Essays on Decision Making under Risk and Uncertainty

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Dong Yan
August 2020

Copyright © by Dong Yan, 2020
All Rights Reserved.

DEDICATION

To my mom, Zhiwei Zong

my father, Hongzhi Yan

For their advice, their faith, and their inspiration.

To my wife, Na Zuo

For all the things great and small.

To my little boy, Xu Yan

For all the happiness and love.

ACKNOWLEDGEMENTS

I would like to express my gratitude and appreciation of those who have supported me during my studies at the University of Tennessee. Especially, I want to thank my dissertation advisor, Christian Vossler. During the last five years, Christian has offered me excellent guidance, mentoring, and suggestions. Without him, my academic career would be much different. I also want to thank my other committee members, including Scott Gilpatric, Charles Sims, Enda Hargaden, and Cary Deck (U of Alabama), as they also provided great suggestions and helped me to finalize my dissertation. In addition, the staff and student colleagues across the university made my five years of study happy and meaningful.

ABSTRACT

This dissertation uses economic theory, in tandem with experiments and other empirical methods, to better understand the role of information – for instance, missing information, conflicting information, and information overload – in determining outcomes in decision settings characterized by risk and uncertainty. In my first chapter, I use theory and experiments to compare market outcomes in a setting where the seller has better information on product quality than the buyer, and examine the effects of introducing a third-party who can credibly relay information on product quality. Under a range of conditions, I find market efficiency is higher when the information intermediaries interact with buyers instead of sellers. Moreover, market outcomes depend on many other factors, including the accuracy of the certification process and the level of competition in the certification market.

My other dissertation chapters examine decision making in settings characterized by ambiguity, where agents may not form unique beliefs about the probabilities associated with potential outcomes either because there is too little information, or too much conflicting information is provided. In the second chapter, I design a novel experiment that takes advantage of information updating to test the predictive power of two popular classes of ambiguity aversion models, and the experimental results suggest that the "kinked" rather than "smooth" specification better predict subjects' behavior. In the third chapter, I conduct MATLAB simulations to study the timing of technology adoption in response to climate change when predictions from multiple climate models are available. Using the data from California's Central Valley, my simulation results show that an increase in the confliction between projections accelerates the adaption speed of new technology, while an increase in the volatility of projections delay the procedure.

TABLE OF CONTENTS

CHAPTER I: PRODUCT QUALITY AND THIRD-PARTY CERTIFICATION IN POTENTIAL LEMONS MARKETS 1

ABSTRACT 2
1.1 INTRODUCTION 3
1.2 THEORETICAL MODEL 8
1.2.1 Benchmark case 9
1.2.2 Seller certification 11
1.2.3 Buyer certification 14
1.2.4 Certification when the certifier moves first 19
1.2.5 Extension 1: Competitive certification market 21
1.2.6 Extension 2: Imperfect certification process 22
1.3 EXPERIMENTAL DESIGN 22
1.3.1 Testable hypotheses and power analysis 25
1.3.2 Participants and procedures 27
1.4 RESULTS 28
1.4.1 Effort 29
1.4.2 Certification uptake 31
1.4.3 Transactions and efficiency 32
1.4.4 Seller misrepresentation 33
1.4.5 Additional findings 34
1.5 CONCLUSION 36
REFERENCES 39

CHAPTER II : AN EXPERIMENTAL INVESTIGATION OF UPDATING UNDER AMBIGUITY 41

ABSTRACT 42
2.1 INTRODUCTION 43
2.2 EXPERIMENT DESIGN 47
2.2.1 Lottery value elicitation 49
2.2.2 BDM training procedures 49
2.2.3 Lottery tasks 50
2.2.4 Survey 53
2.2.5 Participants 53
2.3 THEORETICAL PREDICTIONS AND TESTABLE HYPOTHESES 54
2.3.1 Kinked specification 54
2.3.2 Smooth specification 55
2.3.3 Testable hypotheses 56
2.4 RESULTS 60
2.4.1 Hypothesis tests: kinked versus smooth specifications 61
2.4.2 Hypothesis tests: task complexity 64
2.4.3 Regression models 64
2.5 DISCUSSION 66
REFERENCES 70

CHAPTER III: CLIMATE CHANGE ADAPTATION UNDER DEEP UNCERTAINTY 73

ABSTRACT 74
3.1 INTRODUCTION 75
3.2 OBJECTIVE RISK, AMBIGUITY, AND CLIMATE CHANGE ADAPTATION 80
3.3 AN ILLUSTRATIVE EXAMPLE: IRRIGATED AGRICULTURE IN CALIFORNIA'S CENTRAL VALLEY 89
3.3.1 Irrigated production and water markets 90
3.3.2 Climate change beliefs 94
3.3.3 Results 95
3.4 CONCLUSION 98
REFERENCES 100

APPENDICES 103

APPENDIX A: CERTIFICATION IN LEMONS MARKET 104
Appendix A1: Theory Derivation 104
Appendix A2: Experiment Instructions and Screenshots 108
Appendix A3: Tables 119
APPENDIX B: UPDATING UNDER AMBIGUITY 128
Appendix B1: Theory Derivations 128
Appendix B2: Experiment Instructions 132
Appendix B3: Tables 168
Appendix B4: Figures 175
APPENDIX C: CLIMATE ADAPTATION UNDER DEEP UNCERTAINTY 177
Appendix C1: Theory Derivations 177
Appendix C2: Tables 179
Appendix C3: Figures 181

VITA 184

LIST OF TABLES

Table 1.1. Experiment treatments 119
Table 1.2. Theoretical predictions 120
Table 1.3. Data Description 121
Table 1.4. Treatment effects regressions 122
Table 1.5. Selected hypothesis tests 123
Table 1.6. Conditional certification purchases 124
Table 1.7. Dynamic specification 125
Table 1.8. Quality misrepresentation 127
Table 2.1. Description of lottery tasks 168
Table 2.2 . Variables and descriptive statistics 169
Table 2.3. Lottery valuations, by task (in $) 170
Table 2.4. Hypothesis test results, based on lottery valuations 171
Table 2.5. Linear regression models, bid differences 172
Table 2.6. Multinomial probit models, bid difference categories 173
Table 2.7. Hypothesis test results, excluding participants who poorly understood instructions 174
Table 3.1. Parameter values based on a farm producing alfalfa that is considering switching from flood to center pivot irrigation 179
Table 3.2. Optimal adoption time (t, in years)* 179
Table 3.3. Effect of changing risk and ambiguity on the expected adoption time (t, in years)* 179
Table 3.4. Sensitivity of results to parameter changes (t, in years)* 180

LIST OF FIGURES

Figure 1.1. Subject screen for risk elicitation task 112
Figure 1.2. Subject screen for seller's effort choice 113
Figure 1.3. Subject screen for seller's certification decision 114
Figure 1.4. Subject screen for seller's price offer 115
Figure 1.5. Subject screen for buyer's purchase decision 116
Figure 1.6. Subject screen for seller's round summary 117
Figure 1.7. Subject screen for buyer's round summary 118
Figure 2.1. Individual-level results, kinked versus smooth specifications (Hypotheses 1 to 4) 175
Figure 2.2. Individual-level results, task complexity (Hypothesis 5) 176
Figure 3.1. Streamflow Parameters (million acre feet) 181
Figure 3.2. Parameters values for stochastic process for profit 182
Figure 3.3. Comparison over expected profit trend 183

CHAPTER I: PRODUCT QUALITY AND THIRD-PARTY CERTIFICATION IN POTENTIAL LEMONS MARKETS

Abstract

This paper examines a seller's incentives for investing in product quality when buyers have incomplete information on quality, and either the seller or the buyer can purchase quality certification from a credible third party. When the seller invests in quality before the certifier sets a price, we find that both seller effort and social welfare are higher in a setting where certification is available to the buyer relative to one where it is available to the seller. When the certifier instead moves first in the game, buyer certification continues to incentivize relatively more seller effort, although social welfare is not necessarily higher. In a complementary lab experiment, we find empirical support for some basic implications of the theory: certification improves market outcomes relative to when certification is not available, decreasing the price of certification increases its uptake, and making the certification process error-prone decreases seller effort and social welfare. Comparisons of seller and buyer certification settings suggest that differences are smaller than predicted by theory. Our results also suggest that seller certification is a more robust tool for improving market efficiency.

JEL Classifications: C91, D82, L15, G24

Keywords: Market transparency, Certification, Information and product quality, Asymmetric information, Endogenous quality, Experiments

1.1 Introduction

An extensive literature has explored consequences, such as adverse selection, that arise in markets characterized by asymmetric information between parties engaged in potential transactions. One way to mitigate the information problem is for a certifier to enter the market and provide credible information to the less knowledgeable agent. That both buyers and sellers demand certification is evident from existing markets. For instance, in product markets where sellers have better information on product quality than buyers do, food companies certify their products as organic, consumers pay a business such as Lemon Busters to inspect a used car prior to purchase, and firms engage with credit rating agencies to attract buyers for their bonds. In this paper we focus on the seller's incentives to invest in product quality when a certification service is available, and are specifically interested in comparing markets where the seller has the option to purchase certification, a setting we will refer to as seller certification, with otherwise identical markets where instead buyer certification is available. We use theory and experiments to provide new insight on this topic.

As discussed by Stahl and Strausz (2017), there are fundamental differences between seller and buyer certification settings that, in a model where product quality is exogenously determined, lead to a higher market transparency and social welfare in the seller certification case. Seller certification acts as a signaling device, and leads to transparency as it conveys information to the buyer about quality even when it is not used. In fact, certification is the only credible signal of quality, as the price set by the seller contains no meaningful information absent certification. In the buyer certification case, certification acts as an inspection tool. Here, the seller's price does provide a quality signal, albeit an imperfect one, as the seller is exposed to the possibility that the buyer will pay for certification and then forego the purchase if the product is not high quality. The

certifier is incentivized to price the service in a way that minimizes market transparency, as this decreases the value of the seller's price signal and, in turn, maximizes the demand for certification.

In this study we show theoretically that, despite the decrease in market transparency, buyer certification provides a stronger motivation for sellers to invest in product quality. In fact, for the seller-certification case, product quality is fully revealed in the market, but the certifier captures the value of effort expended toward increased quality, leaving the seller with little incentive to create value through higher quality. Conversely, in the buyer-certification setting where the certifier provides the service at a lower cost to decrease transparency and increase demand, the seller invests more in quality improvement since he captures a higher return on quality investment.

Our analysis complements prior work on seller certification, which also demonstrates that introducing a certifier can improve upon market outcomes, but that the seller's incentive to invest in product quality is nevertheless below the first-best case. This literature dates back to Biglaiser (1993), who demonstrates how an information intermediary (e.g., a certifier) affects market efficiency. Albano and Lizzeri (2001) generalize the model of Lizzeri (1999) to allow for endogenous effort, and find that seller certification increases market efficiency; however, an inefficiency still remains as the certifier will always charge a positive price for the information and reduce the seller's incentive to invest in quality improvement. As another example, Biglaiser and Li (2018) show that when the seller certification process is perfectly revealing, this information mechanism crowds out any imperfect signal, and yields the certifier significant market power. In the extreme case, buyers have perfect information on quality, but sellers have absolutely no incentive to invest in quality. Marinovic, Skrzypacz, and Varas (2017) and Board and Meyer-ter-Vehn (2013) focus on the relationship between certification and reputation. In their model, they assume the firm can both build a reputation and use certification to signal product quality. A

certification trap occurs in that environment as low reputation firms have no incentive to invest in quality as the certifier will take all their surplus.

Our research also contributes to the understanding of tradeoffs between seller and buyer certification. On this note, while there is a sizable literature focused on seller certification (see Dranove and Jin, 2010), only a few papers have compared the two types of processes. Importantly, these studies abstract from the seller's quality investment choice. Durbin (1999) first illustrates that the role of certification depends on the trading party and focuses on how a certifier can maximize profit through the choice of whether to engage with the buyer or seller. Fasten and Hofmann (2010) allow the certifier to trade with both parties, but the focus is on the profit-maximization problem of the certifier. Stahl and Strauz (2017), as described above, allow for strategic interaction between the certifier and engagement party but assume product quality is exogenously determined. In their model, seller certification is more beneficial as it leads to a higher market transparency and social welfare.

Our model considers the case with one potential buyer, one seller, and one certifier. In our baseline model, the seller first makes a costly choice of effort to probabilistically determine the quality of the product, and then presents a take-it-or-leave-it offer to the buyer. Using a similar structure, we consider cases where the seller or buyer has the option to purchase certification. In the seller certification scenario, consistent with Biglaiser and Li (2018), the certifier sets the price of her service after quality is determined (that is, the certifier cannot commit to a price prior to the seller choosing effort). In this case, a profit-maximizing certifier will set the price of the service equal to the price gap between the high quality and low-quality item, leaving the seller no better off in the case where they have a high-quality item to sell. In the buyer certification scenario, knowing that the low-quality seller has an incentive to mislead the buyer, a buyer will purchase

certification with a probability less than one. With a lower certification price, the buyer's random inspection mechanism yields the seller an expected positive benefit for producing a high-quality product, thus motivating effort to increase quality.

We further consider a setting where the certifier moves first by committing to a price. While this has no effect on buyer certification, in the case of seller certification changing the order of the game does motivate the certifier to lower her price to incentivize seller effort which in turn increases demand for certification. Nevertheless, buyer certification continues to better motivate investment into product quality. From a social welfare perspective, the advantage of buyer certification is the additional effort towards quality improvement, while the disadvantage is the loss that occurs when a misleading low-type seller is caught and the product goes unsold. We show that buyer certification unambiguously increases social welfare when the certifier moves second in the game. Whether this is true in the case where the certifier moves first depends on the value of a low-quality product to buyers.

We consider two additional extensions. First, we assume the certification price is set competitively, rather than by a monopolistic certifier. In this case, seller certification leads to relatively higher product quality, as the certifier's rent-seeking behavior becomes limited while the advantage of transparency still exists. Second, we allow errors in the certification process in the sense that the certifier has some probability of classifying a low-quality item as high-quality and vice versa. The presence of errors in the certification process lowers both the price of certification and the seller's effort choice. However, we find the magnitudes of these effects differ across buyer and seller certification, leading to ambiguity in terms of which leads to more desirable market outcomes.

The theoretical model yields several testable hypotheses regarding investment in quality under varying market and institutional conditions, some of which we test in a complementary laboratory experiment. The experimental investigation contributes to an experimental literature on information provision in markets. Cason and Gangadharan (2002) examine whether certification improves market efficiency in a food labelling setting, and find that when compared to cheap talk and reputation building mechanisms, the certification system is a more reliable approach for improving market efficiency. Brown, Falk, and Fehr (2004) study a related setting where firms and workers enter into contracts, and the workers can choose effort levels that deviate from what was contracted. They find that introducing third-party enforcement leads to near-competitive contract terms. As no prior experiment has focused on comparing outcomes from parallel seller and buyer certification settings, we concentrate on this in the experiment. We compare settings where seller certification is predicted to be superior in terms of product quality and/or social welfare with others where buyer certification carries these potential advantages. We further investigate the effects of introducing errors in the certification process, which has not been examined in the prior experimental literature.

The experimental results serve to both confirm and challenge the theoretical model. We find support for some primary hypotheses, including that certification improves market outcomes relative to when certification is not available, decreasing the price of certification increases its uptake, and introducing errors decreases both seller effort and social welfare. However, comparisons of the parallel seller and buyer certification settings suggest that differences are smaller than predicted by theory. Other interesting findings include differences in how market surplus is divided between buyers and sellers across certification settings, and changes in

certification uptake when the inspection process is error prone. Additional details of our results are described later.

1.2 Theoretical model

In our model, there are three players in the market: one seller, one buyer and one certifier. We define player roles below.

Seller: the seller produces one unit of a product at some cost. The quality of the output θ is determined by the seller's effort, and is either high or low (discrete), $\theta \in \{L, H\}$. Define e as the probability that the item is high quality, and $c(e)$ as the cost function. We assume $c(e)$ is twice differentiable, strictly increasing and convex on the unit interval, $e\epsilon[0,1]$. Also, we assume $c'(0) = 0$, which guarantees an interior solution. Absent any certification cost, if the seller exerts effort e and the good is sold at price P, firm profits equal $P - c(e)$.

