Necessity Is the Mother of Invention: Fear and Promise of the Unknown

Alycia Chin, David Hagmann, and George Loewenstein

Author Note

Alycia Chin,*† Economic and Risk Analysis, Public Company Accounting Oversight Board, 1666 K Street, Washington DC 20006, china@pcaobus.org, ORCID ID: 0000-0002-9570-0549; David Hagmann, Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213, hagmann@cmu.edu; George Loewenstein, 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

Acknowledgements

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.

Word count: 4,398

Abstract

In “explore-exploit” situations, decision makers choose between taking a risk on an unknown option (exploration) and taking advantage of a known option (exploitation). We extend the traditional paradigm to a setting in which movement is constrained to adjacent locations, with payoffs correlated among neighboring options. Participants do not face random payoff shocks, but have to decide whether to persist in their exploration, aiming to attain a higher potential payoff, or to trace back their steps to a previously explored option, incurring lower payoffs along the way. We study how explore-exploit-retreat behavior is affected by the presence of losses and predict that people will be (a) more likely to explore when *experiencing* losses, but (b) less likely to explore when they *anticipate* losses. In two studies, we find support for our predictions, with significant differences in exploration and exploitation behavior consistent with attempts to avoid losses. We additionally demonstrate that loss aversion can be adaptive, leading participants in certain environments to higher total earnings.

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

Introduction

A perennial decision problem facing individuals and firms alike is the extent to which they should search for and invest in technologies. New tools promise to enhance our future productivity, but come with considerable upfront cost: because individuals have to learn how to navigate them, there is a loss in productivity as the new technology is initially adopted. It is only once people have become familiar with its use that, we hope, the switching costs are justified. However, rarely do we know in advance whether the promised gains will indeed materialize. Many promises of enhanced productivity turn out to have been overestimated, leading us to realize that we invested in the tool's adaptation for nothing. We may then have no option but to revert to the old way of operating.

Searching in this environment is characterized by three key features: (1) future returns are unknown, (2) there is the potential to incur a cost during a transition period, and (3) there is an additional cost to reverting to the previous process. The previous literature on explaration, reviewed below, focuses on the first key feature. Participants are exposed to a set of options about which they have no information and can freely explore, choosing any option they see fit (CITE). Any positive (or negative) outcome is then a random shock that the decision maker could not have anticipated. In contrast, we propose that exploration in many consequential domains is characterized by a path of payoffs. Returns may initially be low, but decision makers may continue to pursue their exploration because they expect the investment to lead to higher gains in the future. Moreover, while previous experiments allow decision makers to instantly return to a known option, this is rarely available in the real world: as costly as it is to switch to a new technology, it is also costly to undo the change.

Firms choosing to pursue research and development, for example, may have invested in hiring workers and building facilities. They have to decide how long to continue investing resources into the effort and if they should cut their losses and retire the new infrastructure. Researchers, too, face this conundrum when deciding whether to continue pursuing a new line of research. Having invested time in gaining an understanding in the topic, they may not be willing to abandon the idea after a single failed experiment – and returning to previous ideas after months of leaving them dormant comes with costs of its own.

An open question, then, is what influences the decision to pursue these new, risky endeavours and, just as importantly, the decision to abandon the efforts. We propose that the presence and anticipation of losses can play a central role. When the currently chosen option leads to losses, individuals and firms alike are most motivated to seek out alternatives (CITE loss aversion). We may then conclude that losses are beneficial for exploration. However, when people can *anticipate* losses, yet do not incur them, we would expect the opposite effect: the risk of incurring the loss would deter someone from exploring.

The tradeoff embodied in this situation is captured in “explore-exploit” models, in which decision makers choose between exploring unknown options and exploiting known options (Zwick et al. 2003). It is nearly impossible to find optimal solutions to explore-exploit decisions (e.g., Gittins, 1979; Lee, Zhang, Munro & Steyvers, 2011), so decision makers use a variety of solutions to address these problems (see Cohen, McClure & Yu, 2007; Wilson, et al., 2014; and Mehlhorn, et al., 2015 for a review). Moreover, these models assume that every available option can readily be chosen and that payoffs across options are not correlated. Yet,

Exploration, Loss Aversion, and Anticipated vs. Experienced Losses

This paper examines how explore-exploit decisions are affected by loss aversion, the tendency for people to be more sensitive to losses than to gains of the same magnitude (Kahneman & Tversky, 1979). Despite the general importance of loss aversion in human behavior (Benartzi & Thaler, 1995; Genesove & Mayer, 2001; Pope & Schweitzer, 2011; Rick, 2011), there is relatively little research on loss aversion in explore-exploit decisions, as most tasks have concentrated on situations with strictly positive payoffs (CITES; c.f., Krueger, Wilson, & Cohen, 2017; Teodorescu & Erev, 2014). We examine the influence of loss aversion using a framing manipulation, assigning participants to observe an environment with all gains, or an objectively equivalent scenario that involves gains and losses. Since the two situations involve objectively identical payoffs, purely pecuniary considerations would drive subjects to behave similarly in the both situations. In fact, however, they make strikingly different explore-exploit decisions depending on how payoffs are framed.

In examining the influence of losses on exploration behavior, we draw a distinction between *anticipated* and *experienced* losses. Using a time series framework, anticipated losses are defined as negative payoffs that participants may experience from their next action (including exploring or exploiting). In contrast, experienced losses are defined as negative payoffs that participants received from the action they just took. We expect that participants will try to avoid losses regardless of whether they are anticipated or experienced, but this avoidance will manifest in different ways. If participants anticipate that exploration will lead to a loss, they will be less likely to explore. Conversely, if participants have just experienced a loss, they will be less likely to exploit that option again (see Table XX). Importantly, because participants in our experiments have many options to explore, there always remains some uncertainty about whether exploration is worthwhile.

