Received: 29 September 2021 | Revised: 23 September 2022 | Accepted: 17 October 2022

DOI: 10.1002/bdm.2309

RESEARCH ARTICLE

WILEY

Fear and promise of the unknown: How losses discourage and promote exploration

Alycia Chin[1] | David Hagmann[2] | George Loewenstein[3]

[1]Office of the Investor Advocate, Securities and Exchange Commission, Washington, District of Columbia, USA

[2]Department of Management, The Hong Kong University of Science and Technology, Hong Kong

[3]Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

Correspondence
Alycia Chin, Office of the Investor Advocate, Securities and Exchange Commission, Washington, DC 20549, USA.
Email: readlinga@sec.gov

Funding information
This work was supported by funding provided to the third author by Carnegie Mellon University.

Abstract

Many situations involving exploration, such as businesses expanding into new products or locations, expose the explorer to the potential for subjective losses. How does the potential to experience losses during the course of a search affect individuals' appetite for exploration? In three incentivized studies, we manipulate search outcomes by presenting participants either with a gain-only environment or a gain-loss environment. The two environments offer objectively identical incentives for exploration: Using a framing manipulation, we decrease gain-loss payoffs and provide participants an initial endowment to offset the difference. Participants decide how to explore a one-dimensional space, receiving payoffs based on their location each period. We predict and find that participants are motivated to avoid losses, which increases exploration when they are incurring losses but decreases exploration when they face the potential for losses. We conclude that exploration is driven by hope of potential gains, constrained by fear of potential losses, and motivated by avoidance of experienced losses.

KEYWORDS
decisions from experience, explore-exploit, framing, loss aversion

1 | INTRODUCTION

Many important real-world situations involve explore-exploit decisions in situations characterized by local and global optima. For example, consider an organization that accrues modest profits in its local market (a local optimum). That organization might explore opening new locations but abandon those changes if exploration does not quickly increase total profits. In such cases, the organization may fail to discover that a more substantial commitment to expansion would have yielded far superior outcomes. In other words, short-run declines in performance at points close to a local optimum may deter decision-makers from further exploration.

Within the broad set of consequential situations that can be characterized as explore-exploit decisions, one important distinction is between environments in which only gains are possible and environments in which both gains and losses are possible.[1] For example, an organization making substantial profits might see reduced profitability from expansion but not losses. Alternatively, an organization with a smaller profit margin might incur losses from the same expansion. In this paper, we predict (and test) the idea that search behavior in these two situations is likely to differ, even holding constant the returns to exploration.

A large body of research across domains finds that people treat gains differently from losses: people tend to be "loss-averse," meaning that they dislike losses more than they like equivalent magnitude gains (Kahneman & Tversky, 1979). Loss aversion affects how long homeowners keep their homes on the market (Genesove & Mayer, 2001), the risk-adjusted rate of return on stocks (Benartzi & Thaler, 1995), the effort put into tasks by experimental participants (Imas et al., 2016), the performance of professional golf players (Pope & Schweitzer, 2011), and the impact on shopper behavior of incentives to use resusable bags (Homonoff, 2018). Moreover, people

Alycia Chin, the Securities and Exchange Commission disclaims responsibility for any private publication or statement of any SEC employee or commissioner. This article expresses the author's views and does not necessarily reflect those of the commission, the commissioners, or other members of the staff.

[1]A third situation, which we do not examine, occurs when only losses are possible.

exhibit diminishing sensitivity to both gains and losses; a property that the conventional notion of diminishing marginal utility assumes only applies to gains (Kahneman & Tversky, 1979). These features of decision-making are consequential whenever people make decisions under conditions of risk, a situation that includes explore-exploit decisions. Disproportionate aversion to losses has two major consequences: First, when people experience losses but know they can achieve modest gains by returning to an earlier location, they should be especially likely to do so. Second, people may be especially averse to exploration when it could lead to losses. Diminishing sensitivity to losses has one consequence: If people who experience losses choose not to return to an earlier location, they should be risk-seeking, which in the context of search means searching further.

We present three experiments examining the consequences of the potential to experience losses, induced by a framing manipulation, on explore-exploit decisions. Based on the basic features of the prospect theory's value function, we hypothesize and test the following two predictions about exploration in domains characterized by gains and losses. First, people will explore more or be likely to retreat when they are experiencing losses, with the goal of exiting from the loss domain. Second, when people are experiencing gains, but exploration exposes them to the potential for losses, they will explore less. Thus, the extent to which the presence of losses motivates or discourages exploration depends on whether the loss is currently experienced or possible and avoidable.

Applied to the previously discussed example of an expanding firm, we predict that a firm that becomes unprofitable from exploration (experiencing losses) would be more likely to revert to its initial operating mode or to go out further on a limb by taking further risks, relative to a comparable firm that becomes less profitable from expansion (but remains profitable). Furthermore, a firm that faces the potential to become unprofitable from expansion may also be reluctant to explore, thereby avoiding the future experience of losses, relative to a firm that does not face such potential.

Whether the predicted pattern of behavior ultimately results in higher or lower payoffs depends on the actual but *ex ante* unknown, pattern of returns faced by decision-makers. When a superior (local or global) optima is within reach, exploration will often lead to greater payoffs; when it is not, exploration will often decrease payoffs. To examine this prediction, we use different underlying payoff structures and show that the effects of reduced exploration due to loss aversion depend on the specific payoff structure of the environment. When there are gains available from exploration, environments with a potential for losses are likely to yield lower payoffs than environments characterized only by gains.

1.1 | Background on loss aversion and exploration

Research on the influence of losses in explore-exploit decisions has shown that participants are sensitive to losses (e.g., Teodorescu & Erev, 2014) and that experiencing losses leads to more exploration (e.g., Lejarraga et al., 2012). For instance, Lejarraga et al. (2012) review 15 experiments that used the "decisions from sampling" paradigm (e.g., Hertwig et al., 2004; Rakow & Newell, 2010). In this paradigm, decision-makers are first presented with two gambles and can repeatedly choose and observe the outcome of either one (i.e., explore) without affecting their payoffs. After that initial exploration stage, they are asked to make a single decision, and the outcome of that final choice determines their earnings. The review showed that in 12 of 15 studies participants explored more extensively when payoffs were in the loss domain than in the gain domain (choosing a median of 11.25 vs. 8.75 draws).[2] That is, when participants faced the prospect of incurring a loss, they were more likely to explore. Although this finding superficially appears opposite to our prediction, it is not. Because participants facing a loss during the final round are likely more motivated to avoid that outcome than those facing a gain, and initial exploration is not penalized; exploration has little downside. As such, loss aversion, in this very different paradigm, leads to greater exploration during a period when there are no payoffs.

