Decision Changes and Risk: An Experiment

Mitch Addler

University Of Arizona

Economics Department

3rd Year Paper Conference Draft

April 9, 2018

Abstract

I conduct an individual choice experiment in which subjects repeatedly decide on how to allocate funds between a safe and risky asset. The goal is to understand how preferences are discovered regarding risk and if the preference discovery process can be aided. Subjects are randomly given the opportunity to change their decision every round before outcomes are realized. If offering this option triggers reflection or simulates experience, preferences will be discovered more quickly. I compare results for changed decisions to unchanged decisions to see if offering subjects the opportunity to change is welfare-improving.[1]

[1]I am grateful for comments and suggestions from my peers and the faculty in the Economics Department at the University of Arizona. I am especially grateful for helpful suggestions and comments from my advisor, Charles Noussair, and Julian Romero. I would also like to thank the Economic Science Lab at the University of Arizona for funding my experiment. If the reader has any questions or comments, I can be reached at maddler@email.arizona.edu. All mistakes in this paper are my own. This draft is complete in the sense that I present results, but it is very much a preliminary draft, and I plan to significantly add to/change sections of the paper in the near future. I welcome and appreciate all comments.

1 Introduction

Important decisions are not made often. Some of the most significant financial decisions are made only a handful of times in the average person's life. These decisions include buying a home, buying a car, choosing how to invest for retirement, and purchasing insurance. By nature, when people make these decisions they do not have much experience in the decision environment. Lack of experience can lead to mistakes even for the most rational and intelligent among us. Even a person who perfectly understands a decision environment: the benefits and costs of every option, all contingencies and strategies, and all potential outcomes, can end up making a decision that is suboptimal ex post because the person did not fully understand their preferences in that specific decision environment.

The scenario mentioned in the opening paragraph isn't a problem if we make the assumption that all economic agents have stable preferences and are aware of said preferences. Then, the natural economic analysis can occur where a rational agent optimizes with respect to her objective function given some constraint(s). Evidence has been presented, however, that indicates that not all economic agents have stable preferences of which they are fully aware. This evidence has given rise to two schools of thought regarding preferences. One argument, often put forth by psychologists is that preferences are constructed by agents and are highly dependent on context. The other argument, often presented by economists, is that individuals have stable, fundamental preferences, but those preferences have to be discovered through experience and reflection. I discuss the details of this discovered preference argument in Section 2.1.

If we operate under the assumption that preferences are a primitive that needs to be discovered by agents in a decision environment, it becomes clear that decisions made in environments where an agent has little experience may not be optimal given her true stable preferences. The agent is, as an assumption, optimizing, but she may not be optimizing with respect to the right objective function. One natural question to ask is: How do we improve welfare in these situations?

In decision environments like buying a home or deciding on a retirement investment strategy, there is typically a low level of experience. It is also prohibitively costly to acquire experience by making multiple decisions in that environment. So, if there was a way to simulate experience or help agents discover their preferences quickly, there could be welfare gains. One potential option which I test experimentally in this paper is to give subjects an opportunity to change a decision before the outcome of the initial decision is realized. This would be relatively costless to implement in different real-world decision environments. If this opportunity allowed agents to learn about their preferences, then they would optimize with respect to an objective function that is at least closer to their true objective function and likely see a welfare improvement.

In this paper I present an experiment to test the validity of this opportunity to change decisions as a means of speeding up or aiding in the preference discovery process. Subjects make a series of similar decisions and are randomly given the opportunity to change decisions after each round, before outcomes are presented. If this procedure can speed up the process of preference discover and allow subjects to reach stable preferences in this simple environment, the procedure may help with more complex decisions as well.

The remainder of the paper proceeds as follows. Section 2 discusses the motivation for the experiment and the impetus for the research in further detail. Section 3 discusses prior work, theoretical and experimental, that informs the research presented in this paper. Section 4 presents a model which can explain the preference discovery process in this decision environment and that may extend to other environments as well. Section 5 details the experimental design, Section 6 presents hypotheses regarding potential experimental results, Section 7 presents results and discusses the hypotheses, and Section 8 concludes.

2 Motivation

In this section, I describe motivation for the experiment and hypotheses. I begin with the theoretical impetus for the experiment in Section 2.1 and then discuss real world scenarios and welfare implications in Section 2.2.

2.1 The Discovered Preference Hypothesis

The choice of investments for a retirement fund, the purchase of real estate, and the purchase of insurance are all significant economic decisions which are not made frequently. Each of these decisions will likely be made a handful of times over the course of the average person's life. Despite their infrequency, these decisions have a large impact on life. Sometimes the effect is immediate, as is the case of the purchase of a new home. Other times the effect is delayed, in the case of insurance or investment for retirement.

Much of fundamental economic theory is predicated on the assumption that people have preferences over outcomes in all of the aforementioned scenarios. Without this fundamental assumption, much of the applicable economic analyses would be impossible. Utility maximization, for example, would not be applicable to the decision of whether to purchase a new home if there were no fundamental preferences over homes for the economic agent in question. Without the fundamental preferences, one could not represent those preferences with a utility function and perform the workhorse calculations to determine the optimal decision. It seems then that economists must agree that there must exist some fundamental, stable preference over choices/outcomes in a given context or framework. Without this assumption, economic analysis loses its bite.

