Wait, Wait... Don't Tell Me: Repeated Choices With Clustered Feedback

David Hagmann[1], Jason L. Harman[2], and Cleotilde Gonzalez[1]

[1]Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213
[2]Department of Psychology, Louisiana State University, Baton Rouge, LA 70809

December 19, 2015

Abstract

When individuals make repeated choices between two lotteries without having a description of their payoffs, they have to form beliefs based on the observed outcomes of their choices. Previous research finds that choices differ consistently after learning about outcomes compared to having an objective description, an effect termed the description-experience gap. We introduce a new clustered feedback mechanism in which participants receive feedback consisting of individual outcomes of a number of choices at once, rather than observing the outcome immediately after making a decision. Presenting clustered feedback closes the description-experience gap and leads individuals to act as if they had a description available. We also use lotteries with rare outcomes and find that the description-experience gap is greatest when a high payoff is rare, and is closed by clustered feedback, but does not emerge when a high payoff is common.

1 Introduction

Feedback is an instrumental part of decision making. Sales figures signal the success of a manager's marketing campaign, the likelihood of obtaining an interview indicates whether a job seeker is applying to the right job openings, and exam scores suggest aptitude in a particular subject. When things go poorly, feedback allows us to adjust our actions and avoid repeating mistakes. Information about outcomes is therefore instrumentally valuable and a key component of any environment in which we repeatedly make decisions, so that we can respond to what we learn from our past actions.

Intuitively, we might also suspect that getting feedback quicker and more frequently is desirable. For example, the sooner a job seeker learns that she is underqualified for a position, the quicker she can target other opportunities. If we are taking suboptimal actions, we should want to know immediately and make early corrections. However, many realistic processes are fraught with noise, so that a single bad (good) outcome does not necessarily signal a bad (good) decision. For example, merely getting a few rejections from job applications should not discourage someone from applying further. If receiving feedback influences people's tendency to explore options (e.g. because they are quickly discouraged by bad outcomes), then introducing a mechanism that delays feedback can improve ex post outcomes.

Previous research has shown that providing as feedback the sum of multiple outcomes, rather than individual realizations, can increase risk taking, because it shrouds the losses that occured along the way (Gneezy and Potters 1997, Fellner and Sutter 2009). However, in many instances, it is not feasible to protect decision makers from learning about bad outcomes. In this paper, we propose a "clustered feedback" mechanism in which decision makers make multiple decisions prior to observing individual outcomes. We hypothesize that observing multiple outcomes at once provides a more representative impression of available options and discourages extreme sensitivity to bad outcomes.

In a laboratory experiment, we vary whether people obtain feedback immediately after making a decision, or whether they make a few decisions with no feedback before observing the outcomes of their choices. We then apply the clustered feedback mechanism to an environment in which a description of the options is available a priori (and feedback is thus non-informative), and to an environment in which participants have to learn what the options are from their observation of outcomes (such that feedback is instrumental). We hypothesize that obtaining less frequent feedback in a setting without description leads participants to explore a risky option further; and when that risky option has a higher expected value, delaying instrumental feedback can make decision makers better off.

2 Background

Our paper builds on previous research in both experimental economics and experimental psychology. Despite the common interest in repeated decisions and the role of feedback frequency on risky choice, the two literatures have remained disparate. We explore separately previous work on repeated descriptive decisions and repeated experiential decisions.

Descriptive decisions are those involving full information of probabilities and outcomes, but no outcome feedback after a choice. They have been particularly prominent among economists, who are concerned that learning about outcomes affects subsequent risk-taking (confirmed recently in a range of settings by Imas forthcoming). However, decision makers in real world settings rarely have such clear descriptions of the options they are facing. Instead, for decisions they make repeatedly, they learn from their experience and avoid options that worked out badly and pursue options that had good outcomes. Although experiencing an option repeatedly should lead beliefs to converge to the true state, boundedly rational decision makers may not explore sufficiently or may not remember outcomes correctly and thus fail to converge (Hertwig et al. 2004).

The last decade has seen a growth in research studies involving decisions under uncertainty made from experiential information (Barron and Erev 2003). Experiential decisions are those made without having a priori a description of the probabilities and payoffs of the available options. Instead, decision makers learn about their options purely from observing feedback about outcomes of their decisions. One main insight from this research is that people make different choices across these two types of information: decision makers behave as if they underweight rare events when relying on experience, while they overweight rare bad outcomes when relying on description. This difference has become known as the Description-Experience gap (Hertwig et al. 2004, Hertwig and Erev 2009).

Experiments that study the Description-Experience gap have participants explore the gambles without a description, either by observing outcomes without consequences before making one consequential choice (i.e., sampling paradigm); or by making a fixed number of repeated, incentivized decisions (for a comparison of these paradigms, see Gonzalez and Dutt 2011). The results from such experiential binary decisions are then compared to one-shot decisions made with an explicit description (e.g., \$4 with p=.80, or \$0 otherwise vs \$3 with p=1). However, in many naturalistic decision situations, experience and description are combined together, and few studies have investigated the Description-Experience gap when both experience and descriptions are simultaneously present (e.g. Jessup et al. 2008, Lejarraga and Gonzalez 2011).

In Economic decisions for example, it is known that decision makers who decline a single round of a risky gamble given an explicit description (e.g. \$200 if heads comes up or -\$100 if tails comes up) may participate if they could play the same gamble repeatedly (an observation first made by Samuelson 1963). Thus, they vary not only the learning mechanism, but also switch between repeated and single choices. Jessup et al. (2008) introduce the "repeated descriptive paradigm" in which decision makers repeatedly choose from two options with description and observing the outcome of every decision. They ask participants to make 120 incentivized decisions between a safe option and a risky option and manipulate the payoff variance of the risky option (within-subjects) and whether participants observe the outcome after every decision (between-subjects) or whether they receive no feedback until the end of the experiment. When small gains are common in the risky option, they find that feedback reduces the proportion of risky choices (a somewhat counter-intuitive result that is contradicted by Langer and Weber 2005, that we will review momentarily).

In prospect theory, decision makers' utility is affected more by losses than by equivalently sized gains (Tversky and Kahneman 1992). How often losses or gains are incurred, however, depends on how frequently outcomes are evaluated. Read et al. (1999) observe that in risky choices, people often pay attention to individual outcomes rather than integrating multiple related outcomes. Frequent evaluations therefore increase the effect of loss aversion and – to the extent that loss aversion leads to suboptimal decisions – lead to worse outcomes.