In contrast to the decisions from the sampling paradigm in which only the final choice counts, in other studies relevant to the gain-loss distinction, *every* choice affects participants' earnings. In one paradigm reported in two papers (Lejarraga & Hertwig, 2017; Rakow et al., 2015), researchers give participants two options to choose from and define exploration as the number of times participants switch from one option to another. In this setting, participants who choose between options that only incur losses "explore" more, in that they switch back and forth between the two options more frequently. In a second paradigm in which every choice affects payoffs (Yechiam, Zahavi, et al., 2015), participants are presented with a risky option (a lottery with two outcomes) and a safe option that always returns the same amount. The experiment varies whether both options' payoffs are in the loss or the gain domain. Participants are more likely to explore, by switching between the two options (vs. repeatedly drawing from the same option), when in a task with losses. Finally, a third line of research explores how variation in the amount of information provided to participants affects exploration when exploring the payoffs of lotteries with positive and negative outcomes (Krueger et al., 2017). Participants are shown two buttons representing lotteries with unknown payoffs and probabilities. They are then restricted to selecting and observing the outcomes of each button an unequal number of times (e.g., Button A once and Button B thrice). Following those initial trials, participants make an additional six choices over the two lotteries, freely selecting however they prefer. Participants who observed only losses in their initial forced trials were subsequently more likely to choose the option that had fewer forced draws (compared to participants who observed only gains). That is, participants who observed losses were more likely to "explore" the relatively unknown option. Thus, overall, this literature appears to show that losses spur increased exploration, as measured by sampling before making a single consequential choice (Lejarraga et al., 2012), alternating between uncertain options (Yechiam, Zahavi, et al., 2015), or selecting a gamble with relatively uncertain payoffs (Krueger et al., 2017).

[2]Notably, the review also finds that a substantial share of respondents did not exhibit this asymmetry or were asymmetric in the opposite direction.

1.2 | Experimental task

Our paper considers a novel variant on explore-exploit problems. Unlike the studies mentioned above, which each present only two options, in our "Grain Game" paradigm, participants are placed in a one-dimensional environment that allows them to explore up to 70 options and receive payoffs on every choice (see Figure 1). Participants play the role of a farmer who traverses an environment and plants seeds; the extent of movement in each period is limited to a single step. Upon initializing the experiment, participants are assigned to a starting position that is automatically planted. As such, they see a single piece of "grain," with its corresponding payoff. All other locations initially have a question mark icon. As shown in Figure 1, as participants move, the most recent payoff for each location is displayed on the screen, whereas unexplored locations continue to display the question mark icon. If participants move close to the edge of the screen, the graphics shift to reveal new locations.

On each of the 70 turns, participants choose to plant in one of the three locations: the same location as the previous turn (defined as a decision to "exploit"), one space to the left or one space to the right. We define "exploration" as moving to a previously unexplored location and "retreat" as returning to a previously visited location.

Much like a geographic contour such as a mountain range, payoffs at adjacent points in the environment are highly correlated. That is, someone receiving a payoff of five is much more likely to receive a payoff of four in a neighboring location than a payoff of zero. Moreover, if participants have chosen a location in a previous round, they visually observe that prior outcome. If they have not previously chosen a given location, they have no experience of that location's payoffs and see a question mark icon. Each time participants select a given location, payoffs are determined by a fixed value that differs across locations plus noise of ±2 points. For the majority of locations, there is a .56 chance of receiving the expected value, a .15 chance of receiving ±1 from the expected value, and a .07 chance of receiving ±2 from the expected value. In the standard gain-only setting, the total range of points is [0, 25]. In the gain-loss setting, we shift all the gain-only payoffs down by 10 and provide an offsetting initial endowment, producing payoffs with a range [−10, 15]. The environments always include possible payoffs in the lower end of the range so that those in the gain-loss setting are likely to confront losses.

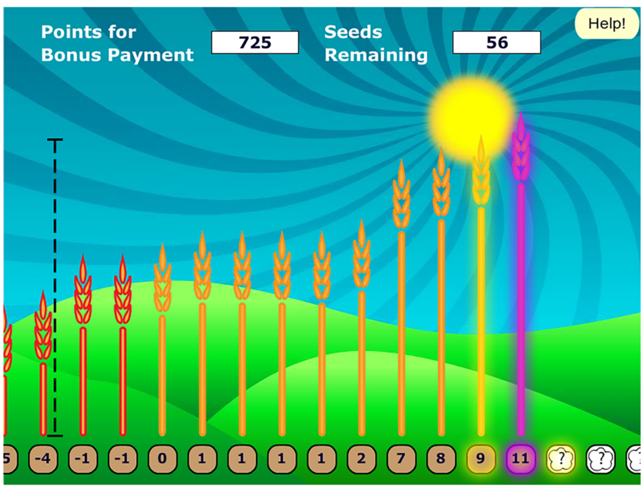

FIGURE 1 Screenshot of the Grain Game. *Note*: This figure shows one section of the gain-loss condition of the Grain Game. The top of the screen shows accumulated points and remaining turns (seeds remaining). The black dashed bar indicates the payoff range. The most recent payoffs received from explored locations are given in boxes below the grain, and unexplored positions are designated by question marks. The participant's current position is highlighted in purple; in this turn, they can exploit that position, retreat to the left, or explore to the right. The Help button displays instructions. In the gain-only condition, payoffs were increased by 10, such that the current location would have returned a payoff of 21. See Supporting Information for image of gain-only condition. [Colour figure can be viewed at wileyonlinelibrary.com]

For locations with expected values near the endpoints of the range in each condition, the probabilities associated with certain payoffs are collapsed. For instance, for a gain-only location with an expected value of 25 (the maximum possible), the chance of receiving 25 on any given turn was .78 (=.56 + .15 + .07), the chance of receiving 24 was .15, and the chance of receiving 23 was .07. If participants repeatedly chose a given location (e.g., by exploiting), the noise regenerates, and the graphics update to show the most recent payoff. As such, payoffs can vary with repeated selection, and there is some uncertainty associated with each position.

Our exploration task is characterized by three key features: (1) returns are *ex ante* unknown, (2) returns from options are locally correlated, and (3) movement is constrained to adjacent positions. Whereas the explore-exploit literature has considered a wide range of variations (see Mehlhorn et al., 2015 for a review), correlated payoffs in the context of exploration have, to the best of our knowledge, rarely been considered with human decision-makers (e.g., Wu et al., 2018) or animals, such as bumblebees foraging flower patches (Real, 1992). Moreover, few existing paradigms constrain movement, as we do, a condition that directly corresponds to search behavior in nature (e.g., a bumblebee can more easily search nearby flowers vs. those further away) and in the economy (e.g., a company can transition from gas-powered vehicles to electric cars but only gradually).[3] We capture cases in which ongoing investments are required to discover new solutions and there are costs for reverting.