If there must exist some stable, underlying preference over outcomes and actions in any given economic environment, then why do we see phenomena that seem to invalidate this assumption? Examples include the Allais paradox as well as other instances of "preference reversals" that have been studied thoroughly in economic experiments. These instances will be discussed further in Section 3. To reconcile the observed phenomena and the underlying need for existence of stable preferences, Plott's Discovered Preference Hypothesis (henceforth DPH) presents the idea that economic agents have some underlying preferences but they must discover those preferences when presented with an unfamiliar economic scenario [12]. The DPH does not attempt to model how preferences are explicitly discovered, but Plott argues that the discovery of preferences occurs through experience and reflection. Experience is obtained through making decisions and seeing outcomes within a given context. The more similar each decision is, the quicker the discovery of preferences for that particular scenario. Reflection refers to the way an economic agent integrates any feedback she receives from decisions. This feedback can be internal - how she feels about the action or the outcome, or external - how do other agents/players react to the decision. These two elements: experience and reflection, are typically present in any experiment that is not a one-round, one-shot game. In my experimental design, which I discuss in detail in Section 5, There are multiple rounds of similar decisions, so experience and reflection should be present, allowing subjects to discover, at least partially, their stable preferences for this particular decision environment.

In this experimental setting, I employ a procedure that I hypothesize will accelerate the preference discovery process, but this procedure does not fit nicely into either experience or reflection, as both require outcomes of decisions to inform the preference discovery process. The procedure I use is to randomly allow some subjects to change their initial decision in a given round of the experiment *before* the outcome of the decision is determined. The idea is that asking a subject if they would like to change a decision which they have just made, without seeing an outcome, will trigger reflection and introspection. If a subject is presented with the option to change a decision and her preferences are already stable, she has no incentive to change the decision. If a subject is still discovering her preferences over decisions and gets the opportunity to change, she is likely to exercise it if the option triggers some deeper thought about the situation and/or potential outcomes. The option to change a decision will randomly be offered at all rounds of the experiment, so I will be able to see if the procedure helps subjects discover their preferences quicker and if subjects exercise the option less as they learn more about the decision environment and their preferences.

2.2 Change is Costly

The major economic decisions mentioned in the opening paragraph of Section 2.1 do not occur very often, so it is likely that many economic agents facing such decisions do so without much prior experience. If we operate under the DPH, it is extremely likely that people making these major decisions like house purchases and investment portfolios have not discovered their stable preferences yet. This increases the chance that they make a decision that is not optimal given their stable preferences. This should not be surprising, as one would expect less experienced individuals to make more mistakes than experienced ones. There is a subtle difference between an agent who makes a decision that is suboptimal because she doesn't fully understand her preferences and one who makes a mistake due to inexperience. The latter could simply be someone who chooses a dominated option because of inattentiveness or lack of understanding. An agent who takes a decision but doesn't understand her preferences is still optimizing correctly, she is just optimizing the wrong objective function.

Both types of mistakes, those coming from optimizing with respect to the wrong objective function and those coming from not optimizing for other cognitive limitations, are important to understand. Furthermore, a social welfare optimizer like a government would want to try and eliminate these mistakes. The purpose of this paper is to understand the type of mistakes that come from optimizing the wrong objective function. One can imagine a scenario where a person researches many different homes and their values, examines market trends, and is confident that they home they want to purchase is a fair deal. They make the purchase. Only after moving in do they realize that they wanted a bigger house. They still optimized, but did not fully understand their preferences over home size so optimized with respect to the wrong objective function. Now, if there were no switching costs or transaction costs, the new homeowner could transition to a new house that suits her preferences without loss. However, change is costly. Searching for a new home would be costly time-wise, selling the house she just purchased would take time and effort, and there is no guarantee she would break even in the transactions involved in selling the home. Knowing that there are these transaction costs, a benevolent government might want to find a way to help people understand their true preferences over homes before people commit to a purchase. I do not wish to speculate on how this might be accomplished, but the experiment presented in this paper presents one possible procedure that could help speed up the process of preference discovery: offering the opportunity to change a decision before outcomes are realized and the decision becomes final.

Another scenario where the combination of transaction costs and not-fully-discovered preferences can result in welfare losses is in saving for retirement. If a person seeks to avoid paying fees by having a broker or advisor handle their investment portfolio decisions and opts to handle these decisions on her own, problems can arise. If she has not faced a decision like a portfolio allocation problem in the past, she likely has to discover her preferences for risk. If she decides on the portfolio and later learns that she has a preference for higher risk investments, she has to sell current holdings and buy new holdings, and that includes transaction costs. This is the type of scenario that the experiment in this paper seeks to mimic. If there is relatively costless way to speed up the preference discovery process, welfare may be improved. The procedure used in this experiment could be easily implemented in several different contexts and if results show that it does help the preference discovery process, policies that implement such procedures could be welfare-improving.

3 Existing Literature

In this section I present an overview of the existing literature to which this paper contributes. I first focus on the theoretical foundation in Section 3.1. In Section 3.2 I summarize several experiments that have been conducted regarding risk which have informed my experimental design.

3.1 Preference Discovery/Formation

In the discussion regarding an individual's fundamental preferences, there are two main schools of thought. In one school, there exist some fundamental preferences for a decision maker, though the decision maker may not fully understand those preferences at the outset. This is the position that most economists take. The notion that economic agents have existing, stable preference that must be discovered through experience is known as the Discovered Preference Hypothesis (Plott, 1996) [12]. Operating under this DPH is essential for economists to explain various "anomalies" regarding preferences that are observed in a laboratory environment while still being able to apply models of optimization to individual behavior. Plott argues that the preference discovery process can be divided into phases which explain some of the anomalous behavior observed in different contexts like the famous Allais Paradox from Allais 1953 [1]. The preference discovery process is divided into three stages according to Plott. The first stage is when an individual has little experience with a decision and the decision can exhibit myopia or a seeming lack of strategy. The second stage is when an individual becomes aware of the decision environment and begins to strategize. The third stage is when an individual incorporates the rationality and strategy of other agents in the same environment and understands how their decisions will interact with her own. In the context of the current paper, subjects will be exhibiting behavior that is indicative of the first and/or second stages of preference discovery. The third stage is not applicable because the decisions of others have no effect on an individual's payoffs in this experiment.