Table XX. Predictions about the effects of losses on explore-exploit behavior.

Situation 1		Situation 2		Predicted difference between situations
Experienced payoff	Anticipated from exploration	Experienced payoff	Anticipated from exploration	
+	+	+	+/-	Less exploration in Situation 2
+	+/-	-	+/-	Less exploitation in Situation 2

Note. +/- denotes cases where there is uncertainty about whether payoffs will be positive or negative.

Existing research on explore-exploit decisions has used a variety of paradigms to examine the effects of losses. Perhaps the most popular is the "decisions from sampling" paradigm (e.g., Hertwig, Barron, Weber, & Erev, 2004; Rakow & Newell, 2010), in which decision makers are presented with two options and can sample repeatedly from either (i.e., explore) before making a consequential choice for pay. A review of experiments using this paradigm showed that in 12 of 15 studies, participants explored more extensively when payoffs were in the loss domain (Lejarraga, Hertwig, & Gonzales, 2012). However, these experiments did not use mixed payoffs, which means that participants could not avoid losses through exploration. Additionally, in this paradigm participants can costlessly gather information through exploration; it is unclear whether increased exploration in the presence of losses would occur in a situation where exploration is consequential.

Two other experiments address explore-exploit decisions using two-option paradigms. In Krueger, Wilson, and Cohen (2017), participants receive unequal amounts of information about two options at the beginning of the experiment (i.e., one payoff for one option and three payoffs for the other) before making a choice of their own. The results show that participants are more likely to choose the relatively unknown option (explore) when in a condition with only losses, versus a condition with only gains, perhaps because of increased uncertainty regarding the payoffs associated with the unknown option. In Yechiam, Zahavi, and Arditi (2015), participants choose between two unlabeled options. In the background, one option is set to be risky, with even probabilities of receiving 0 or 2, and one safe option receives 1; in in loss conditions, these values are negative. The results show that participants are more likely to explore the two options by switching between them (versus repeatedly draw from the same option) when in a task with losses. Thus, overall this literature appears to show that losses spur increased exploration, as measured by drawing from a relatively uncertain option (Krueger, Wilson, & Cohen, 2017), or switching between options (Yechiam, Zahavi, & Arditi, 2015). However, as with the sampling paradigm, participants have only two options to choose, leading to limited uncertainty about whether exploration is worthwhile.

Perhaps the most closely related work is a series of studies in which participants explored a grid environment containing either 120 or 144 options (Teodorescu & Erev, 2014). For this research, participants "explored" by choosing a new grid location or "exploited" by choosing a previously explored location. In one study (Study 2), participants were assigned to environments in which exploration always led to gains or another in which exploration always led to losses; exploitation always paid zero. Using this setup, the authors found that participants explored less when in an environment of losses, tending instead to exploit options that yielded a fixed, zero payoff. In additional work (Study 1), the authors used grids that contained both gains and losses, and assigned one of those outcomes to be more probable than the other (Teodorescu & Erev, 2014). The results showed that participants acted in accordance with the most probable outcome, exploring more when gains were prevalent and less when losses were prevalent. These results are consistent with an account in which participants developed expectations about the payoffs associated with future choices (in our parlance, anticipated losses), and reduced exploration accordingly. Thus, this paper provides evidence that exploration decisions are affected by anticipated losses.

Adding Ecological Validity: Multiple Options and Retreat Behavior

We study explore-exploit behavior using a novel computer task that asks participants to explore a large environment (as suggested by Rakow & Newell, 2010). To date, only one paper uses an environment with more than two outcomes to study exploration and losses; the authors note that “models that best predict binary decisions in the choice prediction competitions do not provide good predictions of behavior in the current multi-alternative setting” (Teodorescu & Erev, 2014, p. 1020). As such, we help contribute to the literature by studying the effects of losses on exploration behavior using a large environment with more than two options.

The large environment also allows us to examine how explore-exploit decisions vary when participant behavior is constrained. Specifically, in our setup participants who want to return to previously-explored locations must “retreat” and retrace their steps, rather than being able to choose any option at any time. Prior research often allows participants to move in an unconstrained fashion (e.g., Teodorescu & Erev, 2014; CITE), but we believe that many real-life situations are better approximated by our task. For instance, a business that reverses course – switching employees from working on existing products to new R&D and back again – may find that there are switching costs from employees who require retraining on technical subjects, rather than immediate payoffs from each project. Additionally, and perhaps most literally, physical exploration imposes costs for turning back. Westward-bound pioneers in the early history of the

United States might spend weeks or months returning to the east coast, rather than instantly teleporting home. While large environments and switching costs have been discussed in the explore-exploit literature (Mehlhorn., et al., 2015), to our knowledge, no other research has implemented a task that increases ecological validity in these ways.

Overview of the Current Research

We use a novel research task to study two research questions about loss aversion and exploration behavior. We ask (a) what are the effects of anticipated losses on exploration behavior, and (b) what are the effects of experienced losses on exploration behavior? In response to the first question, we predicted that participants who anticipate losses from exploration will be deterred from exploring, consistent with behavior displayed in Teodorescu and Erev (2014). In response to the second question, we predicted that participants who have just experienced a loss will move to a different location to avoid repeated exploitation, using either further exploration or retreat behavior. As noted, to our knowledge no previous research has implemented retreat behavior, so this prediction has not been directly examined. However, it is consistent with experiments showing that participants receiving negative payoffs are more likely to switch their sampling (Yechiam, Zahavi, & Arditi, 2015). Together, our predictions reflect the view that exploration is driven by hope of positive anticipated payoffs, constrained by fear of negative anticipated payoffs, and motivated by avoidance of experienced negative payoffs.

