

Dynamic Risk Preferences Under Realized and Paper Outcomes

Kirby Nielsen^a

March 18, 2019

Abstract

We conduct a large-scale test of dynamic risk preferences. While most of the literature on dynamic risk restricts attention to positively skewed gambles, subjects in our experiment tend to choose negatively skewed lotteries. This allows us to study dynamic risk preferences in an environment that the literature has not analyzed. We find evidence of the reinforcement effect—individuals take on more risk after a gain and take on less risk after a loss. Furthermore, we exogenously vary whether these outcomes are “realized” or on paper, according to the distinction put forth by Imas (2016). We find little difference in the responses to realized and paper outcomes in environments of negatively skewed risk.

KEYWORDS: dynamic risk; reinforcement effect; history-dependent preferences; prospect theory; realization effect

JEL CODES: C91, D81

I. INTRODUCTION

Many economically-relevant situations involve sequential decisions under uncertainty. As a result, it is important to understand how early gains or losses affect subsequent choices. Losses may induce more cautious behavior and *reductions* in risk taking, known as the “reinforcement effect” for losses. On the other hand, individuals may exhibit “loss chasing,” or *increased* risk seeking with the hope of recovering the early loss. Empirical evidence is mixed—some papers have found risk to increase after a loss (Andrade and Iyer, 2009; Langer and Weber, 2008; Verbruggen et al., 2017; Schneider et al., 2016; Clark, 2010; Smith et al., 2009) while others found risk reduction (Shiv et al., 2006; Liu et al., 2010). Theoretical predictions are mixed, as well. Models in Barberis et al. (2001), Dillenberger and Rozen (2014), and Tserenjigmid (2017) predict that individuals take on less risk after a loss, while those in Shefrin and Statman (1985) and Weber and Camerer (1998) predict the opposite.

In this paper, we conduct a large-scale test of individuals’ responses to prior gains and losses. We analyze behavior from nearly 2,000 participants from Amazon’s Mechanical Turk platform. We use a variation of the

^awww.kirbynielsen.com. Department of Economics, Stanford University, 579 Serra Mall, Stanford, California 94305, USA. I thank Puja Bhattacharya, Owen Clady, Lucas Coffman, Alex Imas, and especially Paul J. Healy for helpful comments and conversations. I also thank Itay Sisso for help with his app to flag suspicious survey responses. I’m thankful to the editor and three anonymous referees for their role in greatly improving the paper. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

bomb risk elicitation task (BRET) (Crosetto and Filippin, 2013) to analyze the difference in changes in risk behavior as a response to earlier outcomes. In this task, subjects choose to open a number of boxes, one of which contains a “bomb.” Opening more boxes gives a higher potential reward if the boxes don’t contain the bomb, but increases the probability of selecting the bomb. We describe the BRET in detail in Section II. The BRET is an easy and intuitive way to measure risk behavior.

The BRET also allows for testing changes in risk preferences in a relatively understudied environment. Previous experiments on dynamic risk preferences have focused primarily on lotteries with fixed gain/loss probabilities, using positively skewed or 50-50 gambles (Andrade and Iyer, 2009; Langer and Weber, 2008; Shiv et al., 2006).¹ In contrast, the BRET allows for choosing both positively and negatively skewed gambles, its gambles span a large range of expected value, and the probability of a loss increases as an individual takes riskier decisions. Furthermore, in our environment, lotteries with the highest expected value are those with *negative* skew. As a result, most individuals choose negatively skewed gambles from the menu of available lotteries. To the best of our knowledge, ours is the first paper to study dynamic risk-taking in an environment of negatively skewed gambles. While the literature has documented loss chasing under positively skewed risk, confirming various theoretical predictions, predictions under negatively skewed risk can go in the opposite direction.

In addition, we vary exogenously whether these previous gains or losses are “realized.” Building on earlier observations by Shefrin and Statman (1985) and the model of Barberis and Xiong (2012), Imas (2016) proposes a theory in which prior losses decrease risk taking when those losses are actually realized, but increase risk taking when the losses are only on paper.² He refers to this as the *realization effect*. Imas demonstrates the realization effect through a series of experiments involving risky investment decisions with both realized and paper losses. The realization effect theory applies directly to loss averse individuals choosing whether to accept or reject a positively skewed gamble. In this paper, we ask whether we see a difference in risk-taking behavior following realized versus paper gains and losses in environments where individuals largely choose *negatively* skewed gambles. This is outside the domain of the realization effect theory, and the realization effect has been shown only in environments with positively skewed risk, so we test whether the difference between realized and paper outcomes holds under negatively skewed risk as well.

Overall, we find strong evidence for the “reinforcement effect,” where individuals increase risk taking following a gain and decrease risk taking following a loss. We don’t find any interaction of this effect with how losses are realized—individuals respond to realized and paper outcomes similarly in our experiment. However,

¹Skewness is defined formally in Section II. Intuitively, positively skewed gambles are “long shot” gambles with a low probability of winning a large amount.

²For example, a day of betting at the horse track constitutes a series of paper gains/losses until the individual settles his account at the end of the day and “realizes” his final earnings or losings.

the realization effect theory applies directly to individuals choosing over positively skewed gambles, while individuals in our experiment mostly choose negatively skewed gambles. One cannot apply predictions from Imas (2016) directly to our environment, as the types of risk under consideration differ across the two. We conduct an individual-level parameter estimation to fit the data to the realization effect model. We find that a model which assumes realized losses fits the experimental data better than one with paper losses, regardless of the experimental manipulation of realization. Overall, our results give evidence confirming history-dependent risk preferences with both types of outcomes. In an environment where most individuals are choosing negatively skewed lotteries, we expect to see the reinforcement effect equally under both realized *and* paper outcomes.

II. EXPERIMENTAL DESIGN

We conduct the experiment online using a survey programmed in Qualtrics and subjects from the general population on Amazon’s Mechanical Turk, an online labor platform. Subjects are required to be over 18 years of age, residents of the United States, and are screened for basic Mechanical Turk competency.³ In total, we analyze data from 1,953 survey responses.⁴ The survey takes about 8 minutes to complete and average total payments are just over \$0.50 per subject including a \$0.10 show-up fee, which is very typical for the Amazon Mechanical Turk workforce.