This effect is particularly salient in the stock market. Mehra and Prescott (1985) show that the return to stocks is excessively large compared to bonds even after adjusting for the difference in risk. Benartzi and Thaler (1995) explore the role of loss aversion and myopia in the stock market setting and show that the premium return on stocks can be explained by a model with loss aversion and annual portfolio evaluations. In other words, overreacting to losses in combination with frequently observing outcomes can lead decision makers to prefer safer options with lower long term returns. These findings hold up in a range of extensions, showing that delaying and aggregating feedback increases the rate at which investments in risky options are made (for a review, see Hopfensitz 2009). Myopia does not necessarily discourage risk-seeking, however. Langer and Weber (2005) show that in a setting in which large losses are rare and small gains are common, frequent evaluations may in fact make a risky option more appealing.

Timing of feedback has been shown experimentally to affect choices when descriptions are present. Gneezy and Potters (1997) find that participants prefer to invest in a bond-like asset (a safe option) when they frequently observe feedback about its performance, but a stock-like asset (a risky option) when feedback is given less frequently. Since stocks have a higher expected return, delayed feedback improves earnings by the end of the experiment. Interestingly, decision makers appear to not be sophisticated about the effect of feedback, as they prefer immediate feedback when given the choice (Fellner and Sutter 2009).

While the effect of myopia has been extensively studied in decisions with descriptions, there have been no such studies in decisions without a description to the best of our knowledge. Existing studies of repeated decisions from experience provide immediate feedback only, possibly encouraging myopia. Moreover, immediate feedback may be quite uncommon in the real world, where delayed and clustered feedback characterizes much of the economically consequential settings (e.g. a job applicant who submits multiple applications before learning about any outcomes).

We extend the literature in three ways. First, we compare repeated consequential decisions under uncertainty without a description (as in Barron and Erev 2003) and with an explicit description (as in Jessup et al. 2008). In both cases, participants obtain feedback about individual outcomes throughout the experiment. Second, we introduce a clustered feedback mechanism that shows individual outcome realizations, but with some delay. Participants make a number of decisions without observing any outcomes, before seeing a table displaying at once each of the outcomes of their choices. We compare decisions with clustered feedback against decisions made when feedback is obtained immediately after each choice. Finally, we test the effect of the payoff variance of the risky gamble across conditions. Participants always face a safe option with a sure payoff of 4 and a binary lottery with an expected value of 5. We vary whether the high outcome is common (6.25 with 80% probability, 0 otherwise), both outcomes are equally likely (10 with $p = 0.5$, 0 otherwise), or the high outcome is rare (25 with $p = 0.2$, 0 otherwise).

Our experiment introduces a realistic feedback mechanism into experiential decisions and tests whether delaying feedback when a description is absent may lead to decisions that resemble those made in the presence of a description. Moreover, because we obtain time series data, we can look at the effect of obtaining feedback for the first time after a series of decisions. We hypothesize that in the absence of a description, early negative outcomes discourage further exploration of a risky option and prevent an accurate belief from being formed. Clustered feedback is one way of shielding decision makers from such early negative feedback, leading them to explore an option for a longer period and giving them a better chance of forming beliefs matching the true state.

3 Experimental Design

We recruit 1,216 participants from Amazon Mechanical Turk (54.4% male, mean age 32.7) and incentivize all of their decisions. Participants are asked to make 110 sequential and incentivized choices between two stationary options, presented on the screen as two buttons labeled "Option A" and "Option B." One option represents a safe payoff and the other a binary lottery[1]. Their total earnings are determined by the sum of all outcomes from their choices and ranged from 81 cents to $2.01, with a mean of $1.24. The study took participants approximately seven minutes to complete. The full instructions and interface are presented in Appendix B.

The experimental design compares decisions over three dimensions (all between-subjects). In the first dimension, we compare conditions with or without an explicit description of the available

[1]The order of the safe and risky options is counterbalanced.

options. In the "description" condition, outcomes and probabilities of each of the options are stated explicitly in the instructions and are present on the buttons' labels during every choice made. In the "no description" condition, participants learn outcomes and the frequency of these outcomes only from feedback of the selections they make.

On a second dimension, we vary whether participants observe the realization of their choice immediately after choosing one of the options ("Immediate Feedback," as in the traditional repeated choice paradigm) or whether they only obtain feedback after every 10 trials ("Clustered Feedback"). In the latter condition, participants make 10 choices without observing any outcomes. After 10 trials, they are shown a table with the individual outcomes of their last 10 choices.[2]

The last dimension is the payoff structure of the risky option. In the low variance lottery, the risky option returns 6.25 with $p = 0.8$ and 0 otherwise (denoted $(6.25, 0.8; 0)$). The medium variance lottery consists of the gamble $(10, 0.5; 0)$ and the high variance lottery offers the gamble $(25, 0.2; 0)$. The binary lottery has an expected value of 5 and the safe option a sure outcome of 4. Note that participants are most likely to observe small gains (with respect to the safe option's return of 4) in the low variance setting, while they are most likely to observe a loss in the high variance setting.

Consider a participant who always chooses the risky option. Simulation shows that in the low variance condition, the 95% confidence interval of the sum of their realizations is $[500, 600]$. Their probability of obtaining less than 440 (obtained by always choosing the safe option) is approximately 0.005% (1 in 20,000). For decision makers who are not myopic and who have a description available, we should thus expect all of them to always choose the risky option. On the other hand, if they are myopic and loss averse, many may instead opt for the safe option.

Increasing the payoff variance increases the size of the confidence interval and increases thus the probability of earning less with the risky option than with the safe option. The corresponding intervals for the medium and high variance lotteries are $[450, 650]$ and $[350, 750]$, respectively, and the probabilities of doing worse than with the safe option are 2.24% and 14.14%, respectively. Thus, we would expect the popularity of the risky option to decline with the payoff variance given risk averse agents (without the need to introduce loss aversion).

4 Results

We begin our analysis by looking at the proportion of risky options chosen across all periods, shown in Figure 1. The left panel presents the results for the immediate feedback condition and the right panel presents the clustered feedback condition. Each panel compares the proportion of

[2] In a pilot, we tested block sizes of 5, 10, and 20 and found no significant differences between them. The data from the pilot are available from the authors upon request.