Studies 1a and 1b

Our first two studies were designed primarily to test the prediction that anticipated losses would cause participants to exhibit more conservative, exploitative behavior. We created a novel laboratory task, the "Grain Game," in which participants explored a one-dimensional environment and earned payoffs that depended on their location. To study the welfare effects of exploration, we varied whether long-run exploration was harmful (Study 1a) or beneficial (Study 1b). For both studies, all participants began on the left edge of the environment and did not observe the environment prior to playing. We describe below the experimental task used, with slight modifications, throughout our experiments.

Study 1a

Method

Participants. This study and all subsequent studies in this paper were approved by the Institutional Review Board at Carnegie Mellon University. We recruited 140 U.S. residents online using Mechanical Turk, based on a power analysis suggesting that we needed at least 48 participants per condition (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 57.9% were male.

Experimental task. In "The Grain Game," participants played the role of a farmer who chose where to plant crops (Figure 1).[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 [icon]." They then received instructions on how to move through the game using

[1] Due to programming error, we failed to record the last choice made by participants in Studies 1a and 1b. The first location was automatically assigned, so analyses are based on the remaining 68 choices. Study 2 uses 69 choices.

the keyboard, and that they would receive payment based on their performance (see Supplementary Materials for screenshots of these instructions with complete wording).

To ensure they understood the game and would know the range of payoffs to expect, participants 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 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.

During the game, payoffs in each location were determined by a fixed value that differed across locations and noise of up to two points. For the majority of locations, there was 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. For locations with expected values near the endpoints of the range in each condition, probabilities were collapsed. For instance, for a 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.

On each of 70 turns, participants chose to plant in one of 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 defined “exploration” as moving to a previously unexplored location and “retreat” as returning to a previously visited location. Notably, this restrained movement differs from previously reviewed work in that participants had to retrace their steps to return to earlier locations. As shown in Figure 1, recent payoffs were displayed on the screen and unexplored locations were indicated with a question mark icon. At the end of the study session, participants received payment based on accumulated points (4 points = \$0.01).[2]

Figure 1. Screenshot of the Grain Game.

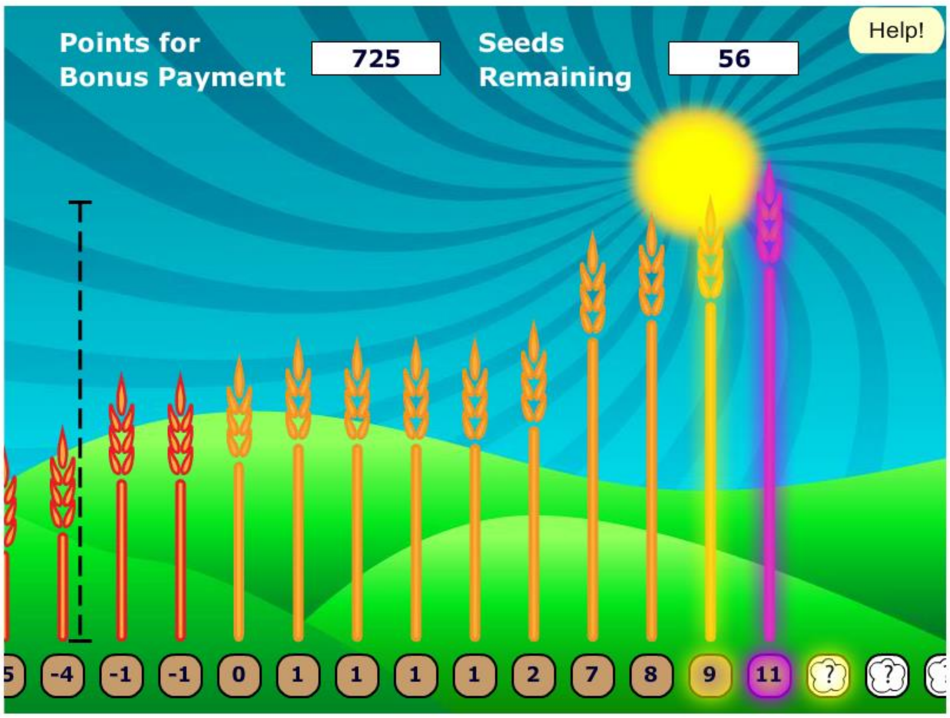

Note: This figure shows the gain-loss condition of the Grain Game. The top of the screen shows accumulated points and remaining turns. The black dashed bar indicates the payoff range. The most recent payoffs received from explored locations are given in boxes below the grain and unexplored positions are designated by question marks. The participant's current position is highlighted in purple; this turn, they can exploit that position, retreat to the left, or explore to the right. The Help button displays instructions. See Supplementary Materials for image of gain-only condition.

Gain-only and gain-loss conditions. To assess the effects of loss aversion, we assigned participants to one of two conditions that differed in the framing of payoffs. Participants in the

[2] 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.

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 an extra 700 points at the beginning of the task. Payoffs were represented graphically by the height of the grain displayed on the screen (Figure 1).

Environment. We created a version of the Grain Game for Study 1a in which long-run exploration was not worthwhile because the highest payoff was available after only eight moves (Figure 2; dotted line labeled "Environment").[3] Gain-loss participants received negative payoffs starting at position 15, after which the highest expected payoff was one point.

Figure 2. Extent of exploration and environment used in Study 1a.