We adapt the BRET to test the effect of previous gains and losses on subsequent risk taking. We endow subjects with \$0.40 and present them with a 5x5 grid of boxes. All but one of the 25 boxes are empty, while one randomly-selected box contains a “bomb.” Subjects choose how many boxes to open with the implicit goal of not opening the box containing the bomb. If a subject opens the box that contains the bomb, he loses \$0.10 for that round.⁵ If he opens only empty boxes in a round, he gains \$ c per box opened. In the “low stakes” condition ($n = 829$), subjects gain \$0.01 per box opened ($c = 0.01$). In the “high stakes” condition ($n = 1, 124$), we double the reward to \$0.02 ($c = 0.02$). To mitigate any effects of confusion and learning, on every decision screen we inform subjects of the probability of selecting a bomb for every possible number of boxes opened, along with how much they would earn if they do not select the bomb and how much they would lose if they do select the bomb. Subjects are also required to answer correctly a number of comprehension questions before starting the experiment.

³We excluded participants who had never completed a task on Mechanical Turk, and we restricted our sample to participants who had completed fewer than 1,000 tasks. We also excluded participants with low approval ratings from previous Mechanical Turk tasks.

⁴A drawback of online data collection is that we cannot perfectly verify survey responses. “Spam” accounts and user-written programs to automatically detect and submit surveys introduce fake data. We filter data using an app which flags suspicious IP Address and/or GPS Locations (Prims et al., 2018). One batch of MTurk surveys was sent out fully anonymous and therefore does not have IP/GPS data, so we cannot filter that portion of the sample. We also dropped responses with duration below 100 seconds, leaving us with a total sample of 1,953 out of the original 2,352 responses.

⁵In the original formulation of the BRET, subjects did not start with an endowment and selecting the bomb “wipe(d) out the subjects’ earnings,” but does not result in an actual loss (Crosetto and Filippin, 2013). We modify this in order to study changes in risk following actual losses, rather than losses based on the reference point of the potential gain.

Subjects select all boxes they wish to open and then learn which box contains the bomb and whether they had selected it.⁶ After revealing the bomb, we inform subjects of their earnings from that round and then they proceed to the next decision screen. After a “practice round” to gain familiarity with the interface, all subjects participate in a total of four rounds. As such, a participant can lose his entire endowment if he clicks the bomb in every round, and can gain a maximum of \$24c per round otherwise. This four round structure follows the experiments in Imas (2016), aside from the fact that we use the BRET while Imas uses a risky investment decision.

We study the effect of realized versus paper gains and losses using a between-subject experimental manipulation. Participants are randomly assigned to either the *Realization* or *Paper* treatment. In the instructions of the *Realization* treatment, we tell subjects that their earnings will be finalized after Round 3. That is, we will transfer their winnings to their account or subtract losses away from their starting balance. At the time of realization, subjects are notified of their current earnings or losses and are asked to type “Close Account” to make the transfer and finalize their balance. After this, they proceed to Round 4 to make choices in the final round in the same way described above. In the *Paper* treatment, we eliminate this account-closing stage but subjects are still notified of their current earnings before proceeding to Round 4.⁷ Rounds 1–3 are identical across treatments. A full copy of instructions and screenshots can be found in the Appendix.

II.I. Skewness

Denote a lottery, L , which gives x^i with probability p^i by $L = (x^1, p^1; x^2, p^2; \dots; x^n, p^n)$. Given the structure of the BRET, a subject chooses his most preferred lottery from a set of 26 available lotteries, one for each number of boxes he could open.⁸ These include negatively skewed and positively skewed alternatives.

We define skewness as a lottery’s third central moment,

$$\mu_3 = \sum_i^n p^i (x^i - \mathbb{E}x)^3.$$

Under this definition, in both the high- and low-stakes conditions, opening 12 or fewer boxes constitutes choosing a negatively skewed lottery and opening more than 12 boxes constitutes choosing a positively skewed lottery.⁹

The available lotteries differ also in expected value. In the low stakes condition, expected earnings under

⁶Note, we do not use a sequential version of the BRET where boxes are “opened” upon selection, more akin to the “balloon task” (Lejuez et al., 2002). We did not want subjects updating their plans within a round, and we wanted to avoid the truncated observations that would result from selecting the bomb early on.

⁷Our implementation of “realization” mirrors the Robustness Check of Study 2 in Imas (2016).

⁸For example, in the Low Stakes BRET, the available lotteries are $\{(\$0, 1), (\$0.01, 0.96; -\$0.10, 0.04), (\$0.02, 0.92; -\$0.10, 0.08), \dots, (\$0.24, 0.04; -\$0.10, 0.96), (-\$0.10, 1)\}$

⁹Since we do not have a 50-50 gamble, the only “unskewed” choice is to open 0 or 25 boxes.

risk neutrality are maximized by opening 7 boxes.¹⁰ On the other hand, expected earnings in the high stakes BRET are maximized by opening 10 boxes. See the Appendix for a full list of probabilities and potential earnings.¹¹ In addition, Figures II and III in the Appendix plot skew and expected value as a function of number of boxes opened in both the low and high stakes conditions.

We run both low and high stakes conditions to see a larger range of lottery choices, spanning a range of “riskiness” and skew. In our environment, expected earnings are maximized in choosing negatively skewed lotteries. As a result, we expect to see most individuals choosing negatively skewed gambles. As discussed, the literature on dynamic risk preferences has focused on positively skewed gambles, and evidence on both loss chasing and the reinforcement effect is largely concentrated in this environment. Our experimental structure allows us to study dynamic risk preferences in a new environment. Furthermore, this is outside the domain of the theory in Imas (2016), and we test whether the difference between realized and paper outcomes holds under negatively skewed choices.

III. THEORETICAL PREDICTIONS

Expected utility theory would predict individuals’ risk taking behavior not change significantly on the heels of a gain *or* a loss, absent a drastic change in wealth. However, a number of recent behavioral theories describe history-dependent risk preferences which would predict changes in risk taking. In addition, the realization effect theory (Imas, 2016) predicts differences in behavior depending on whether these losses or gains are realized. We adapt these predictions to our experimental environment.

III.I. History-Dependent Risk Preferences

The “reinforcement effect” describes the phenomenon where a decision maker’s risk behavior reinforces his prior history—he becomes more risk seeking after a good history and becomes less risk seeking after a bad history. The first evidence of the reinforcement effect is often attributed to Thaler and Johnson (1990), who documented the reinforcing behavior using simple lotteries in a laboratory experiment. Various theoretical models capture this changing risk behavior, either directly or indirectly.

Barberis et al. (2001) study an environment where investors receive utility directly from fluctuations in their financial wealth. In general, investors are loss averse over wealth fluctuations, but the degree of loss

¹⁰Expected earnings are the same opening 7 or 8 boxes.