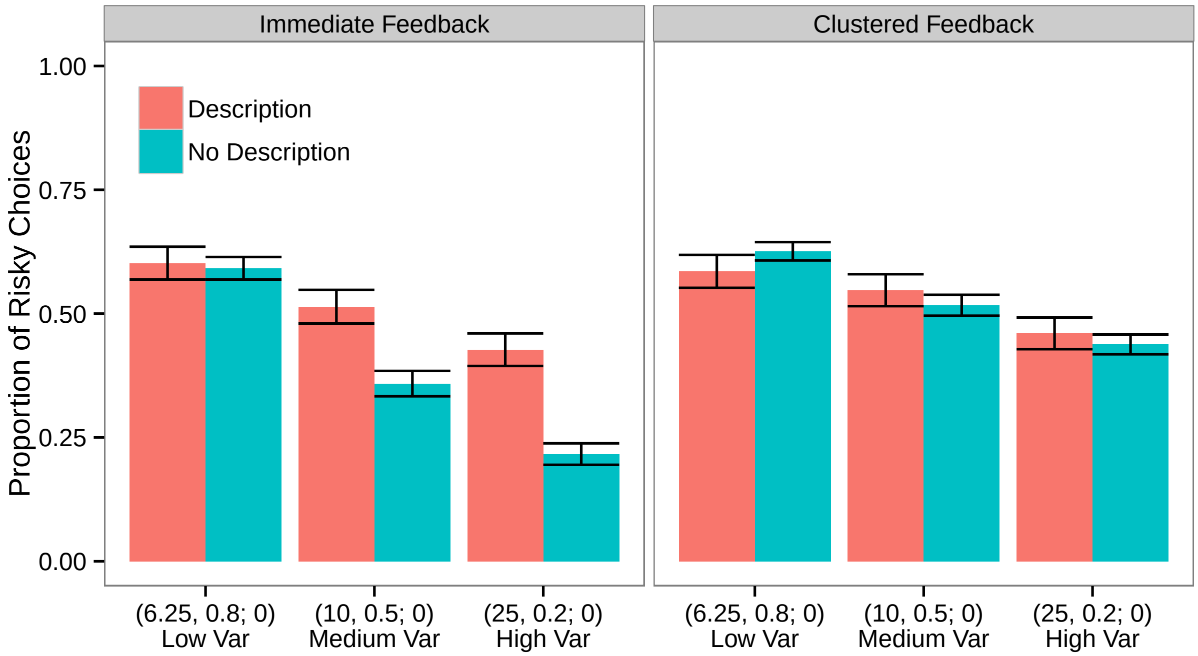

Figure 1: Proportion of risky options chosen by feedback, presence of description, and lottery. There is a significant difference between description and experience with immediate feedback and medium and high variance lotteries, but clustered feedback closes this gap. Error bars show one standard error above and below the mean.

risky choices in the presence or absence of a description, in each of the three lotteries. We can see immediately that in no condition does the proportion of risky choices approach 100% and in fact never exceeds 63%. This suggests that many decision makers appear to evaluate outcomes myopically.

Result 1: The description-experience gap is increasing in payoff variance with immediate feedback

In the left panel of Figure 1, we observe a difference between the description and no description conditions that increases with the payoff variance. When the most common outcome in the risky option is higher than the safe outcome (6.25 with 80% probability vs. 4 guaranteed), there is no difference between the two conditions $(F(1, 205) = 0.07, p = 0.80)$; that is, there is no gap in the presence of immediate feedback between repeated decisions with or without a description. However, in the other two lotteries, we observe a substantial and increasing difference between the conditions with and without a description. In the medium variance lottery, where the high and low outcomes in the risky option occur with equal probability, the absence of a description reduces the proportion of risky choices from 51.4% with description to 35.9% without (a 30% decrease, $F(1, 199) = 13.39, p < 0.001$). In the high variance lottery, the decrease is from 42.7%

in the presence of description to 21.6% in its absence ($F(1, 204) = 28.13, p < 0.001$). Thus, the description-experience gap is robust to repeated consequential choice settings with or without a description, but is dependent on the gamble's payoff profile: the rarer the higher outcome in the risky option is (i.e. the greater that the payoff variance is), the more the absence of an explicit description leads participants towards a safe choice with a lower expected value.

Table 1: Regressions on the probability of choosing the risky option (using a linear probability model).

	(6.25, 0.8; 0)	(10, 0.5; 0)	(25, 0.2; 0)			
No Description	-0.010	-0.009	-0.155^{***}	-0.165^{***}	-0.211^{***}	-0.220^{***}
	(0.041)	(0.041)	(0.042)	(0.043)	(0.040)	(0.040)
Male		0.024		0.067		0.039
		(0.042)		(0.043)		(0.040)
Age		0.0003		0.002		-0.004^{*}
		(0.002)		(0.002)		(0.002)
Constant	0.602^{***}	0.588^{***}	0.514^{***}	0.481^{***}	0.427^{***}	0.408^{***}
	(0.028)	(0.038)	(0.030)	(0.037)	(0.028)	(0.035)
Observations	22,770	22,770	22,110	22,110	22,660	22,660
Note:						*p<0.05; **p<0.01; ***p<0.001

We further test this finding by looking at each individual decision, rather than the proportion of choices. Table 1 presents a linear probability model where the dependent variable is 1 if the risky option was chosen and 0 otherwise. This allows for easily interpretable coefficients and we show in the appendix that the results hold with a logistic regression as well. The regressions include random effects at the individual level and confirm that our findings hold even when controlling for demographics.

Result 2: Clustered feedback closes the description-experience gap.

The right panel of Figure 1 shows the proportion of risky choices when participants obtained clustered feedback. Here, there are no observable differences between decisions made with and without a description ($F(1, 600) = 0.03, p = 0.87$). Moreover, comparing the left and right panels directly, we find no significant effects of feedback frequency on the proportion of risky choices when a description is present ($F(1, 609) = 0.38, p = 0.54$).

To show that the difference between the decisions with and without description is greater when feedback is provided immediately than when clustered feedback is provided, we have to estimate a difference-in-difference model. We do so using the following linear probability model:

$$Y = \beta_0 + \beta_1 D + \beta_2 C + \beta_3 (D \cdot C) + \epsilon$$

Where Y is choosing the risky option (estimated separately for the three lotteries), D is a dummy variable equal to 1 if description is present, and C is a dummy variable equal to 1 for clustered feedback. The estimates, with and without additional controls, are shown in Table 2. Note that the interaction term is significant in the medium and high variance lotteries, showing that clustered feedback closes the description-experience gap.