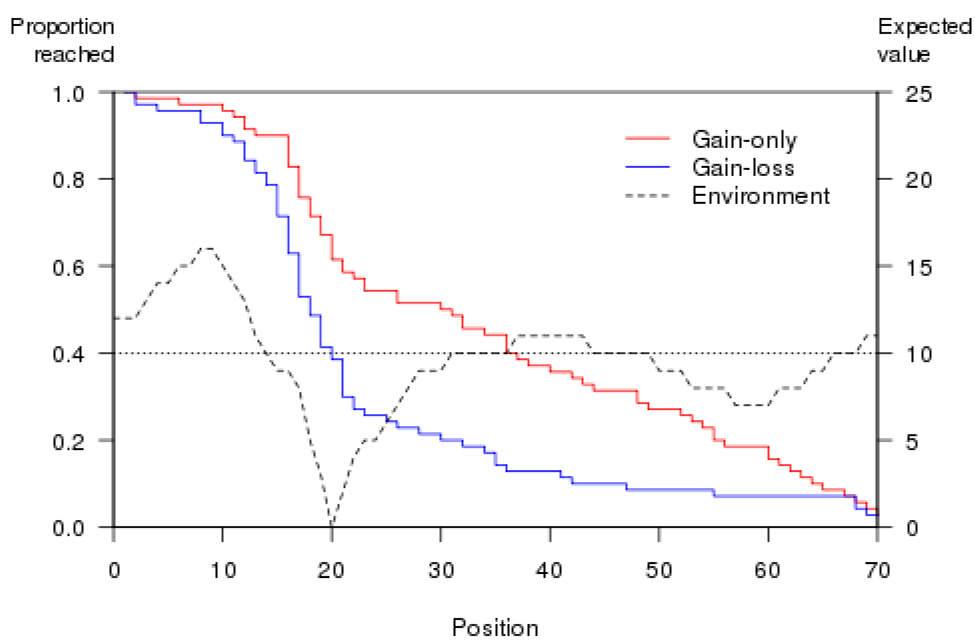

Note. This chart shows the proportion of participants in each condition who reached each position and the underlying environment. The highest adjusted payoff (16 points) is obtained at positions 8 and 9. Participants in the gain-loss condition are expected to first experience losses at position 15, below the dotted line.

Results

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). 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* = 19.7) and the average gain-loss participant explored 21.3 times (*SD* = 16.0; $t(138) = 3.7$, $p < .001$).

Explore-Exploit-Retreat Behavior. To study the effects of anticipated and experienced losses, we analyzed the actions participants took each turn. To do so, we ran the following multinomial logit regression:

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

Specifically, 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 (who knew that negative payoffs were possible) and zero for gain-only participants. *Adjusted payoff* represents the expected 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 a threshold of 10, and would therefore be negative in the gain-loss condition. Our predictions were that (a) the coefficient on β_1 would be negative for exploration, representing a reluctance to explore in the gain-loss condition due to anticipated losses, and (b) β_4 would be negative for exploitation, suggesting a reluctance to exploit after experiencing a loss.

The multinomial logit results are shown in Table 1 with exploitation as the omitted action. As predicted, given our expectations about anticipated 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 44.3 percent less likely to explore than to exploit each turn. They were also 11.5 percent less likely to retreat. Contrary to our expectations, there was no difference in behavior between conditions when participants experienced losses (i.e., coefficient on *gain-loss x below threshold*).

The coefficient on *adjusted payoff* indicates that every one point increase in expected payoffs was associated with a 13.8 percent decrease in exploration regardless of condition, suggesting that participants were more likely to repeatedly exploit high value locations. There was no change in behavior associated with receiving a payoff below the threshold.

Payoffs. The environment in this study was designed to be unfavorable to 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 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, losses discouraged (ex-post) suboptimal exploration and hence made decision-makers better off.

Table 1. Multinomial logit regression results predicting participant actions.

	Study 1a		Study 1b		Study 2	
	Explore	Retreat	Explore	Retreat	Explore	Retreat
Gain-Loss condition	0.557*** [0.500, 0.621]	0.884* [0.786, 0.995]	0.535*** [0.473, 0.606]	0.443*** [0.384, 0.512]	0.849*** [0.814, 0.886]	1.115*** [1.066, 1.165]
Adjusted payoff	0.862***	1.018	0.958***	0.934***	0.898***	0.956***

[t-1]	[0.844, 0.880]	[0.993, 1.044]	[0.945, 0.972]	[0.919, 0.950]	[0.894, 0.903]	[0.951, 0.960]
Below threshold [t-1]	0.916 [0.744, 1.128]	1.219 [0.937, 1.587]	1.720*** [1.388, 2.131]	0.904 [0.707, 1.157]	0.937 [0.887, 1.001]	0.997 [0.906, 1.052]
Gain-Loss x Below threshold [t-1]	0.797 [0.633, 1.003]	0.903 [0.676, 1.206]	0.855 [0.679, 1.078]	1.993*** [1.540, 2.579]	1.772*** [1.641, 1.915]	1.301*** [1.192, 1.420]
Intercept	7.554*** [5.657, 10.088]	0.451 [0.317, 0.642]	1.862*** [1.496, 2.317]	1.832*** [1.426, 2.355]	3.906*** [3.609, 4.227]	1.126** [1.035, 1.225]
Log Likelihood	-9734		-8370		-77,026	
N (actions)	9520		8024		73,692	
N (participants)	140		118		1,068	

Note. The table presents exponentiated coefficients and 95% confidence intervals. 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. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Study 1b

This study used the same design as Study 1a, but we presented participants with an environment that rewarded instead of punished exploration.

Method

Participants. We recruited 118 participants via Amazon Mechanical Turk, based on a power analysis suggesting that we needed at least 43 participants per condition (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.

Stimuli. The payoff function, in which exploration was rewarded, is shown in Figure 3. Participants in the gain-loss condition received negative payoffs from position 10 to 26. All participants who persisted past these low payoffs received the maximum possible payoff at position 39.