Buyer: the buyer neither observes the seller's effort nor the quality of the good prior to purchasing the item. However, the buyer has a prior belief about the seller's effort level, $\tilde{e}$, and updates his belief based on the price, and any information provided by a certifier. The buyer's utility depends on the quality of the product. Define $v = v_h$ if the product is high quality and $v = v_l$ if the product is low quality. Then, absent any certification cost, the buyer's expected payoff from purchasing the item is $E(v|P, \tilde{e}) - P$. If a purchase is made, the buyer knows with certainty the quality of the good.

Certifier: the certifier only benefits from selling its services. The certifier has an inspection technology, which is costless and possibly prone to errors. The certifier makes a "take-it-or-leave-it" price offer, P_c, for the service and provides a product certification only when the quality is determined to be high. The certifier either offers the service directly to the buyer or the seller, depending on the setting.

In the analysis that follows, we first define the benchmark model as the case without any certifier. Next, under the assumption that the certification service is perfectly-revealing, we will compare seller certification and buyer certification when the certifier moves after the seller to illustrate the difference between the two mechanisms. We then consider the case where certifier moves first and demonstrate that the advantage of buyer certification remains. In addition, we extend our model in two ways. First, we examine the case where the certification market is competitive, and then, we consider the case of an error-prone certification process.

1.2.1 Benchmark case

In his classic article, Akerlof (1970) argues that – when the product quality is unknown to potential buyers – this can result in a market where only low-quality products ("lemons") are sold, and eventually all high-quality products are priced out of the market. As it is not obvious that a similar result will hold in a setting where sellers can expend costly effort to improve product quality, we first analyze a simple product market without a certifier. The game has three stages:

1. The seller chooses effort e and then learns product quality θ.
2. The seller makes a "take-it-or-leave-it" offer to the buyer at price P.
3. Observing price P, the buyer updates his belief about product quality and decides whether to purchase the product.

Our model is of a dynamic game with incomplete information, and we rely on the Perfect Bayesian Equilibrium (PBE) equilibrium concept. Allowing for a mixed strategy, let $\sigma_\theta(P)$ denote the probability that a seller with quality θ offers the item with price *P*, and it satisfies the condition that:

(1) $\int_{v_l}^{v_h} \sigma_\theta(P) = 1, for\ \theta \in \{L, H\}$,

where we assume the reasonable range of prices to be between the buyer's valuations of the high- and low-quality product.

The buyer updates his belief after observing the price signal using Bayes' rule. Since the only signal that a buyer has is the price, we define $\mu(P)$ as the buyer's belief that the item is high quality as follows:

(2) $$\mu(P) = \frac{\tilde{e}\sigma_h(P)}{\tilde{e}\sigma_h(P)+(1-\tilde{e})\sigma_l(P)}$$

In other words, when facing price P, the buyer will form her belief based on the prior and the chance that a seller will offer a low (high) quality product at a price P. With the updated belief, the buyer will take the offer if $P \leq \mu(P)v_h + (1-\mu(P))v_l$, and reject it if $P > \mu(P)v_h + (1-\mu(P))v_l$.

Knowing that the buyer has a belief about the expected value based on the price signal, the seller decides the optimal price based on the actual quality. Though we allow the seller to use a mixed strategy, however, it turns out that offering the good at $P^* = E(v) = \mu(P)v_h + (1-\mu(P))v_l$ dominants any other price for a high-quality seller. Any price higher than this will deter the buyer while any price lower than this will not maximize profit. For a low-type seller, he follows the same strategy as the high type seller, and offers the item at $P^* = E(v)$, since a different action will reveal his type to the buyer. Thus, we have $\sigma_h(P^*) = \sigma_l(P^*) = 1$, $\mu(P^*) = \tilde{e}$, and $E(v) = \tilde{e}v_h + (1-\tilde{e})v_l$. Notice here the optimal selling price does not depend on the effort choice of seller. Thus, the seller will set the price just equal to $\tilde{e}v_h + (1-\tilde{e})v_l$ and then choose effort in order to maximize profit

(3) $$\max_{e\in[0,1]} P^* - c(e).$$

Substituting in the value of P^*, the equivalent problem is

(3') $\max_{e\in[0,1]} \tilde{e}v_h + (1-\tilde{e})v_l - c(e)$.

Clearly, the solution to this problem is for the seller to exert minimal effort, which in turn minimizes cost. In equilibrium, the buyer's belief $\tilde{e}$ is correct and $\tilde{e} = 0$. No high-quality item is produced in the market. Conditional on the assumptions made, the solution for the problem is unique.

Lemma 1 When there is no certifier in the market, the optimal effort choice by the seller is $e^* = 0$, i.e. the seller chooses minimum effort, and all items are traded as the value of low-quality.

1.2.2 Seller certification

Now we allow the certifier to enter the market and interact with the seller. The existence of the certifier gives the seller a chance to signal product quality. We assume that certification perfectly reveals quality. If certification is a purchased, a high-quality item will be certified whereas a low-quality item will not. With this additional interaction, the seller certification game has five stages:

1. The seller decides his effort level e, and then learns the quality θ.
2. The certifier determines the certification price P_c.
3. The seller decides whether to purchase certification. If the service is purchased, an item with high quality will convey a certification as a perfect quality signal to the buyer.
4. The seller sets the price P and makes a "take-it-or-leave-it" offer to the buyer.
5. Observing price P and whether the product is certified, the buyer updates his belief about product quality and decides whether to purchase the product.

Based on the certification condition, let us define $\sigma_\theta^c(P)$ as the probability a seller with θ type item offers a certified item at price P, and $\sigma_\theta^u(P)$ as the probability a seller with θ type item

offers an uncertified item at price P. Again, the seller's strategy must satisfy the condition that for each type of seller, the joint probability between choosing a selling price and taking certification equals to one. Since we assume certification provides a perfect signal, after observing the seller's action, the buyer will treat any item with certification as high quality, and thus the maximum acceptable price is $P_H = E^c(v) = v_h$. For any item without certification, the buyer will form the belief $\mu(P)$ as

$$(4) \qquad \mu(P) = \frac{\tilde{e}\sigma_h^u(P)}{\tilde{e}\sigma_h^u(P)+(1-\tilde{e})\sigma_l^u(P)}.$$

which implies $P_L = E^u(v) = \mu(P)v_h + (1-\mu(P))v_l$. We first assume the price of certification, P_c, is exogenously given and define the equilibrium as follows:

Lemma 2 With P_c exogenously determined, under seller certification we have that

1. For any $P_c \leq P_H - P_L$, a high type seller will always purchase certification, and a low type seller will not. A PBE exists for which the expected profit for the certifier is $P_c e$ and the product will be purchased with certainty.
2. For any $P_c > P_H - P_L$ neither the high type nor low type seller will purchase certification. The equilibrium will be the same as in the benchmark case.

Knowing the buyer takes the certification as a perfect signal and high type seller will always purchase certification when $P_c \leq P_H - P_L$, the certifier needs to solve their profit maximization problem. To maximize profit while incentivizing the seller to purchase certification, the certifier will set the price P_c such that the profit of the seller is the same regardless of the quality of the item, which implies $P_c = P_H - P_L$. Under this price we have $\sigma_h^u(P) = 0$, since it is always optimal for the high type seller to purchase the certification. We also know that a low-type seller will never purchase certification since the inspection is perfect, $\sigma_l^u(P) = 1$. Thus,

(4') $\mu(P) = \frac{\tilde{e}\sigma_h^u(P)}{\tilde{e}\sigma_h^u(P)+(1-\tilde{e})\sigma_l^u(P)} = 0,$

(5) $P_L = E^u(q) = \mu(P)v_h + (1-\mu(P))v_l = v_l,$

and

(6) $P_c = P_H - P_L = v_h - v_l.$

With these conditions, the profit-maximizing problem for the seller becomes

(7) $\max_{e\in[0,1]} e(P_H - P_c) + (1-e)P_L - c(e).$

After substituting in the optimal P_c, the problem above is equivalent to

(7') $\max_{e\in[0,1]} v_l - c(e).$

Based on our assumptions, the unique solution for this problem is $e^* = 0$, as in the benchmark case. We summarize our first result in the following proposition.

Proposition 1 The market with available seller certification has a unique equilibrium. The high-quality seller certifies with certainty, and the low-quality seller does not certify. The buyer purchases a certified product at price v_h and an uncertified item at price v_l. The certifier sets the price of certification to be $P_c = v_h - v_l$, and the seller devotes minimal effort to improving the quality. The product is traded with certainty.

Intuitively, although the existence of the certifier provides perfect information into the market, the rent seeking nature of the certifier will also reduce the incentive for the seller to devote additional effort. Note here that the equilibrium for the seller certification case is fundamentally different from the benchmark case in the sense that the equilibrium is a perfect sorting equilibrium and a high-quality product is recognized by the market. While this distinction may seem trivial, consider a slight modification where we instead assume that expending minimal effort, $e = 0$, yields a high-quality item with a probability greater than zero. Thus, a high-quality product can

arise, and when produced the seller will purchase certification. This, in turn, will increase social welfare relative to the benchmark case.

1.2.3 Buyer certification

As with the case of seller certification, the buyer certification has five stages:

1. The seller decides his effort level e and learns the quality θ following production.
2. The certifier determines the certification price P_c.
3. The seller sets a product price P and makes a "take-it-or-leave-it" offer to the buyer.
4. Based on the price of certification and the product price, the buyer decides whether to purchase a certification service. If certification is purchased, the buyer knows the quality for sure.
5. The buyer decides whether to purchase the product.

Allowing for a mixed strategy, we define $\sigma_\theta(P)$ as the probability that a seller with quality θ offers the item with price P, and $\sigma_\theta(P)$ satisfies the condition that

(8) $\int_{v_l}^{v_h} \sigma_\theta(P) = 1, for\ \theta \in \{L, H\}$.

In the absence of certification, the only signal available to the buyer is the offered price. We define $\mu(P)$ as the buyer's belief that the quality of the item is high based on the price signal. Using Bayes' rule, faced with P, the buyer will first consider the probability that a seller with type θ will sell the item at the given price, yielding the belief

(9) $\mu(P) = \frac{\tilde{e}\sigma_h(P)}{\tilde{e}\sigma_h(P)+(1-\tilde{e})\sigma_l(P)}$.

With the price and belief, the buyer has three possible actions:

1. Action s_b: the buyer does not purchase certification, and buys the item directly. The expected payoff is $U(s_b|P,\mu) = \mu(P)v_h + \big(1-\mu(P)\big)v_l - P$.

2. Action s_n: the buyer does not purchase certification, and does not buy the item. This action yields the buyer zero payoff.
3. Action s_h: the buyer purchases certification and purchases the product only if it is high quality. The expected payoff is $(s_h|P,\mu) = \mu(P)(v_h - P) - P_c$

There are three other possible actions, including to pay for the certification service and purchase the item if the quality is low, buy certification but never purchase the item, and pay for certification and always buy the item. Since these actions are payoff dominated, we can safely exclude them from consideration. We can calculate the range of values for $P, \mu,$ and P_c that allows one action to yield a higher payoff than the others. The action s_n is optimal when

(10) $\mathrm{P} > \mu(P)v_h + \big(1 - \mu(P)\big)v_l \quad \text{and} \quad P_c > \mu(P)(v_h - P).$

Following the same logic, s_b is optimal when

(11) $\mathrm{P} \leq \mu(P)v_h + \big(1 - \mu(P)\big)v_l \quad \text{and} \quad P_c \geq (1 - \mu)(P - v_l).$

And last, s_h is optimal when

(12) $P_c \leq \mu(P)(v_h - P) \quad \text{and} \quad P_c \leq (1 - \mu)(P - v_l).$

To fully characterize the equilibrium, we will divide the discussion into three parts. First, assuming a mixed strategy equilibrium, we identify the optimal action set for each player in the PBE. Second, conditional on consistency in beliefs we identify the probability of each action. Last, given the mixed strategy PBE, we will show that there is no meaningful pure strategy PBE in this game.[1] These three steps are sufficient to demonstrate that the equilibrium satisfies the requirement for a PBE.

[1] There is a pure strategy equilibrium in which the seller choses e = 0 and buyer buys as the low-quality price. We ignore this self-fulfilling case in the following analysis.

In a mixed strategy equilibrium, a buyer must be indifferent among s_n, s_b and s_h, which implies:

(13) $P_c = \mu(P)(v_h - P)$,

(14) $P_c = (1-\mu)(P - v_l)$, and

(15) $P = \mu(P)v_h + \big(1-\mu(P)\big)v_l$.

Solving the system yields

(16) $\tilde{P} \equiv \frac{\left(v_h+v_l+\sqrt{\Delta v(\Delta v-4P_c)}\right)}{2}$ and $\tilde{\mu} \equiv \frac{\left(1+\sqrt{1-4P_c/\Delta v}\right)}{2}$,

where Δv is the difference in buyer valuations for high and low-quality items. If the seller sets the price at $\tilde{P}$ and the buyer has the belief $\tilde{\mu}$, then the buyer will be indifferent among these three actions. Since the certifier wants to maximize profit, they will always set the price of certification such that the buyer purchases it. As long as $P_c \leq \Delta v/4$, the buyer will have a demand for certification.

To fully characterize the equilibrium, we need to calculate the optimal strategy of the seller conditional on the realized quality. When $P = \tilde{P}$ and $\mu = \tilde{\mu}$, a high type seller will set $P = \tilde{P}$, since any higher price will result in the buyer not purchasing the item. Any lower price is not optimal since it would only decrease profit while not changing the purchase probability. For a low-type seller, he will face a semi-inspection game. If he offers a price other than $\tilde{P}$, he will be recognized as a low type under an updated belief from the buyer. Thus, if he does not ask for $\tilde{P}$, the maximum price he can ask is v_l. If he asks for $\tilde{P}$ and then he will receive 0 when the buyer chooses certification, or $\tilde{P}$ if the buyer chooses not to. Defining $\sigma^*(s_h)$ as the probability a buyer will check the quality when facing $\tilde{P}$, and $\sigma_l^*\big(\tilde{P}\big)$ as the probability that a low type seller wants to mislead the buyer into thinking his product is high quality, we state Lemma 2 as follows:

Lemma 2 For a given P_c, under buyer certification we have that

1) For any $P_c \leq \Delta v/4$, a high type seller will always sell the product at $\tilde{P}$. A low type seller will sell the product at $\tilde{P}$ with probability $\sigma_l^*(\tilde{P})$, and P_L with probability $1-\sigma_l^*(\tilde{P})$. Faced with $P = \tilde{P}$, the buyer will purchase certification with probability $\sigma^*(s_h)$, and the buyer will never purchase certification when $P = P_L$.
2) For any $P_c > \Delta v/4$, the buyer will never purchase certification. The equilibrium will be the same as the lemon market described in the benchmark model.

For the buyer, $\sigma^*(s_h)$ makes the low type seller indifferent between offering to sell at $\tilde{P}$ or v_l. After calculation, we have that

(17) $\sigma^*(s_h) = \frac{\tilde{P}-v_l}{\tilde{P}}$.

For $\sigma_l^*(\tilde{P})$, knowing the high type seller will always offer the item at price $\tilde{P}$, it follows that

(18) $\mu(\tilde{P}) = \frac{\tilde{e}\sigma_h(\tilde{P})}{\tilde{e}\sigma_h(\tilde{P})+(1-\tilde{e})\sigma_l(\tilde{P})} = \frac{\tilde{e}}{\tilde{e}+(1-\tilde{e})\sigma_l(\tilde{P})}$.

We also have

(19) $\mu(\tilde{P}) = \tilde{\mu}$.