¹¹Mechanical Turk poses a challenge for studying losses under positively skewed risk in a task like the BRET. Calibrating parameters so that the expected earnings are maximized in opening many boxes requires c to be very large. This is atypical for payments on Mechanical Turk, so subjects tend to open very few boxes to secure a large payment, a type of “satisficing.” As a result, very few subjects experience losses. In a small pilot session with $c = 50$, subjects opened 2 boxes per round on average, even though expected earnings under risk neutrality are maximized by opening 12 boxes.

aversion depends on previous portfolio performance. After an increase in wealth, the investor becomes less loss averse. On the other hand, he becomes more loss averse after a decrease in wealth. As a result, investors will change their risk taking behavior after a gain or loss. After a gain, they become more risk seeking, and become less risk seeking after a loss.

Dillenberger and Rozen (2014) put forth a model where the decision maker's risk preference can be affected by his history of gains and losses, or "elations" and "disappointments." In their model, disappointment and elation are determined endogenously as a result of prior outcomes. This history, and the sequence of disappointments and elations, directly affects the decision maker's utility. The authors prove equivalence with the reinforcement effect, where the decision maker becomes more risk averse after a disappointment than after an elation.

Tserenjigmid (2017) identifies another channel through which the reinforcement effect can arise. He generalizes discounted expected utility in a dynamic choice environment. Furthermore, he shows that the reinforcement effect results from individuals having history dependent risk preferences and satisfying monotonicity with respect to first order stochastic dominance.

On the other hand, the "disposition effect" can predict the opposite. Introduced by Shefrin and Statman (1985), the disposition effect describes the tendency of investors to sell stocks that have gained in value and keep stocks that have dropped in value. Shefrin and Statman (1985) explain the disposition effect using prospect theory. Investors set the reference point of a stock at the initial purchase price. A stock that increases in value relative to the purchase price is coded as a gain, while a stock that decreases in value relative to the purchase price is coded as a loss. Prospect theory individuals are risk averse over gains and risk seeking over losses, so the disposition effect suggests the opposite of the reinforcement effect—individuals appear more risk averse after a gain and less risk averse after a loss. Weber and Camerer (1998) develop a model which captures these aspects and predictions of the disposition effect.

Taken together, the theoretical literature gives little consensus. In some environments, theories predict individuals will take less risk after a loss, and in others individuals will take more risk after a loss. Motivated by the theoretical and empirical inconsistency, the realization effect (Imas, 2016) seeks to reconcile these predictions and clarify the environments where risk will increase or decrease.

III.II. The Realization Effect

The full realization effect theory is presented formally in the Appendix of Imas (2016), and we direct the interested reader there for details. The model is based on Cumulative Prospect Theory (CPT) (Tversky and Kahneman, 1992).

Denote a lottery, L , which gives x^i with probability p^i by $L = (x^1, p^1; x^2, p^2; \dots; x^n, p^n)$. For simplicity we assume $x^i > x^{i+1} \forall i$. The decision maker evaluates a lottery as

$$U(L) = \sum_{i=1}^n \pi^i V(x^i | r) \quad (1)$$

where r is the reference point and π^i are decision weights. Under CPT, decision weights are generated from a probability weighting function, $w(\cdot)$, and are given by $\pi^i = w(p^i + \dots + p^n) - w(p^{i+1} + \dots + p^n)$ for $x^i \geq r$ and $\pi^i = w(p^1 + \dots + p^i) - w(p^1 + \dots + p^{i-1})$ for $x^i < r$. The utility function may exhibit loss aversion, parameterized by $\lambda > 1$. $V(x^i | r)$ is defined as

$$V(x^i | r) = \begin{cases} v(x^i - r) & \text{if } x^i \geq r, \\ -\lambda v(-(x^i - r)) & \text{if } x^i < r, \end{cases} \quad (2)$$

where $v(\cdot)$ is a concave function.

In the environment Imas (2016) considers, a decision maker is offered a positively skewed lottery, $L = (x^g, p; x^l, 1-p)$, where $p < 0.5$, $x^g > 0 > x^l$, and $x^g > |x^l|$. Imas assumes the reference point to be the status quo, $r = 0$. The decision maker can accept or reject the lottery in period one. If he accepts the lottery, he observes the lottery outcome before proceeding to period two, where he makes the same decision to accept or reject L once again. If he rejects the lottery in period one, he is not offered a lottery in period two. The primary question, both theoretical and empirical, is whether the decision maker will become more or less likely to accept the lottery in period two after accepting and losing the lottery in period one. The realization effect theory assumes that realized and paper outcomes are bracketed differently within this framework, and therefore predictions depend on whether the loss from L in period one is “realized” or not. The theory predicts that the decision maker will become more likely to accept L if the loss in period one is not realized, and will become less likely to accept L if the period one loss is realized.

The realization effect operates through three primary channels: choice bracketing, loss aversion, and probability weighting. *Choice bracketing* determines the manner in which risky decisions are evaluated, and is the main addition of the theory to standard cumulative prospect theory (CPT) (Tversky and Kahneman, 1992). The realization effect theory assumes that mental accounts remain open after paper gains or losses, but accounts close after realized gains or losses. As a result, subsequent lotteries are bracketed and evaluated *together* with paper gains or losses, but are evaluated *separately* from realized gains or losses. In other words, a decision maker who suffers a paper loss and then considers accepting a second lottery evaluates the second lottery’s prospects in conjunction with the previous loss. However, a decision maker who suffers a realized loss does

not integrate this loss with the prospects of future lottery decisions.¹²

Second, *loss aversion* works to give extra attraction to lotteries that offer the opportunity to erase previous paper losses. A loss averse individual experiences a sharp decrease in utility when he “realizes” outcomes below his reference point. After experiencing a paper loss, the individual gains a heightened attraction to gambles that offer a large reward which could cancel out the paper loss. Since the paper loss has not yet been realized, this new gamble offers the opportunity to avoid realizing any loss altogether, which is particularly appealing to a loss averse decision maker. Lotteries following realized losses don’t offer the possibility of erasing the negative payment since the loss has already been realized. As a result, loss aversion along with the assumed manner of choice bracketing provide differing risk-taking incentives following realized and paper losses, leading to differences in predicted choices.

Finally, *probability weighting* affects how individuals view “tail” events. Individuals tend to overweight low probability events, such as a long-shot chance to win a large sum. Probability weighting skews the individual’s perception of winning the lottery, making the possibility appear more likely. As a result, probability weighting provides incentive for individuals to favor positively skewed lotteries, where the potential upside is larger than the potential downside and where the probability of a gain is less than the probability of a loss. Loss aversion provides an incentive to chase this long-shot lottery after a *paper loss*, since it offers the decision maker the possibility that he will never have to realize the loss. Realized losses, on the other hand, provide no such incentive. As a result, probability weighting combined with loss aversion induces individuals to become more willing to chase long-shot risks after paper, compared to realized, losses.