The second school of thought is that preferences are constructed by individuals as they gain experience in a given context. This argument is presented by Slovic et al. 1990 [13] and Payne et al. [11] and argues that there is a process by which preferences are formed. In this vein of thinking, preferences are not a primitive of a problem and are formed endogenously based on many factors, including the context of a problem. This argument has become popular in the psychology literature in recent decades. This school of thought makes it more difficult to explain the existence and observation of stable preferences, however. Attempts have been made by psychologists to explain how a subject could be constructing preferences, but when they are faced with a more complex situation, they simplify preferences so that they appear stable [2].

The argument between the two sides essentially boiled down to whether or not preferences can be taken as given when examining a choice problem. My current work does not make an attempt to argue one side versus another, as that would require a different experimental design. Instead, I am operating under the DPH and trying to better understand how preferences are discovered, as well as test a procedure that could help expedite the preference discovery process.

3.2 Experiments on Risk

Risk and subjects' attitudes towards it have been studied extensively in economic experiments. Various risk elicitation tasks have been developed to help experimenters gauge subjects' feelings and preferences towards risk. These tasks include the Bomb Risk Elicitation Task (BRET) (Crosetto and Filippin, 2013) [6], the Holt-Laury task (Holt and Laury, 2002) [10], and a method where subjects choose how much money to allocate between a safe and risky asset (Gneezy and Potters, 1997) [9]. All of these tasks vary in their complexity and method of revealing subjects' preferences for risk. More risk elicitation tasks exist, but these are three examples of those prominent in the literature on risk preferences.

The BRET consists of a decision problem where subjects must choose how many boxes on a screen to open. For each box they open, they receive a small payment. However, there is a "bomb" in one of the boxes and if the bomb is opened, the subjects lose their payment. In the Holt-Laury task, there is a series of binary decisions between lotteries presented to subjects. One set of lotteries is increasing in risk, so there is a "switching point" which provides a measure of subjects' risk aversion.

The final task, the task used in Gneezy and Potters (1997) [9], is the risk elicitation task I employ in this paper. The task consists of an allocation decision by subjects. Two goods/investments/assets are presented to subjects. One of them involves some amount of risk, i.e. a chance that payoff for that good is low/zero and a chance the payoff is high. The other good has a lower return but has no risk associated with it - the return has probability 1. This task is attractive for my experimental design for a few reasons, which I will discuss in Section 5.

These various risk elicitation tasks have been used extensively to examine different behavioral phenomena regarding risk. Charness and Gneezy (2012) examine differences in gender when it comes to taking risks in an investment environment [5]. Charness and Gneezy (2010) look at the different behavioral components which may affect risk aversion like myopic loss aversion, illusion of control, and ambiguity aversion [4]. Dohmen et al. (2010) examines the relationship between cognitive ability and risk aversion, but this field still has room for growth [7]. I aim to contribute to this literature by examining how risk preferences are discovered by experimental subjects and if the discovery process can be accelerated through interventions like giving subjects the opportunity to change decisions.

4 Modeling Preference Discovery

The goal of this paper is to explore whether giving subjects the opportunity to change a decision before revealing an outcome will help them learn about their risk preferences. Up until this point, I have not explicitly stated how I believe these risk preferences will be discovered by subjects. I am agnostic towards providing an explicit model of preference discovery/learning, but I would be remiss if I did not mention and describe a model that I think would fit well. Typically, thew literature on learning focuses on learning as it relates to strategies, not preferences. This could be troublesome if we were considering complicated preferences over social outcomes, perhaps, but in the simple context of the individual decision environment I use and describe in section 5, many of the learning models can be applied. If one recasts the individual choice problem as a game between a subject and chance or nature, one can think of a person's preferences over risk as preferences over strategies in this game, where subjects are trying to maximize their individual utility functions. I do not attempt to model the utility of subjects, as it is sufficient to see whether a stable strategy arises in the choice environment subjects face.

One specific model of learning can be directly applied to the individual choice environment I use in the experimental design. That model is the Experience-Weighted Attraction Model from Camerer and Ho (1999) [3]. This model combines the two previously separate models of leaning: learning from experience and belief-based learning into one more general setup. This is an attractive feature because in an individual choice environment involving risk, it is natural to assume that learning about preferences will come from direct experience: making decisions and seeing outcomes, as well as from subjects imagining what *would have* happened had they made a different decision in a given choice environment. The model captures the idea of the "law of actual effect" and the "law of simulated effect." Decisions that are taken are taken more often in the future if they cause positive outcomes (actual effect), and foregone decisions that would have yielded positive outcomes are more likely to be taken in the future (simulated effect).

The model is very rich and flexible parametrically. See Camerer and Ho (1999) for an explanation of all parameters and their intuitive meaning. Rather than restate the model here, I describe how it could be implemented in the context of this experiment. Subjects will make a series of decisions regarding risk. In each round, subjects view the outcome of their decision. This is where updating can occur. Subjects see their decision and the outcome, and they can update their preferences (or discover them) based on the value provided to them by the outcome of the decision they took, as well as the outcome that would have occurred, had they taken a different decision. This allows subjects to update the "attractions" of each option, which are values indicating how likely the subject is to take that decision again in the future. The attractions inform some sort of probabilistic choice rule, like one that follows a logit or probit form.

One difficulty in using such a model with this experiment is that the decision is different in each round. Parameters are changed and a strategy which would have dominated in one round may be inferior in the next, if a subject does not notice differences in parameters. So, the model of learning would have to be able to break down the problem into a series of attributes which describe a particular choice problem. This would then allow for a subject to gain an understanding of her optimal decision in a given choice problem as it relates to a number of factors. Such factors might include the expected value of an asset, the variance associated with the random outcome, the probability of success, etc. If one were to specify a decision rule in terms of the factors associated with a given situation, that is, if one could completely describe the choice environment as a function of several parameters, one could directly apply multiple learning models.