Table 2: Difference-in-difference analysis showing the description-experience gap is closed by clustered feedback.

	(6.25, 0.8; 0)	(10, 0.5; 0)	(25, 0.2; 0)
Description Present	0.010	0.155***	0.211***
	(0.039)	(0.041)	(0.038)
Clustered Feedback	0.034	0.158***	0.221***
	(0.039)	(0.041)	(0.039)
Description x Clustered Feedback	−0.051	−0.125*	−0.188***
	(0.055)	(0.058)	(0.055)
Constant	0.592***	0.359***	0.216***
	(0.028)	(0.029)	(0.027)
Observations	45,210	44,110	44,440
Note:	*p<0.05; **p<0.01; ***p<0.001		

Result 3: The description-experience gap emerges early and proportions of risky choices stabilize quickly.

Since participants make 110 repeated and incentivized decisions, we can also examine their behavior over time. The time series of the proportion of risky choices is shown in Figure 2, along with a linear regression line. In the immediate feedback setting (left column), the time series with and without description look the same in the low variance condition, but a gap appears within the first ten trials in the medium and (to a greater extent) the high variance conditions. After an initial drop in the proportion of risky choices with immediate feedback, the lines continue approximately parallel to one another. In the clustered feedback condition, we do not see such an initial drop in the early trials.

Table 3 and Table 4 explore time trends with immediate and clustered feedback using linear

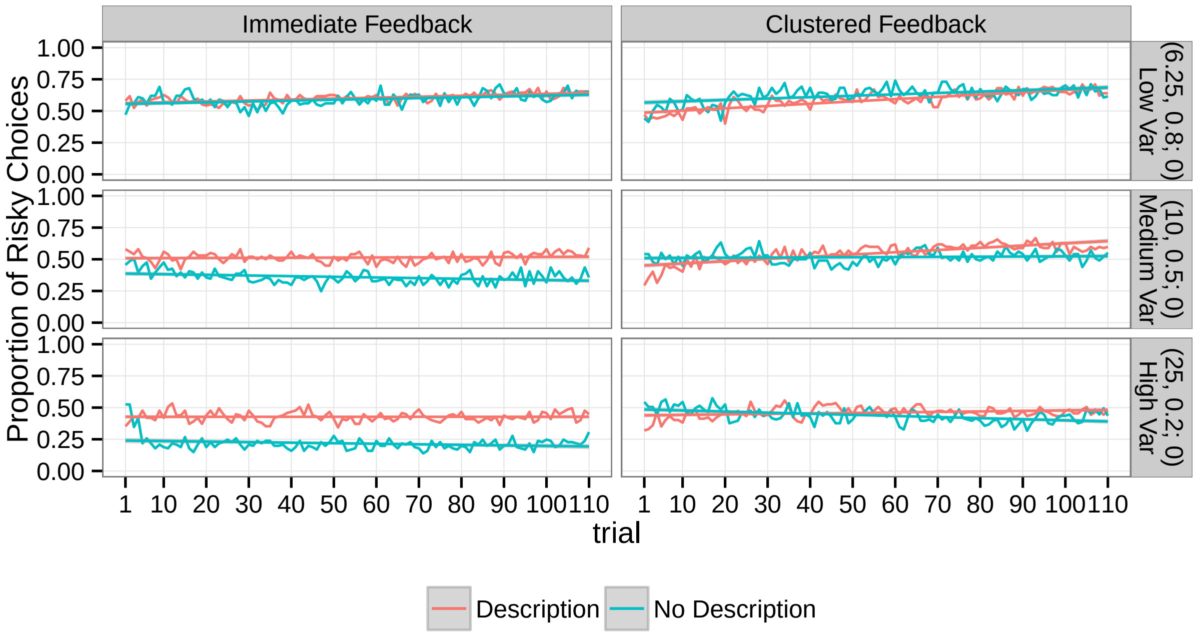

Figure 2: The proportion of risky choices in each trial with a linear regression line and 95% confidence region superimposed. The trend over time varies depending on the lottery and the feedback frequency.

probability models. Looking first at Table 3, we see that the main effect of description persists in the medium and high variance lotteries when controlling for a time trend. In those lotteries, we moreover see that the proportion of risky choices decreases in the absence of a description, but remains unchanged with description. Although the coefficients on the trial parameter are small, the magnitude becomes economically significant over the course of 110 trials. A difference of 0.001 per trial means the proportion of risky choices increases by 12 percentage points from the first to the last trial. Once we include an additional dummy for the first block (the first 10 trials), we see that the time trend is entirely driven by those first 10 decisions, further supporting our hypothesis that early decisions are driving the description-experience gap.

Table 4 shows the same regressions for the clustered feedback conditions. Interestingly, the proportion of risky choices is increasing when a description is presence across all lotteries (the description and trial interaction). Even though the outcome feedback provides no instrumental information, it appears that the risky option becomes more appealing over time – something we did not see with immediate feedback. Moreover, this effect persists even with a dummy for the first 10 trials (the first block) during which no feedback is obtained. In contrast to the immediate feedback condition, we see that the risky option is initially chosen more frequently in the absence of a description. Over time, however, this difference vanishes.

Table 3: Difference-in-difference analysis comparing the time trends between the presence and absence of a description with immediate feedback.

	(6.25, 0.8; 0)	(6.25, 0.8; 0)	(10, 0.5; 0)	(10, 0.5; 0)	(25, 0.2; 0)	(25, 0.2; 0)
Description Present	0.006	0.009	0.121**	0.141**	0.187***	0.226***
	(0.042)	(0.042)	(0.043)	(0.044)	(0.041)	(0.041)
Trial	0.001***	0.001***	−0.001***	−0.0001	−0.0004***	0.00003
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Description x Trial	0.0001	0.00004	0.001***	0.0004	0.0004**	−0.0001
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)	(0.0002)
First Block		0.039*		0.083***		0.102***
		(0.015)		(0.015)		(0.014)
Description x First Block		−0.010		−0.060**		−0.113***
		(0.021)		(0.021)		(0.019)
Constant	0.555***	0.542***	0.387***	0.359***	0.240***	0.205***
	(0.030)	(0.030)	(0.031)	(0.031)	(0.029)	(0.029)
Observations	22,770	22,770	22,110	22,110	22,660	22,660

Note: *p<0.05; **p<0.01; ***p<0.001

Table 4: Difference-in-difference analysis comparing the time trends between the presence and absence of a description with clustered feedback.