We study dynamic risk preferences under both types of outcomes. While the realization effect makes predictions only in specific domains, we take the insight from Imas (2016) and study realized versus paper outcomes more broadly. We ask how risk behavior changes with prior outcomes, and whether this depends on whether the prior outcomes are “realized” or not. We study a negatively skewed environment, which is outside the scope of the realization effect theory and may result in very different decisions in dynamic risk environments.

IV. RESULTS

In the Appendix, we report summary statistics from the data, including average earnings and average number of boxes opened in each round by treatments.

Figure I plots the change in number of boxes opened from the previous round, separately for the two treatments, depending on whether the outcome in the previous round was a gain or a loss. The change in number of

¹²For example, consider a lottery that offers a chance to win \$500 at the risk of losing \$100. Informally, a decision maker who loses in the first round but does not realize this loss considers the second lottery as offering him the chance to finish with earnings of \$400 or -\$200, integrating his first loss into the lottery outcomes. A decision maker who realizes the first loss does not perform this integration and brackets the second lottery separately.

boxes opened is coded as the number of boxes opened in Round t minus the number of boxes opened in Round $t - 1$.¹³ Therefore, a positive change indicates an increase in risk taking while a negative change indicates a decrease in risk taking.

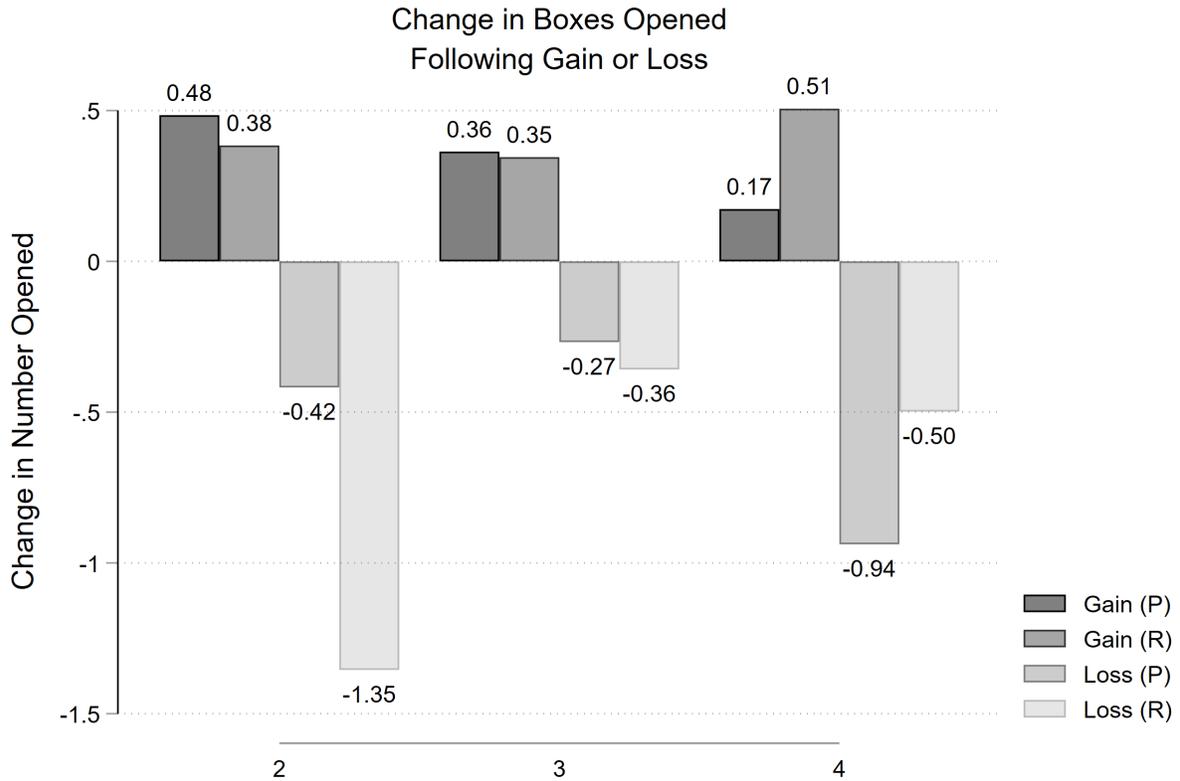


Figure I: Overall change in number of boxes following paper and realized gains and losses in each round. (P) indicates *Paper* treatment and (R) indicates *Realization* treatment.

The main result from Figure I is immediate—individuals take on *more* risk after a gain and take on *less* risk after a loss. All of the binary comparisons within treatment show significant differences ($p < 0.0064$, Wilcoxon Rank-Sum test). For each category, we run Fisher-Pitman permutation tests against a null hypothesis that the change in risk is 0. All but one of the tests are significant at conventional levels—the change in risk after a gain is significantly greater than zero and the change in risk after a loss is significantly less than zero ($p = 0.12$ for the change in Round 3 after a loss in the Realization treatment, and $p < 0.053$ for all other tests). Thus, our data overwhelmingly support the reinforcement effect (Barberis et al., 2001; Dillenberger and Rozen, 2014; Tserenjigmid, 2017).¹⁴

¹³In the original BRET paper, Crosetto and Filippin (2013) report that “the degree of risk aversion negatively correlates with the choice (of number of boxes to open).” Therefore, the choice of x represents a subject’s degree of risk aversion in a given round, and we take $x - x'$ as a simple measure of how this changes from one round to the next.

¹⁴Thaler and Johnson (1990) find the *opposite* of the reinforcement effect prediction when the potential gain of the lottery can make

To analyze the effect of prior outcomes on risk taking more generally, we run a regression predicting the change in number of boxes opened. We run two sets of specifications. In the first (columns 1 and 2), we split the sample by whether the individual has a positive or negative balance and include a dummy for losing the lottery in the previous round. In the second (columns 3 and 4), we split the sample by whether the individual experienced a previous gain or loss and include a dummy for having a negative balance. In all regressions, we include a round variable to control for pure time trends. We look at choices in rounds 2-4, and cluster errors at the subject level.¹⁵

	Change in Number of Boxes Opened			
	Negative Balance	Positive Balance	Previous Gain	Previous Loss
Loss (t-1)	-2.251*** (0.352)	-0.500*** (0.168)		
Negative Balance			0.989*** (0.261)	-0.795*** (0.243)
Round	-0.119 (0.173)	-0.107** (0.0511)	-0.0930* (0.0512)	-0.175 (0.164)
Constant	1.693** (0.665)	0.651*** (0.148)	0.611*** (0.147)	0.395 (0.593)
Observations	1,011	4,848	4,433	1,426

Standard errors clustered at the subject level
*** p<0.01, ** p<0.05, * p<0.1

Table I: Regression predicting change in number of boxes opened from one round to the next, based on previous outcome and current balance.