1.3 | Theory and predictions

In examining the influence of losses on exploration behavior, we draw a distinction between *potential* and *experienced* losses. Potential losses are defined as negative payoffs that participants believe that they could experience from their next action. Potential losses are not possible in gain-only environments, which only contain payoffs of zero or more. Participants in mixed gain-loss environments are most likely to fear losses when near a location with a payoff of zero; because payoffs at adjacent points in our environment are correlated and have limited noise, the potential for a loss can be inferred from a positive, but near-

[3] One exception that considers movement constraints is Yechiam, Rakow et al. (2015).

zero, payoff. In contrast, *experienced* losses are negative payoffs that participants receive from the action they just took, whether that was exploring, exploiting, or retreating to a previously explored option.

Participants in the gain-only condition not only are informed that payoffs are zero or higher each turn but also may recognize that payoffs are sufficiently positive that any single-position shift (which is all that is allowed) is unlikely to lead to a negative payoff. We would expect that participants in the gain-only condition would be risk-averse when it comes to exploration, given that they are in the domain of gains, where the value function from the prospect theory is concave.

Participants in the gain-loss condition, who are experiencing gains but are close enough to the zero point that they face the potential of losses, will also be risk-averse because of the curvature of the value function. However, for these participants, a second, even more important, factor is likely to contribute to increased risk-aversion: the potential for losses, coupled with loss aversion. We therefore predict that those who are in the domain of gains will be more risk-averse—i.e., less likely to explore—if they are in the gain-loss condition than if they are in the gain-only condition (where they know no losses can be incurred as a result of exploration).

We expect that participants in the gain-loss condition who are experiencing losses in their current position will be motivated to avoid them in the future. As a consequence, they should be more risk-seeking—i.e., apt to explore—than those who are either experiencing gains with no potential for losses or experiencing gains and anticipate losses from exploration. We also expect that such participants will be more likely to "retreat" or return to a prior choice, if they had previously experienced gains and started to experience losses as a result of exploration. The behavior in both cases results because (1) the curvature of the value function implies that participants will be naturally risk-seeking in the domain of losses and (2) loss aversion predicts that they will be anxious to climb out of, or avoid, a loss situation.

Combining the analysis of these three situations generates straightforward predictions: Individuals who are experiencing losses in the gain-loss condition will be the most likely to shift positions, either by exploring or by retreating to a location in which they did not experience losses. Those in the gain-loss condition who are experiencing gains will be least likely to explore. Finally, those in the gain-only condition will be in between in their propensity to explore.

Although the only conditions we examine empirically are those described here, similar lines of reasoning would lead to parallel predictions in different search settings, including those in which (1) payoffs are uncorrelated in adjacent locations, (2) unrestricted movement is possible, or (3) retreat is not permitted. Our novel environment seeks to represent exploration as it occurs in a wide range of real-world settings in which discoveries and insights most commonly emerge from incremental exploration rather than from unanticipatable random and potentially large shocks that could occur with uncorrelated payoffs.

Notably, our paper represents a deviation from standard literature in the explore-exploit tradition. Multiple papers using the decisions from experience and sampling paradigms concentrate on the *probabilities* associated with risky choices, and the fact that participants often appear to under weigh rare events—in contrast to prediction made under the prospect theory that people often overweigh these events (e.g., Hertwig et al., 2004; Hertwig & Erev, 2009; Newell et al., 2015; Yechiam, Rakow, et al., 2015; see also Liang et al., 2019). In these situations, a natural and fruitful way to analyze behavior is through models that concentrate on learning, as participants need to learn about payoffs through repeated draws. Here, in contrast, we concentrate on predictions derived from the prospect theory's *value function*, which includes a reference point at zero and loss aversion. This concentration on the value function leads us to emphasize how risk-seeking changes over positive and negative experienced and potential payoffs.

2 | STUDY 1A

Our first two studies were designed primarily to test the prediction that potential losses would cause participants to exhibit more risk-averse exploitative behavior. In the "Grain Game," as already described, participants explore a one-dimensional environment and earn payoffs that depend on their location. To study the consequences of the potential occurrence of losses, we varied whether long-run exploration would reduce (Study 1A) or increase (Study 1B) total payoffs by varying global and local optima in the underlying payoff structures. Our theory predicts that the potential to incur losses will decrease exploration and, hence, that the gain-loss framing will increase payoffs in the Study 1A environment and decrease payoffs in the Study 1B environment. To simplify analysis of behavior, in both studies, all participants began on the left edge of a one-dimensional environment. This feature allows us to quantify exploration as the maximum distance from the left edge that the individual advances (explores).

2.1 | Method

2.1.1 | Participants

This study and all subsequent studies in this paper were approved by the Institutional Review Board at Carnegie Mellon University. We recruited 140 US residents online using Mechanical Turk, based on a power analysis suggesting that we needed at least 48 participants for each of the two conditions (we used a *t*-test assuming that the gain-only group would explore to Location 35 and the gain-loss group would explore to Location 25, with $\sigma = 15$, $\alpha = .05$, and $\beta = .90$). The average age was 31.4 years ($SD = 9.3$), and 58% were male.

2.1.2 | Experimental task

In "The Grain Game," participants played the role of a farmer who chooses where to plant crops (Figure 1). Participants were told that "In this game, you will play the role of a farmer who just bought a 70-acre field. Unfortunately, you don't know which parts of your field will produce the most grain. Each turn, you have two choices. You can plant a seed in a spot where you have not planted before ['?' icon], or in a part of the field that you have already tried [‘planted’ icon].” They then received instructions on how to play the game using their keyboard and were told that they would receive payment based on their performance (see Supporting Information for screenshots of these instructions with complete wording).

To ensure that participants understood the game and would know the range of payoffs to expect, they next answered attention check questions asking them to report the high and low values that they could receive (“What is the highest/lowest number of points you can get each turn?”) and the size of the field they had available. Those who incorrectly answered either of the first two questions received a warning (e.g., “Try again! You can get more/fewer points” or “Try again! You can't get that many/few points”) and had two subsequent attempts to answer. After three incorrect responses, the correct answers were displayed (e.g., “The maximum number of points you can get is 25”). Participants had to input these values before being allowed to move to the next part of the experiment. As we believed these instructions were sufficient for educating participants about possible payoff amounts, we do not exclude any participants who completed the game. Participants went on to make 70 consequential decisions.[4] At the end of the study session, participants received payment based on accumulated points (four points = $0.01).[5]

2.1.3 | Experimental conditions: Gain-only and gain-loss

To assess the effects of the potential for and/or experience of losses, we used a two condition between-subjects experiment. Specifically, we assigned participants to one of the two conditions that differed in the framing of payoffs. Participants in the *gain-only* condition were in an environment with payoffs ranging from 0 to 25 points per turn. Those in the *gain-loss* condition had payoffs that were lower by 10 points (resulting in a range of −10 to 15 points) but, to equate potential payoffs, received extra points at the beginning of the task. Thus, although participants were randomly assigned to a different point system, the expected value of the experiment (conditional on the same behavior) was the same. Payoffs were represented graphically by the height of the grain displayed on the screen (Figure 1).