	(6.25, 0.8; 0)	(6.25, 0.8; 0)	(10, 0.5; 0)	(10, 0.5; 0)	(25, 0.2; 0)	(25, 0.2; 0)
Description Present	-0.080^{*}	-0.092^{*}	-0.060	-0.028	-0.047	-0.002
	(0.039)	(0.040)	(0.040)	(0.041)	(0.039)	(0.039)
Trial	0.001^{***}	0.001^{***}	0.0001	0.0001	-0.001^{***}	-0.001^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Description x Trial	0.001^{***}	0.001^{***}	0.002^{***}	0.001^{***}	0.001^{***}	0.001^{**}
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
First Block		-0.090^{***}		-0.010		0.051^{**}
		(0.015)		(0.016)		(0.016)
Description x First Block		0.033		-0.092^{***}		-0.133^{***}
		(0.022)		(0.023)		(0.023)
Constant	0.565^{***}	0.596^{***}	0.509^{***}	0.512^{***}	0.486^{***}	0.469^{***}
	(0.027)	(0.028)	(0.029)	(0.029)	(0.027)	(0.028)
Observations	22,440	22,440	22,000	22,000	21,780	21,780

Note: $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

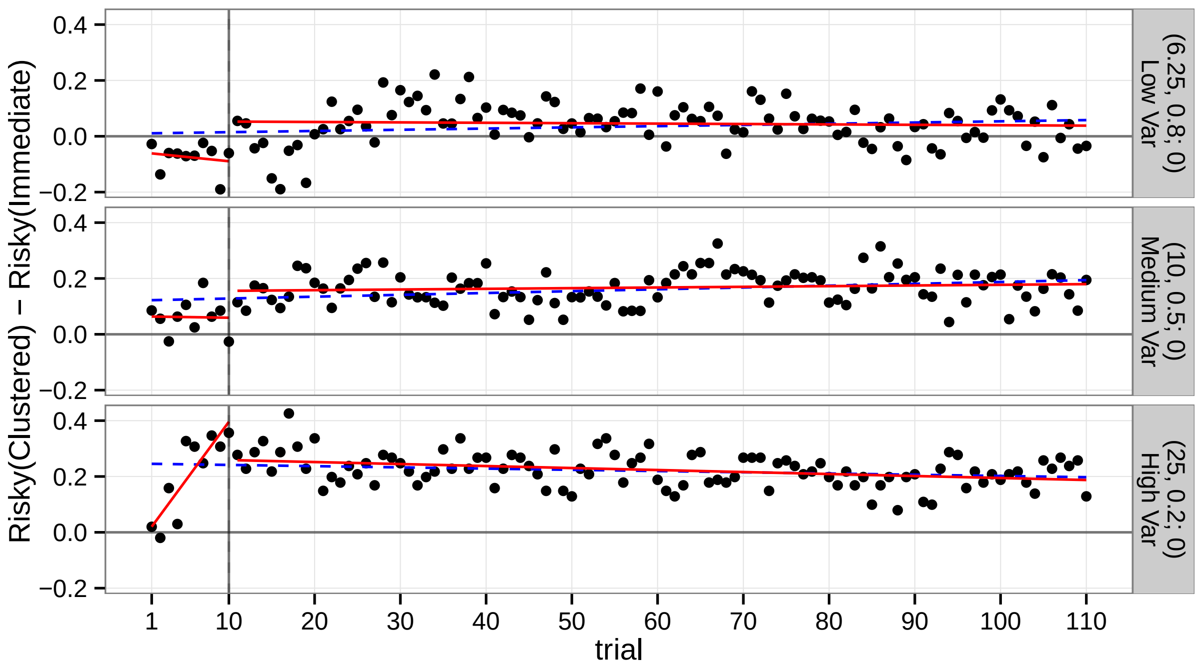

Figure 3: Difference in the proportion of risky choices between the clustered and immediate feedback conditions (without description) over time. A value above zero means the risky option is more frequently chosen with clustered feedback. Also shown is a single regression line (dashed blue) as well as two separate regressions for the first 10 trials and the next 100 trials (red).

We now examine a potential mechanism for why clustered feedback closes the description-experience gap. With no description, participants in the clustered feedback condition do not obtain any information about the two options until after the first 10 trials. It is only for their 11th decision that they know the outcomes of the decisions they have previously made. We can thus examine the difference in risky choices between the immediate and clustered feedback conditions and look for a discontinuity between the 10th and 11th trial. Figure 3 plots the difference in proportion of risky choices between the clustered and immediate feedback in each trial of the no description condition. A value above zero indicates that those with clustered feedback are more likely to choose the risky option than those with immediate feedback. Also plotted is a linear regression line for the whole sample (dashed blue line) and two separate regressions with a break point between trials 10 and 11 (solid red lines). If there were no systematic difference between the first ten trials and the remaining 100, we would expect the red lines to overlap fully with the blue line. When the red lines have different slopes and there is a discontinuity between the two lines, this suggests a systematic difference before and after the breakpoint (trial 10).

In the low variance condition (top panels), we see that those in the immediate feedback condition select the risky option more often than those in the clustered feedback condition before trial 10 (the first red line segment is below 0 and decreasing). Once those in the clustered feedback condition observe the past outcomes (after trial 10), they switch to the risky option and pick it at a consistently (and slightly) higher rate than those in the immediate feedback condition for the rest of the task (the second red line segment is flat and above 0). In the medium variance condition (middle panel), we see a similar pattern, but both line segments are above 0, indicating that the risky option is chosen more frequently throughout with clustered feedback than with immediate feedback. Finally, in the high variance condition (bottom panel), we see that participants with clustered feedback are much more likely (nearly 40 percentage points) to pick the risky option in trial 10 (the last before feedback is received) and, for the rest of the task, choose the risky option approximately 20 percentage points more frequently. These visual findings are confirmed by Chow tests, which show that for each lottery, the differences in proportion of risky choices between the conditions with clustered and with immediate feedback differ significantly in the first 10 trials compared to the remaining 100 ($p < 0.001$ for all Chow tests).[3]

Result 4: Repeated Decisions With Descriptions Differ From One-Time Decisions

Our implementation of a setting with description differs from much of the literature in that the decision is made repeatedly and that feedback about outcomes is observed. We showed that the risky option is chosen more frequently in later trials, even though feedback provides no new instrumental information. This highlights a concern with previous work on the description-experience gap, which confounded the presence or absence of a description with whether a decision was made

[3]We explore additional breakpoints in Appendix A.