Figure 3. Extent of exploration and environment used in Study 1b.

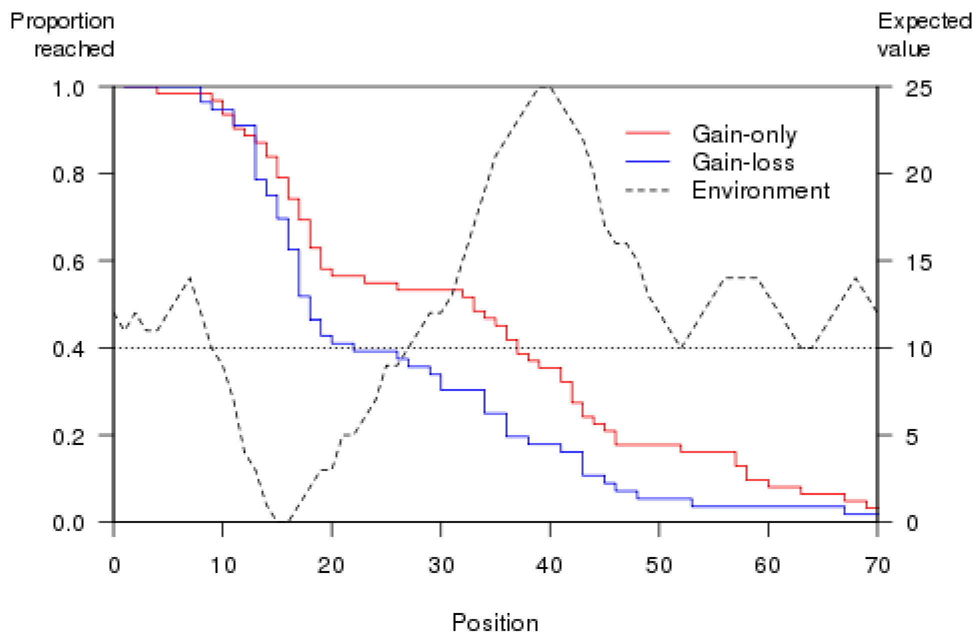

Note. This chart shows the proportion of participants in each condition who reached each position and the underlying environment. The highest adjusted payoff (25 points) is obtained at position 39. Participants in the gain-loss condition are expected to first experience losses at position 9.

Results

Overall exploration. Following our analysis from Study 1a, we first present 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 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 ($t(116) = 2.34$, $p = .02$).

Explore-Exploit-Retreat Behavior. We analyzed decisions to explore, exploit, and retreat using a multinomial logit regression as in Study 1a. The results show that, consistent with expectations about anticipated losses, participants in the gain-loss condition were less likely to explore than those in the gain-only condition (Table 1). These participants were also less likely to retreat. Consistent with our expectations about experienced losses, those in the gain-loss condition were nearly twice as likely to *retreat* (relative to exploiting) as those in the gain-only condition after receiving a payoff below the threshold (i.e., coefficient on *gain-loss x below threshold*). In other words, experiencing losses led participants to retreat rather than encouraging them to explore further.

In addition to these effects of interest, 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 an expected payoff below the threshold was associated with a higher probability of exploring, but the magnitude of this effect did not differ for those in the gain-only and gain-loss conditions.

Payoffs. The environment 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, 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), who explored less.

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 they anticipated receiving negative payoffs. In both Studies 1a and 1b, we found support for this hypothesis, as those in the gain-loss condition were less likely to explore than those who were in an environment without negative values. Second, we hypothesized that people would be less likely to exploit after experiencing a loss. Consistent with this hypothesis, negative payoffs were associated with less exploitation (and increased retreat behavior) in Study 1b; however, this pattern was not statistically significant in Study 1a. Given this inconsistency, it is possible that the effects of experienced losses are more sensitive to the specific environment than those of anticipated losses. To address this possibility, we created a large landscape for participants to explore in Study 2 that would introduce additional variation into the environmental context.

Study 2

Method

Participants. We recruited 1,068 U.S. residents online using Mechanical Turk. The average age was 31 years (*SD* = 10.0) and 61% were male. We recruited a large sample because we anticipated an increase in variation given the large environment that we were using.

Stimuli. We created a version of the Grain Game using an environment containing 700 locations. To increase ecological validity, the contours of the underlying environment were based off of hiking elevation maps from hikes in Virginia (www.hikingupward.com). 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 fix participant expectations, all participants viewed a graph of potential payoffs before starting the game (Figure 4) and were told they would start in a random position.

The graph initially showed a subset of payoffs, but participants could pan to the right or left using prominent arrows to see the entire environment.[4]

Figure 4. Preview of payoffs, Study 2.