In the first specification, the negative coefficient on lagged loss confirms that individuals reduce risk taking after experiencing a loss, compared to a gain, in the previous round. This is true both for individuals with a negative balance and for individuals with a positive balance, although the coefficient is much larger for those with a negative balance. In the second specification, we find that the effect of current balance depends on prior outcomes. For those who won the lottery in period $t - 1$, having a negative balance leads to a further increase in risk taking relative to those with a positive balance. On the other hand, for those who lost the lottery in period $t - 1$, a negative balance indicates a reduction in risk relative to those with a positive balance.

Result 1. *Individuals exhibit the reinforcement effect, increasing risk taking after a gain and decreasing risk taking after a loss.*

up the entire loss. This is the case for nearly all of our subjects—only 13 out of 2,066 subjects cannot make up their loss at the start of Round 4. All of the results hold similarly if we drop these 13 subjects.

¹⁵Since we are looking at the effect of outcomes and choices in the *previous* round, we look only at rounds 2-4.

IV.I. Realized vs. Paper

The question remains as to whether subjects treat realized gains and losses differently from those on paper. The two treatments differ only after Round 3—before Round 3, all gains or losses in both treatments were on paper. However, sophisticated agents could anticipate the realization and alter behavior accordingly. Thus, we test whether the change in risk following a gain or loss in any round differs across treatments.

As seen in Figure I, the treatment comparisons show little overarching pattern. There is a significant difference between realized and paper losses in Round 2, where individuals decrease risk taking significantly more in the Realization treatment ($n = 177$ Paper, $n = 186$ Realization, $p = 0.0043$). There is no significant difference in any of the other comparisons.

We run a regression predicting the change in risk taking after realized and paper losses for all rounds and separately for Round 4 only as this is when the experimental manipulation occurred. We confirm a very modest, if any, effect of realization. When analyzing all rounds, the significantly negative coefficient on a loss in the previous round confirms the reinforcement effect. However, we find no effect of realization. Both the coefficient on realization and on the interaction between realization and a previous loss are small and insignificant.

	Change in Number of Boxes Opened	
	All	Round 4
Loss (t-1)	-0.819*** (0.185)	-1.074*** (0.295)
Realization	0.0668 (0.0753)	0.338* (0.178)
Loss (t-1) * Realization	-0.172 (0.235)	0.0949 (0.413)
Negative Balance	-0.155 (0.145)	-0.0954 (0.271)
Round	-0.0401 (0.0512)	
Constant	0.470*** (0.157)	0.178* (0.106)
Observations	5,859	1,953

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table II: Regression predicting change in number of boxes opened from one round to the next, based on whether outcomes were realized or on paper.

Looking only at Round 4, where the experimental manipulation takes place, we find a similar result. Individuals decrease risk taking after losses, and there is a small effect of realization ($p = 0.058$) which mitigates this effect. While there may be a very modest difference between realized and paper losses in the treatment round, the effect moves *opposite* the realization effect. The standard realization effect predicts that individuals will decrease risk taking after realized losses and increase risk taking after paper losses, but our results

suggest, if anything, the opposite.¹⁶

Result 2. *Individuals exhibit the reinforcement effect equally in environments of both realized and paper outcomes.*

However, we cannot apply the predictions from Imas (2016) uniformly to our environment since the bomb risk task does not fit into the setup of the model. As a result, predictions in the BRET may depend on individual risk and loss aversion parameters, as well as on current balance and number of boxes opened previously. That is, while the prediction for loss averse individuals choosing positively skewed lotteries is clear—individuals in the *Realization* treatment should decrease risk-taking after a loss and individuals in the *Paper* treatment should increase risk-taking after a loss—predictions under negatively skewed risk are ambiguous. We use the realization effect model to derive specific predictions for whether risk taking will increase or decrease from Round 3 to Round.

We assume a “myopic” version of the model from Imas (2016).¹⁷ Let \mathcal{L} be the set of 26 possible lotteries in the BRET. We assume an individual i ’s choice of $L_i^t \in \mathcal{L}$ in Round t maximizes his utility according to the model outlined in equations (1) and (2). We run a maximum-likelihood estimation procedure to estimate individual preferences. We assume a multinomial logit choice model,

$$P(L) = \frac{e^{\theta U(L)}}{\sum_{\hat{L} \in \mathcal{L}} e^{\theta U(\hat{L})}} \quad (3)$$

We estimate preferences using a power utility function, $v(x) = x^\alpha$, and the Prelec (1998) probability weighting function, $w(p) = e^{-(-\ln p)^\beta}$. For each subject, we use choices from Rounds 1–3 to estimate individual-specific values for α , λ , and θ , fixing $\beta = 0.60$. Given individual i ’s choices of L_i^1, L_i^2, L_i^3 , the maximum-likelihood procedure chooses values of α , λ , and θ using a grid search to maximize the likelihood function, $f(L_i^1, L_i^2, L_i^3 | \alpha, \lambda, \theta, \beta) = \sum_{t=1}^3 \ln(P(L_i^t | \alpha, \lambda, \theta, \beta))$. Possible values for α ranged from 0 to 2 in 0.1-step uniform increments. Values for λ ranged from 0 to 6 in 0.25-step uniform increments. There were 10 possible values for θ , ranging from 0 to 10 in non-uniform increments.¹⁸

Using decisions from Rounds 1–3 and actual gain/loss outcomes, we estimate individual-specific α , λ , and θ parameters using the model specified above. Given there was no exogenous realization in these rounds, we incorporate all gains and losses as paper gains/losses. Using the estimated parameters, we then estimate the

¹⁶Note, this also reduces any concern that our “realization” manipulation was not strong enough. If this had been the case, we would expect both treatments to fall under the predictions of paper losses, meaning we would see individuals *increase* risk taking after a loss.

¹⁷Deriving predictions for the non-myopic case is very complicated in our environment, where individuals choose over a menu of lotteries with many possible balances and decision paths. The non-myopic case is beyond the scope of this paper, and it is possible that predictions would differ for non-myopic agents.

¹⁸Tversky and Kahneman (1992) find estimates of $\alpha = 0.88$ and $\beta = 0.61$, Camerer and Ho (1994) find $\alpha = 0.32$ and $\beta = 0.56$, and Wu and Gonzalez (1996) find $\alpha = 0.52$ and $\beta = 0.74$. von Gaudecker et al. (2012) estimate loss aversion for multiple populations. They estimate $\lambda = 1.5$ for student populations, though it varies with the stakes of the decision. Only 8% of the population has $\lambda > 5$.

likelihood of the subject’s choice in Round 4 under two specifications: 1. assuming the balance until round 4 had been “realized,” and 2. assuming the balance was still on paper. Then, for each subject, we test whether the model better fits assuming outcomes had been realized or on paper.