In order for the beliefs to be consistent we must have

(20) $\sigma_l^*(\tilde{P}) = \frac{\tilde{e}(1-\tilde{\mu})}{\tilde{\mu}(1-\tilde{e})}$,

where $\sigma_l^*(\tilde{P})$ is the probability that a low type seller offers the item at $\tilde{P}$.

Based on the derivations above, the seller's profit-maximization problem is

(21) $\max_{e\in[0,1]} e\tilde{P} + (1-e)\sigma_l^*(\tilde{P})(1-\sigma^*(s_h))\tilde{P} + (1-e)(1-\sigma_l^*(\tilde{P}))q_l - c(e)$,

which can be simplified by substituting in the expression of $\sigma_l^*(\tilde{P})$ and $\sigma^*(s_h)$ to yield

(21') $\max_{e\in[0,1]} e\tilde{P} + (1-e)v_l - c(e)$.

This maximization problem makes intuitive sense. In equilibrium, a high-type seller will always claim $\tilde{P}$ and successfully sell the product, while the expected profit for a low-type seller is always v_l as the buyer will adjust the chance of certification to make a low-type seller indifferent between attempting to mislead the buyer or not.

Taking the derivative of the revenue and cost functions with respect to e separately, we have

(22) $MB = \tilde{P} - v_l$ and $MC = c'(e)$.

Combining this result with our assumptions about the cost function, we can easily show that there exists an optimal effort e^* that is strictly positive and solves the maximization problem above. The intuition here is that, when a seller knows that the buyer might purchase certification, he will put more effort into product development since quality has a higher chance of being recognized. Compared with the result of seller certification, clearly a market with available buyer certification leads to a higher quality improvement.

Though this result may seem counterintuitive, there is a clear story here. For the seller certification case, the only incentive compatible way to reveal quality is through certification. At the same time, any price claim without the reinforcement of certification is not credible. Such a mechanism provides the certifier with significant market power and allows the certifier to obtain extra rent from their service. On the other hand, for the buyer certification case, the action to claim a price of $\tilde{P}$ also serves as an information device. When a seller self-claims to be a high type, he sends the market a signal, because he incurs an opportunity cost in expectation if his product is, in fact, low quality. The existence of the second source of information limits the rent-seeking ability of the certifier, thus in turn increasing the seller's incentive to invest in quality improvement.

Stahl and Strausz (2017) find that a market with available seller certification yields a higher social welfare as the product is always sold. With endogenous effort, we find seller certification is still trade promoting compared with buyer certification, but the overall social welfare may become ambiguous as the effort level decreases. Here we show the social welfare is always higher under buyer certification when the seller moves first in the game. The advantage of buyer certification is the additional effort towards quality improvement, while the loss is the untraded item when a misleading low-type seller is caught. We can express the comparison as follows

(23) $e_b \Delta v - \sigma_l^*(\tilde{P})\sigma^*(s_h)v_l,$

where e_b is the equilibrium effort level in the case of buyer certification. Substitute in the value of each term, and in equilibrium the expression above reduces to:

(24) $e_b \Delta v[1 - \frac{(1-\mu)}{\tilde{P}} v_l],$

which is strictly positive as $\tilde{P} > v_l$ and $\mu < 1$.

Proposition 2: If the certification price is set after the seller observes product quality, a market with a certification service available to the buyer leads to a higher seller effort level and higher social welfare when compared with a market where instead certification is available to the seller. The probability a product is traded is strictly less than 1.

1.2.4 Certification when the certifier moves first

In this section, we focus on a case where the certifier sets their price prior to when the seller chooses effort. With seller certification available, knowing that the seller will make his effort choice after observing the certification price, the certifier has the incentive to reduce the certification price to incentivize the seller to expend effort, which increases the chance of purchasing certification. On the other hand, the buyer's optimization problem when buyer

certification is available is unaltered by the order of the game; therefore, the certification price and seller's optimal effort are the same as before.

To facilitate a cardinal comparison across certification mechanisms, we assume a flexible form for the cost of effort function; in particular, we assume this has the exponential form, $c(e) = e^{\lambda}$, where $\lambda > 1$ is a parameter that characterizes the convexity of the cost function. For the seller certification case, using backward induction, the optimal certification price and marginal benefit of effort are as follows

(25) $P_c = \frac{(\lambda-1)\Delta v}{\lambda}$ and $MB_s = \Delta v - P_c = \frac{\Delta v}{\lambda}$.

With buyer certification, based on the discussion in section 2.3, we can write the certifier's profit maximization problem as

(26) $\max_{P_c} [e_b(P_c) + (1 - e_b(P_c))\sigma_l^*(\tilde{P}(P_c))]\sigma^*(s_h(P_c))P_c$.

By choosing P_c, the certifier simultaneously influences many aspects of behavior. A higher price of certification makes the certification less favorable for the seller. However, if the chance the buyer purchases certification falls, this also increases the chance that a low-type seller may distort his price in an attempt to mislead the buyer. The solution to (26) is

(27) $P_c^* = -\frac{\Delta v(-2n\lambda-\lambda+n+1)\sqrt{(n+1)(\frac{1}{4}+n(\lambda-\frac{1}{2})^2)}+(n+1)(\lambda^2-(4n+1)\lambda+n+1)}{2\lambda^2}$,

where $n = \frac{v_l}{\Delta v}$, represents the relative value of a low-quality product. The difference between the marginal benefit of effort in the buyer and seller certification cases is equal to

(28) $MB_b - MB_s = \frac{\Delta v}{2} + \frac{\sqrt{\Delta v(\Delta v-4P_c^*)}}{2} - \frac{\Delta v}{\lambda}$,

which is always positive, and increases with n. Intuitively, when the relative value of the low-quality item increases, the value of the price signal in the buyer certification case increases as well,

which in turn leads to a lower certification cost and higher seller effort. Here, we provide our third proposition, a detailed proof of which is provided in the appendix.

Proposition 3: With the cost function $c(e) = e^{\lambda}$ and $\lambda > 1$, the seller's effort is higher in a market with available buyer certification relative to a market with available seller certification. The difference in effort across the certification mechanisms increases with the relative value of the low-quality product, $\frac{v_l}{v_l+v_h}$.

Whether buyer or seller certification leads to higher social welfare is ambiguous. As discussed in Section 2.3, the social welfare differences depend on whether the expected value of the additional effort induced by buyer certification outweighs the expected loss that occurs when a misleading low-type seller is caught, and the product is not sold. When the value of a low-quality product is sufficiently low, the additional effort induced by buyer certification is no longer large enough to offset the expected losses from misleading sellers.

1.2.5 Extension 1: Competitive certification market

In the previous sections we focused on a monopolistic certifier. At the other extreme, the certification market may be either competitive or under some regulation, and be one where the certifier charges the same price regardless of the engaging party. In this case, we can easily show that seller certification leads to higher seller effort. As the cost function and the price of the low-quality item are the same in the two cases, we compare the two mechanisms based on the expected price of the high-quality item. To set up the comparison we have

$$(29) \quad (P_H - P_C) - \tilde{P} = v_h - P_C - \frac{\left(v_h+v_l+\sqrt{\Delta v(\Delta v-4P_C)}\right)}{2},$$

which equals to zero when $P_C = 0$, and is strictly increasing in P_C. Intuitively, the constraint on the certification price limits the certifier's ability to obtain rents in both cases. However, the

transparency advantage of seller certification provides the seller a higher incentive to devote effort, as a seller of a high-quality item can extract all available surplus without providing the buyer any compensation for potentially misleading behavior.

1.2.6 Extension 2: Imperfect certification process

Last, we focus on the case where the certification process is error prone. Let τ denote the accuracy rate, or the probability that a high (low) quality product will be correctly deemed by the certifier to be of high (low) quality. We assume that errors are symmetric, in the sense there is a 1-τ probability that a low-quality item is certified to be high quality, and a 1-τ probability that a high-quality is determined to be low quality. We assume the probability is common knowledge to all parties. When the certifier sets the price after the seller chooses effort, regardless of whether the inspection is perfect or not, the certifier will always price so that they capture all seller surplus, leaving the seller with zero incentive to improve the quality. Therefore, Proposition 2 applies.

When the certifier sets her price before the seller chooses effort, it turns out that when the error rate, $1-\tau$, together with value of the low-quality item becomes large, the seller certification mechanism will lead to higher effort. This is because under buyer certification, when a high-quality product is certified to be low quality, the seller faces an additional opportunity cost, as he is not able to sell the item. On the other hand, under seller certification, when the high-quality product is certified to be low quality, the seller can still sell it at the price of a low-quality item. Support for these claims can be found in the Appendix.

1.3 Experimental design

The theory models provide several testable hypotheses. As a starting point for understanding whether the models are good approximations for actual behavior, we use a lab experiment with human subjects to test a subset of them. While mathematically tractable, a

sophisticated game with three players introduces considerable complexities. To reduce this complexity, we streamline the experimental design by automating the role of the certifier. Specifically, we hold fixed certification prices within treatments. As such, the experimental game conceptually captures a setting where either the certifier moves first, or otherwise one where the certification market is competitive or regulated. One additional advantage of this design choice is that it allows us to compare otherwise identical buyer and seller certification markets while holding the price of certification fixed, a case described in Section 2.5

The experimental treatments are summarized in Table 1.1, and Table 1.2 presents theoretical predictions for outcomes of interest.[2] The experimental design varies as between-subject treatment variables whether seller or buyer certification is available, and whether the certification process perfectly or imperfectly reveals product quality. As a within-subject treatment variable we vary whether the certification price, P_c, is the theoretically optimal certification price derived from the model in Section 2.4 or instead an alternative price. This alternative price in the case of seller (buyer) certification is the optimal price for the parallel buyer (seller) certification treatment. We introduce as a control a standard lemons market, a market where certification is not available. We now provide details of the experimental game.

The experimental game involves a single seller and a single buyer. The seller makes an effort choice, e, by selecting an integer from 0 and 10, and this choice translates linearly into the probability the product will be of high quality – each additional unit of effort increases this chance by 10 percentage points and 10 units of effort guarantees the product is high quality. The cost of effort function is $c(e) = e^2$, and costs range from 0 to 100 lab dollars. The buyer has valuations of $v_l = 250$ and $v_h = 500$ lab dollars, respectively, for a low-quality and a high-quality item. The

[2] All tables and figures are located in the appendix.

seller makes a take-it-or-leave it price offer to the buyer, and there is no outside opportunity to sell the product.

When seller certification is available, the seller has the option to purchase certification, at a known price, after product quality is revealed. The outcome from the certification process is known to the seller prior to making a price offer. When buyer certification is available, the buyer has the option to purchase certification, at a known price, after receiving the seller's price offer. In either certification setting, when a product is certified to be of high quality, this information is revealed to the buyer before he is asked to make a purchase decision.

The certification process can be perfect or imperfect. Under perfect certification, there is a 100% chance that a high-quality product, and a 0% chance that a low-quality product, will be certified as high quality to the buyer. Under imperfect certification, we set $\tau = 0.7$, and these percentages change to 70% and 30% respectively; i.e., there is now a 30% error rate. With seller certification, unless the product is certified to be of high quality, no information about quality is made explicit to the buyer. To be clear, the buyer does not know for sure whether the seller simply elected not to purchase certification or instead whether certification was purchased but the outcome was unfavorable. Under buyer certification, the buyer learns the outcome from the certification process. In any case, the error rates are known to both buyers and sellers.

A seller's earnings are equal to the sales price (if any) minus production cost and any certification cost. A buyer's earnings are the difference between the product's valuation and sales price (if a purchase is made), less any certification cost. In addition, both the buyer and the seller receive 100 lab dollars in fixed income in each decision period. This fixed income helps to make earnings positive over the course of the experiment, given that losses are possible; e.g., a buyer may mistakenly overpay for a low quality product.

1.3.1 Testable hypotheses and power analysis

Our experimental design allows for tests of several hypotheses related to seller effort, the probability certification is purchased, and social welfare. For ease of interpretation we express social welfare (i.e., the sum of seller, buyer, and certifier profits) in terms of the percentage of possible welfare achieved; i.e., as an efficiency measure. Maximum social welfare equals 400, and is attained when the seller chooses the maximal effort of 10 (which guarantees a high-quality product, and costs 100), and the product is successfully sold to the buyer (who values it at 500). Note that any revenue collected by the (simulated) certifier does not alter welfare calculations, as this represents a transfer between parties. Below we summarize as hypotheses key predictions from the theory:

> **Hypothesis 1**: With the exception of buyer certification under the alternative certification price, both effort and social welfare are higher when certification is available relative to the no certification setting.
>
> **Hypothesis 2**: Under perfect inspection and the monopolistic certification price, buyer certification leads to higher seller effort relative to the seller certification case.
>
> **Hypothesis 3**: Under imperfect inspection and the monopolistic certification price, seller certification leads to higher seller effort relative to the buyer certification case.
>
> **Hypothesis 4**: Under perfect inspection and holding constant the certification price, seller certification will lead to higher seller effort relative to the buyer certification case.
>
> **Hypothesis 5**: Seller effort and social welfare are lower when the certification process is imperfect (i.e., subject to error).

Hypothesis 1 is the key prediction that certification induces higher seller effort and leads to a higher social welfare relative to the standard lemons market. As we have devised scenarios, for comparison purposes, that price the certifier out of the market, in a few treatments the expectation is instead that the market with certification does no better than our control setting. Hypothesis 2 compares the two treatments with perfect inspection and a monopolistic certification price, and this prediction follows directly from Proposition 3. Hypothesis 3 and 5 follow from the

analysis in Section 2.6, where we consider imperfect inspection. Hypothesis 4 compares buyer and seller certification while holding fixed the certification price. Based on our model, when applying the alternative price into the seller certification model, the seller will maximize profit by choosing the highest effort level since the marginal profit is always higher than the marginal cost. On the other hand, when faced with the alternative price, the buyer will never purchase certification as the price is higher than the potential loss from overpaying for a low-quality product. With the exception of Hypothesis 3, these predictions are not specific to the parameters chosen in the experiment.

To determine sample sizes, we conducted a paid pilot experiment with 22 participants. During the pilot, participants first encountered the control condition for ten rounds and then switched to the buyer certification with perfect inspection treatment for another ten rounds. Each set of 10 rounds was preceded by two practice rounds. Aside from the fact that two between-subjects treatments were included in the same session, the procedures followed those described in Section 3.2. In our power calculations, we assume that the estimated within and between-subject variances from the pilot session are representative of all treatments.[3] Moreover, we assume that tests are based on a linear regression model with standard errors clustered at the participant-level. With these assumptions in mind, we settled on sample sizes of 34 for each of the five between-subject treatments. This allows us to detect a minimum effect size of 1.4 effort units for a between-subjects test at 80% power, using a 5% significance level. To place this into perspective, considering all cases where treatment effects are expected to arise, the smallest predicted treatment

[3] Interestingly, both the within and between-subjects variances of the effort and market efficiency variables were virtually identical across the two treatments. This suggest that estimated variances from the pilot serve as a reasonable proxy for the remaining treatments, as assumed in the power calculations.

effect is 2 units of effort. The minimum detectable effect sizes are 15% and 22% percentage points, respectively, for the certification purchase and social welfare (efficiency) measures.

1.3.2 Participants and procedures

One hundred and seventy-six undergraduate students participated in experiment sessions conducted in the UT Experimental Economics Laboratory during the Summer and Fall of 2019. Participants were drawn from a large pool of University of Tennessee students registered as potential participants in economics experiments. The pool resembles the general population of students with respect to gender, age, and academic college. Overall, there are 10 sessions, and between 16 and 22 participants in each. Experiment sessions lasted approximately 90 minutes and individual earnings averaged $23.

Decisions were made via networked computers using a program coded with the software z-Tree (Fischbacher, 2007). The software collected all decisions and made all earnings calculations. Written instructions were provided to each participant, which were read aloud by the same moderator in each session. Prior to the certification experiment, participants faced a risk elicitation procedure modelled after Holt and Laury (2002). In the certification experiment, there were two practice rounds that gave participants experience in both the buyer and seller roles. This was followed by 20 rounds, and participants were paid based on the outcomes in each round. Earnings from the certification experiment were converted at the rate of 200 lab dollars to 1 US dollar. At the conclusion of the experiment, participants completed a short questionnaire which elicited basic demographics and further included the 10-item Big-Five personality instrument of Gosling, Rentfrow, and Swann (2003). Participants were paid privately and in cash at the end of the session. Representative instructions and computer screenshots are provided in the Appendix.