2.1.4 | Underlying payoff structure

Across the 70 locations available to participants, payoffs were correlated. In particular, the underlying payoff structure (shown in Figure 2, right axis, line labeled “Environment”) shows that participants who explored would experience an initial rise in payoffs, through Position 8. After that, payoffs declined before rising again to reach a moderate level for the remainder of the field. This payoff structure was created to establish a situation in which long-run exploration was not worthwhile because the global optimum was near the participants' starting position. Specifically, the highest payoff was available after only eight moves from the left-hand side of the environment.

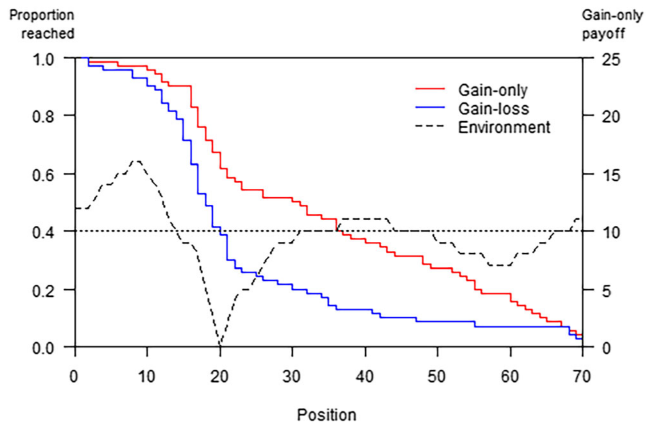

FIGURE 2 Extent of exploration and environment used in Study 1A. *Note*: This chart shows the proportion of participants in each condition who reached each position (red and blue lines) and the underlying payoff structure of the environment (black dashed line). The highest payoff (16 points in the gain-only condition or 6 points in the gain-loss condition) is obtained at Positions 8 and 9. Participants in the gain-loss condition first experience an option with negative expected payoffs at Position 15, below the dotted line. [Colour figure can be viewed at wileyonlinelibrary.com]

2.2 | Results

2.2.1 | Overall exploration behavior

To examine whether exploration varied between the two conditions, we created a hazard plot showing the proportion of participants in each condition who explored to a given location (Figure 2, left axis). All participants started at the left-most position, so overall exploration was captured by movement to the right. The figure shows that participants in the gain-loss condition were much less likely to explore beyond Position 15, where they first started to experience losses. The average gain-only participant explored 32.5 times ($SD = 19.8$) and the average gain-loss participant explored 21.3 times ($SD = 16.1$; $t[138] = 3.7$, $p < .001$).

2.2.2 | Explore-exploit-retreat behavior

To study the effects of potential and experienced losses, we analyzed the actions that the participants took each turn. Actions were classified as *exploring* a new location, *retreating* to a previously explored location, or *exploiting* the current location. To analyze those actions, we ran the following multinomial logit regression:

[4] Due to a programming error, we failed to record one choice made by participants. The regressions in Table 1 also include a lag term, with the result that we are modeling 68 explore-exploit-retreat decisions.

[5] For exploratory purposes, we also asked participants about their thought process while playing the game, the purpose of the game, and their interest in exploring (six statements). We concentrate on behavior rather than these responses.

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

Specifically, participant i at time t takes action $j \in \{exploit, explore, retreat\}$. *Gain-loss* is an indicator variable set to one for gain-loss participants and zero for gain-only participants. *Adjusted payoff* represents the value of the payoff from the previous turn, scaled to the gain-only range of [0, 25] by adding 10 to every payoff from the gain-loss condition. *Below threshold* is an indicator variable representing whether the adjusted payoff from the previous turn was below 10 points and would therefore be negative in the gain-loss condition. We used clustered standard errors at the individual level to account for correlation in each participant's actions.

Our main predictions were that (a) the coefficient on β_1 (*gain-loss*) would be negative for exploration, representing a reluctance to explore in the gain-loss condition due to potential losses and (b) β_4 (*gain-loss* × *below threshold*) would be negative for exploitation, suggesting a reluctance to exploit after experiencing a loss. We also expected that participants would be more likely to exploit when payoffs were high, leading to a positive coefficient for β_2 (*adjusted payoff*).

The multinomial logit results are shown in Table 1 with exploitation as the omitted action. The coefficients are exponentiated, meaning that values above one represent cases where the predicted action is more likely. Conversely, values less than one represent cases where the action is less likely. As predicted, given our expectations about potential losses, participants in the gain-loss condition were less likely to explore than those in the gain-only condition. Specifically, the coefficient on *gain-loss* suggests that gain-loss participants were less likely to explore than to exploit on each turn. Contrary to our expectations, there was no difference in behavior between conditions when participants experienced losses (i.e., coefficient on *gain-loss* × *below threshold*).

The coefficient on *adjusted payoff* indicates that every one point increase in payoffs received was associated with a decrease in exploration and retreat regardless of condition, suggesting that participants were more likely to repeatedly exploit high-value locations rather than move to a new location. Participants below the 10-point threshold were particularly unlikely to explore.

2.2.3 | Points received

The underlying payoff structure in this study was designed to be unfavorable to exploration because the global optima were relatively close to the participants' starting location. 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 on average, received significantly fewer points ($M = 816.3$, $SD = 163.4$) than those in the gain-loss condition who explored less on average ($M = 881.8$, $SD = 172.5$; $t[138] = -2.30$, $p = .02$). That is, the presence of losses in the environment discouraged ex post suboptimal exploration and made decision-makers better off.

3 | STUDY 1B

This study used the same design as Study 1A with one key change. Here, we presented participants with an underlying payoff environment that rewarded, instead of punished, exploration because reaching the global optimum required long-run exploration. This setup allows us to test whether, in line with our theory, the potential to incur losses decreases exploration even if exploration would lead to higher payoffs.

TABLE 1 Multinomial logit regression results predicting participant actions

	Study 1A		Study 1B		Study 2	
	Explore	Retreat	Explore	Retreat	Explore	Retreat
Gain-loss condition	0.472** (0.137)	0.738 (0.296)	0.413** (0.115)	0.403* (0.174)	1.068 (0.180)	1.312 (0.241)
Adjusted payoff [t-1]	0.916*** (0.012)	0.957** (0.015)	0.851*** (0.008)	0.915*** (0.011)	0.830*** (0.005)	0.861*** (0.006)
Below threshold [t-1]	0.568*** (0.075)	0.872 (0.149)	0.525*** (0.073)	1.087 (0.179)	0.770*** (0.051)	0.812** (0.062)
Gain-loss condition × below threshold [t-1]	0.772 (0.121)	1.415 (0.288)	0.957 (0.143)	1.346 (0.234)	1.808*** (0.148)	1.391*** (0.134)
Constant	5.312*** (1.435)	0.782 (0.280)	12.722*** (3.030)	1.609 (0.557)	13.580*** (1.963)	3.702*** (0.594)
Num.Obs.	9,520		8,024		40,868	
RMSE	0.38		0.39		0.38	

Notes: The table presents exponentiated coefficients and standard errors. Participant actions were classified as *exploring* a new location, *retreating* to a previously explored location, or *exploiting* the current location. *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 from the prior turn was below the threshold of 10.
*$p < .05$.**$p < .01$.***$p < .001$.