We estimate an average $\alpha = 0.54$ and $\lambda = 2.08$, consistent with previous literature. For a vast majority (77%) of subjects, the realized model fits the data better. That is, most subjects’ round 4 behavior is better-predicted assuming the balance until round 4 had been realized, rather than still on paper. There is no difference between the actual experimental treatments. As shown in the first two columns of Table III, the realized model better explains the data for 77.55% of subjects in the experimentally-manipulated *Realization* treatment, and for a nearly-identical 77.35% of subjects in the *Paper* treatment.¹⁹ Binomial tests confirm statistical significance of the realized model fitting better in both the *Realization* and *Paper* treatments ($p = 0.0000$).²⁰

	Rounds 1–3 Paper		Rounds 1–3 Realized	
	Realized Fit	Paper Fit	Realized Fit	Paper Fit
<i>Realization</i> Treatment	77.55%	22.45%	83.39%	16.61%
<i>Paper</i> Treatment	77.35%	22.65%	85.41%	14.59%

Table III: Percentage of subjects for whom the Realized vs. Paper model fits better, by experimental treatment. First two columns assume outcomes in Rounds 1–3 are on paper while second two columns assume they were realized.

We run a similar procedure assuming individuals treat outcomes from Rounds 1–3 as realized gains/losses, rather than paper as above. We find similar results. We estimate an average $\alpha = 0.36$ and $\lambda = 1.29$. Again, for a vast majority of subjects (84%), the realized model fits the data better and there is no difference between the actual experimental treatments. The second two columns of Table III present the results by treatment. The realized model better explains the data for 85.41% of subjects in the experimentally-manipulated *Realization* treatment, and 83.39% of subjects in the *Paper* treatment.

Comparing across the two specifications, the model assuming individuals realize all gains or losses in Rounds 1–3 fits better in estimating individual parameters. The average loglikelihood is significantly higher assuming realized outcomes (-7.715 vs. -7.174, binomial test $p = 0.000$).²¹ Overall, the data is most consistent with assuming individuals treat all outcomes as realized. This leaves an interesting avenue for future work to investigate endogenous realization.

¹⁹We find that the paper model is more likely to fit better for individuals with higher estimated α parameters, and the realized model is more likely to fit better for individuals with lower estimated λ parameters. This was not a prior hypothesis, and we encourage future research to explore more of these relationships.

²⁰We confirm this using Vuong’s closeness test for non-nested models, which gives a z-score of -31.03 in favor of the model of realized outcomes.

²¹Vuong’s test gives a z-score of -160.60 in favor of the model of realized outcomes.

Overall, the model predicts individuals in both treatments will take on less risk after a loss, on average, in our environment.²² This gives new insight into the application of realized losses to negatively skewed environments. In this domain, realizing losses relative to keeping them on paper is unlikely to generate much difference in risk taking, since individuals will take on less risk after both types of outcomes. This is in stark contrast to dynamic risk in environments of positively skewed risk, where individuals have been shown to take on more risk after paper losses.

V. DISCUSSION

We demonstrate the reinforcement effect in a large scale experiment on Amazon’s Mechanical Turk. Using the bomb risk elicitation task, we find that individuals take on more risk after a gain and take on less risk after a loss. We exogenously vary whether these gains and losses are “realized” or are just on paper, according to the distinction put forth in the realization effect theory of Imas (2016). We find no difference in the response to realized and paper outcomes in our environment. We find that a model of realizing outcomes better explains the data in both treatments, regardless of exogenous experimental realization. In our domain, which is outside the scope of the theory, we see little evidence of a realization effect.

The realization effect relies heavily on positively skewed risk. In many environments, like in our experiment, positively skewed choices are those that are relatively riskier. This is especially true of financial risks, such as choosing over a menu of potential stocks or investment projects which vary in their risk/reward trade-off. Previous papers on loss chasing and the realization effect use lotteries with fixed probabilities of gains and losses, essentially “forcing” individuals into an environment of positively skewed risk. In environments where the risk is more akin to the BRET, where risk determines the probability of a gain or a loss, individuals can choose the skewness of their preferred risk along with choosing the expected value. In environments where the optimal choice is a negatively skewed gamble as in our experiment, we are unlikely to see behavioral differences between risk preferences under realized and paper outcomes.

In a new working paper, Merkle et al. (2018) test the realization effect using the same investment task as in Imas (2016), but where subjects can choose how much to invest into a negatively skewed lottery. They vary whether these outcomes are realized or on paper. Consistent with our results, they find no realization effect. These results further support the current findings regarding skewness of risk. When negatively skewed gambles are the most desirable, or are the only gambles offered, there is unlikely to be a difference in risk taking between realized and paper outcomes.

These results help shed light on the boundaries of the realization effect and the environments in which

²²We estimate Round 4 choices based on parameter estimates and Round 3 outcomes.

we can expect to see loss chasing rather than reinforced losses. For positively skewed risks, such as familiar state/national lotteries, penny stocks, and insurance among others, we might expect differences between realized and paper outcomes. The literature suggests these are environments where we are most likely to see loss chasing. On the other hand, investing in the stock market is generally a negatively skewed risk. In this environment, our results suggest we are unlikely to see loss chasing and differences between realized and paper outcomes. Instead, we expect to see individuals reinforce their outcomes, taking on more risk after a gain and less risk after a loss. While there are many factors which vary between risky environments, our results suggest that skewness of risk is one dimension that plays a role in dynamic risk choices.