In each decision round, players are anonymously and randomly matched into groups of two, with one player in a group assuming the seller role and the other player the buyer role. Participants are re-matched in each period; i.e., we use a "strangers" design. To facilitate learning and to minimize cognitive burden, a player maintains a particular role (buyer or seller) for five consecutive rounds before switching to the alternative role.[4] In treatments with certification, there are two possible certification prices which vary within session. Specifically, at the beginning of the experiment half of the participants are assigned the monopolistic price, and the other half are assigned the alternative price. These prices stay in effect for ten rounds, and then prices are switched for the remaining ten rounds. This way, when a participant assumes the role of seller (buyer), he faces each of the two certification prices exactly half of the time.

1.4 Results

Table 1.3 provides descriptive statistics for variables used in the data analysis. In terms of participant characteristics, the risk elicitation task suggests that 67.5% can be classified as risk-averse, 42% are female, 62% have participated in a prior (unrelated) experiment, and about 39% of participants have experience in sales. On average, across all treatment settings and decision rounds, the seller chooses 3.7 units of effort and sells the product with a 66.3% chance. Conditional on the certification service being available, certification is purchased 39.3% of the time. The average efficiency is 50.6%.

Table 1.4 presents estimates from a set of linear regression models associated with several outcome measures of interest: seller effort, whether certification was purchased, whether the product was sold, and social welfare. The included explanatory variables are the set of treatment

[4] To help reduce order effects due to the within-session price variation, within a session, half of the participants first face the monopolistic certification price while the other face the alternative certification price.

indicators, the coefficients on which measure (unconditional) mean differences in the outcome relative to the control condition. Standard errors are clustered at the participant level. In the Appendix, we extend these models to include control variables, which does not alter the main conclusions drawn.

Overall, a casual examination of the regression results in comparison to the theoretical point predictions in Table 1.2 suggests that differences in experimental outcomes are not as dramatic as theory predicts. In general, differences in effort are less stark across treatments, certification is not purchased often enough, too many products go unsold, and efficiency is less than predicted. Many of the basic, directional hypotheses are nevertheless supported by the experimental data.

1.4.1 Effort

Based on the parameters of the experimental design, point predictions of seller effort span the full range of the choice set. Theory predicts zero effort in the control and some certification treatments with alternative prices, and high levels of effort in the monopolistic certifier treatments with a perfect certification process. As illustrated in model (1) in Table 1.4, mean effort in the control lemons market is 2.7 units, which is much higher than predicted. This higher than expected effort could stem from other-regarding preferences, or otherwise attempts to produce a high-quality product mostly by "luck" in the hopes of being able to sell it for a reasonable profit. Consistent with theory, effort in the BPA and BIA treatments are not statistically different from the control. Also, in support of Hypothesis 1, for the six treatments where certification is predicted to increase, effort is in fact statistically higher relative to the control. The rank-ordering of the treatments in terms of mean effort is also roughly consistent with theoretical predictions.

Tests of the other hypotheses are presented in Table 1.5, based on t-tests of the estimated coefficients presented in Table 1.4. Hypothesis 2 claims that under prefect certification and monopolistic certifier pricing, buyer certification markets yield higher effort than markets with available seller certification. The direction of the estimated effect is consistent with theory, but is small in magnitude and not statistically significant. Hypothesis 3 claims that under imperfect certification, seller certification will lead to a higher effort level under the selected parameter values due to the additional opportunity cost of delivering the high-quality item in the buyer-side mechanism. However, again we fail to find a significant effect, and our point prediction of the difference is also small in magnitude.

Hypothesis 4 predicts that if we hold the certification price fixed, seller certification leads to a higher effort relative to buyer certification. The experiment reveals that the seller side mechanism motivates an extra 1.5 units of effort under the higher certification price, and an extra 0.8 units under the low certification price. While the former effect is statistically significant, the latter effect is in the right direction but is not statistically different than zero. Under Hypothesis 5, making the certification process error prone reduces seller effort. This hypothesis is confirmed for both buyer and seller certification markets. In particular, the change in certification quality decreases effort by 1.5 units under seller certification and 1.4 units for buyer certification markets.

The regressions also show that the market responds to a change in certification price, and in the expected direction. We find a significant decrease in effort for the perfect certification comparisons when the higher certification price is in effect. For the imperfect certification treatment, though the direction of change is consistent with theory, we fail to find a significant difference.

1.4.2 Certification uptake

A treatment effects regression of the certification purchase decision is presented as model (2) in Table 1.4. As certification is not available in the control, the intercept is excluded from the regression such that the coefficients measure the mean purchase frequencies for each treatment. The rank-ordering of observed certification purchase probabilities is consistent with theoretical predictions, with the exception that the SIM treatment leads to the highest certification usage. As theory predicts that certification purchases should be based on other information available, in Table 1.6 we summarize conditional purchase frequencies. In markets where certification is available to the seller, we display the purchase frequencies conditional on the quality of the item. Theory predicts that the seller always certifies a high-quality item and never certifies a low-quality one. Although there are some differences from these stark predictions, the experimental results are rather close to predictions in the perfect inspection cases. When certification is perfectly-revealing and the product is high quality, the service is purchased 92% and 98% of the time, respectively, under monopolistic and alternative pricing. When the product is low quality, certification is purchased just 1% and 7% of the time with these prices, respectively.

When the certification becomes error-prone, we see a significant drop in the usage of the certification service when product quality is high, and an increase in uptake when quality is low. Under monopolistic pricing, the usage is 74% (high quality) and 18% (low quality). Under alternative pricing, these figures are 90% and 60%. Especially in the latter case, where the certification price is rather low, there is clear evidence that sellers with a low-quality product attempt to mispresent their product quality by purchasing the service and hoping that there is an error in their favor. Albano and Lizzeri (2001) argue that a monopolistic certifier might have the incentive to increase the certification error rate in order to attract more consumers for their service, which is supported by our lab result.

In Table 1.6, we present buyer certification purchase frequencies conditional on the product price offered by the seller. According to the theory, the buyer checks quality only if the price is higher than the value of the low-quality item. Though some participants still purchase certification when the price is lower than 250, generally speaking, the data are consistent with theoretical predictions. We also notice that the frequency of certification purchases significantly decreases when the price of the service is high, which also aligns with theory. From the data, we notice that buyers underutilize certification when the process is error prone and the seller's price offer exceeds 250, which might explain our earlier findings with respect to seller effort under monopolistic certification pricing. A lower chance of certification means a lower chance that a high-quality product is misclassified as low, and thus, the opportunity cost of effort in the buyer certification setting is much lower than what theory predicts. We also notice that buyers overuse the certification service when the process is perfectly revealing, which contrasts the finding in the imperfect inspection case. One explanation is that some participants select certification only in cases where it maximizes their lowest possible earnings. With a perfect certification process, purchasing certification serves as insurance, which increases the potential lowest payoff to the price of certification. On the other hand, under imperfect inspection, purchasing certification lowers the potential lowest possible payoff as the chance of being misled is still positive after inspection. Thus, a participant may not purchase certification as it reduces the payoff under the worst-case scenario.

1.4.3 Transactions and efficiency

Model (3) and (4) in Table 1.4 coincide with the product sold and social welfare measures, respectively. Our model, as well as Stahl and Strausz (2017), predicts that seller certification is trade-promoting relative to the buyer certification setting. This claim is largely supported by the

data. The transactions frequency is statistically equal to the control for each of the buyer certification treatments. Three of the four seller certification treatments yield statistically different, and higher transaction frequencies.

In support of Hypothesis 1, we find empirical evidence of a higher social welfare in all certification treatments where theory predicts this to be true. In the control group, though the seller's effort is higher than what is predicted theoretically, with only a 60% chance that a transaction occurs, efficiency is 41.1%. The two perfect seller certification treatments yield the highest efficiency, and these figures are significantly higher than the buyer perfect certification with monopolistic pricing. For the two buyer certification treatments with alternative pricing, we fail to find any significant changes compared with the control group, which is consistent with theory.

We find no statistical difference in the transaction rate when comparing the control with buyer perfect certification. Thus, the social welfare improvement of buyer perfect certification is driven by the improvement in quality. Also, together with the fact that we fail to detect a significant difference in effort between the seller and buyer perfect certification under monopolistic pricing, the higher social welfare in seller certification treatments is driven by the higher chance that a product is sold.

1.4.4 Seller misrepresentation

To gain additional insight into market behavior, we further analyze seller quality misrepresentation, which we define as any case where the seller offers a low-quality item at a price higher than 250 (i.e., the buyer's induced valuation for a low-quality item). It is logical to interpret such behavior as an attempt to mislead the buyer into thinking the product is high quality. Table 1.7 summarizes the misrepresentation probability and average offering price according to product

quality and treatment. Certification treatments have lower misrepresentation frequencies when compared to the control, even for treatments where theory predicts certification will never actually be purchased. When certification perfectly reveals quality, seller certification significantly reduces the misrepresentation behavior compared with the buyer certification treatments, and this result is consistent with the theoretical prediction. However, we fail to find a similar difference when certification is error prone. In addition, we fail to reject the hypothesis that the misrepresentation probability is the same across all buyer certification treatments.

Our results also indicate that making certification error prone reduces the mean offer price for the high-quality item in all cases. We further find that imperfect certification increases the price of the low-quality item in the seller certification case, and fail to find a similar effect among all buyer certification treatments. Based on our data, under seller certification with imperfect inspection, a lower certification price will significantly increase the price of the low-quality item together with the misrepresentation probability. This finding is consistent with the result of effort input and usage of certification for the changes of seller imperfect certification treatment.

1.4.5 Additional findings

To further explore welfare implications, we include model (5) in Table 1.4, using as the outcome variable an alternative measure of market efficiency. In particular, this measure is identical to our social welfare measure but excludes any certifier profits. This measure thus provides insight into the question of whether the introduction of a certifier to an existing lemons market actually makes buyers and sellers better off. The answer to this question is that "it depends". Only in the case of available seller certification and perfect inspection technology is market efficiency higher relative to the control.

The models in Table 1.7 add demographic characteristics to the treatment effect models. Prior conclusions drawn are robust to the presence/absence of these controls. Moreover, we find that some characteristics are correlated with observed outcomes. Females purchase the certification service more frequently. Subjects who have previously participated in an economics lab experiment, and those with sales experience (from outside the lab) invest less in product quality. Using results from the Big-5 personality survey, people with a higher “openness” measure devote more effort when in the seller role and are more likely to use certification.

To study possible behavioral dynamics, in Table 1.7 we further extend the model specification to include treatment-specific time trends, and further allow the mean outcomes to differ across the first and second half of the experiment. The latter set of interactions are motivated by the fact that in certification treatments participants face two certification prices, with the price switching after the tenth period. In the regression models, the indicator "Lasthalf" equals 1 for observations tied to the last half of the experiment. The variable “Period” is the period number (1 to 20). In terms of seller effort, we find that effort decreases in the control as well as for the two buyer certification treatments when an alternative certification price is in effect. In these three treatments, our theory predicts zero effort. Thus, this decrease in effort indicates that participants may be learning to play the theoretically optimal strategy as they become more experienced. In contrast, the effort levels in the seller certification treatments do not present any significant change over periods, which potentially indicates subjects determine what they feel is their best strategy more quickly. Further, we find that after transitioning to the second half of the session (i.e., after experiencing a certification price change), the seller's effort level in the BPM treatment significantly falls, and by a large magnitude. In contrast, the effort level in the BPA treatment significantly increases. This result suggests that subjects' experience in the first half of the

experiment spills over to impact behavior after certification prices change. The drop in effort in the BPM treatment potentially explains the failure to reject Hypothesis 2 from the prior analysis based on averaging across all periods of the experiment.

In terms of certification uptake, we find that in the BIPM treatment, the usage decreases over time together with a significant restart effect in the last ten rounds. For the likelihood of purchase, we find that over time the number of transactions decreases for the control group, and there is a positive restart effect in the last ten rounds. In the two buyer certification treatments under monopolistic certifier pricing, we find the chance of transaction significantly decreases in the last half of the experiment, while transactions increase over time in the buyer perfect certification treatments.

1.5 Conclusion

In this study, we use theory and experiments to examine a seller's incentives for investing in product quality when buyers have incomplete information on quality, and either the seller or the buyer can purchase quality certification from a credible third party. While prior work examines investment in product quality in the case of seller certification, ours is the first paper to focus on the buyer certification case, and to make comparisons between otherwise equivalent buyer and seller certification markets. When the certification process perfectly-reveals quality, and the certification market is monopolistic, the buyer certification mechanism leads to a higher quality. Under the seller mechanism, where the seller has the option to purchase quality certification from a third party, the existence of a perfect certification service crowds out the value of information provided by the seller's imperfect quality signal (i.e., price), thus yielding the certifier significant bargaining power (Biglaiser and Li, 2018). In contrast, under buyer certification, where the buyer has the option to purchase certification, there is no information crowding-out effect. Due to this

difference, the seller mechanism leads to a much higher price for the certification service, which in turn decreases the seller's incentive to devote resources to quality improvement. Similar to Stahl and Strausz (2017), our model also shows that buyer certification brings less transparency to the market. However, with the advantage in quality improvement, we identify many cases where the buyer mechanism leads to higher social welfare, which contradicts the prior finding in the literature when quality is instead assumed to be exogenous.

We conduct a complementary lab experiment to check the validity of the theory. For both buyer and seller certification, we confirm that increasing the certification price decreases product quality, as does making the certification process error prone. However, when comparing buyer and seller certification mechanisms, empirical differences are smaller than theory predicts, and are statistically insignificant in three of four comparisons. A difference in complexity may explain this result as, with available buyer certification, buyers have to form beliefs about the probability the product is high quality based on the seller's price offer. We also find some evidence of behavioral dynamics that differ across the buyer and seller certification settings, with some dynamics tied to the fact that participants face a relatively high and a relatively low certification price within the same session. It is possible that when participants first engage in markets with available buyer certification where the certification price is high, buyers become overly pessimistic and, relative to what theory predicts, associate a seller's price offer with a relatively low probability the item is high quality. Based on our theory, a lower belief about seller effort reduces the buyer's incentive to purchase and in turn reduces seller's future effort choices even after the certification price is lowered. The data provides some suggestive evidence that this pessimism does not appear to wane with additional experience, and spills over to the lower certification price setting. This same

behavioral pattern does not arise under seller certification, suggesting that introducing seller certification is more likely to be effective when the existing market conveys little transparency.

Although our theory also predicts that the imperfect seller certification mechanism brings a higher effort level when the product of the error rate and the value of low-quality item becomes larger, our lab results fail to detect significant differences under the parameter values of experiment design. Compared with the theoretical prediction, our data shows that the usage of certification is much lower. The lower usage of certification decreases the chance that a high-quality item is misclassified as low, thus reducing the expected cost of effort under buyer certification. Without any further exploration, we think the under-use of certification is driven by an objective to maximize the lowest payoff. Under imperfect inspection, purchasing certification lowers the potential lowest possible payoff as the chance of being misled is still positive after inspection. Thus, a participant may not purchase certification as it reduces the payoff under the worst-case scenario.

In the experiment, in order to compare buyer and seller certification mechanisms while holding prices fixed, we simplified the decision setting by exogenously setting certification prices. Extending the design to have participants play the role of certifiers may lead to interesting behavioral insights. For instance, the collapse of AAA-rated structured financial products during the years 2007 and 2008 has brought attention again to the certifier moral hazard issue. Bolton, Freixas, and Shapiro (2012) show the problems of the current credit rating system not only include the agency's conflict of understating risk to attract business, but also the fact that firms can "window shop" for the most favorable rating. Our design can easily include (multiple) certifiers to explore the validity of these results.