If one does not wish to explicitly model the learning process of subjects, a very simple model of utility which incorporates random shocks for different decisions could suffice. Suppose subject i has a utility function, $u_i(\cdot)$, that represents her true, stable preferences over some set of alternatives. In the context of this paper, those alternatives would be investment portfolios or endowment allocations. However, it is natural to think that a subject is inexperienced in this particular environment, so she may need to discover her true preferences. She has some sense of what they are. That is, she sees some realization of a utility function, $\tilde{u}_i(\cdot)$, which is close to her true utility function. As she gains experience in the choice environment, the function $\tilde{u}_i(\cdot)$ converges to $u_i(\cdot)$. One very simple example of such functions would be $u_i(x,r) = x - v_i(r)$ and $\tilde{u}_i(x,r) = x - \tilde{v}_i(r,\epsilon)$ where x is the monetary payoff from a decision, r is some measure of risk, $v_i(r,\epsilon)$ is how the subject reacts to the risk, and ϵ_i is a random variable with mean 0 and variance σ_i^2. Using this very simple representation of utility in a risky decision environment, one can model learning by simply arguing that, with experience, $\tilde{v}_i(r,\epsilon)$ converges to $v_i(r)$. One might impose functional form assumptions on $v(\cdot)$ and ϵ so that this converges with probability 1 after some level of experience. Again, This example of a specific utility function was meant to show how one could model learning in a choice environment without explicitly modeling the mechanisms which guide the learning.

5 Experimental Design

Subjects participated in a computerized individual choice experiment which was coded and administered using Z-Tree [8]. The experiment consisted of 10 rounds of decisions. Each decision was identical in structure, but parameters changed from round to round. Subjects were paid according to their decision and the outcome of one randomly selected round. The structure of one round was as follows.

Subjects were endowed with 100 ECU. The conversion rate was 10 ECU = 1 USD. In a decision round, the subjects were presented with information about two goods: Good A and Good B. This information included a price per unit, a value, and a probability. The price per unit for both goods was 1 ECU and this price did not vary across rounds. The value was the per unit monetary value of each good. The value of Good A and Good B changed across rounds. The probability was the probability that the good was actually worth its value. This probability was always 1 for Good A, making it a riskless choice, while the probability for Good B was always less than 1. Purchasing units of Good B always entailed some level of risk.

Subjects were asked to consider the information presented to them and decide how to allocate their 100 ECU between the two goods. This can be thought of as a portfolio allocation problem. There is a "safe" asset, Good A with a definite return, and a "risky" asset, Good B, which has higher return but a risk of losing the investment. Subjects could choose any allocation they desired, as long as they spent all 100 ECU acquiring their portfolio and purchased only whole units of Good A and B. See Appendix A for a complete set of experimental instructions.

Table 1 shows the parameters for each decision made by subjects. The order of these decisions remained the same for all subjects, to avoid confounding order effects across subjects. The columns titled "EV maximizing?" indicate which god dominates in each decision in terms of expected value. This indicates that a risk neutral subject should invest fully in the indicated good in each round. These parameters were chosen with a focus on making the decision nontrivial for subjects. I wanted it to require thought on the part of subjects. The difficulty of the decision coupled with unfamiliarity with the decision environment would make ideal conditions for preference discovery/learning to occur.

In any given round, after subjects decided on how to allocate their endowment between the two goods, there was a chance that they would be given the opportunity to change their decision. The

Decision Parameters

Round	EV maximizing?	Good A value	Good B value	Good B Probability	Good B EV	EV Maximizing?
1	Yes	1.2	2.7	0.4	1.08	No
2	No	0.8	2.3	0.6	1.38	Yes
3	No	1	2.5	0.5	1.25	Yes
4	No	0.6	2.1	0.7	1.47	Yes
5	Yes	1.4	2.9	0.3	0.87	No
6	Yes	1.2	2.9	0.3	0.87	No
7	No	0.8	2.1	0.7	1.47	Yes
8	No	1	2.5	0.5	1.25	Yes
9	No	0.6	2.3	0.6	1.38	Yes
10	Yes	1.4	2.7	0.4	1.08	No

Table 1: The parameters for each of the 10 rounds of the experiment, along with an indication of which good dominates in terms of expected value.

probability that a subject would be given this opportunity was 0.5. This probability was constant across all rounds, and was not communicated to subjects. Subjects were simply told that there is a chance they will be given the opportunity to change their decision. This adds a level of ambiguity to the experiment, but was done in order to incentivize subjects to treat each decision as if it were their only decision, while still allowing for the potential to change a decision. This helps mitigate any frivolous behavior in a decision round by avoiding a scenario where subject knows, or thinks she knows, that she will have the opportunity to change her decision. It stops subjects from acting without thought in any given round because they never know whether or not their initial decision will count.

Offering subjects the opportunity to change a decision before the outcome is realized is a procedure that has not been used often in economic experiments. I employ the procedure in this experiment in order to see whether or not this opportunity helps subjects discover their risk preferences more quickly. Traditional theory would say that, if preferences are stable, a decision like the one subjects are making in this experiment should be optimal. Without any additional information, there should be no reason for a subject to alter their decision when given the opportunity to do so. I should note that there is the potential concern that subjects, when presented with the opportunity to change their decision, will seize the opportunity only because they think that's what the experimenter wants. This experimental design allows for me to test for this effect by comparing instances of decision changes in early periods as opposed to later ones. If my hypotheses are correct, subjects who are given the opportunity to change their decision in early rounds will exercise the option more so than those in later rounds, because preferences over risk in this environment are more stable at the end than the beginning. If I see similar rates of changed decisions when comparing rounds 1 and 2 to the rates in rounds 9 and 10, I can see whether the driving force is the experimenter demand effect or the discovery of risk preferences.