REFERENCES

- Andrade, E. B., Iyer, G., 2009. Planned versus actual betting in sequential games. *Journal of Marketing Research* 46 (3), 372–383.
- Barberis, N., Huang, M., Santos, T., 2001. Prospect theory and asset prices. *The Quarterly Journal of Economics* 116 (1), 1–53.
- Barberis, N., Xiong, W., 2012. Realization utility. *Journal of Financial Economics* 104 (2), 251–271.
- Camerer, C. F., Ho, T.-H., 1994. Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty* 8, 167–196.
- Clark, L., 2010. Decision-making during gambling: An integration of cognitive and psychobiological approaches. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences* 365, 319–330.
- Crosetto, P., Filippin, A., 2013. The “bomb” risk elicitation task. *Journal of Risk and Uncertainty* 47, 31–65.
- Dillenberger, D., Rozen, K., 2014. History-dependent risk attitude. *Journal of Economic Theory* 157, 445–477.
- Imas, A., 2016. The realization effect: Risk-taking after realized versus paper losses. *American Economic Review* 106 (8), 2086–2109.
- Langer, T., Weber, M., 2008. Does commitment or feedback influence myopic loss aversion? *Journal of Economic Behavior & Organization* 67 (3–4), 810–819.
- Lejuez, C., Read, J., Kahler, C., Richards, J., Ramsey, S., Stuart, G., Strong, D., Brown, R., 2002. Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (bart). *Journal of Experimental Psychology: Applied* 8 (2), 75–84.
- Liu, Y.-J., Tsai, C.-L., Wang, M.-C., Zhu, N., 2010. Prior consequences and subsequent risk taking: New field evidence from the taiwan futures exchange. *Management Science* 56 (4), 606–620.
- Merkle, C., Müller-Dethard, J., Weber, M., 2018. Closing a mental account: The realization effect for gains and losses. SSRN Working Paper.
- Prelec, D., 1998. The probability weighting function. *Econometrica* 66, 497–527.
- Prims, J., Sisso, I., Bai, H., 2018. Suspicious ip online flagging tool.
URL <https://itaysisso.shinyapps.io/Bots>

- Schneider, S. L., Kauffman, S., Ranieri, A., 2016. The effects of surrounding positive and negative experiences on risk taking. *Judgment and Decision Making* 11, 424–440.
- Shefrin, H., Statman, M., 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40 (3), 777–790.
- Shiv, B., Loewenstein, G., Bechara, A., 2006. The dark side of emotion in decision-making: When individuals with decreased emotional reactions make more advantageous decisions. *Cognitive Brain Research* 23 (1), 85–92.
- Smith, G., Levere, M., Kurtzman, R., 2009. Poker player behavior after big wins and big losses. *Management Science* 55, 1547–1555.
- Thaler, R., Johnson, E., 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science* 36 (6), 643–660.
- Tserenjigmid, G., 2017. History-dependent risk aversion and the reinforcement effect. Working Paper.
- Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5 (4), 297–323.
- Verbruggen, F., Chambers, C. D., Lawrence, N. S., McLaren, I. P. L., 2017. Winning and losing: Effects on impulsive action. *Journal of Experimental Psychology* 43, 147–168.
- von Gaudecker, H.-M., van Soest, A., Wengstr om, E., 2012. Experts in experiments: How selection matters for estimated distributions of risk preferences. *Journal of Risk and Uncertainty* 45 (2), 159–190.
- Weber, M., Camerer, C. F., 1998. The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior and Organization* 33, 167–184.
- Wu, G., Gonzalez, R., 1996. Curvature of the probability weighting function. *Management Science* 42, 1676–1690.

I. PROBABILITIES

Number Boxes Opened	Probability of Bomb	Probability No Bomb	Low Stakes		High Stakes	
			Gain	Loss	Gain	Loss
1	0.04	0.96	\$0.01	-\$0.10	\$0.02	-\$0.10
2	0.08	0.92	\$0.02	-\$0.10	\$0.04	-\$0.10
3	0.12	0.88	\$0.03	-\$0.10	\$0.06	-\$0.10
4	0.16	0.84	\$0.04	-\$0.10	\$0.08	-\$0.10
5	0.2	0.8	\$0.05	-\$0.10	\$0.10	-\$0.10
6	0.24	0.76	\$0.06	-\$0.10	\$0.12	-\$0.10
7	0.28	0.72	\$0.07	-\$0.10	\$0.14	-\$0.10
8	0.32	0.68	\$0.08	-\$0.10	\$0.16	-\$0.10
9	0.36	0.64	\$0.09	-\$0.10	\$0.18	-\$0.10
10	0.4	0.6	\$0.10	-\$0.10	\$0.20	-\$0.10
11	0.44	0.56	\$0.11	-\$0.10	\$0.22	-\$0.10
12	0.48	0.52	\$0.12	-\$0.10	\$0.24	-\$0.10
13	0.52	0.48	\$0.13	-\$0.10	\$0.26	-\$0.10
14	0.56	0.44	\$0.14	-\$0.10	\$0.28	-\$0.10
15	0.6	0.4	\$0.15	-\$0.10	\$0.30	-\$0.10
16	0.64	0.36	\$0.16	-\$0.10	\$0.32	-\$0.10
17	0.68	0.32	\$0.17	-\$0.10	\$0.34	-\$0.10
18	0.72	0.28	\$0.18	-\$0.10	\$0.36	-\$0.10
19	0.76	0.24	\$0.19	-\$0.10	\$0.38	-\$0.10
20	0.8	0.2	\$0.20	-\$0.10	\$0.40	-\$0.10
21	0.84	0.16	\$0.21	-\$0.10	\$0.42	-\$0.10
22	0.88	0.12	\$0.22	-\$0.10	\$0.44	-\$0.10
23	0.92	0.08	\$0.23	-\$0.10	\$0.46	-\$0.10
24	0.96	0.04	\$0.24	-\$0.10	\$0.48	-\$0.10
25	1	0	\$0.25	-\$0.10	\$0.50	-\$0.10