References

Akerlof, G. A. (1970). The market for lemons: Quality and the market mechanism. *Quarterly Journal Economics*, 84, 488-500.

Albano, G. L., & Lizzeri, A. (2001). Strategic certification and provision of quality. *International Economic Review*, 42(1), 267-283.

Biglaiser, G. (1993). Middlemen as experts. The RAND Journal of Economics, 212-223.

Biglaiser, G., & Li, F. (2018). Middlemen: the good, the bad, and the ugly. *The RAND Journal of Economics*, 49(1), 3-22.

Board, S., & Meyer-ter-Vehn, M. (2013). Reputation for quality. *Econometrica*, 81(6), 2381-2462.

Bolton, P., Freixas, X., & Shapiro, J. (2012). The credit ratings game. *The Journal of Finance*, 67(1), 85-111.

Brown, M., Falk, A., & Fehr, E. (2004). Relational contracts and the nature of market interactions. *Econometrica*, 72(3), 747-780.

Cason, T. N., & Gangadharan, L. (2002). Environmental labeling and incomplete consumer information in laboratory markets. *Journal of Environmental Economics and Management*, 43(1), 113-134.

Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4), 935-63.

Durbin, E. (1999). McDonald's or the Michelin Guide? Revealing Quality through Private-Sector Certification. Mimeo, Olin School of Business, Washington University.

Fasten, E. R., & Hofmann, D. (2010). Two-sided certification: The market for rating agencies (No. 338). SFB/TR 15 Discussion Paper.

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

Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality* 37(6), 504-528.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.

Lizzeri, A. (1999). Information revelation and certification intermediaries. *The RAND Journal of Economics*, 214-231.

Marinovic, I., Skrzypacz, A., & Varas, F. (2018). Dynamic certification and reputation for quality. *American Economic Journal: Microeconomics*, 10(2), 58-82.

Spulber, D. F. (1996). Market microstructure and intermediation. *Journal of Economic Perspectives*, 10(3), 135-152.

Stahl, K., & Strausz, R. (2017). Certification and market transparency. *The Review of Economic Studies*, 84(4), 1842-1868.

CHAPTER II : AN EXPERIMENTAL INVESTIGATION OF UPDATING UNDER AMBIGUITY

Abstract

We formulate new hypotheses that take advantage of information updating in order to discriminate between the two major specifications of multi-prior ambiguity models: "kinked" and "smooth". In particular, across comparable decision settings, we examine the effects of adding or trimming out certain priors, updating the weight on particular beliefs, changing the payoff for a single potential state, and modifying the distribution within certain priors. Our results show that the kinked specification does well in consistently predicting choices from 68% of participants, and the smooth specification predicts well for just 10%. We find evidence that people may use a compound lottery as one of their priors, subjects are insensitive to information that the best prior is more likely, and people place lower values on ambiguous lotteries that are relatively more complex. Our experimental methods are likely to be useful in other contexts, as they allow for simple tests of decision-making under ambiguity without placing restrictions on the weights participants place on priors, or reliance on comparisons to decision-making under risk.

JEL Classifications: C91, D81, D83

Keywords: Uncertainty, ambiguity aversion, risk aversion, updating, maxmin expected utility, α-maxmin expected utility, recursive expected utility, lab experiment, self-protection.

2.1 Introduction

With the subjective expected utility (SEU) model, Savage (1954) argues that, under uncertainty, a decision maker maximizes utility by attaching a subjective probability distribution (i.e., priors) to all potential states of the world. With several inventive experiments, Ellsberg (1961) shows that people hold a non-neutral attitude between risk (known probability) and uncertainty (unknown probability), and labels this phenomenon "ambiguity aversion". Ambiguity aversion models have been applied to a wide range of settings, including participation in financial markets (Easley and O'Hara 2009), medical treatment decisions (Berger, Bleichrodt, and Eeckhoudt 2013), the take-up of genetic tests (Hoy, Peter, and Richter 2013), and climate change policy (Millner, Dietz, and Heal 2013). Even though experimental and theoretical research has been informative, it remains an open question as to which theoretical model, if any, adequately characterizes behavior under ambiguity. The importance of this issue is exemplified by several theoretical analyses that find conclusions are sensitive to the assumed ambiguity aversion model (e.g., Chambers and Melkonyan 2017), and overall mixed results from experiments.

In this study, we design novel experimental tests that allow us to discriminate better between the two main classes of multi-prior ambiguity models, commonly referred to as "kinked" and "smooth". The new tests take advantage of the fact that the two classes of models make stark predictions on how preferences change in response to new information. Under the kinked specification, people adjust utility based on the changes in best and worst priors, while smooth ambiguity assumes that people take every potential prior into consideration. Across decision tasks, we add or trim out certain priors, manipulate the weight placed on a particular prior, change the payoff of a single potential state, and modify the distribution within certain priors. Simple

differences in valuations across related tasks are used to test theoretical predictions, giving rise to powerful tests while minimizing auxiliary theoretical and statistical assumptions.

Our research contributes to two different streams of literature. First, we complement the handful of recent experiments that provide insight on the ability of ambiguity theories to predict actual behavior. Halevy (2007) and Adellaoui et al. (2011) use an Ellsberg-type experiment to distinguish among several theories by comparing valuations across four lotteries characterized either by objective risk, ambiguity, or two-stage lotteries involving objective risk. Halevy (2007) finds that the choice patterns of about 70% of participants can be explained by one of two smooth model specifications. He further finds that most participants able to reduce compound lotteries are highly likely to be ambiguity neutral, which is a prediction consistent with the SEU model. With a more general design that varies the probabilities of the winning event, Adellaoui et al. (2011) instead find only weak associations between ambiguity attitudes and compound risk attitudes.

Hey, Lotito, and Maffilletti (2010) and Hey and Pace (2014) use a Bingo Blower that contains many balls of three colors that are visible to participants but constantly moving. Across between-subject treatments, the number of balls increases (although their proportions did not), which the authors argue increases the level of ambiguity. Participants complete a large number of decision tasks, distinguished by the payouts associated with the ball colors. Findings from both studies suggest that simple models such as SEU and α-maximin expected utility (α-MEU; a kinked specification) are better at predicting behavior than more sophisticated ones, such as Choquet expected utility (CEU). Ahn et. al. (2014) ask participants to make financial allocation decisions over risky and ambiguous assets, and analyze the data based on parametric representations of the SEU, α-MEU, and recursive expected utility (REU; a smooth specification) models. The authors

find that they are unable to reject the SEU model for 60% of participants, while the behavior of others may be explained by models of ambiguity or pessimism/optimism.

Our paper also contributes to recent research that focuses on how new information alters decision making under ambiguity. Liang (2020), in a setting where participants bet on a horse in a virtual horserace, exogenously varies the credibility of information on the likely winner, and focuses on how valuations change in response to the information. In contrast, in our design the information provided to participants is always accurate. Epstein and Halevy (2019) also focus on how people interpret information. In particular, they conduct an experiment where they allow participants to draw a ball before they bet on the Ellsberg urn. However, the value of this draw is ambiguous as the total number of balls is unclear. In our study, information provided relates to a specific element of an Anscomb-Aumann act, and has a limited effect on estimating the possibility of each state; in contrast, the information in Epstein and Halevy (2019) provides a basis of updating the probability of each state.

Relative to prior work, our experimental design has several advantages. First, our design is flexible in that we can carefully vary the number of priors, the payoffs associated with states, the distribution of a single prior, and the weight over priors. In contrast, the Bingo Blower approach for instance only allows one to change payoffs. The design of Cubitt et al. (2020) bears similarity to our own as the authors also use a combination of cards and balls to create an environment with ambiguous weight over priors. However, they only use two colors of balls and do not associate colors with different payoffs. Within our design, we can create a situation where the expected value of two lotteries is the same across two lottery tasks. However, people characterized by smooth specification will prefer one lottery over the other for a hedging reason, while MEU

maximizers will prefer the other lottery since the best prior is more likely to occur.[5] In the Cubbitt et al. (2020) design, it is not possible to make one prior more credible without changing the expected value of the lottery. Second, relative to Halevy (2007), Ahn et al. (2014) and Cubitt et al. (2020), our identification strategy does not rely on within subject comparisons with risky but unambiguous lotteries. The presence of unambiguous lotteries is likely to introduce context and reference points, with the potential to alter preferences for ambiguous lotteries. Thus, our design is likely to enhance external validity. Third, as the contents of all of the urns in our experiment are hidden, there is no potential for participants to form reasonable, objective estimates over the probabilities of outcomes. Fourth, there is a direct connection between the underlying theories we are testing and experimental treatments, giving rise to simple comparisons of valuations across related tasks, and avoiding the need for sophisticated econometric methods. Last, similar to Halevy (2007), we elicit valuations for different lotteries using the Becker-DeGroot-Marschak (BDM) mechanism (Becker, et al., 1964). Relative to other elicitation approaches, such as binary choices, this mechanism collects more precise information about preferences. As this mechanism is unfamiliar to participants, we provide extensive training, as advocated by Plott and Zeiler (2005) and Cason and Plott (2014).

Our results suggest that the vast majority of our participants display behavior consistent with the kinked specification. To be clear, using differences in valuations across tasks within six comparison sets, we find that 68% of participants adhere to the directional predictions of the kinked specification based on four or more comparison sets, in contrast to just 10% for the smooth specification. We also find supportive, but not universal, evidence that preferences for ambiguity are independent of the payoffs for a specific state, which is one of the most important assumptions

[1] A similar idea has been tested by Yang and Yao (2017).

underlying kinked ambiguity theories. With our experimental design, we have also uncovered interesting results that cannot be easily reconciled by existing theories. First, we find subjects might use a compound lottery as a prior, while the previous theories treat all priors as single lotteries. Second, by framing a subset of mathematically equivalent lotteries in two different ways, we find that elicited preferences differ based on the complexity of the decision task. In particular, participants tend to place higher values on simpler lotteries. Finally, our data analysis reveals that those with lower college GPAs are more likely to deviate from the prediction of the kinked specification, and further that those with higher GPAs are more likely to value mathematically equivalent lotteries equally. This suggests that the ability of a theory to predict behavior in settings characterized by ambiguity may be tied to critical thinking and learning skills. We further find statistical correlations that imply that those who are emotionally less stable, based on a Big-5 personality instrument, are more likely to adhere to predictions from the smooth ambiguity specification.

2.2 Experiment Design

Table 2.1 presents the 14 lottery tasks in the experiment, representing six comparison sets. Tasks identified by the same number belong to the same comparison set. In each task, participants place bids to play a particular lottery, which here refers to the opportunity to randomly draw one ball from an urn, and receive an amount of money based on the color of the ball drawn. Values associated with lotteries are elicited using the BDM mechanism, which under mild assumptions is incentive compatible in the sense that bidding one's true value for the lottery is a dominant strategy (Azrieli, Chambers, and Healy 2018).

Though the BDM mechanism has been widely used in experiments to elicit valuations, both empirical and theoretical concerns have been raised. For instance, Cason and Plott (2014)

provide evidence that the BDM mechanism can give rise to misconceptions, although there is suggestive evidence that experience and exposure to mistakes can help mitigate the effects of misconceptions. As Baillon et al. (2015) point out, ambiguity averse subjects may be able to use the randomness of the BDM mechanism as a hedge against ambiguity, which in turn potentially violates incentive compatibility. To promote incentive compatibility, in our experiment the uncertainty over the outcome of the BDM is first resolved (i.e., whether the lottery is successfully purchased and if so at what price) before the lottery itself is played out. Further, although we ask participants to bid on several lotteries, only one task is randomly selected to be played out. Overall, we take the potential limitations BDM as given, and as discussed below we attempt to minimize the potential influence that the elicitation mechanism has on our findings through our experimental procedures.

An experiment session proceeds as follows. The experiment is conducted using paper and pencil. Participants are provided with written instructions, an example of which is provided in Appendix B. The moderator reads the instructions aloud and answers any questions about the procedures. Prior to the lottery tasks, subjects complete a double multiple price list (DMPL) (Gneezy, Imas, and List 2015), which elicits risk and ambiguity preferences. This is followed by extensive training on the BDM mechanism. Participants then complete 14 lottery tasks.

After all tasks are completed, earnings are determined for the DMPL and lottery tasks. It is common knowledge that after all lottery tasks are completed, just one of the tasks is binding in the sense that it will be played for real. For transparency, all relevant processes are determined by participant volunteers: selection of the binding task; the random BDM price; the draw of a card to determine what urn is in effect (when necessary); and the selection of a ball from the urn. Random

processes associated with earnings for the DMPL are similarly determined by volunteer actions. The session ends with a questionnaire.

2.2.1 Lottery value elicitation

The BDM procedure for eliciting valuations (i.e. reservation prices) for lotteries works as follows. Participants place a bid for the right to play the lottery. Then, a price is randomly determined. If the bid is higher than the randomly determined price, the participant successfully makes a purchase (plays the lottery) and pays the random price; otherwise, no transaction takes place. During the experiment, there is no constraint on bidding – participants can bid any amount. To avoid anchoring associated with the upper bound of the price distribution (see Cason and Plott 2014), participants are informed that the price distribution is a uniform distribution from $0.00 to a "maximum price". We state that the maximum price can be as high as the payoff associated with the highest-valued ball. As rational valuations for any of the lotteries are between the minimum and (expected) maximum possible price, the BDM mechanism remains incentive compatible.

2.2.2 BDM Training Procedures

As the BDM is not a familiar elicitation procedure to most, we provide participants with extensive training, following the recommendations of Plott and Zeiler (2005) and Cason and Plott (2014). In particular, our instructions closely follow those of Butler and Vossler (2018). The training is divided into three steps. In the first, participants are provided with instructions that clearly describe the purchase procedure, provide examples, and include language to disassociate bidding in this context with bidding in auctions. In the second step, subjects are asked to work through a set of practice examples and one additional question that elicits whether people can identify the dominate strategy. To motivate thoughtful responses, participants are paid for correct

answers. After all answers are checked, the moderator discusses each answer and highlights how misconceptions (deviations from the dominant strategy) can decrease earnings.

In the last step, participants work through two training rounds. In the first, they bid on the opportunity to draw a ball from an urn containing three black balls (each worth $2 if drawn), which represents a simple version of the ambiguous lottery tasks. In the second, they bid on an urn that contains one black ball (worth $2 if drawn), one red ball ($3) and one white ball ($4). In each round, although only one bid is placed, the BDM procedure is played out multiple times, using a narrow price distribution that is unknown prior to bidding. This procedure thus motivates learning by giving rise to scenarios where over or underbidding leads to losses or foregone earnings. The instructions for the lottery tasks further emphasize the dominant strategy; i.e., the instructions state, "It will be in your best interest to place a bid equal to the highest price you are willing to pay". Thus, even if the training procedures did not clarify the incentives of the BDM mechanism, simply believing that this advice is true may be sufficient to motivate truthful demand revelation. Importantly, as the elicitation mechanism does not change across tasks, it is reasonable to interpret differences in bids across tasks as a meaningful metric from which to test ambiguity theories.

2.2.3 Lottery tasks

Each task involves urn(s) containing at most three colors of balls: black, red and white. Across all tasks, the black ball is tied to the low prize, the red ball represents the medium prize, and the white ball is associated with the highest payoff. Most of our urns contain only three balls. We choose this structure because the number of priors increases rapidly as balls are added when there are three possible states. In the Halvey (2007) two-color setting, the author attempted to limit the maximum number of priors to be 11 to reduce task complexity. In our design, for the most complex tasks, the maximum number of priors is ten.

For each valuation task, subjects are presented with both written and graphical depictions of the urn(s). The picture contains most of the valuable information, including the total number of balls, potential colors and some other information. The tasks are grouped into two parts based on procedural differences. Part I includes the first two comparison sets (six tasks), and each task therein involves one possible urn. Part II includes sets 3 to 6 (eight tasks), and each task therein involves several possible urns. To implement this, a card is first drawn from a deck to determine what urn a ball will be drawn from. The number of cards is known, but only limited information is provided in terms of the likely values on the cards (i.e., the probability of selecting a particular urn). Within each part, related tasks are encountered in pairs. Each task is individually labelled, and appears on a separate page of the instruction packet. To control for order effects, the order of the two parts is randomized across sessions. Moreover, the order of the task pairs, as well as which task within a pair is first seen, is randomized across participants within a session.