After subjects have made their initial decision and before subjects are possibly given the opportunity to change that decision, I impose a 30 second waiting screen. This is done to ensure that subjects are not simply speeding through all decisions in order to complete the experiment and receive payment as soon as possible. I also want to encourage thought and reflection about the decisions they have made, as this can help subjects discover their preferences in this experimental environment. Once the waiting period is complete, subjects are potentially presented with the opportunity to change their decision. They are reminded of their decision as well as all of the other information initially provided. They then enter their final decision and proceed to the next screen, where they are informed about the outcome of the risk associated with Good B and reminded of their final decision. Subjects are also informed of their ECU earnings for the round, should the round be randomly selected to determine their payment.

Once the outcome of a round is revealed, subjects can continue on to the next decision. They are presented with different parameters for Good A and Good B and are asked to make another decision. This process repeats for a total of 10 rounds. Then, a final screen presents subjects with which round was randomly selected to determine payment as well as their earnings for that round. Finally, a questionnaire was distributed to all subjects which asked them about their gender, age, and ethnicity, as well as the following questions:

- When you were given the opportunity to change a decision, were you pleased?
- Did being asked if you would like to change a decision affect how you thought about later rounds?
- Did you become more confident in your decisions as the experiment progressed?

I chose to ask such questions because I want to know if subjects became aware of any learning they did throughout the course of the experiment. I am also interested if subjects knew that asking them if they would like change their decision changed/didn't change their thought process.

6 Hypotheses

In this section, I present several hypotheses and briefly explain the intuition and foundation for each.

Hypothesis 1 (H1): *Offering subjects the opportunity to change their decision will speed up the risk preference discovery process.*

The central procedural element of this paper is that subjects will get the opportunity to change their decision without seeing the outcome of their initial decision. If preferences are stable, subjects would not exercise the option. If preferences are still being discovered, subjects may decide to change their decision if the opportunity to change prompted some reflection about their preferences. If this is the case, subjects who are given the opportunity to change their decision should discover their stable preference quicker than subjects who do not get the opportunity, or who get the opportunity less often.

Hypothesis 2 (H2): *There will be a negative correlation between instances of decision changes and experience in the decision environment.*

This hypothesis is formed from the DPH, discussed at length in Sections 3.1 and 2.1. Subjects who are still in the process of discovering their preferences will desire to change their decision more often when they have little experience in the decision environment.

Hypothesis 3 (H3): *Monetary payoffs will, on average, be higher when subjects change decisions than when subjects do not.*

This intuition behind this hypothesis is that the opportunity to change a decision will trigger reflection and/or deeper thought about the choice. This reflection could lead subjects in the direction of expected value maximization and receiving higher monetary payoffs (on average).

Hypothesis 4 (H4): *Subjects will invest less in the risky asset after observing a bad outcome of risk in the preceding round.*

This hypothesis is based on myopic loss aversion. Subjects are able to evaluate their decisions every period, and they see the outcomes of their choices. It is natural to suspect that subjects who see a bad outcome associated with risk will shy away from risk in future periods, despite the fact that the risk changes from round to round and each round is independent.

7 Results

Data was collected from three experimental sessions in the Economic Science Laboratory at the University of Arizona.There were a total of 58 subjects who participated in the experiment, and each subject completed 10 rounds of decisions, yielding 580 total observations. In this section, all numbers regarding payoffs and decisions are left in terms of ECU (exchange rate: 10 ECU = 1 USD).

Table 2 presents summary statistics for several variables in the data. Average payoff for all subjects across all rounds was 121.8 ECU, the average amount invested in the risky asset, Good B, was 39.7 ECU, and the average time to make a decision in any given round was 13.8 seconds. The last row in Table 2 is a variable I constructed in the data which represents the difference between

	mean	sd	min	max
Round payoff	121.8	58.9	0	263.5
Investment B	39.7	33.0	0	100
Choice time	13.8	16.79	1	157
$EV_{max} - EV_i$	15.3	16.7	0	87
N	580			

Table 2: Summary statistics for several variables of interest in the data.

the maximum expected value a subject could get in a round, EV_{max}, and the expected value of the decision they actually made, EV_i. The larger this difference, the worse the decision (in strictly monetary terms). The average difference was 15.3 ECU. This difference, $EV_{max} - EV_i$ is helpful because it shows how far a decision is from a benchmark strategy: expected value maximization.

7.1 Decision Changes

The main goal of this paper was to see if offering subjects the opportunity to change decisions would speed up the preference discovery process. In order to explore this, I first examined what might lead subjects to take such an opportunity when I offer it to them. Table 3 presents results from probit regressions where the dependent variable is whether or not a subject exercised the opportunity to change a decision. In all cases, regression coefficients are presented along with standard errors in parentheses. Standard errors are clustered at the subject level. There are 278 total observations because that is how many times, out of 580, subjects were randomly given the opportunity to change a decision. Column (1) is a simple, uncontrolled test of Hypothesis 2. The lack of statistical significance of the coefficient on the round of the experiment suggests that the number of decisions a subject completed had a negligible impact on their desire to change a decision. Column (2) introduces a variable, B prob, which is the probability that the risky asset receives the high return. This was included to test for whether it is the really risky decisions that subjects desire to change. The coefficient on B prob is highly statistically significant and negative, indicating that subjects were actually more likely to change a decision when it involved a less risky asset. Column (3) adds choice time, which is the amount of time, in seconds, subjects took before making their initial decision in a round. Again, the coefficient is statistically significant and negative, indicating that subjects who took a long time on their initial decision did not wish to change the decision. This indicates that those who exercise the option to change decisions may have made their initial decisions too hastily. Columns (4) and (5) use an alternative measure of risk, the variance of the risky asset.