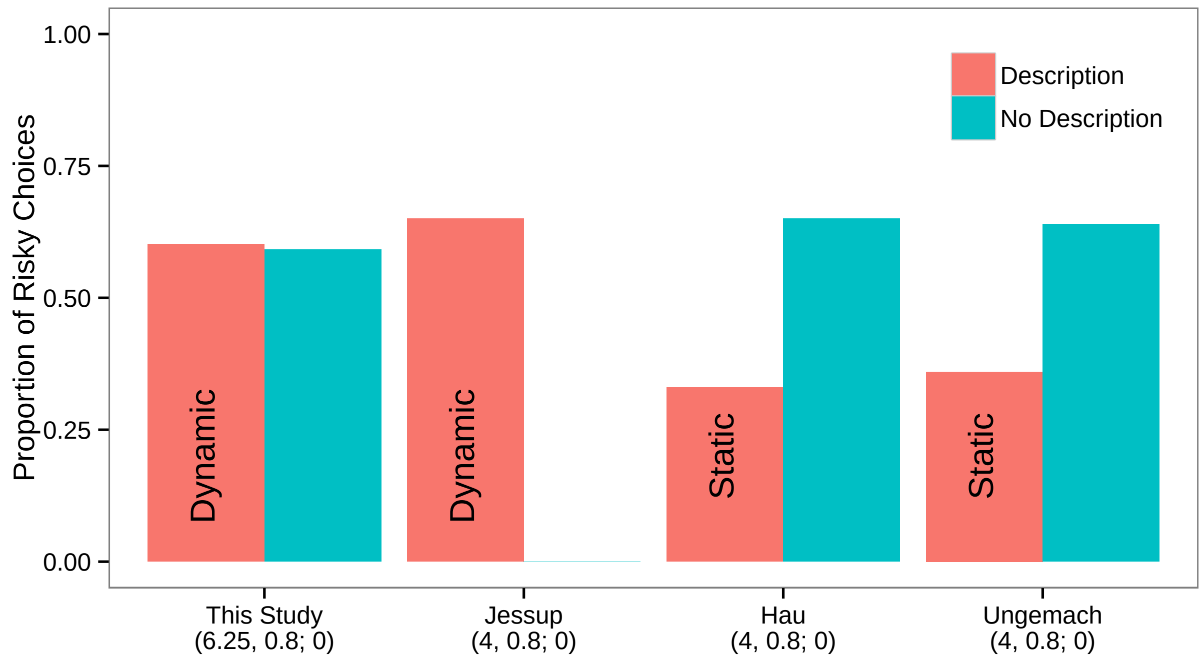

Figure 4: Comparison to other studies on repeated choice in the presence and absence of description.

once or repeatedly, respectively. In Figure 4 we compare the proportion of risky choices participants in the low variance lottery with immediate feedback made against three other studies using similar lotteries (Hau et al. 2008, Ungemach et al. 2009, Jessup et al. 2008).[4] Note that Jessup et al. use repeated decisions with description, as do we, but do not include a condition without description.

While the conditions without a description appear the same in all three studies that include it, there is a substantial gap between repeated decisions with immediate feedback (dynamic description) in this study and in that of Jessup et al. and one-time decisions with description (static) used by Hau et al. and Ungemach et al.. In fact, comparing the dynamic description condition with the no description conditions suggests that merely having participants make decisions with description repeatedly is sufficient to close the description-experience gap. This may explain why we did not find this gap in our low-variance condition.

[4]We thank Orestis Kopsacheilis for suggesting this comparison. The percentage from Jessup et al. (2008) is estimated based on Figure 1 as they do not report the proportion of risky choices numerically.

5 Discussion

Previous work has established that decision makers choose differently if they have a description of available options than if they have to learn about them from observing outcomes of repeated choices. We introduce a novel clustered feedback mechanism in which decision makers repeatedly choose between two options without obtaining any feedback before learning about the individual realizations of their choices. We find that when no description is available, those who receive clustered feedback choose a risky option more often than those who receive immediate feedback and make the same choices as participants who have a description of the options available. Our results show that the effect is driven by greater exploration in the periods prior to obtaining any outcome feedback: those who learn of outcomes immediately quickly give up on the risky option and thus fail to learn its true properties.

Our experiment also provides insights into the kind of tasks that are likely to produce a description-experience gap when feedback is provided immediately. In a setting that is most likely to produce a small gain compared to the safe option, we find no significant difference between the presence or absence of a description. However, as the low outcome becomes more common and the high outcome becomes rarer, an increasing difference between the two settings emerges.

These findings apply in many real world settings, where decision makers do not have available an objective description of their options' payoffs. When they have to draw on their experience, insufficient exploration of available options may lead them to make ex post suboptimal choices. A job seeker, for example, may give up too quickly on finding employment in a particular line of work after a small number of applications. Our results suggest that delaying feedback can encourage more exploration and help decision makers obtain a more accurate perception of the true state. That can help them avoid costly mistakes from forgoing more rewarding options.

Although our focus in this paper has been on individuals making repeated decisions, we can map the same idea to domains where many people make decisions after observing feedback from others. Consider a decision to buy a product online after reading reviews from other consumers. An item for which 995 out of 1,000 reviews are favorable may be a great purchase. However, suppose the first 5 consumers had been displeased, leaving the product with 5 out of 5 unfavorable reviews. Other prospective buyers would reasonably have avoided this purchase, never learning how satisfied they would have been with this item. There would not be another 995 reviewers to leave positive feedback. Such a website may help consumers by not disclosing ratings based on very few reviews. This may hold even more so when consumers can rank the available options by reviews, such that a single bad review may condemn a product to the last pages, making it unlikely that anyone would discover it in their search results.

References

Barron, Greg, Ido Erev. 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making* **16**(3) 215–233.

Benartzi, Shlomo, Richard H. Thaler. 1995. Myopic Loss Aversion and the Equity Premium Puzzle. *The Quarterly Journal of Economics* **110**(1) 73–92.

Fellner, Gerlinde, Matthias Sutter. 2009. Causes, Consequences, and Cures of Myopic Loss AversionAn Experimental Investigation*. *The Economic Journal* **119**(537) 900–916.

Gneezy, Uri, Jan Potters. 1997. An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics* **112**(2) 631–645.

Gonzalez, Cleotilde, Varun Dutt. 2011. Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review* **118**(4) 523–551.