3.1 | Method

3.1.1 | Participants

We recruited 118 participants via Amazon Mechanical Turk, based on a power analysis suggesting that we needed at least 43 participants for each of the two conditions (we used a *t*-test assuming that the gain-only group would explore to Location 42 and the gain-loss group would explore to Location 28, with $\sigma = 20$, $\alpha = .05$, and $\beta = .90$). They were 30.7 ($SD = 8.7$) years old on average, and 56% were male.

3.1.2 | Experimental conditions: Gain-only and gain-loss

Similar to Study 1A, Study 1B was a two condition, between-subjects, experiment in which participants were randomly assigned to a gain-only or gain-loss framing. The instructions and attention check questions were identical to Study 1A.

3.1.3 | Underlying payoff structure

The payoff function is shown in Figure 3 (right axis, line labeled "Environment"). In this study, long-run exploration was rewarded, because the local optimum for participants starting the game was significantly lower than the global optimum. In particular, participants received the maximum possible payoff at Position 39; however, to realize these payoffs, they had to first traverse through an area of low rewards. Participants in the gain-loss condition received negative payoffs in expectation from Positions 10 to 26.

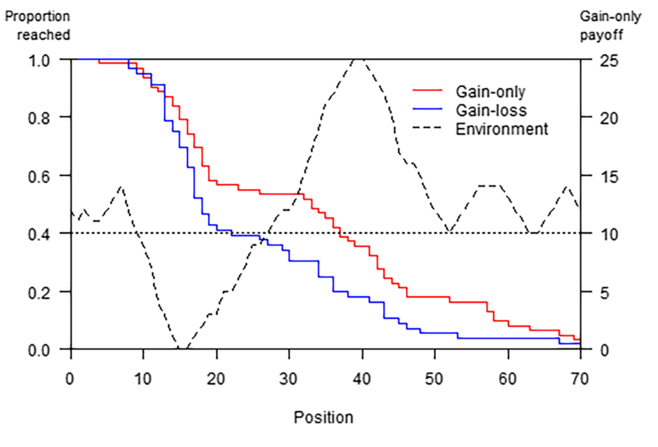

FIGURE 3 Extent of exploration and environment used in Study 1B. *Note*: This chart shows the proportion of participants in each condition who reached each position (red and blue lines) and the underlying payoff structure of the environment (black dashed line). The highest payoffs (25 points for the gain-only condition and 15 points for the gain-loss condition) are obtained at Position 39. Participants in the gain-loss condition first experience losses in expectation at Position 9. [Colour figure can be viewed at wileyonlinelibrary.com]

3.2 | Results

3.2.1 | Overall exploration

Following our analysis from Study 1A, we generated a hazard plot of exploration (Figure 3). As shown, participants in the gain-loss condition explored less than those in the gain-only condition. The average distance explored was 23.0 ($SD = 14.7$) for those in the gain-loss condition and 30.3 ($SD = 18.3$) for those in the gain-only condition ($t[116] = 2.34$, $p = .02$).

3.2.2 | Explore-exploit-retreat behavior

We classified participant actions as explore, exploit, or retreat and analyzed these actions using a multinomial logit regression with the same form as in Study 1A (Table 1). The results show that participants in the gain-loss condition were less likely to explore than those in the gain-only condition, holding all else equal (coefficient on *gain-loss*); this pattern is consistent with the idea that gain-loss participants were reluctant to explore because of potential losses. The odds that a participant would retreat were also lower in the gain-loss condition.

In addition, the coefficients on *adjusted payoff* suggest that obtaining a high payoff was correlated with increased exploitation and diminished exploration and retreat behavior, demonstrating that participants sought to repeatedly select locations with relatively high payoffs. Additionally, receiving an expected payoff below the threshold was associated with a lower probability of exploring, but the magnitude of this effect did not differ for those in the gain-only versus gain-loss conditions.

3.2.3 | Points received

The underlying payoff structure in this study was designed to be favorable to exploration. 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 farther, received more points ($M = 810.0$, $SD = 185.7$) than those in the gain-loss condition ($M = 740.9$, $SD = 187.9$; $t[116] = 2.01$, $p = .047$).

3.3 | Discussion of Studies 1A and 1B

We predicted that losses would affect exploration in two ways. First, we expected that people would be less likely to explore when they faced the potential to receive negative payoffs. Studies 1A and 1B both provide support for this hypothesis, as those in the gain-loss condition were less likely to explore than those who were in an environment where only gains were possible. That is, the possibility of incurring losses reduced exploration, and this pattern persisted regardless of whether exploration was ultimately rewarding. Second, we hypothesized that people would be less likely to exploit after experiencing a loss, due to loss aversion. We find no statistically significant evidence for this hypothesis in Study 1, which leaves it as our main focus in Study 2.

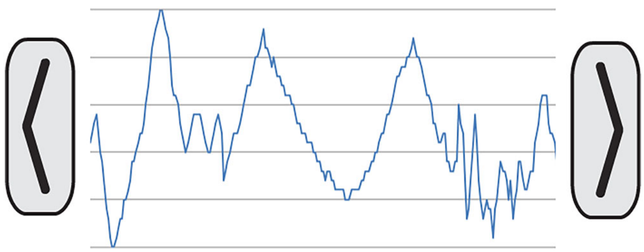

FIGURE 4 Preview of underlying payoff structure used in Study 2. *Note*: All participants viewed this graph prior to starting the study. The arrows on the left and right could be used to view the entire environment. As a result, participants were aware of the range of potential payoffs as well as how frequently they occurred. [Colour figure can be viewed at wileyonlinelibrary.com]

4 | STUDY 2

For Study 2, we used a much larger landscape, initiated participants to a random starting location (rather than at the left-most edge), and increased our sample size substantially. These changes allowed us to increase the variation in underlying payoff profiles that participants experienced, addressing concerns that the effects in Studies 1A and 1B were unique to these two specific payoff structures. Additionally, the larger sample allowed us to more robustly test our hypothesis that experiencing losses would cause participants to move.

4.1 | Method

4.1.1 | Participants

We recruited 601 US residents online using Mechanical Turk. The average age was 32 years ($SD = 10.1$), and 61% were male. We recruited a larger sample because we anticipated an increase in variation of decisions given the large environment and the random initial starting location that we used in this study.