It is also natural to ask whether or not the number of times a subject has been offered the opportunity to change a decision has any bearing on whether she exercises the option or not. One could argue that being asked too much would cause a person to ignore the opportunity. This not

Table 3: Changing a Decision: Probit Regressions

	(1)	(2)	(3)	(4)	(5)
round	-0.141	-0.151	-0.141	-0.139	-0.121
	(0.119)	(0.132)	(0.138)	(0.111)	(0.129)
B prob		-3.070***	-2.826***		
		(0.742)	(0.585)		
choice time			-0.0254*		-0.0240
			(0.0121)		(0.0145)
B variance				5.496***	5.382***
				(0.706)	(0.650)
constant	-1.473**	-0.0538	0.0298	-10.73***	-10.39***
	(0.451)	(0.789)	(0.765)	(0.854)	(0.780)
N	278	278	278	278	278

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

ion is captured by the inclusion of the # of opps variable. This represents the number of times a subject has been prompted with a decision change opportunity, at that particular round in the experiment. Table 4 presents more probit regression results with the dependent variable being a changed decision. I also introduce A value, the monetary value associated with the riskless asset as a dependent variable. One could argue that if someone makes an initially risky decision, they might be tempted to choose a safer one if the safe option is more valuable. Column (5) includes the most variables in the regression and all variables are statistically significant. Coefficient signs suggest that having a more valuable riskless option as well as a riskier Good B would lead subjects to change decision more. I also obtain a statistically significantly negative coefficient on # of opps, which seems to provide support the idea that being asked the same question repeatedly might lead one to disregard it all together.

Since regression (5) from Table 4 has the most statistically significant coefficients out of the models presented so far, it is helpful to look instead at the marginal effects associated with changes in the level of the independent variables. Table 5 presents the marginal effects for the independent variables from regression (5) in Table 4. All marginal effects are calculated at the means of the variables in question. All statistical significance disappears when looking at these marginal effects. Under alternate conditions (not calculating marginal effects at the means) there is still some statistical significance, but that table is not included in the paper. Results are available on request. The lack of significance of these margninal effects could stem from a variety of factors. One factor is that these variables all take on fairly small values (probabilities are less than 1 for example) and they are close together. This could lead to the strange regression result coupled with

Table 4: Changing a Decision: Probit Regressions Continued

	(1)	(2)	(3)	(4)	(5)
round	-0.193				
	(0.123)				
A value	2.784***		2.950***	0.0961	22.05***
	(0.363)		(0.393)	(1.393)	(6.646)
# of opps		-0.325	-0.488	-0.357	-1.287***
		(0.224)	(0.290)	(0.189)	(0.364)
B variance				5.424*	
				(2.528)	
choice time					-0.0715*
					(0.0354)
B prob					31.59**
					(10.28)
constant	-4.473***	-1.335**	-4.400***	-10.50***	-38.56**
	(0.198)	(0.453)	(0.374)	(2.439)	(11.74)
N	278	278	278	278	278

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

insignificant marginal effects.

It appears then, that there is not much about the decision environment that can predict whether or not a subject will change her decision. Perhaps there is something about the specific subject, or the particular initial decision of that subject which leads her to exercises the option to change her decision. Decision changes occur in a nontrivial number of cases. Out of 278 total opportunities to change decisions being offered to subjects, 75 decision changes were made. On average, when subjects changed their decision, they increased investments in the risky asset by 6.5 ECU. The largest change was a shift from buying 0 units of the risky good to buying 100 units of the risky good. The reverse also occurred at one point (with a different subject). Perhaps the repeated nature of the experiment is desensitizing subjects to risk and encouraging a strategy of expected value maximization. Since subjects are presented with multiple decisions about risk, one after another, perhaps they become desensitized to risk. This could also be a function of the outcomes they have seen. After each round, subjects were notified whether or not the risky good, Good B, had a good outcome or a bad outcome. Could subjects be myopic, in which case a subject who sees a bad risky outcome would avoid risk in the next round. Or, a subject who sees a good outcome of risk may be more likely to take a risky choice in the subsequent round. Behavior like this would go along with an outcome based learning model. If risk pays off, a subject might be more likely to take risk in the future.

Table 5: Probit Regression: Marginal Effects

	(5)
choice time	-3.73e-09
	(1.81e-08)
# of opps	-6.72e-08
	(3.21e-07)
A value	1.15e-06
	(5.43e-06)
B prob	1.65e-06
	(7.77e-06)
N	278

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7.2 Expected Value of Decisions

In this section, I examine the expected monetary payoffs resulting from subjects' decisions and attempt to explain certain behaviors observed in the data. I conducted various OLS regressions on the dependent variable EV of decision, the expected value of a subject's decision. The results for these regressions are presented in Table 6. Independent variables include round, the round of the experiment (takes values 1-10), EV_{max}, which is the expected value of the expected-value-maximizing strategy, and changed decision (a binary variable that takes a value of 1 if a subject changed a decision). Results in all three regressions in Table 6 support Hypothesis 3, if only weakly. Using coefficients in regression (3), on average, the expected value of subjects' decisions increased by .591 ECU, when controlling for the maximum possible expected value. It also appears that changing a decision has no statistically significant effect on the expected value of a subject's final decision.

I ran an alternative analysis where I constructed a new variable, $EV_{max} - EV_i$ that represents the difference between the subjects' actual decision and the decision that is expected value maximizing. The larger $EV_{max} - EV_i$ is, the worse off the decision, in expected value terms. I use this constructed variable as the dependent variable in a series of regressions presented in Table 7. The goal with these regressions is to understand what drives subjects' decisions away from the expected value maximizing one. Is there some feature of the decision environment which drives choices away from expected value maximization? It is immediately evident from the magnitude and statistical significance level of the coefficient on the A value variable that when the value of the safe asset, Good A increases, people make decisions that more closely match expected value maximization. This is fairly easily explained by my parameter choices. If you look back at Table 1 in Section 5, one can see that in some cases, the safe asset dominates the risky one in expected value terms.