Hau, Robin, Timothy J. Pleskac, Jürgen Kiefer, Ralph Hertwig. 2008. The descriptionexperience gap in risky choice: the role of sample size and experienced probabilities. *Journal of Behavioral Decision Making* **21**(5) 493–518.

Hertwig, Ralph, Greg Barron, Elke U. Weber, Ido Erev. 2004. Decisions from experience and the effect of rare events in risky choice. *Psychological Science* **15**(8) 534–539.

Hertwig, Ralph, Ido Erev. 2009. The descriptionexperience gap in risky choice. *Trends in Cognitive Sciences* **13**(12) 517–523.

Hopfensitz, Astrid. 2009. Previous Outcomes and Reference Dependence: A Meta Study of Repeated Investment Tasks with Restricted Feedback. TSE Working Paper 09-087, Toulouse School of Economics (TSE).

Imas, Alex. forthcoming. The Realization Effect: Risk-Taking after Realized versus Paper Losses. *American Economic Review* .

Jessup, Ryan K., Anthony J. Bishara, Jerome R. Busemeyer. 2008. Feedback Produces Divergence From Prospect Theory in Descriptive Choice. *Psychological Science* **19**(10) 1015–1022.

Langer, Thomas, Martin Weber. 2005. Myopic prospect theory vs. myopic loss aversion: how general is the phenomenon? *Journal of Economic Behavior & Organization* **56**(1) 25–38.

Lejarraga, Toms, Cleotilde Gonzalez. 2011. Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes* **116**(2) 286–295.

Mehra, Rajnish, Edward C. Prescott. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* **15**(2) 145–161.

Read, Daniel, George Loewenstein, Matthew Rabin, Gideon Keren, David Laibson. 1999. Choice Bracketing. Baruch Fischhoff, Charles F. Manski, eds., *Elicitation of Preferences*. Springer Netherlands, 171–202.

Samuelson, Paul A. 1963. Risk and uncertainty-a fallacy of large numbers. *Scientia* **98**(612) 108–&.

Tversky, Amos, Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* **5**(4) 297–323.

Ungemach, Christoph, Nick Chater, Neil Stewart. 2009. Are Probabilities Overweighted or Underweighted When Rare Outcomes Are Experienced (Rarely)? *Psychological Science* **20**(4) 473–479.

A Supplemental Analysis

A.1 Logistic Regression on Probability of Choosing the Risky Option

Table 5: Logit models on the probability of choosing the risky option.

	(6.25, 0.8; 0)	(6.25, 0.8; 0)	(10, 0.5; 0)	(10, 0.5; 0)	(25, 0.2; 0)	(25, 0.2; 0)
No Description	−0.509	−0.501	-1.231^{**}	-1.342^{***}	-1.772^{***}	-1.843^{***}
	(0.357)	(0.357)	(0.377)	(0.373)	(0.361)	(0.361)
Male		0.332		0.860^{*}		0.343
		(0.368)		(0.379)		(0.364)
Age		−0.005		0.021		-0.035^{*}
		(0.017)		(0.017)		(0.018)
Constant	0.997^{***}	0.809^{*}	0.371	−0.056	−0.248	−0.419
	(0.252)	(0.332)	(0.269)	(0.326)	(0.254)	(0.314)
Observations	22,770	22,770	22,110	22,110	22,660	22,660

Note: $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

A.2 Other Break Points

We can further plot the F-statistic for the Chow test as a function of the break point. Figure 5 does this for every break point between 5 and 105 and for each lottery in the environment without description. We see that in the high variance lottery, the optimal break point (where the F-statistic reaches its maximum value) is indeed after 10 trials. In the low variance condition, there is a small peak after 10 trials, but a much larger one after 20. It appears that people in the clustered feedback condition are willing to explore for a second block of feedback before deciding on an option. Finally, in the medium variance condition, we see that any breakpoint between 10 and 30 could model the data substantially better than a single regression line, and do so better than higher breakpoints. Together with Figure 3, this provides strong evidence that clustering feedback influences search behavior early in the process and promotes exploration, which in turn leads to a greater proportion of decision makers choosing the ex post better option.

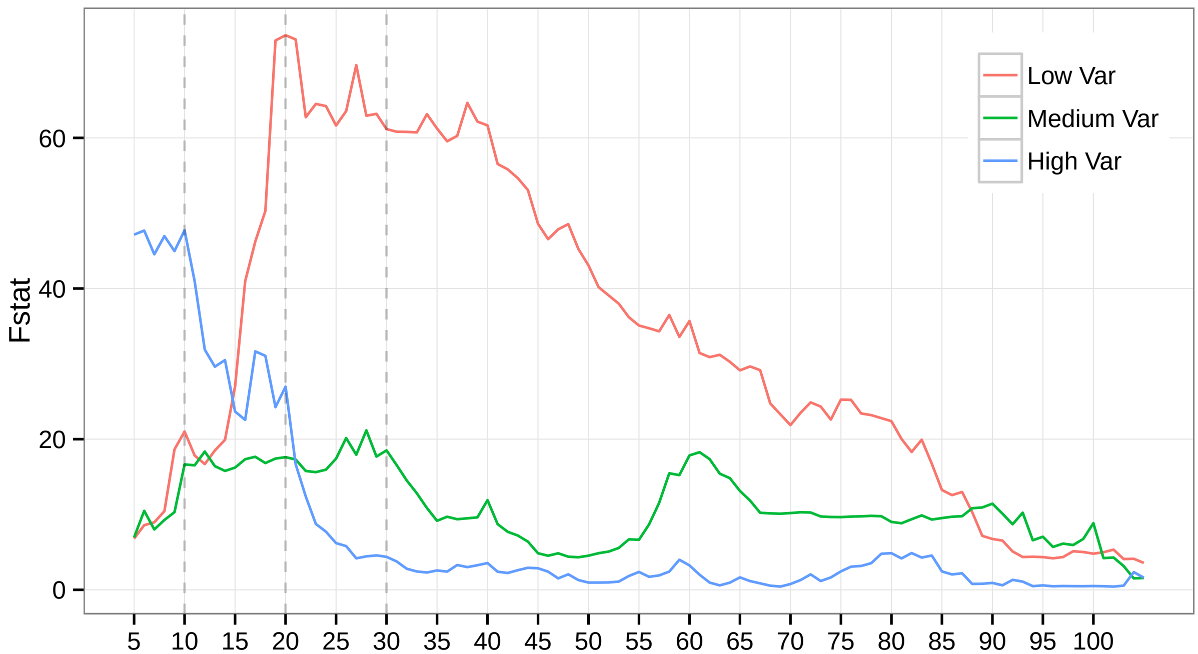

Figure 5: F-Statistic of the Chow test for breakpoints in trials 5 through 105. Dashed gray lines show trials 10, 20, and 30.