4.1.2 | Underlying payoff structure

We created a version of the Grain Game using an environment containing 700 locations. To increase ecological validity, we used correlation occurring in a natural setting: the contours of the underlying environment were based off of hiking elevation maps from hikes in Virginia.[6] Each participant was randomly assigned to a starting position in this environment that was at least 70 positions removed from either edge, to ensure that participants could explore in either direction as far as they wanted. To provide participants with fixed expectations about the nature of the underlying payoffs, all participants viewed a graph of the landscape before starting the game (Figure 4) and were told they would start in a random position on this graph.[7] That is, they could learn that there were multiple local optima and get a sense of how frequent different payoffs were. As in Studies 1A and 1B, participants in the gain-loss condition received an initial endowment, and payoffs for each option were shifted downward by 10 points.

4.2 | Results

4.2.1 | Overall exploration

To examine patterns in exploration behavior, we calculated the probability of exploration for each condition-payoff combination, as plotted in Figure 5. In this figure, we again normalize payoffs to the range of [0, 25] by adding 10 points to the outcomes in the gain-loss condition. The negative slope in the left panel demonstrates that in both conditions, participants were less likely to explore when receiving higher payoffs, showing that they were sensitive to rewards. More importantly, a comparison of average behavior across the two conditions shows a difference in exploration patterns. To the left of the red dotted line, gain-loss participants experienced losses, whereas gain-only participants received low positive payoffs. Over this range, gain-loss participants displayed a higher propensity to explore, consistent with our prediction that experiencing losses encourages people to change locations. In contrast, to the right of the red dotted line, gain-loss participants were receiving payoffs just above zero and could anticipate that exploration might lead to negative payoffs. Here, as expected, gain-loss participants were less likely to explore than gain-only participants.

Another way to visualize this interaction is by calculating the difference in the proportion of exploration decisions across conditions. In the right panel of Figure 5, we subtract the proportion of exploration decisions in the gain-only condition from those in the gain-loss condition for each payoff. This panel shows directly that participants in the gain-loss condition are more likely to explore when receiving negative payoffs. With one exception, each negative payoff shows an increase in the probability of exploration between one and 13 percentage points. Additionally, gain-loss participants are less likely to explore than gain-only participants when above the zero-payoff threshold. At the highest payoffs in each condition, where it is unlikely that losses could be incurred from further exploration, there is no meaningful difference between conditions.

4.2.2 | Explore-exploit-retreat behavior

As in previous studies, we ran a multinomial logit on each participant action (Table 1). Consistent with our prediction about the effects of experienced losses and in line with the primary focus of this study, the

[6]Adapted from www.hikingupward.com.

[7]We asked the same exploratory questions as in Studies 1A and 1B. We focus on experimental effects rather than these questions. Additionally, we asked participants whether they copied the graph shown at the beginning of the study. Only four participants reported that they did so, and we included them in our analyses.

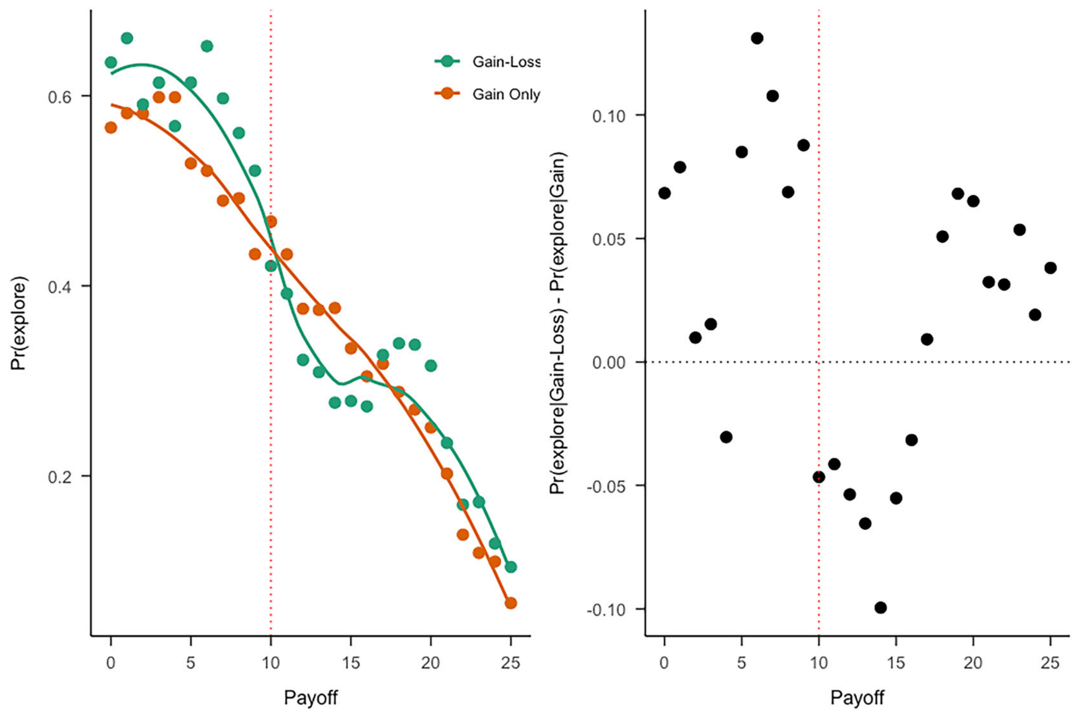

FIGURE 5 Probability of exploration by condition, Study 2. *Note*: The left panel shows the probability of exploring by condition. The lines show the best fitting smoothed spline. The right panel shows the difference in probability of exploration between the two conditions, with values above zero on the y-axis, representing higher exploration rates for gain-loss participants. In both panels, payoffs are rescaled to the gain-only range of [0, 25]. For payoffs less than 10 (i.e., to the left of the red dotted lines), participants in the gain-loss condition incur losses. [Colour figure can be viewed at wileyonlinelibrary.com]

FIGURE 6 Evolution of points received by condition and turn, Study 2. *Note*: This graph shows the points received by participants in each of the two conditions, divided into five, 14-period panels. Adjusted payoffs are shown such that +10 represents a payoff of 0 in the gain-loss condition. *D* and *p*-values represent Kolmogorov–Smirnov tests and corresponding significance values for differences between the gain-only and the gain-loss conditions for each group of 14 periods. [Colour figure can be viewed at wileyonlinelibrary.com]