Table 6: Expected Value of Decisions: OLS

	(1) EV of decision	(2) EV of decision	(3) EV of decision
round	0.880***	0.457**	0.591*
	(0.145)	(0.155)	(0.246)
EV_{max}		0.436***	0.497***
		(0.0730)	(0.105)
changed decision			3.398
			(4.618)
constant	113.8***	57.69***	49.67***
	(1.097)	(8.620)	(12.95)
N	580	580	278

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This would make it easy for both a risk neutral and risk averse person to invest wholly in the safe asset in those cases. That is what I believe drives the result of the coefficient on A value. When we restrict attention to the subsample of rounds where subjects were offered decision change opportunities, we see that changing a decision had no statistically significant effect on the outcome variable in question.

Table 7: Distance from Expected Value Maximization: OLS

	(1) $EV_{max} - EV_i$	(2) $EV_{max} - EV_i$	(3) $EV_{max} - EV_i$
round	0.0901	0.0140	
	(0.145)	(0.153)	
EV maximum		0.209**	0.0679
		(0.0690)	(0.0909)
A value		-26.20***	-17.89*
		(4.781)	(7.792)
B value		-1.612	-11.14
		(3.967)	(7.145)
changed decision			0.464
			(5.856)
constant	14.85***	17.42	51.02*
	(1.097)	(13.89)	(19.45)
N	580	580	278

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Another phenomenon that can be examined in the data is one that is directly related to reinforcement learning: myopia. While they seem unrelated at first, reinforcement learning can sometimes reflect a certain level of myopia. If a subject takes an action and receives a positive outcome from that action, she is more likely to take that action again. But, what if she got incredibly lucky? The subject would then be myopic if she does not consider the strategic elements of an action or recognize the risk associated with a decision. It's almost as if a myopic individual would take an action, put the action into a black box, and see the result pop out the other side. She'd then take another action based only off of the feedback she has: the previous outcome. The fallacious logic of "something good just happened, so it might happen again!" is not new in economic environments, and it can be applied again here.

Table 8 presents results from various OLS regressions where the dependent variable is the amount invested in the risky asset, Good B. To examine the effect of past outcomes on current decisions, I include two variables: "last outcome good" and "# of good outcomes." "Last outcome good" takes a value of 1 if the riskiness associated with Good B in the previous round had a good result, i.e. Good B had a high return. "# of good outcomes" is a count of how many of the previous rounds had good outcomes for the risky asset for that particular subject. A value and B value are defined as before. One relationship worth explaining is that of the coefficients on A value and B value in columns (2) and (4). Both coefficients are statistically significantly negative. The sign on A value is to be expected; one would think that if the safe asset was more valuable, one would invest less on average in a risky asset. The sign on B value is somewhat perplexing unless someone remembers the details of the experiment. When Good B is very valuable, it is also very risky (the probability of receiving the value is lower). When Good B is less valuable, the probability of receiving its value is higher. So, the negative coefficient on B value tells us that subjects are more willing to take the risk on Good B when the probability of success is higher.

Table 8: Amount Invested in Risky Asset: OLS

	(1)	(2)	(3)	(4)
last outcome good	-1.778	0.819		
	(2.471)	(2.169)		
A value		-39.56***		-42.21***
		(5.362)		(5.295)
B value		-38.42***		-35.39***
		(5.344)		(5.228)
# of good outcomes			2.319**	2.079**
			(0.730)	(0.723)
constant	40.57***	175.0***	33.12***	164.5***
	(2.338)	(12.28)	(2.800)	(12.38)
N	580	580	580	580

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Now that I've explained a potential regressional oddity, turn attention to the coefficients on "last outcome good" and "# of good outcomes." One might expect that the most recent result would stick in subjects' minds and inform their next decision. If that were the case, we would expect a statistically significant coefficient on "last outcome good," which we do not see. What we do see is a statistically significant and positive coefficient on "# of good outcomes." The coefficient is highly statistically significant which is interesting because I did not provide subjects with any sort of history of previous rounds. If subjects were to consciously keep track of good outcomes in order to inform decisions, they would have to do it on their own. It would seem that most subjects wouldn't do this, however. Subjects might be subconsciously keeping track of good outcomes and it informs their decision-making process. This would go along with the model of attraction based learning because subjects would see risky actions paying off and be more likely to take risky future actions as a result. The effect is small, however, and can be interpreted as an average of 2 extra ECU being invested in the risky asset for every good outcome the subject observed so far. This can speak to Hypothesis 4. Subjects seem to not respond immediately to good outcomes in previous rounds, but respond to the accumulation of good outcomes over the course of the experiment to inform their decisions.

I have yet to speak to Hypothesis 1 and am not in a position to do so in the current draft. I have only had the data for four days and have not had time to perform a deeper analysis of the individual trends which may be present in the data. Speaking to Hypothesis 1 might also require me to explicitly model the learning behavior/discovery of preferences, which I have not included in this draft. I hope to do so in future drafts.

8 Conclusion

In this experiment I offer subjects the opportunity to change decisions as a means of helping them discover their preferences in an unfamiliar choice environment. This option was exercised by roughly 27% of subjects who were given the opportunity. Some subjects changed drastically, others made minor adjustments. The round in which they were asked seemed to have a negligible effect on the decision to change, indicating that either subjects did not adequately discover preferences over the course of the experiment, or that the opportunity to change a decision is not exercised because preferences need to be discovered. Evidence points to the idea that subjects make mistakes, some minor and some not. Offering subject the opportunity to change their decision may be welfare improving for a portion of the population. It is still unclear why subjects would exercise the option to change a decision when no additional information is presented.

Evidence suggests that subjects update their strategies in risky individual choices based on the prevalence of good outcomes of risk in the history of play, even though no history was provided for them. Subjects are not persuaded to invest more in a risky asset due to immediately previous favorable results, but the aggregation of past successes involving risk. The results of this experiment suggest that it may be welfare improving to allow people to costlessly change recent decisions before implementing them, as this can avoid mistakes that need to be corrected.Though this paper can not yet speak to the reasons why someone would change a decision, it does show that in a nontrivial number of cases, the option to change is exercised by subjects in a laboratory environment.It should be noted that this was a simple choice environment. Only one decision had to be made each round, and each round was similar in structure. If 27% of decisions were changed when the opportunity was given, the numbers might be much more significant in more strategic or more complicated environments like games or markets.

