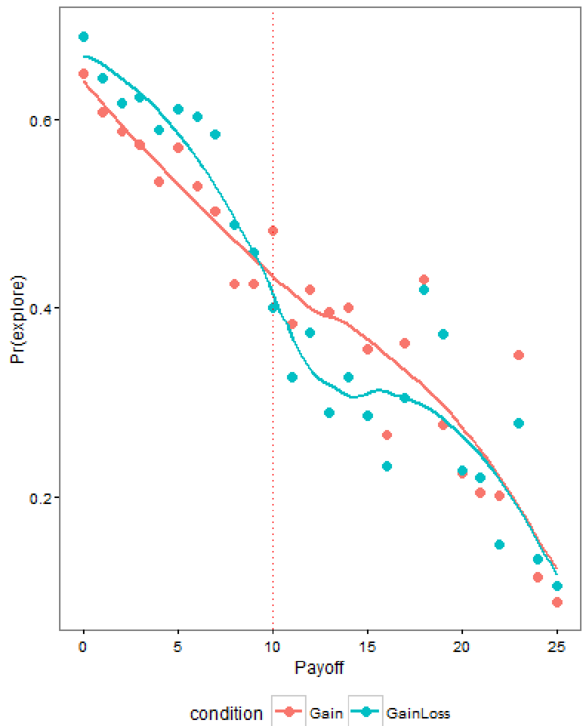

Note. This chart shows the probability of exploring as a function of the adjusted payoff for each experimental condition. Payoffs are rescaled to the gain-only range of [0, 25]. The lines show the best fitting smoothed spline. Below the threshold (i.e., to the left of the red dotted line), participants in the gain-loss condition incur losses.

Figure 8. Difference in exploration probability across conditions (Study 2).

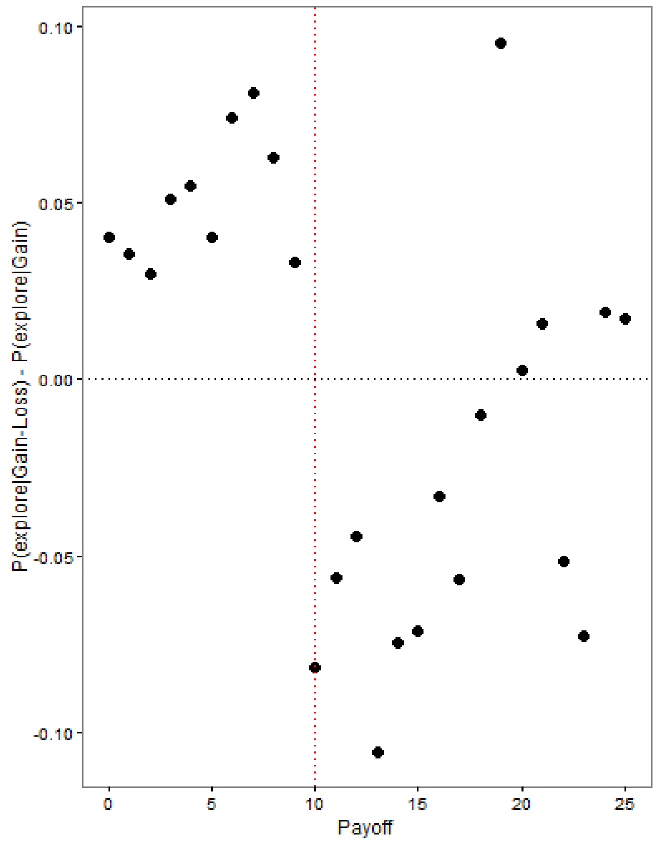

Note. This chart shows the difference in probability of exploration between the gain-loss and gain-only conditions. Values above 0 indicate increased exploration for those in the gain-loss condition. To the left of the red dotted line, participants in the gain-loss condition incur losses.