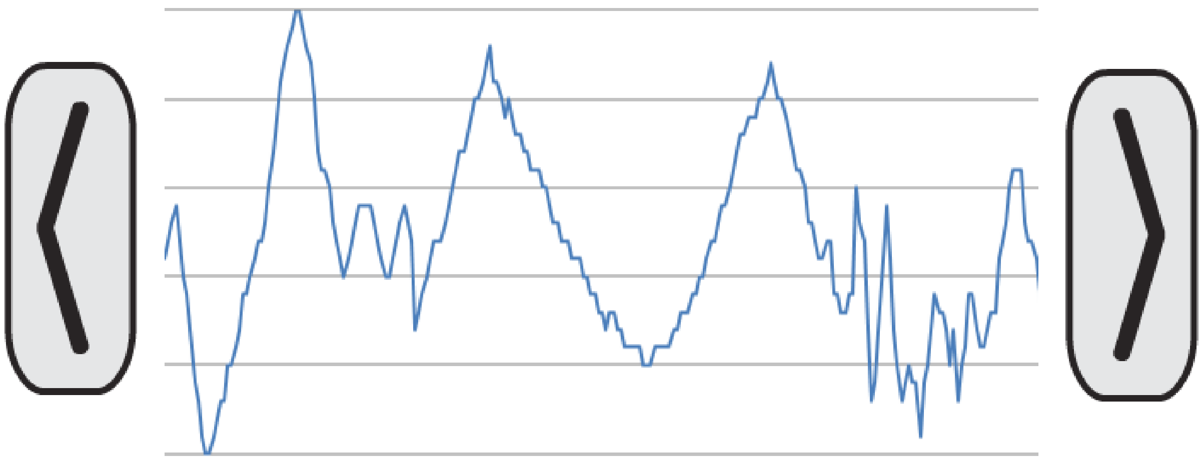

Note. All participants viewed this graph prior to starting the study. The arrows on the left and right could be used to view the entire environment.

Results

Overall exploration. To examine patterns in exploration behavior, we calculated the probability of exploration for each condition and payoff combination, using adjusted payoffs of [0, 25] by adding 10 to every gain-loss payoff (Figure 5). The negative slope in the left panel demonstrates that in both conditions, participants were less likely to explore when receiving higher payoffs. More importantly, a comparison of the two conditions shows an interaction between the condition and exploration behavior. Gain-loss participants were less likely to explore when payoffs were just above the threshold (to the right of the red dotted line), in a situation where they could anticipate, but were not receiving, negative payoffs. In contrast, gain-

[4] We asked the same exploratory questions as in Studies 1a and 1b. Additionally, we asked participants whether they copied the graph shown at the beginning of the study. Again, we focus on experimental effects rather than these questions.

loss participants were more likely to explore than gain-only participants when receiving negative payoffs (to the left of the red dotted line), and were experiencing losses.

Another way to visualize this interaction is by comparing the proportion of exploration decisions across the both 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. This panel again shows that participants in the gain-loss condition are more likely to explore when receiving negative payoffs, and less likely to explore when above the threshold.

Figure 5. Probability of exploration by condition, Study 2.

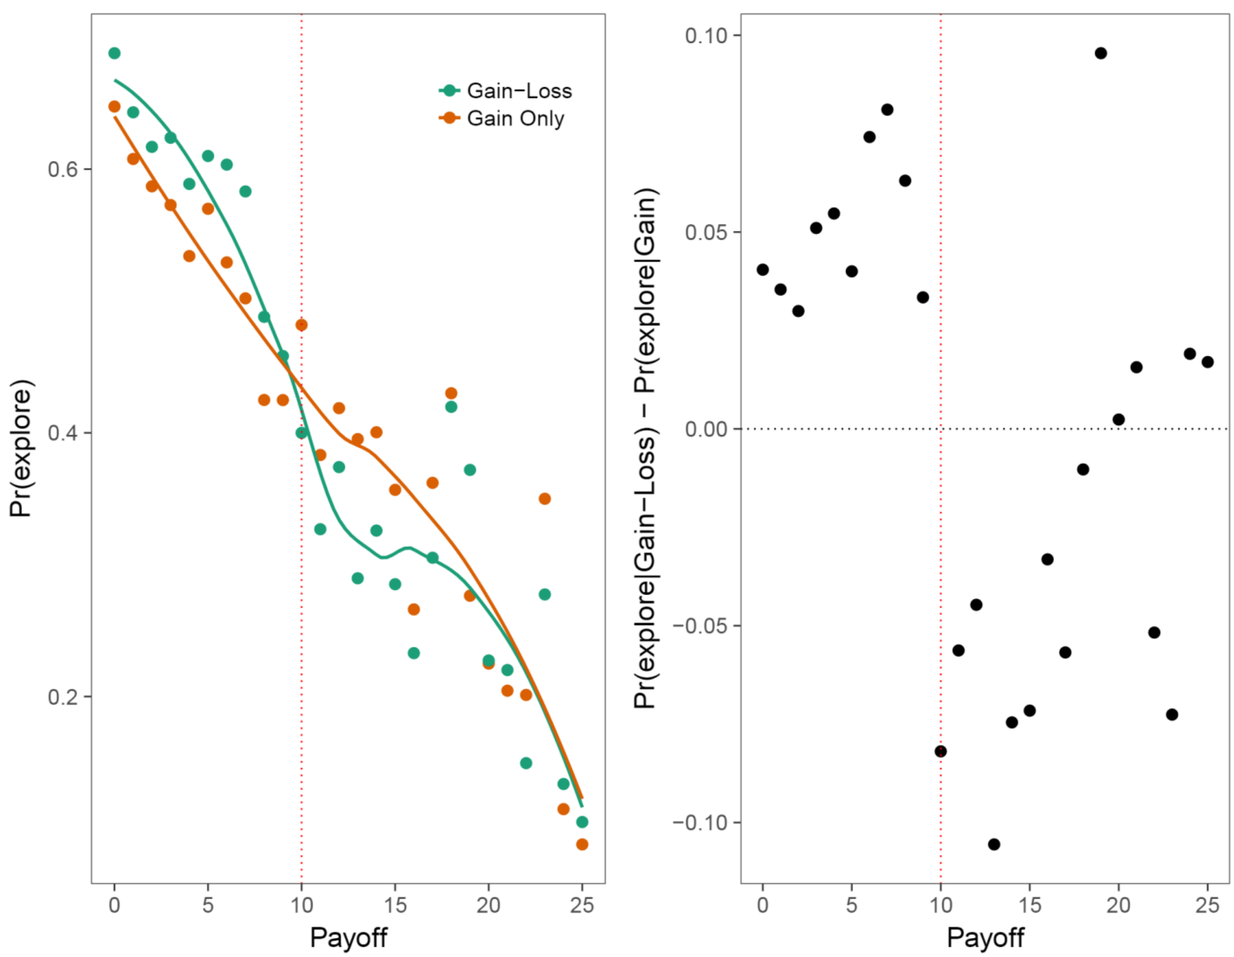

Note. The left panel shows the probability of exploring by condition. The lines show the best fitting smoothed spline. The right panel shows the difference in probability of exploration between the two conditions, with values above zero on the y-axis representing higher exploration rates for gain-loss participants. In both panels, payoffs are rescaled to the gain-only range of [0, 25]. Below the threshold (i.e., to the left of the red dotted lines), participants in the gain-loss condition incur losses.