The four tasks in the first comparison set allow us to identify the value of self-protection when the payoff of a state is altered. In particular, we implement self-protection by eliminating the possibility that the lowest-value (black) ball is drawn. We describe this process as follows:

> IF this urn contains one or more black balls, these balls will be replaced with red or white balls before a ball is drawn.

To be clear, we do not ask participants to value self-protection directly, and instead infer this value based on the difference in bids across tasks. We prefer this approach, as it allows us to keep the endowment (i.e., owning nothing) and framing consistent across all tasks in the experiment. To alter the payoff of a state, we vary the value of the highest-value (white ball) across tasks. To help insure *ceteris paribus* comparisons, participants are told that the same urn will be used for all four tasks.

The second comparison set allows us to test the effects of adding a prior. Task 2A only has two priors (all black or all white balls), whereas Task 2B has three (all black, all red, or all white balls). Tasks in set 3 and 4 identify the effect of changing the distribution within priors, and tasks in set 5 and 6 identify the effect of changing the weight on a specific prior. To make *ceteris paribus* comparisons across tasks within each of these sets, participants are told that the same deck of cards will be used within each set; i.e., the chances that a particular urn is selected does not change across paired tasks.

The lotteries in set 3 are identical to those in set 4, and the lotteries in set 5 are identical to those in 6. The main difference is how the identical lotteries are framed. In particular, for tasks in set 3 and 5, participants are told:

> IF a black ball is selected from the urn, you would repeat the procedure – draw a different card, and then draw a ball from the selected urn – until a red or white ball is chosen.

In contrast, the lotteries in sets 4 and 6 eliminate any mention of black balls. Therefore, we hypothesize that any difference in valuations across identical (but differently framed) lotteries may be attributable to decision complexity.

Our experimental tasks represent a new approach for implementing decisions involving ambiguity. Here, we briefly contrast this with the popular second-order probability method. With the second-order probability method, by changing the number of priors, researchers can easily modify the ambiguity level. In addition, employing a uniform distribution allows researchers to assume subjects place equal weight on each prior, which greatly simplifies the theoretical analysis. However, a shortcoming of this method is that it results in subjects valuing a compound lottery instead of a situation truly characterized by ambiguity.

In our experiment, instead of using a uniform distribution, participants are told, e.g., "there are four different urns, numbered 1, 2, 3 and 4. You would first draw a card from a stack of TWELVE cards to select what urn is used." This modification allows researchers to introduce unknown probabilities, and the ambiguity level can be altered by providing (incomplete) information on the distribution of card values. To make subsequent theoretical analysis tractable, we hold fixed the unknown probabilities within a comparison pair (i.e. participants are told that the same deck of cards would be used for either task). We then maintain the assumption that subjects place the same weights on the different priors within the task pair.

2.2.4 Survey

The experiment ends with a two-page questionnaire that elicits demographic information such as age and gender. Participants self-report their level of understanding regarding experiment procedures, and how well they were compensated using 5-point Likert scale questions. In addition, we include the 10-item Big-Five personality instrument of Gosling, Rentfrow, and Swann (2003). Previous research suggests that there is considerable heterogeneity in ambiguity preferences, with respect to both the level of ambiguity aversion and what theories explain behavior. By collecting information about personality traits, we hope to provide insight on whether personality may explain some of this heterogeneity.

2.2.5 Participants

The experiment was conducted at a designated experimental laboratory at a large public university. One-hundred and seventeen undergraduate students, recruited from a variety of majors, participated in the experiment during July 2017 and January 2018. There are data from four sessions, each lasting approximately 90 minutes. Earnings for the experiment are determined by a $10 lump sum for completing all tasks and the questionnaire, the payout from the DMPL, earnings

from the BDM training procedures, and the randomly selected lottery task. Earnings were paid in private at the end of the experimental session. Participants earned between $15 and $32.75, with a mean of $21.91.

2.3 Theoretical Predictions and Testable Hypotheses

This section focuses on the theoretical predictions provided by the kinked and smooth specifications. The kinked specification can be derived from a variety of models, including MEU, CEU, contraction expected utility, and α-MEU (see Ahn et al. 2014). The smooth ambiguity model is developed by Klibanoff, Marinacci, and Mukerji (2005), which is alternatively labeled as REU in Halevy (2007) and Ahn et al. (2014). Following Ahn et al. (2014), we chose the α-MEU as the representative model for the kinked specification. Below we provide a theoretical framework and derive testable hypotheses from the experiment.

2.3.1 Kinked specification

Ghirardato, Maccheroni, and Marinacci (2004) derived α-MEU by relaxing the uncertainty aversion assumption of the standard MEU model (Gilboa and Schmeidler 1989). As shown in Ahn et al. (2014), α-MEU can be derived from different classes of preferences, and in each setting the value of α depends on the set of priors. Applying the theory to our experiment decision tasks, let $s \in \{B, W, R\}$ denote the possible ball colors in an urn (i.e., "states"), let p_s denote the probability of each color, and let x_s denote the payoff associated with the color. In the case that the p_s are known, the expected utility from a lottery involving a single urn can be represented as

(1) $\quad v = \sum_{s\in\{B,W,R\}} p_s u(x_s).$

In a setting characterized by ambiguity, an agent will not know the objective distribution of colors. Instead, she assigns subjective probabilities to the payoffs. In a multiple priors model, the decision maker does not simply assign point estimates of the prior probabilities. Instead, she assigns a

subjective weight over prior probabilities. Denote ρ as a subjective probability distribution (a prior) over payoffs in the subject's mind and Δ as the set of all possible priors. The utility of a person exhibiting α-MEU preferences becomes:

(2) $$V = \alpha \min_{\rho \epsilon \Delta} v(\rho) + [1 - \alpha] \max_{\rho \epsilon \Delta} v(\rho),$$

where $0 \leq \alpha \leq 1$, represents the level of ambiguity aversion. When $\alpha = 1$, the subject is extremely ambiguity averse, while $\alpha = 0$ indicates the subject is ambiguity loving. As suggested by (2), α-MEU assumes that an agent's utility is a weighted average of the expected utilities associated with the worst and best priors.

2.3.2 Smooth specification

The REU model assumes that people place a subjective weight on *every* prior. Define w_ρ as the subjective weight placed on a prior. Then, the expected utility for a person with smooth ambiguity with discrete priors can be expressed as

(3) $$V = \sum_{\rho \epsilon \Delta} w_\rho \varphi[v(\rho)],$$

where, under ambiguity aversion, $\varphi(.)$ is a concave function over the expected utility of a single prior. As pointed out by Halevy (2007), the concavity of $\varphi(.)$ is equivalent to an aversion to the potential mean spreads.

The smooth ambiguity model is motivated by a single model, and places less structure on preferences relative to alternative models. For example, compared with ROCL model by Segal (1988, 1990), REU allows a subjective probability weight (instead of an objective weight) on different priors, together with different preferences over the first and second stage of a lottery. One potential issue with REU is, when a flexible functional form is assumed for $\varphi(.)$, ambiguity preferences are difficult to distinguish from the changes in the weights on different priors. Thus, in prior work researchers either induce an equal weight on different priors (Halevy 2007) or instead

make functional form assumptions to separately distinguish ambiguity aversion from weights on priors (Ahn et al. 2014). As stated previously, we do not impose uniform weights on priors. Instead, we frame related decision tasks in a way that holds fixed the weights within a set of tasks. This allows us to test hypotheses while maintaining ambiguity.

2.3.3 Testable Hypotheses

We formulate hypotheses based on four situations: changing the payoff for a specific state; adding a new prior; modifying the distribution within some given priors, and updating the weight on a specific prior. In all cases, by exploiting differences in valuations due to information updating, the theories lead to distinct predictions. In the case of updating the weight on a specific prior, it is not clear how best to derive predictions from the α-MEU model as this setting has not been previously considered. So we consider two cases, one where α is assumed not to vary across related tasks, and a second where the value of α varies when an agent believes the best (or the worst) outcome is more likely to occur. Last, embedded in our design are pairs of mathematically equivalent tasks that vary only in terms of framing, and we consider a fifth set of hypotheses based on decision complexity. Unless otherwise indicated, in the derivations below we assume that subjects are ambiguity averse, and the value of α is held fixed when making predictions based on a related set of tasks.

2.3.3.1 The value of self-protection when the payoff of a state is altered

The difference in valuations for tasks 1A and 1A′ connotes the value of self-protection. That is, the tasks are identical with the exception that, if the urn contains one or more (low paying) black balls, in 1A′ any black ball will be replaced with either a red or a white ball. Tasks 1B and 1B′ parallel 1A and 1A′, with the exception that the value of the (high paying) white ball is increased from $15 to $30.

As suggested by equation (2) and (3), the two models differ in how people value an ambiguous event. α-MEU assumes the agent's utility function can be expressed as a weighted average of the expected utility from the worst and best prior. This specification implies that when subjects have at least three priors with a strict preference ordering, any changes to the best prior should not influence the utility gain from trimming out the worst one. The intuition here is that with the additive functional form, changing the best prior cancels out when calculating the utility gain from self-protection. On the other hand, the value of self-protection can be easily influenced by a change in the prize level under REU, since it assumes people update their subjective weights after trimming out certain beliefs. The differences across specifications gives rise to the following hypothesis. Theory derivations for this and the other hypotheses are provided in Appendix A.

Hypothesis 1: For an agent exhibiting α-MEU preferences, $V_{1A'} - V_{1A} = V_{1B'} - V_{1B}$. For an agent exhibiting REU preferences, $V_{1A'} - V_{1A} < V_{1B'} - V_{1B}$.

2.3.3.2 Adding a new prior

To distinguish between α-MEU and REU, Halevy (2007) studied whether a lower potential variance in expected payoff would increase subject valuations. Following the same logic, we design tasks 2A and 2B in such a way that the only difference is that 2B has an additional prior: 3 red balls. However, to maintain ambiguity, we did not place any restriction on the weight of each prior. As mentioned in Halevy (2007), one of the testable differences is the additional concave functional form proposed in the smooth specification. We assume subjects update based on the re-weightings of Bayes' rule by Hanany and Klibanoff (2009). Based on the discussion by Halevy (2007), our second hypothesis could also be derived from second-order preferences, which means subjects prefer the urn with a potential lower variance when the expected mean is the same (Halevy

and Feltkamp 2005), or an issue preference, which means people might prefer one issue (risk) to the other (ambiguity) (Ergin and Gul 2009).

Hypothesis 2: For an agent exhibiting α-MEU preferences, $V_{2A} = V_{2B}$. For an agent exhibiting REU preferences, $V_{2A} < V_{2B}$.

2.3.3.3 Changing the distribution within priors

All multi-prior ambiguity models have two components. The first one is a prior that maps potential outcomes with a probability distribution; the second one is the weights on the priors, which reflect the relative possibility for the occurrence of each prior. This structure provides an explicit theoretical representation; however, it can become complicated when one attempts to define the set of priors in a decision setting. Suppose a person faces 1A and for simplicity, she has three priors: all black, all red, and all white. Applying α-MEU, we can express her utility as

$$V = \alpha[u(5)] + (1 - \alpha)[u(15)].$$

Now suppose she believes with certainty that the chance the urn contains all red balls or all black balls is the same. In this case there are two ways to express her utility. First, since the probabilities of the best and the worst priors are still unknown, we can express the utility as before. However, if she takes advantage of the new information and uses a compound lottery as one prior, then her utility becomes

$$V' = \alpha[u(5)] + (1 - \alpha)[u((15 + 10)/2)].$$

In our experiment, the only difference between 3A and 3B is as follows. In 3A, Urn 2 involves a two-thirds chance of drawing a white ball and a one-third chance of red ball, while Urn 3 involves a one-third chance of a white ball and two-thirds chance of red ball. For 3B, both Urn 2 and Urn 3 give an equal chance of drawing a red or white ball. Agents are told that the chance

of taking a draw from Urn 2 or Urn 3 is the same, and that both occur with probability greater than zero. As a result, 3A is a convex combination between priors relative to 3B.

If an agent follows the α-MEU model and only uses a single lottery as prior, she prefers 3A since the best prior has a higher expected payoff. For a person with α-MEU preferences, who uses a compound lottery as prior, she should be indifferent between lotteries 3A and 3B. In contrast, if the decision maker follows the REU model, she prefers the one with a lower mean spread, i.e., she should derive more utility from lottery 3B. The same hypotheses hold when comparing 4A and 4B. Indeed, when we trim out the black balls in tasks 3A and 3B, and recalculate the outcome distribution in each potential prior, these are precisely the same lotteries as 4A and 4B, respectively.

Hypothesis 3: For an agent exhibiting REU preferences, (a) $V_{3A} < V_{3B}$ and (b) $V_{4A} < V_{4B}$. For an agent exhibiting α-MEU preferences, who uses a single lottery as a prior, (a) $V_{3A} > V_{3B}$ and (b) $V_{4A} > V_{4B}$. For an agent exhibiting α-MEU preferences, who uses a compound lottery as a prior, (a) $V_{3A} = V_{3B}$ and (b) $V_{4A} = V_{4B}$.

2.3.3.4 Changing the weight on a specific prior

We now focus on a situation where two lotteries are equivalent except for the fact that the best prior is more likely to happen in one of them. The lotteries 5A (6A) and 5B (6B) have the same best and worst priors: drawing a white ball for sure and drawing a red ball for sure. In 6A and 6B, the lotteries are the same except that, in constructing the latter, two of the red balls in Urn 2 are moved to Urn 3 in exchange for two white balls. A sophisticated bidder should notice that the best prior is more likely to occur in 6A since the only possible outcome of a draw from Urn 3 or Urn 4 would be a white ball. Hence, what we label as an agent with “sophisticated α-MEU" preferences prefers 6A to 6B. We label an agent who does not update across the two lotteries as holding "naïve α-MEU" preferences. At the same time, the REU model predicts that agents should

prefer 6B to 6A since the former task has a lower variance. The same hypotheses hold when comparing 5A and 5B. Indeed, when we trim out the black balls in tasks 5A and 5B, and recalculate the outcome distribution in each potential prior, these are precisely the same lotteries as 6A and 6B, respectively.

Hypothesis 4: For an agent exhibiting naïve α-MEU preferences, (a) $V_{5A} = V_{5B}$ and (b) $V_{6A} = V_{6B}$. For an agent exhibiting sophisticated α-MEU preferences, (a) $V_{5A} > V_{5B}$ and (b) $V_{6A} > V_{6B}$. For an agent exhibiting REU preferences, (a) $V_{5A} < V_{5B}$ and (b) $V_{6A} < V_{6B}$.

2.3.3.5 Preferences and decision complexity

One common assumption is that preferences are stable across decision settings. However, this assumption has been challenged by behavioral economists (e.g., Smith 2010). As lotteries 3A (5A) and 3B (5B) are equivalent to 4A (6A) and 4B (6B), respectively, aside from framing, our design provides the opportunity to examine whether people have the same preferences over identical lotteries.

Hypothesis 5: Decision complexity has no effect on valuations. In particular, (a) $V_{3A} = V_{4A}$, (b) $V_{3B} = V_{4B}$, (c) $V_{5A} = V_{6A}$, and (d) $V_{5B} = V_{6B}$.

2.4 Results

Table 2.2 provides summary statistics related to the participants, and Table 2.3 provides statistics associated with each of the lottery task valuations. Table 2.4 presents hypothesis test results, and Figures 1 and 2 provide histograms of participant-level differences related to the various hypotheses, using bin widths of 25¢. As relayed from Table 2.2, the average participant is approximately 21 years old, 42% of our sample is female, and 40% have either a part-time or a full-time job. According to decisions in the DMPL, about 71% of our subjects may be classified

as risk-averse, while 88% of our subjects may be classified as ambiguity-averse.[6] The correlation coefficient between risk and ambiguity attitudes is 0.15, which is close to the results of Dimmock et al. (2015a, 2015b).