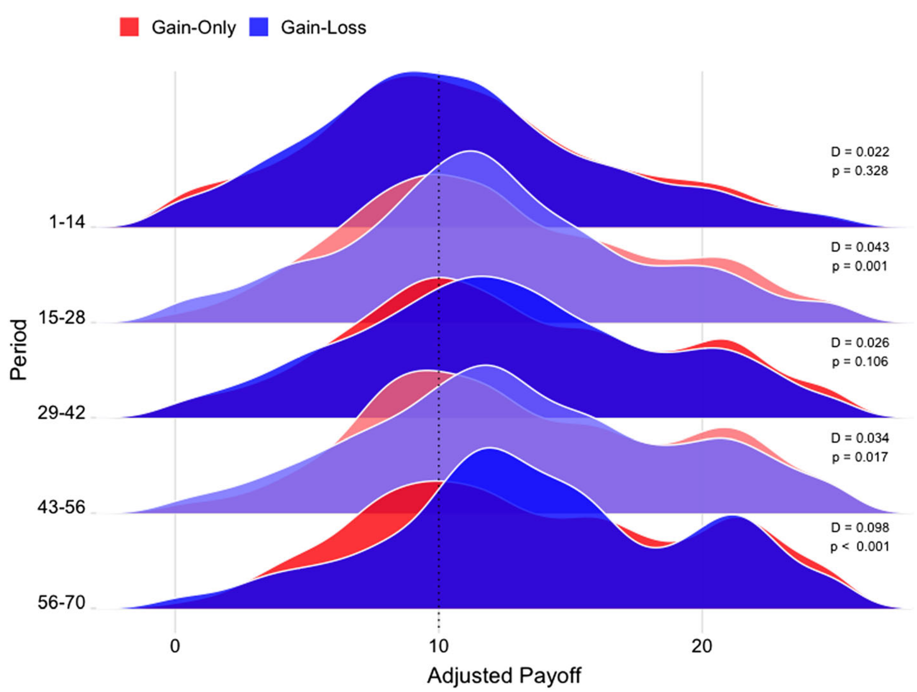

interaction between *gain-loss* and *below threshold* shows that choices following low payoffs differ across conditions. Participants in the gain-loss condition who are receiving negative payoffs are both more likely to explore and to retreat (i.e., less likely to exploit) than those in the gain-only condition who are receiving low, but positive, payoffs.

The regression coefficients on *adjusted payoff* indicate that participants in both conditions were less likely to move to a different location when receiving higher payoffs; for each additional point in earnings associated with a prior location, participants were less likely to explore and less likely to retreat. Finally, there was no significant main effect of being in the gain-loss environment.

4.2.3 | Points received

There was no difference in average payoffs between the two conditions ($p = .57$). Unlike in our previous studies, the underlying payoff structure was not designed to be favorable or unfavorable to exploration. Instead, it was possible to reach different global and local optima depending on one's randomly assigned starting position.

4.2.4 | Changes over time

In line with literature on decisions from experience (e.g., Wu et al., 2018; Yechiam & Busemeyer, 2005), we also explored changes over time. To do so, we divided the 70-period game into five equal groups of 14 periods each and graphed the distribution of points received for participants in the two conditions (see Figure 6). The figure shows that, at the beginning of the game, there is no difference in the distribution of points received across the two conditions, as expected by random assignment. However, as the game progresses, there is a growing difference between the distribution of points received by gain-only and gain-loss respondents. Specifically, the gain-only respondents are more likely to receive points immediately below the adjusted payoff threshold of +10 (i.e., 0 in the gain-loss condition), as shown in the bottom panel. This difference may arise because loss-averse participants in the gain-loss condition explore or retreat to avoid such payoffs, whereas gain-only participants do not experience low positive payoffs with the same levels of aversion.

5 | GENERAL DISCUSSION

Many important life decisions require an initial investment before realizing significant rewards. Decision-makers who choose to explore these paths may experience initial losses, leading them to wonder whether continued investments are worthwhile or whether they should revert to their old ways. We study these kinds of exploration decisions using an exploration task with constrained movement in which participants can explore up to 70 options. Overall, our results show that in situations with unknown potential outcomes in which there is uncertainty about the payoffs associated with exploration, *potential* losses lead to less exploration (Studies 1A and 1B) and *experienced* losses lead to less exploitation (Study 2). These findings are consistent with prospect theory's value function, according to which decision-makers are loss-averse and risk-averse in the gains domain and risk-seeking in the loss domain. In the end, our two primary findings are that people try to avoid losses through both exploration and exploitation and modify their behavior depending on their expectations about what each action might yield. As such, we conclude that exploration is driven by hope of anticipated gains, fear of potential losses, and avoidance of experienced losses.

5.1 | Implications for explore-exploit literature

Paradigms that present participants with only two options are common in the explore-exploit literature (e.g., Hertwig et al., 2004; Rakow & Newell, 2010; see also Lejarraga et al., 2012). By allowing participants to explore more than two options, our work builds on Yechiam, Rakow et al. (2015), who asked participants to explore a 10×13 two-dimensional grid that was visually divided into relatively risky and relatively safe halves. In that paper, as in ours, the participants' movement was restricted to immediately adjacent locations. The authors analyzed choices on the border between the risky and the safe areas to show a slight decrease in risk-taking following the occurrence of a rare loss. However, the environment only contained losses, which limited the authors' ability to examine loss aversion as we do. In contrast, we use a framing manipulation, allowing us to compare the decisions of participants who know they cannot face losses to those who face the potential for losses.

Using an environment with more than two options allowed us to classify the participants' movement as explore, exploit, or retreat behavior. We believe that most scholars would consider both exploit behavior and retreat behavior to indicate risk aversion, because both require reselecting a known location. In contrast, exploration is indicative of risk-seeking. As such, one contribution of our paper is to show that people experiencing losses will move on and can choose either risk-averse (retreat) behavior or risk-seeking (exploration) behavior to do so. It would be worthwhile for future research to determine when losses lead people to choose each course of action and whether these decisions are moderated by prior experience or individual differences.

In our experiment, when a participant experiences a payoff, it is displayed on the screen for the remainder of the experiment. As a result, the task requires little working memory. It is possible that hiding the display of prior payoffs would lead to increased retreat behavior, if participants could not recall a previous outcome or wondered if that outcome might have changed (recall that in our setting, the underlying value of each location was fixed with noise). A future experiment could test whether exploration increases when prior choices are hidden, by asking participants to make left- and right-movement choices without showing a payoff gradient.

5.2 | Implications for literature on losses and loss aversion

A natural question arising from our work is whether the loss aversion-driven patterns of explore-exploit behavior that we observe are adaptive. Our first two studies serve as demonstrations that loss aversion can be helpful or harmful depending on whether the environment rewards exploration. Study 1A, which used an environment unfavorable to exploration, contributes to an expanding literature documenting situations in which biases help people make better decisions. Kahneman and Lovallo (1993), for example, showed that loss aversion and risk aversion can be beneficial to managers who are overconfident, 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. In contrast, as a result of its different payoff structure, Study 1B yields the more conventional result that loss aversion leads to lower payoffs. Finally, Study 2 has no overall difference in payoffs between those who can and those who cannot experience losses. Thus, whereas our studies do not demonstrate that loss aversion is generally helpful, our results indicate that it is not always a mistake.