Explore-Exploit-Retreat Behavior. As in previous studies, we ran a multinomial logit on each participant action. Consistent with Studies 1a and 1b, and with our prediction about the effects of anticipated losses, those in the gain-loss condition were less likely to explore than those in the gain-only condition (Table 1). Additionally, consistent with our prediction about the effects of experienced losses, the interaction between *gain-loss condition* and *below threshold* shows that choices after 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, participants were 10 percent less likely to explore and 4 percent less likely to retreat. Finally, there was no significant main effect of receiving a payoff below the threshold.

Payoffs. There was no difference in average payoffs between the two conditions (p = .41). Unlike in our previous study, the environments were not designed to be favorable or unfavorable to exploration. Instead, payoffs depended on the initial starting location.

General Discussion

In this research, we examine the effects of loss aversion on explore-exploit-retreat decisions, asking how anticipated and experienced losses affect these decisions. Overall, we find that in situations with many potential outcomes, where there is uncertainty about the payoffs associated with exploration, *anticipated* losses lead to less exploration (Studies 1a, 1b, and 2) and *experienced* losses lead to less exploitation (Studies 1b and 2). Distinguishing between anticipated and experienced losses becomes possible when using a novel computer task that features a large environment (suggested by Rakow & Newell, 2010), because participants who experience losses can maintain hope that as yet unexplored options will yield positive payoffs. In contrast, paradigms that require many sampling decisions over two options may not be able to find such patterns, as there may be little uncertainty about the payoffs associated with each option. Together, our two primary findings suggest that people try to avoid losses through both exploration and exploitation; as such, this work is consistent with literature on loss aversion that states that negative payoffs are particularly painful (Kahneman & Tversky, 1979), and empirical research showing that participants explore less in environments with negative payoffs (Teodorescu & Erev, 2014).

Our choice to use an exploration task with many options also allows us to constrain participants' behavior, requiring them to retrace their steps to return to previously explored locations, and giving rise to a new classification of actions as explore, exploit, or retreat behavior. Prior literature has interpreted 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 considered risky and exploiting considered safe. In a risk-based framework, "retreat" behavior may also be viewed as a safe strategy, since it entails returning to a known location. It is therefore interesting to see that in our work, both exploring and retreating (i.e., "risky" and "safe" behaviors) arise from experiencing losses.

A natural question arising from our work is whether loss aversion is adaptive for explore-exploit decisions. Our first two studies serve as existence proofs showing that loss aversion can be helpful or harmful depending on whether the environment rewards exploration. However, because these two studies featured specific environments, our work falls short of a generalized claim about the benefits of loss aversion for exploration. Future research could attempt to create measures of environmental favorability, to measure when exploration would be beneficial, and to see how such environments vary across different contexts or populations. Even without such measures, our work contributes to an expanding literature documenting situations in which biases can 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 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). Consistent with these papers, our Study 1a, which used an environment unfavorable to exploration, found that loss aversion can help people receive higher payoffs.

In the end, our results suggest that losses have a systematic effect on explore-exploit behavior. Experiencing a loss may inspire change, while small gains may induce complacency and a desire to exploit the known option. In settings in which someone may want to induce risk-taking and exploration (e.g., a private equity fund investing in a firm with an established product), artificially creating losses may motivate new initiatives. As Johnson (2011) 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. It may be fruitful to wonder whether other seemingly unfavorable conditions can also have beneficial side effects on decisions.

References

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. doi: 10.1098/rstb.2007.2098

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. doi: 10.1162/003355301753265561

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

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

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. doi: 10.1287/mnsc.39.1.17

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

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108,* 480-498. doi: 10.1037//0033-2909.108.3.480

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. doi: 10.1016/j.cogsys.2010.07.007

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

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. doi: 10.1037/dec0000033

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): 129-57.

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. doi: 10.1016/j.jcps.2010.04.004

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. doi: 10.1037/0033-2909.103.2.193

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. 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. doi: 10.1037/a0038199

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. doi: 10.3758/s13423-014-0749-4

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. doi: 10.1037/a0029383

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. doi: 10.1287/mksc.22.4.503.24909

Supplementary Materials

Instructions

These are selected screenshots from the instructions of the Grain Game.