As one indication of comprehension of experimental procedures, 83% of participants correctly identified the dominant strategy of the BDM in the context of a numerical example. As a second indication, we compare bids from task 1A to 1B, and bids from task 1A′ to 1B′. These tasks are identical except that the payoff associated with the white ball is increased from $15 to $30. Therefore, if any weight is placed on the best prior, valuations in the second of each pair of tasks should be higher. Valuations for task 1B and 1B′ are $2.72 (Std. Err. = 0.36) and $3.33 (0.41) higher than 1A and 1A′, respectively. Less than 2% of decisions correspond with higher valuations for the objectively less favorable lotteries. Last, 86% of respondents rated their understanding of instructions as either a 4 or 5 on a 5-point Likert scale, where 5 indicates "I understood very well".

2.4.1 Hypothesis tests: kinked versus smooth specifications

Panel (a) of Table 2.4 provides test statistics associated with the four hypotheses derived to distinguish between the kinked and smooth ambiguity models. We present both paired t-tests and nonparametric Wilcoxon signed-ranks tests. Given how the hypotheses are formulated, the smooth model predicts that differences in bids should be negative in every case. The kinked specification is supported by differences of zero, and in some cases by a positive difference. To provide some basic information on heterogeneity, Figure 1 displays histograms of the individual-level differences associated with each hypothesis test. Moreover, the last three columns of Table

[2] Unfortunately, the ambiguity MPL used in two sessions contained errors, and we are unable to characterize ambiguity in these sessions. Missing values prevents us from using this ambiguity measure as a control variable in the regression analysis.

2.4 present the percentage of respondents with negative, zero, and positive differences as they relate to the null hypotheses.

As can be gleaned from Figure 1, the most frequent outcome associated with each comparison is a difference of zero. In fact, considering all comparisons, 53% of the differences are (exactly) zero. This evidence supports the notion that the kinked specification has good predictive power. Some observations close to zero may be attributable to decision errors. If we consider deviations of +/- 50¢, this figure of *approximately* zero differences increases to 64%.

Hypothesis 1 exploits the difference in functional form between two ambiguity models. 1A and 1B are identical to 1A′ and 1B′, respectively, with the exception that latter set of tasks induces self-protection by trimming out the black balls. As stated before, here we treat the difference between the bid on 1A (1B) and 1A′ (1B′) as the value for self-protection and test whether it varies with the payoff associated with highest-value outcome (i.e. a white ball). Using paired t-tests we weakly reject (p=0.08) that the value for self-protection is the same regardless of the value of the white ball. The Wilcoxon test fails to reject equality with p = 0.26. The somewhat contrasting results can be explained by Figure 1(a). Although the frequency of positive and negative deviations is roughly the same, there are a few large and negative differences (i.e., one at -$15 and one at -$20), thus triggering a slight difference in the mean. It is possible that the negative and significant difference in means is partially due to decision error. If we restrict the sample to those who stated they understood the instructions well, Hypothesis 1 is no longer rejected (t = -1.52, p = 0.16).

Hypothesis 2 exploits differences that arise when a new prior is added. There are many more negative differences than positive ones, giving rise to a statistical difference in mean bids between 2A and 2B, and a difference in the median bids. The aggregate results thus support the

smooth specification. At the participant level, an approximately equal number of pairwise choices can be explained by the kinked (44%) and smooth (43%) specifications.

For Hypothesis 3, our tests rely on changing the distribution within particular priors. The two pairwise comparisons reveal that the decisions of very few participants (less than 20%) can be explained by the smooth specification. At the aggregate level, there are no statistical differences when comparing 4A and 4B, but a positive and significant difference when comparing 3A and 3B. The high frequency of zero differences (59% and 62%) may support the conjecture that people use a compound lottery as a prior. Most of the multi-prior ambiguity models, including α-MEU and REU, are based on Ansombe and Aumann's (1963) joint objective-subjective approach, where utility is derived from a distribution of subjective weights over objective lotteries. Even though this structure does not force the objective lotteries to be one-stage, the existence of a compound lottery would challenge the mapping between the subjective weight and associated objective priors, as the set of priors would be unclear. Our finding also provides an example that runs counter to Segal (1988, 1990), who showed that under the assumption of time neutrality and compound independence, ambiguity aversion could be modeled as a failure of the reduction of compound lotteries. That most participants reveal indifference between 3A and 3B, and likewise between 4A and 4B, provides suggestive evidence of the successful reduction of compound lotteries.

The last hypothesis, which exploits differences due to changing the weight on a single prior, provides mixed support for either the kinked or the smooth specifications. At the participant level, the majority of outcomes adhere to the predictions of the kinked specification. However, there are a fair percentage of subjects (40%) that reveal a preference for 5B over 5A, leading to the null hypothesis to be rejected and providing aggregate support of the smooth specification. The smooth specification, however, is a poor predictor of the observed differences between 6A and 6B. With

very few positive differences, this lends support of the naïve α-MEU model in contrast to the sophisticated α-MEU model.

Examining within-subject behavior, 79 of 117 participants (68%) adhere to the directional predictions of the kinked specification in at least four of the six comparisons. Just 12 participants (10%) revealed negative differences in at least four of the comparison sets. The remaining 22% of people cannot be easily classified. Of these, ten participants (9%) exhibit preferences consistent with the kinked specification in three cases and the smooth specification in three cases.

2.4.2 Hypothesis tests: task complexity

Panel (b) of Table 2.4 provides test statistics associated with four pairwise comparisons corresponding to the null hypothesis that task complexity does not matter. Figure 2 contains histograms for each comparison using participant-level data. If we assume that the kinked specification approximately characterizes behavior, then the important maintained assumption in these comparisons is that the value of α is fixed. If we instead rely on the smooth specification, we need the maintained assumption that the weights are held fixed. Overall, the results suggest a revealed preference for the relatively simple lotteries. These differences are significant in the aggregate for three of the four pairwise comparisons. The most striking difference occurs when comparing 5A to 6A, where nearly half of participants reveal a preference of 6A over 5A. Moreover, the average person with this directional preference values the simpler task 6A by $2.18, which is a 23% increase relative to 5A. One possible explanation for the results is that respondents are more ambiguity averse (e.g., α increases) in more complex environments.

2.4.3 Regression models

Table 2.5 presents linear regressions, where the dependent variable is the difference in bids as they correspond with the earlier hypothesis tests. These differences are stacked to create a panel

data set; i.e., the first outcome for a subject is the difference associated with H1, the second outcome is the difference associated with H2, etc. Included as control variables are the Big-5 personality measures, gender, a measure of risk from the risk MPL, GPA, and an indicator for participants who are majoring in a STEM field. As those in a STEM field should have a higher aptitude for quantitative analysis, we hypothesize that this may explain heterogeneity in the data. Also included are controls for order effects, including whether Part I tasks were encountered first in the session, whether the #A task preceded the #B task in a comparison set in Part I, and whether the #A task preceded the #B task in a Part II comparison set. Finally, we include a full set of hypothesis indicators. All control variables are demeaned, which means that the coefficients on the hypothesis indicators are interpretable as measures of the mean differences associated with the various hypotheses. Standard errors are clustered by participant.

Model 1 and Model 2 present results based on the set of theory model comparisons (i.e., H1-H4) and complexity comparisons (H5), respectively. The coefficients on the hypothesis indicators largely confirm the earlier results from paired t-tests: mean differences are statistically different from zero, at the 5% level, for H2, H3a, H4a, H5b, and H5c. Thus, this serves as a robustness check on our earlier results. Based on Model 1, the bid difference is decreasing in risk aversion, and is lower for those in a STEM major. Based on Model 2, mean differences due to complexity are increasing with GPA. Overall, the control variables do very little to explain the variation in the data.

Table 2.6 presents the results from multinomial probit models. The dependent variable, similar to that of the prior regressions, is based on bid differences. Here, we create a categorical dependent variable with three categories: negative deviation, zero deviation, and positive deviation.

The zero deviation category serves as the reference category. The control variables are the same as in the linear regressions.

Model 3 suggests that the probability of deviating from the zero difference prediction, in either direction, decreases with GPA and emotional stability. Deviating in the positive direction decreases with risk aversion and increases with extraversion. For Hypothesis 3 and 4, positive deviations may also be consistent with the kinked specification. As a robustness check, we estimated a binary probit model, where the dependent variable is whether the deviation is consistent with the kinked specification. The results tell a similar story: those with higher GPAs and a higher emotional stability are more likely to make decisions that adhere to the qualitative predictions of the kinked specification. Turning to Model 4, those with a higher GPA are less likely to alter valuations for mathematically identical lotteries based on task complexity. Those who are more extraverted are more likely to place a higher valuation on the less complex task, and those who are emotionally stable are less likely to bid in this manner.

2.5 Discussion

In this paper, we formulate new hypotheses that take advantage of information updating to distinguish between two major specifications of ambiguity models: “kinked” and “smooth”. To test these hypotheses, we introduce a novel experimental design that allows one precise control over important attributes of multi-prior ambiguity models. In particular, identification is achieved by varying the number of priors, the payoffs associated with states, the distribution of particular priors, or the weight over priors. Our identification strategies do not rely on comparisons between ambiguous and risky lotteries, nor sophisticated econometric methods.

Our results show that most participants display behavior consistent with the kinked specification, and only 10% of participants provide evidence supporting smooth ambiguity models.

In addition, we further provide fodder for future theory development. We provide some evidence that participants may have priors in the form of a compound lottery. Most existing theoretical models only treat priors as simple lotteries. Our evidence also shows that, assuming participants follow the kinked specification, the degree of ambiguity aversion is independent of changing the value of one specific state, and is unaffected by changing the credibility (weight) of the best prior. Last, by framing mathematically equivalent, ambiguous lotteries in two different ways, we find evidence that valuations change due to complexity; in particular, participants tend to place higher values on more simply framed lotteries.

Ghirardato, Maccheroni, and Marinacci (2004) show that ambiguity attitudes depend on the relationship between states and outcomes. However, there is little discussion about how the mapping is structured and which factors influence the level of ambiguity. Our experimental evidence suggests that, consistent with the kinked theory prediction, people's ambiguity attitudes will not change when the expected value of the lottery, together with the worst and best prior, stays constant. Our results further support the notion that the ambiguity level is independent of the outcome of the states. This is not necessarily a prediction of the kinked specification, but consistent with our evidence. To the extent that the ambiguity level is approximately constant across related settings, researchers can elicit the value of α in one setting and apply it to analyze another.

Our results generally support the notion that the relatively simple ambiguity models are more likely to predict behavior, which is consistent with some prior experimental work. The identification strategy we use to test our second hypothesis is similar to Halevy (2007), although in our case we do not place any restrictions on subjective weights. Our results are fairly similar in that, in this comparison, there is some support that subjects may have second-order beliefs.

However, results from other hypothesis tests suggests that the smooth specification is not a robust predictor.

Some prior experimental evidence favors SEU theory, and so it is natural to ask whether this model predicts well the patterns in our data. For our first and second hypotheses, the SEU prediction coincides with that of the smooth specification, whereas the prediction for the third and fourth is a null effect. If we applied the same, simple rule for characterizing whether behavior is consistent with a particular theory at the individual level, we would classify 50 of 119 subjects (42%) as following SEU. The vast majority of these participants (41 of 50) can also be classified as following the kinked specification. Recalling that we can classify 68% with the kinked specification, we can conclude that – in our experiment – the kinked specification does better than both the smooth ambiguity and SEU models in predicting behavior.

There are nevertheless patterns in the data that are not easily explained by the theories we consider. One alternative explanation we explore here is motivated by observations from participant decision sheets. Decisions in our experiment were made with paper and pencil, and there is written evidence that some participants may have based valuations on the simple counting of the number of balls associated with each color. In our econometric analysis, we did uncover a negative correlation between GPA and deviations from the kinked specification prediction, which is suggestive evidence that some participants may be using simple heuristics to guide decisions. Tied to Hypothesis 3a, a simple count of the balls (noting that black balls are non-binding), yields 6 red and 3 white balls for lottery 3A, versus 5 red balls and 2 white balls for 3B. If one assumes each of these balls has an equal chance of selection, people should value 3A over 3B. This is what we find based on average valuations, and moreover 31% of individuals valued 3A over 3B. Using a similar heuristic suggests that, tied to Hypothesis 4a, people would value 5B over 5A. We again

do find average valuations in this direction and 40% of individuals value 5B over 5A. We note that in other tests of Hypotheses 3 and 4, the ratio of red to white balls is exactly the same. Interestingly, the aggregate results reveal no statistical difference in mean valuations. Moreover, unlike for H3a and H4a, for H3b and H4b the number of non-zero deviations is smaller, and there are similar fractions of negative versus positive deviations. Turning to tests of decision complexity, applying the heuristic predicts that $V_{3A} = V_{4A}$, $V_{3B} < V_{4B}$, $V_{5A} < V_{6A}$, and $V_{5B} = V_{6B}$. Interestingly, these predictions match the aggregate results well, and helps to explain why we find a high number of negative deviations at the participant-level for H5b and H5d relative to H5a and H5c. Whether such simple heuristics can explain patterns in the data from related experiments remains an open question.

References

Abdellaoui, M., Klibanoff, P., & Placido, L. (2015). Experiments on compound risk in relation to simple risk and to ambiguity. *Management Science* 61(6), 1306-1322.

Ahn, D., Choi, S., Gale, D., & Kariv, S. (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics* 5(2), 195-223.

Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics* 34(1), 199-205.

Azrieli, Y., Chambers, C. P., & Healy, P. J. (2018). Incentives in Experiments: A Theoretical Analysis. *Journal of Political Economy* 126(4), 1472-1503.

Baillon, A., & Bleichrodt, H. (2015). Testing ambiguity models through the measurement of probabilities for gains and losses. *American Economic Journal: Microeconomics*, 7(2), 77-100.

Baillon, A., Halevy, Y., & Li, C. (2014). Experimental elicitation of ambiguity attitude using the random incentive system. University of British Columbia working paper.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science* 9(3), 226-232.

Berger, L., Bleichrodt, H., & Eeckhoudt, L. (2013). Treatment decisions under ambiguity. *Journal of Health Economics* 32(3), 559-569.

Butler, J. M. & Vossler, C. A. (2018). What is an Unregulated and Potentially Misleading Label Worth? The case of "Natural"-Labelled Groceries. *Environmental and Resource Economics* 70(2), 545-564.

Cason, T. N. & Plott, C. R. (2014). Misconceptions and game form recognition: challenges to theories of revealed preference and framing. *Journal of Political Economy* 122(6), 1235-1270.

Chambers, R. G. & Melkonyan, T. (2017). Ambiguity, reasoned determination, and climate-change policy. *Journal of Environmental Economics and Management* 81, 74-92.

Cubitt, R., Van De Kuilen, G., & Mukerji, S. (2020). Discriminating between models of ambiguity attitude: A qualitative test. *Journal of the European Economic Association*, 18(2), 708-749.

Dimmock, S. G., Kouwenberg, R., Mitchell, O. S., & Peijnenburg, K. (2015a). Ambiguity aversion and household portfolio choice puzzles: Empirical evidence. *Journal of Financial Economics* 119(3), 559-577.

Dimmock, S. G., Kouwenberg, R., & Wakker, P. P. (2015b). Ambiguity attitudes in a large representative sample. *Management Science* 62(5), 1363-1380.

Easley, D., & O'Hara, M. (2009). Ambiguity and nonparticipation: The role of regulation. *The Review of Financial Studies* 22(5), 1817-1843.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics* 75(4), 643-669.

Ergin, H., & Gul, F. (2009). A theory of subjective compound lotteries. *Journal of Economic Theory* 144(3), 899-929.

Epstein, L. G., & Halevy, Y. (2019). Hard-to-interpret signals (No. tecipa-634).

Ghirardato, P., Maccheroni, F., & Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118(2), 133-173.

Gilboa, Itzhak, and David Schmeidler. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141-153.