Table IV: Gain and loss probabilities and payoffs

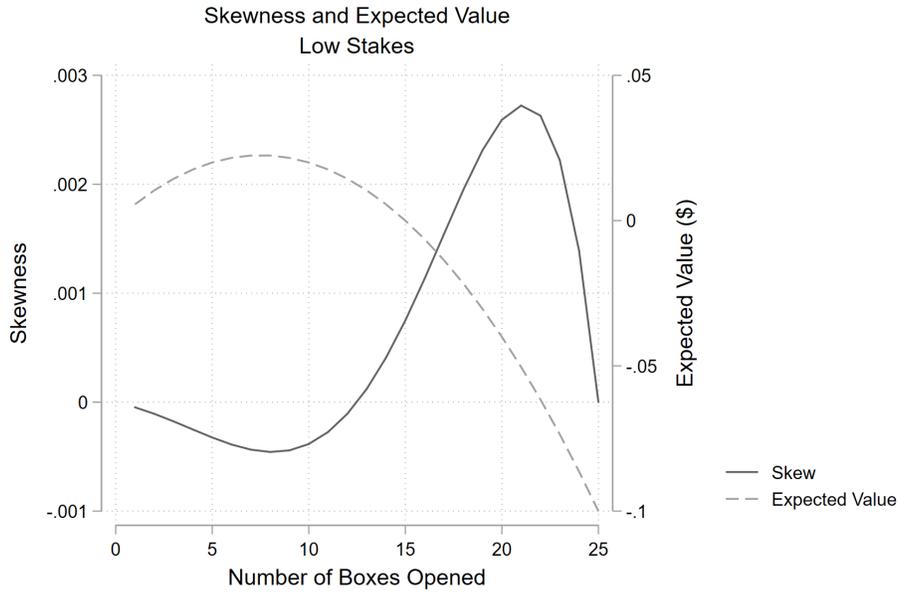


Figure II: Low Stakes Skew and Expected Value

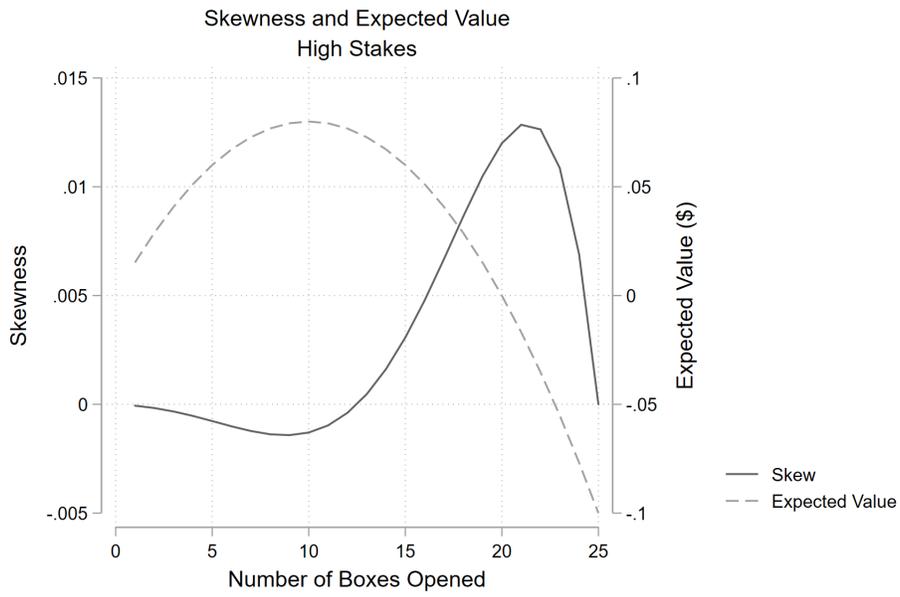


Figure III: High Stakes Skew and Expected Value

II. SUMMARY STATISTICS

Paper Treatment: Low Stakes			
	Average # Boxes Opened	Percent Clicked Bomb	Average Earnings
Round 1	6.26	16.00%	3.17
Round 2	6.56	22.00%	2.26
Round 3	6.57	22.56%	1.87
Round 4	6.54	29.33%	0.69

Paper Treatment: High Stakes			
	Average # Boxes Opened	Percent Clicked Bomb	Average Earnings
Round 1	6.58	17.38%	7.81
Round 2	6.94	25.66%	6.31
Round 3	7.30	27.98%	6.18
Round 4	7.12	33.28%	4.36

Realization Treatment: Low Stakes			
	Average # Boxes Opened	Percent Clicked Bomb	Average Earnings
Round 1	6.71	14.25%	3.98
Round 2	7.09	26.91%	1.88
Round 3	7.20	26.65%	2.02
Round 4	7.47	34.30%	0.49

Realization Treatment: High Stakes			
	Average # Boxes Opened	Percent Clicked Bomb	Average Earnings
Round 1	8.11	25.38%	8.31
Round 2	7.88	31.35%	6.64
Round 3	8.04	32.31%	6.61
Round 4	8.20	36.15%	5.40

Instructions

Welcome! Thank you for participating in this HIT.

In addition to your \$0.10 for participating, you have now been given a starting balance of \$0.40. Any additional money you gain will be added to this starting balance, and losses will be subtracted from this balance.

In this HIT, you will make decisions over 4 rounds. For each round, you will see a 5x5 grid containing 25 boxes, an example of which is shown below.

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20
21	22	23	24	25

Most of these boxes are empty, but 1 of them contains a “bomb.” The placement of the bomb is completely random – each of the 25 boxes is equally likely to contain the bomb.

You can choose to open as many boxes as you like, from 0-25 (Note: This means you are not required to open any boxes if you do not wish to do so). If you open the box that contains the bomb, you will lose \$0.10. However, if you only open empty boxes, you will gain \$0.02 per box opened.

You will click on all the boxes you wish to open, and then progress to the next page to reveal which box contained the bomb. You will NOT see whether a box contains the bomb until you’ve made ALL your choices on which boxes to open and have progressed to the next page. That is, if you’ve clicked on the bomb, you will not learn that immediately. You will see the bomb placement after you’ve opened all of the boxes you want to open and you progress to the next page.

You will make these decisions over 4 separate rounds. *[However, after the 3rd round, we will finalize your earnings up until that point by adding your earnings to your account balance or withdrawing your losses from your account balance. So if you have positive earnings, we will add those to your \$0.40 starting balance, and we will finalize this account balance before Round 4. If you have negative earnings, we will subtract these from your \$0.40 starting balance, and then finalize your account balance before Round 4.*

We will tell you your balance before you start Round 4. Then, we will transfer your earnings to your account (in the case of positive earnings) or away from your account (in the event of losses), and then you will participate in one final round. Your earnings from the last round will be transferred separately at the end of the task. Note: You can still earn money by opening empty boxes in Round 4, and you will still lose \$0.10 from your balance if you open a bomb in Round 4. But in Round 4, we will have already taken money out of your balance (if you've clicked on a bomb) and we will have added money to your balance (if you've opened empty boxes).]

The placement of the bomb is completely random across these 4 rounds, so the bomb is equally likely to be behind any of the boxes in any of the rounds.

Please remember that, if you click on a bomb, you won't be notified of that in "real-time." You'll find out after you have made all your choices and progressed to the next round.

1	2	3	4	5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20
21	22	23	24	25



You chose 1, 4, 5, 8, 10, 12, 13, 18, 23, 25.

The bomb was behind box 4. You will lose \$0.10.

1	2	3		5
6	7	8	9	10
11	12	13	14	15
16	17	18	19	20
21	22	23	24	25

You chose 1, 5, 7, 9, 13.

The bomb was behind box 12. You did not select the bomb! You will gain 10 cents.

1	2	3	4	5
6	7	8	9	10
11		13	14	15
16	17	18	19	20
21	22	23	24	25



You are about to begin Round 4. To establish your earnings before you make your final decisions, the amount you've lost will be subtracted from your \$0.40 starting balance.

Your current earnings: -10 cents

This amount will now be withdrawn from your account and transferred back to the experimenter.

To transfer your losses and finalize your account balance, you must enter "Close Account" below.