I plan more deeply analyze the data at the subject level to understand how decisions evolved over the course of the experiment and hopefully speak more deeply about how risk preferences are discovered and if offering decision changes speeds up the process. I hope to expand the use of this decision change procedure to work involving more complicated environments like games and markets to see if the phenomenon of changing decisions without additional information persists.

References

[1] M. Allais. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica*, 21(4):503–546.

[2] J. R. Bettman, M. F. Luce, and J. W. Payne. Preference Construction and Preference Stability: Putting the Pillow to Rest. *Journal of Consumer Psychology*, 18:170–174, 2008.

[3] C. Camerer and T. H. Ho. Experience Weighted Attraction Learning in Normal Form Games. *Econometrica*, 67(4), July 1999.

[4] G. Charness and U. Gneezy. Portfolio Choice and Risk Attitudes: An Experiment. *Economic Inquiry*, 48(1):133–146, Jan. 2010.

[5] G. Charness and U. Gneezy. Strong Evidence for Gender Differences in Risk Taking. *Journal of Economic Behavior & Organization*, 83(1):50–58, June 2012.

[6] P. Crosetto and A. Filippin. The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, 47(1):31–65, Aug. 2013.

[7] T. Dohmen, A. Falk, D. Huffman, and U. Sunde. Are Risk Aversion and Impatience Related to Cognitive Ability? *American Economic Review*, 100(3):1238–1260, June 2010.

[8] U. Fischbacher. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, June 2007.

[9] U. Gneezy and J. Potters. An Experiment on Risk Taking and Evaluation Periods. *The Quarterly Journal of Economics*, 112(2):631–645, 1997.

[10] C. A. Holt and S. K. Laury. Risk aversion and incentive effects. *The American Economic Review; Nashville*, 92(5):1644–1655, Dec. 2002.

[11] J. W. Payne, J. R. Bettman, and E. J. Johnson. *The Adaptive Decision Maker*. Cambridge University Press, 1993.

[12] C. Plott. Rational Individual Behavior in Markets and Social Choice Processes: the Discovered Preference Hypothesis. In *The Rational Foundations of Economic Behavior*, pages 225–250. St. Martin's Press, 1996.

[13] P. Slovic, D. Griffin, and A. Tversky. Compatibility Effects in Judgement and Choice. In *Insights in Decision Making: A Tribute to Hillel J. Einhorn*. University of Chicago Press, 1990.

Appendix A: Experimental Instructions

Thank you for participating in this experiment on economic decision-making. In this experiment, you will be presented with a series of decisions.

This is an individual experiment. You will not interact with other subjects during the experiment and your decisions will only affect your payment, not the payment of others. This experiment consists of 10 rounds. In each round, you will be presented with information and asked to make a decision.

After you make a decision, there is a chance you will be offered an opportunity to change your decision before the outcome is presented to you. You will not know at the beginning of the round whether or not you will get the opportunity to change your decision in that round.

Your payment in this experiment is based on your decision in one randomly selected round. All 10 rounds have an equal chance to be selected to determine your payment.

Detailed instructions for each round are below. You will perform the following procedure 10 times.

You will be given 100 experimental credits at the start of each round. The conversion rate for experimental credits is always 10 experimental credits = $1. This conversion rate does not change between rounds. These credits do not carry over between rounds. Your decision in each round does not affect any other round.

On your computer screen, you will be given information about two different goods: Good A and Good B. These are the only two goods you can buy and you must spend all 100 experimental credits in each round by purchasing units of Good A and/or Good B. You can only purchase whole units of each good.

One unit of Good A will cost 1 experimental credit.
One unit of Good B will cost 1 experimental credit.

The prices of each good will not change between rounds.

You will be given the following information about each good: a value and a probability. The probability tells you, if you purchase this good, what the chances are that you receive the value of that good.

Here is an example:

Good A: Price: 1 EC	Good B: Price: 1 EC
Value: 0.9 EC	Value: 2.25 EC
Probability: 1	Probability: 0.75

Experimental Instructions

So, when you buy Good A, there is a 100% chance the good is worth 0.9 EC. For good B, there is a 75% chance the good is worth 2.25 EC, while there is a %25 chance that Good B is worth 0 EC.

So, an example of a decision you can make in this round is to buy 50 units of Good A and 50 units of Good B. If this is the decision, you will receive 50 x 0.9 = 45 EC from Good A. For earnings from Good B, there is a %75 chance you receive 2.25 x 50 = 112.5 EC and a %25 chance you receive 0 EC from Good B. Your payoff for each round is the sum of your earnings from Good A and from Good B.

After you make your decision about how to divide your 100 EC between the two goods, there will be a short waiting period. Then, you might get the opportunity to change your decision. After that, the outcome of each round will be presented to you. You will also see your EC earnings for that round, if that round were to be randomly chosen to determine your payment.

You will repeat this process for a total of 10 rounds. In each round, you will see two goods, but the values and probabilities associated with each good may be different. In each round, there is a chance you will be given the opportunity to change your decision.

Again, one round will be randomly selected to determine your payment for the experiment.

Once all rounds are complete, you will see which round was randomly selected to determine your payment, as well as how many EC you earned in that round.

You have been given a piece of blank paper and a pen. You can use them if you want to do so during the experiment. The experimenter will not gather data from this paper. The paper will be destroyed after the experiment concludes, or you may take it with you when you leave the lab.

Once you are finished with all rounds, please wait patiently for other subjects to finish. Once all subjects have completed all rounds, a brief questionnaire with survey questions and demographic questions will be distributed. Your answers to the questions in this questionnaire do not influence your payment. After that is complete, you will be called individually into a separate room and paid privately by the experimenter. Once you receive payment you are free to leave the lab. Thank you for your participation.