A growing literature questions whether loss aversion is a universal phenomenon (e.g., Erev et al., 2008; Gal & Rucker, 2018; Yechiam, 2019). Losses have been found to affect the rate at which participants learn about experimental stimuli (Bereby-Meyer & Erev, 1998), increase how consistent they are when making risky choices (Yechiam & Telpaz, 2013), and increase performance. As such, some scholars now argue that losses affect attention rather than subjective utility about gains and losses (Yechiam, 2019; Yechiam & Hochman, 2013). Unfortunately, in our studies, we do not have data that allow us to observe attention, and therefore we cannot speak directly to this issue. Future work could address attentional effects by asking participants to recreate the environment they had explored and see whether recall is higher among those in gain-loss environments.

If experiencing losses can motivate potentially highly rewarding exploration, there may be strategic incentives to create a setting in which payoffs are framed as small losses. Companies, for example, may become complacent when they are earning profits but would be motivated to innovate and cut costs if those profitability were trending downward—toward losses. In settings in which someone may want to induce risk-taking and exploration (e.g., a company trying to spur new research and development), artificially framing current processes as losses may motivate new initiatives. As Johnson (2010) argues, shortages have a tendency to activate human ingenuity, as expressed in the proverb that "necessity is the mother of invention." Our findings suggest that experiencing losses has a similar effect in motivating exploration.

ACKNOWLEDGMENTS

We thank Efrat Aharonov, Vikram Jambulapati, Melissa Knoll, Amanda Markey, Chris Stocks, Eric M. VanEpps, and participants at the 2016 Society for Judgment and Decision Making meeting for comments. We also thank Steven B. Nash for research assistance.

CONFLICT OF INTEREST

Not applicable.

DATA AVAILABILITY STATEMENT

The data and code to reproduce all statistical analyses and figures in the manuscript are available via OSF: https://osf.io/dktuc/.

ORCID

Alycia Chin https://orcid.org/0000-0002-9570-0549
David Hagmann https://orcid.org/0000-0002-2080-997X
George Loewenstein https://orcid.org/0000-0003-2790-0474

REFERENCES

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, *110*(1), 73–92. https://doi.org/10.2307/2118511

Bereby-Meyer, Y., & Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. *Journal of Mathematical Psychology*, *42*, 266–286. https://doi.org/10.1006/jmps.1998.1214

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. https://doi.org/10.1002/bdm.602

Gal, D., & Rucker, D. D. (2018). The loss of loss aversion: Will it loom larger than its gain? *Journal of Consumer Psychology*, *28*(3), 497–516. https://doi.org/10.1002/jcpy.1047

Genesove, D., & Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics*, *116*(4), 1233–1260. https://doi.org/10.1162/003355301753265561

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. https://doi.org/10.1111/j.0956-7976.2004.00715.x

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Science*, *13*(2), 517–523. https://doi.org/10.1016/j.tics.2009.09.004

Homonoff, T. A. (2018). Can small incentives have large effects? The impact of taxes versus bonuses on the demand for disposable bags. *American Economic Journal: Economic Policy*, *10*(4), 177–210.

Imas, A., Sadoff, S., & Samek, A. (2016). Do people anticipate loss aversion? *Management Science*, *63*(5), 1271–1284. https://doi.org/10.1287/mnsc.2015.2402

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

Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, *39*(1), 17–31. https://doi.org/10.1287/mnsc.39.1.17

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Krueger, P. M., Wilson, R. C., & Cohen, J. D. (2017). Strategies for exploration in the domain of losses. *Judgment and Decision Making*, *12*(2), 104–117.

Lejarraga, T., & Hertwig, R. (2017). How the threat of losses makes people *explore* more than the promise of gains. *Psychonomic Bulletin Review*, *24*, 708–720. https://doi.org/10.3758/s13423-016-1158-7

Lejarraga, T., Hertwig, R., & Gonzales, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, *124*(3), 334–342. https://doi.org/10.1016/j.cognition.2012.06.002

Liang, G., Newell, B. R., Rakow, T., & Yechiam, E. (2019). Further investigations of how rare disaster information affects risk taking: A registered replication report. *Psychonomic Bulletin & Review*, *26*, 1411–1417. https://doi.org/10.3758/s13423-019-01594-w

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191–215. https://doi.org/10.1037/dec0000033

Newell, B. R., Rakow, T., Yechiam, E., & Sambur, M. (2015). Rare disaster information can increase risk-taking. *Nature Climate Change*, *6*, 158–161. https://doi.org/10.1038/nclimate2822

Pope, D. G., & Schweitzer, M. E. (2011). Is Tiger Woods loss averse?: Persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, *101*(1), 129–157. https://doi.org/10.1257/aer.101.1.129

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. https://doi.org/10.1002/bdm.681

Rakow, T., Newell, B. R., & Wright, L. (2015). Forgone but not forgotten: The effects of partial and full feedback in "harsh" and "kind" environments. *Psychonomic Bulletin & Review*, *22*(6), 1807–1813. https://doi.org/10.3758/s13423-015-0848-x

Real, L. A. (1992). Information processing and the evolutionary ecology of cognitive architecture. *The American Naturalist*, *140*, S108–S145. https://doi.org/10.1086/285399

Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., & Damasio, A. R. (2005). Investment behavior and the negative side of emotion. *Psychological Science*, *16*(6), 435–439. https://doi.org/10.1111/j.0956-7976.2005.01553.x

Teodorescu, K., & Erev, I. (2014). On the decision to explore new alternatives: The coexistence of under- and over-exploration. *Journal of Behavioral Decision Making*, *27*(2), 109–123. https://doi.org/10.1002/bdm.1785

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision-spaces. *Nature Human Behavior*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4

Yechiam, E. (2019). Acceptable losses: The debatable origins of loss aversion. *Psychological Research*, *83*(7), 1327–1339. https://doi.org/10.1007/s00426-018-1013-8

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*(3), 387–402. https://doi.org/10.3758/BF03193783

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effect of losses over gains. *Psychological Bulletin*, *139*(2), 497–518. https://doi.org/10.1037/a0029383

Yechiam, E., Rakow, T., & Newell, B. (2015). Super-underweighting of rare events with repeated descriptive summaries. *Journal of Behavioral Decision Making*, *28*(1), 67–75. https://doi.org/10.1002/bdm.1829

Yechiam, E., & Telpaz, A. (2013). Losses induce consistency in risk taking even without loss aversion. *Journal of Behavioral Decision Making*, *26*(1), 31–40.

Yechiam, E., Zahavi, G., & Arditi, E. (2015). Loss restlessness and gain calmness: Durable effects of losses and gains on choice switching. *Psychonomic Bulletin & Review*, *22*, 1096–1103. https://doi.org/10.3758/s13423-014-0749-4

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Chin, A., Hagmann, D., & Loewenstein, G. (2022). Fear and promise of the unknown: How losses discourage and promote exploration. *Journal of Behavioral Decision Making*, e2309. https://doi.org/10.1002/bdm.2309