B Instructions

In this experiment, you will make repeated choices between two options. Each choice you make will result in a payoff. You will make 110 choices involving the same two options. Your payment will be the SUM of all the payoffs of your 110 choices divided by 400. This will be paid to you as a bonus payment within 24 hours of completing the experiment.

Option A returns $`high` 20% of the time and $0 80% of the time. The payoff is randomly determined in each choice and it is independent of your previous choices and previous payoffs.

Option B returns $4 every time you click it. [only in conditions with description]

The two options will remain the same for the entire task. However, the amounts that you receive may be determined by chance. [only in conditions without description]

You will see the payoff earned in each of your choices for a brief period. You will also see the SUM of the payoffs from all of your choices at the end. At that point, your bonus payment will be calculated. [only in conditions with immediate feedback]

You will see the payoffs earned in the previous n choices after every 10 trials. You will also see the SUM of the payoffs from all of your choices at the end. At that point, your bonus payment will be calculated. [only in conditions with clustered feedback]

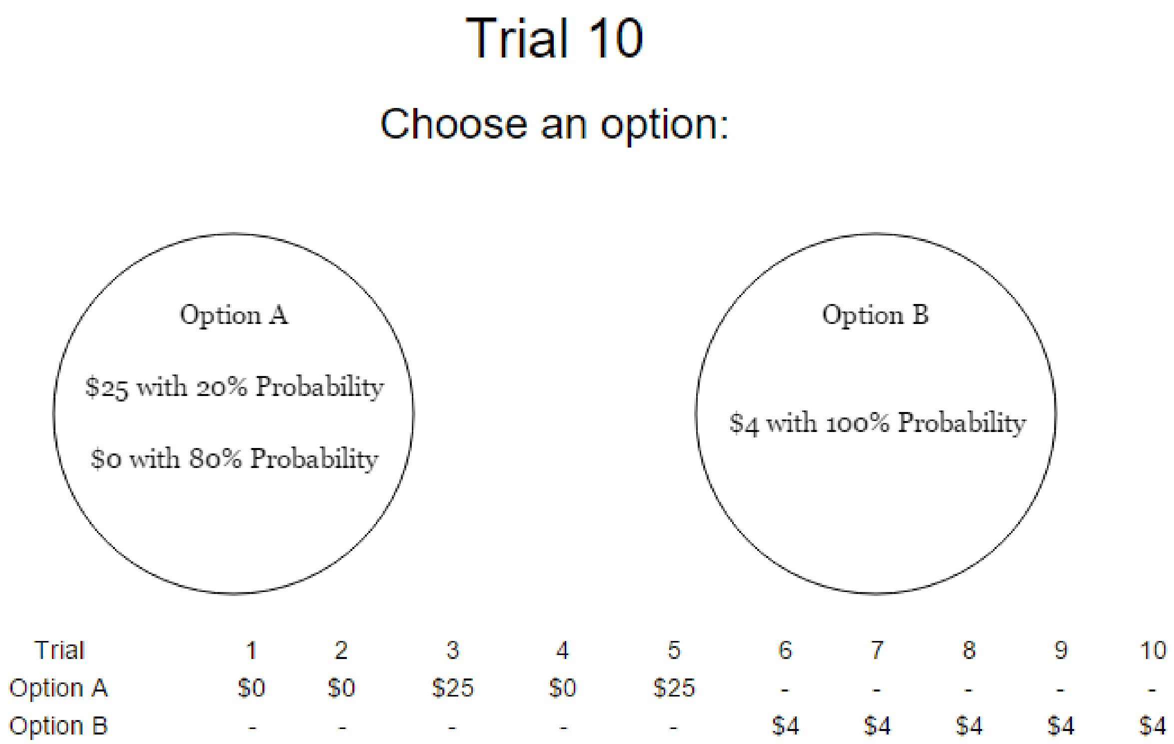

Figure 6: The interface as seen by participants in the clustered feedback conditions with description. After 10 trials, the decision maker observe the 10 realizations from her prior choices for 7 seconds before continuing. Conditions without description omit the lines displaying the payoffs and associated probabilities (leaving only "Option A" and "Option B"), but are otherwise identical. In the immediate feedback condition, participants see only the previous realization for 0.7 seconds. The order of the options (i.e. their corresponding labels) is randomly determined for each participant and remains fixed throughout the task.

[[Task—Figure 6]]

How often do you think you chose Option A and Option B? The two numbers have to add up to 110.

Option A [ ]

Option B [ ]

How often do you think you got a `$high` payoff and a $0 payoff? The two numbers have to add up to `Option X` (whichever was associated with the risky payoff).

`High Payoff` [ ]

$0 [ ]

What is your gender?

[Male / Female]

What is your age?

[ ]

What is the highest level of education that you have completed?

[Some high school / High school / Some college / Associate's degree / Bachelor's degree / Master's degree / Professional or doctoral degree]

What is your monthly household income?

[ ]

Do you have any of the following accounts? Please select all that apply.

[Checking account / Savings account / Credit card / Investment acount / Retirement account (401(k), IRA, etc.)]

Which political party do you most strongly support and/or identify with?

[Democrat / Republican / Independent / Libertarian / Green / Socialist/ Tea Party / Other / None]

How often is your monthly credit card bill paid in full (i.e. no debt remaining)?

[100% of the time (always paid in full) / 75% of the time / 50% of the time / 25% of the time / 0% of the time (never paid in full)]

How often do you get paid?

[Weekly / Bi-weekly (every 2 weeks) / Monthly / It varies / Other]

How often do you deposit money into your person savings account?

[Weekly / Bi-weekly (every 2 weeks) / Monthly / Quarterly (every three months) / Yearly / Never]

How often do you deposit money into a retirement account?

[Weekly / Bi-weekly (every 2 weeks) / Monthly / Quarterly (every three months) / Yearly / Never]

How often do you play the lottery?

[Daily / Weekly / Monthly / Yearly / Rarely (only a few times in my life) / I've never done it]

The SUM of the outcomes from your choices is: `sumPayoffs`

You have earned the following bonus for your participation:

[`sumPayoffs`/400]

This amount will be sent to you as an mTurk bonus payment within 24 hours of completing this study.