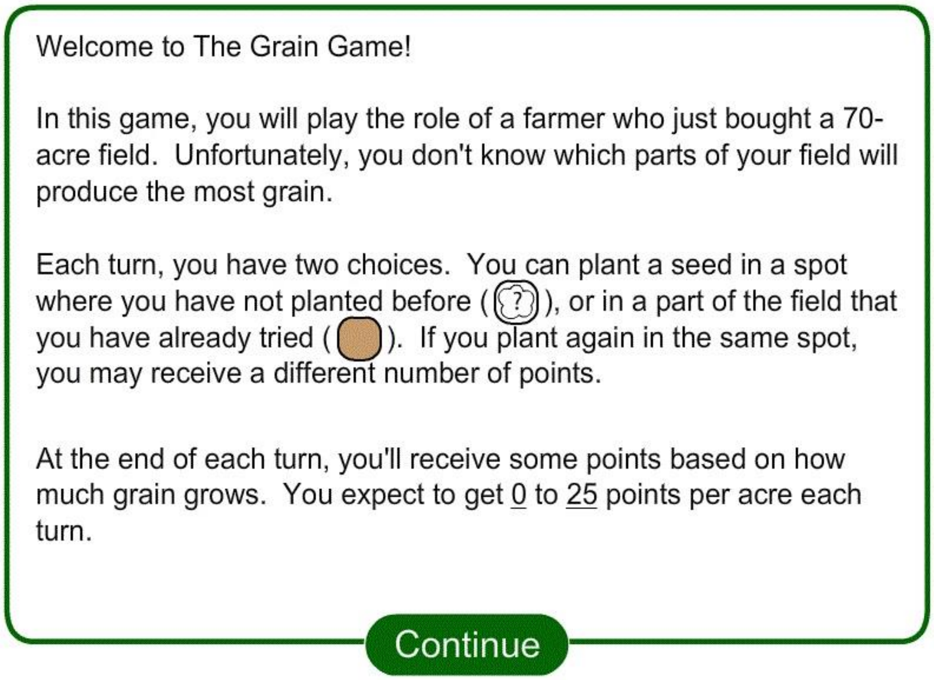

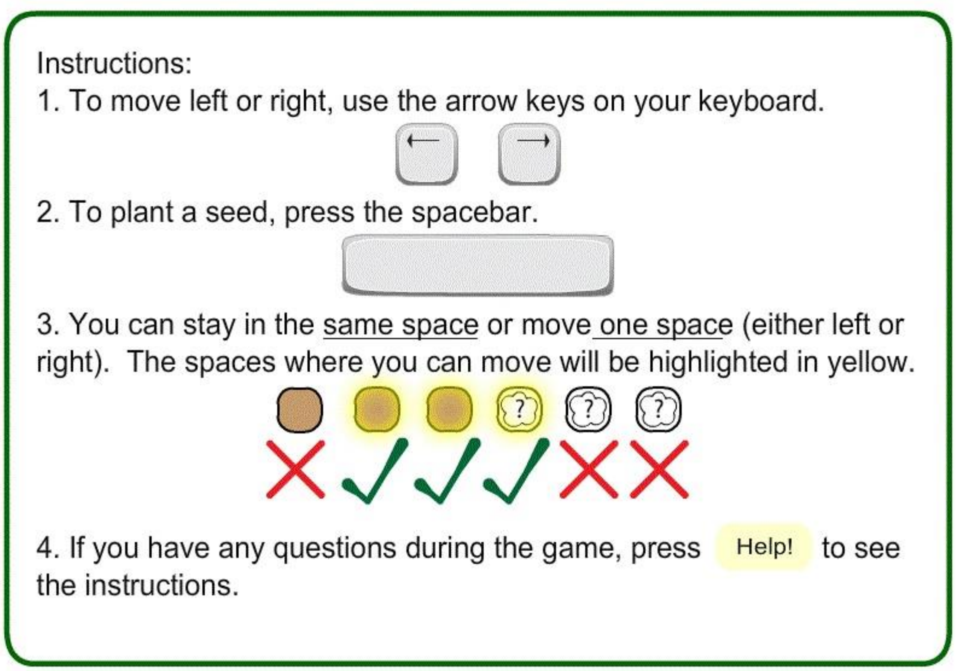

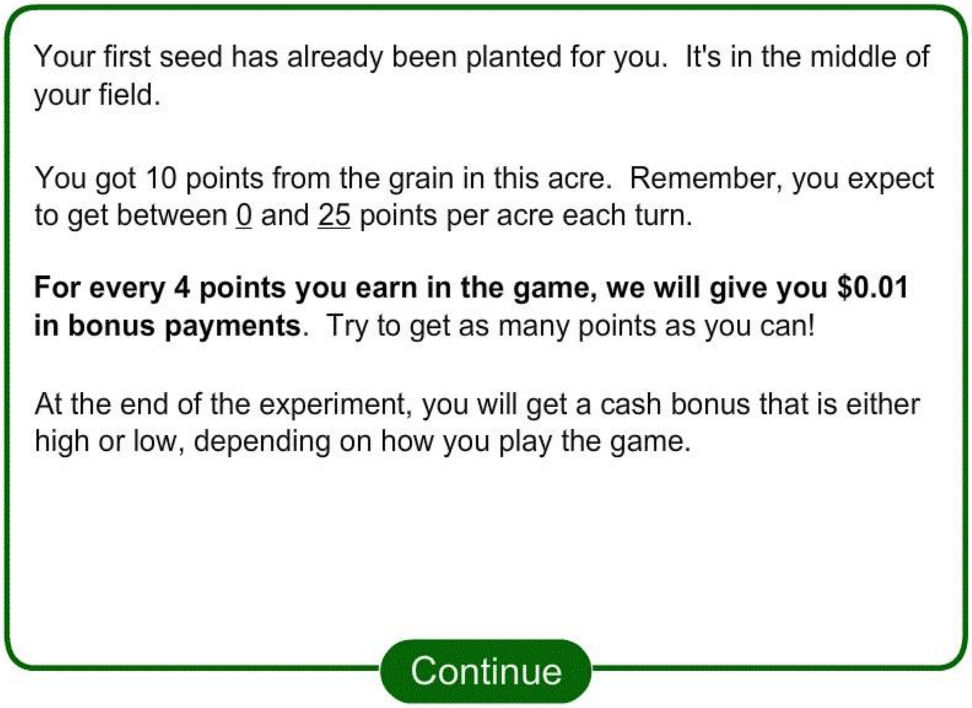

Screenshot of Gain-only condition.