Gneezy, U., Imas, A., and List, J. (2015). Estimating Individual Ambiguity Aversion: A Simple Approach. NBER Working Paper No. 20982.

Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality* 37(6), 504-528.

Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica* 75(2), 503-536.

Halevy, Y., & Feltkamp, V. (2005). A Bayesian approach to uncertainty aversion. *The Review of Economic Studies* 72(2), 449-466.

Hanany, E., & Klibanoff, P. (2009). Updating ambiguity averse preferences. *The BE Journal of Theoretical Economics* 9(1), ISSN (Online) 1935-1704.

Hey, J. D., Lotito, G., & Maffioletti, A. (2010). The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of Risk and Uncertainty* 41(2), 81-111.

Hey, J. D., & Pace, N. (2014). The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity. *Journal of Risk and Uncertainty* 49(1), 1-29.

Holt, C. A. (1986). Preference reversals and the independence axiom. *The American Economic Review*, 76(3), 508-515.

Hoy, M., Peter, R., & Richter, A. (2014). Take-up for genetic tests and ambiguity. *Journal of Risk and Uncertainty* 48(2), 111-133.

Karni, E., & Safra, Z. (1987). " Preference reversal" and the observability of preferences by experimental methods. *Econometrica*, Vol. 55, 675-685.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica* 73(6), 1849-1892.

Liang, Y. (2020). Learning from unknown information sources. Available at SSRN 3314789.

Millner, A., Dietz, S., & Heal G. (2013). Scientific Ambiguity and Climate Policy. *Environmental and Resource Economics* 55(1), 21-46.

Plott, C. R., & Zeiler, K. (2005). The Willingness to Pay–Willingness to Accept Gap, the "Endowment Effect", Subject Misconceptions, and Experimental Procedures for Eliciting Valuations. *American Economic Review* 95(3), 530-545.

Savage, L. J. (1954). *The Foundations of Statistics.* Hoboken, NJ: John Wiley & Sons.

Segal, U. (1987). The Ellsberg paradox and risk aversion: An anticipated utility approach. *International Economic Review* 28(1), 175-202.

Segal, U. (1988). Does the preference reversal phenomenon necessarily contradict the independence axiom? *The American Economic Review* 78(1), 233-236.

Segal, U. (1990). Two-stage lotteries without the reduction axiom. *Econometrica* 58(2), 349-377.

Siniscalchi, M. (2006). A behavioral characterization of plausible priors. *Journal of Economic Theory* 128(1), 1-17.

Yang, C. L., & Yao, L. (2017). Testing ambiguity theories with a mean‐preserving design. *Quantitative Economics*, 8(1), 219-238.

CHAPTER III: CLIMATE CHANGE ADAPTATION UNDER DEEP UNCERTAINTY

Abstract

Climate change adaptation is receiving increased attention due to the growing acceptance that the impacts of climate change are now inevitable. Understanding and predicting adaptation behavior is important from the purview of formulating environmental policy; however, this presents challenges as the benefits of adaptation depend on individual expectations of climate change that are difficult to ascertain. In this paper, we model a setting where a firm makes an adaptation choice based on objective expectations of climate change (climate projections), and is influenced by both risk (climate variability) and ambiguity (inability to assign probabilities to various climate projections). We find that climate variability works to delay adaptation while ambiguity works to hasten adaptation. We then apply our model to the decision to invest in water-saving irrigation technology in California's Central Valley. Using sixty-four down-scaled climate projections for water availability in the Central Valley, we are able to quantify the amount of climate variability and ambiguity facing farmers and solve for the optimal adaptation strategy under different situations. Our results illustrate how downscaled climate projections can be used to generate theoretically-grounded predictions of climate change adaptation in other settings.

JEL Classifications: C61, D81, Q25, Q54

Keywords: Real options, Adaptation gap, Irrigation, Water scarcity, Climate variability

3.1 Introduction

According to the Intergovernmental Panel on Climate Change (IPCC) fifth assessment report, the annual average global mean surface temperature could be regularly 1.5°C higher than its preindustrial level within the next 40 years (Pachauri et al. 2014). The rapid change in temperature generates physical and economic impact, including substantial species extinction, global and regional food insecurity, and consequential constraints on common human activities. To limit the damage, a combination of mitigation policy and adaptation action is required. When a firm considers making an adaptation decision, it must predict the future effects of climate change to assess the benefits and costs of possible adaptation strategies. Unfortunately, the future effects of climate change are highly uncertain. As stated in the IPCC fifth assessment report, "there does not exist at present a single agreed on and robust formal methodology to deliver uncertainty quantification estimates of future changes in all climate variables" (Pachauri et al. 2014). Consequently, climate uncertainty creates a situation where adaptation investments must account for multiple climate projections with little or no ability to assign probabilities to these different projections. The traditional way to model this situation is to assign weights to these climate projections either by assuming the decision-maker applies an equal weight to each projection or updates weights according to Bayes rule. However, since the projections rely on different assumptions, are not independent, and are based on the same historical record, any weighting scheme that results in a single representative projection will ultimately be subjective (Tebaldi and Knutti 2007, Millner et al. 2013).

In this research, we focus on how the uncertainty created by conflicting climate projections impacts the individual adaptation decision. In particular, we focus on a situation where a firm must decide when to adopt a new technology under two kinds of uncertainty. The first one is objective

risk, which is characterized by the weather variability inherent in a single projection. The second one is ambiguity, which appears when multiple climate projections are presented at the same time, making the decision-maker unsure about which one to trust. Both kinds of uncertainty create the possibility of climate adaptation investments that are harmful to the firm (i.e., maladaptation). However, since climate change implies differential effects on objective risk and ambiguity, it is important to understand whether they have the same type of impact on adaptation investments.

To analyze this situation, we combine real options theory with ambiguity models, which focus on how the decision-maker makes a risky adaptation investment when multiple beliefs (climate projections) are valid at the same time. Our theoretical results show that updating expectations to account for climate change does not necessarily alter the timing of adaptation because the impact of risk and ambiguity offset each other. When the future climate becomes more variable, the investment in adaptation becomes riskier and a decision-maker will delay the adaptation process. Delaying allows the decision-maker to gather additional information on the damages avoided by investing in adaptation, and an increase in risk makes the delay more favorable as this increases the potential loss associated with maladaptation. On the other hand, as expectations over the future climate become more ambiguous, this speeds up adaptation. A more ambiguous future does not alter the potential loss associated with maladaptation, but decreases the chance of maladaptation as the agent believes a bad climate state is more likely to happen. Believing that the bad state is more likely to happen reinforces the value of the adaptation investment which creates an incentive to hasten adaptation. The counter-vailing effects of risk and ambiguity suggest that greater awareness of the damages created by climate change may not expedite adaptation investments. We apply the theory to an illustrative example of water conservation in California to show that the effect of risk and ambiguity do work in opposite

directions, leaving the timing of water conservation investments largely unaffected by greater awareness of climate change.

Ambiguity aversion theories have been applied to the climate change mitigation setting, although little research has focused on the effects of ambiguity on adaptation. Given the potential rebound or complementary effect between adaptation and mitigation efforts, it is essential to understand how ambiguity preferences affect adaptation behavior to improve the efficiency of climate mitigation policy. Heal and Millner (2013) point out that there exist two kinds of ambiguity in climate analysis. One is scientific uncertainty, which arises from the inability to precisely predict climate change due to the current limited knowledge of climate change, while the other one is socio-economic ambiguity and arises from the inability to precisely predict the impacts of climate change due to a limited understanding of how people react to a changing climate. Here, the literature suggests that ambiguity has a modest and nuanced effect on climate mitigation policies. Lemoine and Traeger (2012) discuss the impact of uncertainty in the climate tipping point, a temperature threshold causing irreversible changes once passed, and conclude that ambiguity aversion can amplify or dampen the effect of tipping points on optimal mitigation policy. Millner, Dietz, and Heal (2013) combine the smooth ambiguity model (Klibanoff et al. 2005) and DICE models and find increasing ambiguity aversion favors greenhouse gas abatement, but the magnitude is limited. For self-protection behavior, Roux (2018) studies the case of climate insurance and shows the importance of the historical climate record. However, Chambers and Tigran (2017) show that the optimal response to climate change can be sensitive to the particular specification used to model ambiguity.

Considering the effect of ambiguity on adaptation decisions introduces several additional challenges. For instance, adaptation is often characterized by the adoption of new technologies and

investments in protective infrastructure like sea walls and reservoirs that require large sunk costs. These investment decisions are influenced by climate-dependent risk (e.g., annual variability in temperature and precipitation). Such adaptation investments require a careful accounting of the effect of risk, ambiguity, and interactions between the two. An emerging theoretical literature addresses this type of investment by combining ambiguity theory with real options theory. Epstein and Chen (2002) provide the foundation of applying ambiguity theory in a dynamic, stochastic setting characterized by a Brownian motion. Nishimura and Ozaki (2004) show how the Epstein and Chen (2002) approach, paired with Choquet utility, can be applied to a job search problem (optimal stopping problem) to deal with ambiguity in expected salary. Nishimura and Ozaki (2007) use the Maxmin utility model (Gilboa and Schmeidler 2004) to show the influence of ambiguity on the value of an irreversible investment opportunity. Miao and Wang (2011) rely on the recursive multiple-priors utility model developed by Epstein and Wang (1994) and find that the effect of ambiguity differs for a firm's decision to enter a market and exit a market. Specifically, a higher ambiguity level will delay market entry while speeding up an exit.

In this paper, we extend the theoretical model in Nishimura and Ozaki (2007) and apply it to the problem of climate change adaptation. We consider a setting where a firm is considering when to install an adaptation technology that protects it from the effects of future climate change. However, the benefits received by the firm upon the adoption of the technology are uncertain due to future climate variability and ambiguity. Thus, our problem is comparable to Miao and Wang (2011), as the agent is paying a one-time fee to limit the impact of future uncertainty. Modern ambiguity theories can be classified into two groups (Ahn et al. 2014), the kinked specification and the smooth specification. Following Epstein and Ji (2013) and Miao and Wang (2011), our model follows the kinked specification and assumes the decision-maker cares more about the

potential worst belief, and a more ambiguous environment will increase the relative weight on the worst belief (Epstein and Wang 1994). Although the smooth specification has been applied in other climate change contexts (Millner et al. 2013, Jensen and Traeger 2014), the kinked specification is more convenient and conveys its own advantages. To fully define the utility function, the kinked specification only requires the set of priors and the ambiguity attitude parameter. In contrast, the smooth specification requires one to place a weight on each particular prior. Further, the kinked specification setting has been applied in many other settings, including participation in financial markets (Easley and O'Hara 2009) and the uptake of insurance by African farmers (Bryan 2019).

We apply our model to water conservation in California's Central Valley to empirically investigate the degree to which the impact of risk and ambiguity offset one another. Our model considers a farmer who must decide when to incur a sunk cost to adopt a more efficient irrigation technology and has available information from sixty-four climate models. This adaptation investment is risky because annual variation in precipitation makes the benefit of the more efficient technology unpredictable. Thus, there is an opportunity cost of adaptation (a forgone option value) since irreversible investment in adaptation limits the ability to respond to new information and, in turn, generates an incentive to wait (Treasury 2009, Margulis et al. 2010, Chambwera et al. 2014). Ambiguity and risk interact in that several of the climate change projections imply more variable precipitation (i.e., extreme floods and droughts). We solve for the optimal adoption time based on a measure of water availability that characterizes the farmer's efficient adaptation strategy, and investigate the effects of an increase in risk and ambiguity. We find that an increase in risk delays adaptation while an increase in ambiguity hastens adaptation. In our particular example, a 5% increase in risk increases the optimal adoption timing by 23.2% while a 5% increase in ambiguity hastens switching time by 1.23%. When considering the climate change projections, our results

show that the technology adoption decision of an ambiguity neutral farmer is very similar to an extremely ambiguity averse farmer.

3.2 Objective Risk, Ambiguity, and Climate Change Adaptation

To build the intuition underlying the countervailing effects of risk and ambiguity, we first present a simple two-period, two-state model. Consider a firm that uses water as an input to production. Examples include farms that depend on water for irrigation and electric utilities that depend on water for cooling. Period 1 conveys no uncertainty, while period 2 has two possible states, either wetter or drier than historical conditions. The firm must decide when to adopt new water-saving technology (adaptation) to mitigate potential profit losses due to limited access to water. The firm can either adopt the technology in period 1 or wait and adopt in period two if period 2 is characterized by a dry state. Adoption requires a one-time purchase and installation cost, which is a sunk cost. If the firm chooses to invest in the technology in period one, its expected profit can be written as

(1) $$E(\pi) = \pi_1^a - I + \beta\left(p_w \pi_{2,w}^a + (1 - p_w)\pi_{2,d}^a\right).$$

where π stands for the profit, I is the cost of investment, and $\beta = \frac{1}{1+\rho}$ is the discount factor. Superscript a means the technology has been adopted, n means not adopted, 1 and 2 denote the time period, and w and d stand for whether the future state is wetter or drier than the current period. The parameter p_w is the probability the future will be wetter than historical conditions. We assume the technology does not decrease period 2 profits regardless of future climate conditions: $\pi_{2,w}^n \leq \pi_{2,w}^a$ and $\pi_{2,d}^n \leq \pi_{2,d}^a$. If the firm did not adopt in period 1, we assume the firm only adopts the new technology in period 2 if the realized state is dry: $\pi_{2,w}^a - \pi_{2,w}^n < I < \pi_{2,d}^a - \pi_{2,d}^n$. The profit in period 1 does not have the state subscript as we assume no uncertainty.

On the other hand, if the firm waits until the second period, the expected profit is

(2) $$E(\pi') = \pi_1^n + \beta[(1 - p_w)\left(\pi_{2,d}^a - I\right) + p_w \pi_{2,w}^n] \,.$$

which is the discounted weighted average of profits in the dry state (in which case the technology is purchased) and the wet state (no technology is purchased), plus the profit based on using the traditional technology in period 1. Delaying technology adoption (i.e., not adopting in period 1) will be optimal when $E(\pi') > E(\pi)$, which implies

(3) $$\beta\, p_w\left(I + \pi_{2,w}^n - \pi_{2,w}^a\right) > \pi_1^a - \pi_1^n - I.$$

In other words, delaying will be optimal if the net benefit of adopting today (the right-hand side) is small enough.

Here we define risk as a mean-preserving spread. Thus, an increase in risk means the wet state becomes wetter while the dry state becomes even drier. This change increases the difference in profits from adaptation in the dry state, $\pi_{2,d}^a - \pi_{2,d}^n$, and decreases this difference in the wet state, $\pi_{2,w}^a - \pi_{2,w}^n$. An increase in risk thus increasing the value of the left-hand side of (3), which makes it more favorable to delay the adaptation decision. Since the cost of maladaptation in the wet state is expressed as $I - (\pi_{2,w}^a - \pi_{2,w}^n)$, equation (3) has an intuitive explanation. The increase in risk leads to a wetter wet state, thus increasing the cost of maladaptation. As the possibility of the wet state hasn't changed, the expected cost of maladaptation increases, and waiting becomes more favorable.

However, it is unrealistic to assume the firm knows with certainty the probability of the future wet and dry states. To relax this assumption, we assume the firm might hold multiple beliefs on p_w, which might come from different climate models. We define Δ as the set of all possible beliefs. Although we will relax this in the more general specification, for convenience, we assume

here that the firm is extremely ambiguity averse and focuses only on the worst case. Then we represent his expected discounted profit from adopting in the first period as

(4) $$E(\pi) = \pi_1^a - I + \beta \min_{p_w \in \Delta} \left(p_w \pi_{2,w}^a + (1 - p_w)\pi_{2,d}^a\right),$$

where $\pi_{2,w}^a > \pi_{2,d}^a$. If the firm postpones technology adoption, the discounted expected profit will be

(5) $$E(\pi) = \pi_1^n + \beta \min_{p_w \in \Delta} p_w \left[(1 - p_w)\left(\pi_{2,d}^a - I\right) + p_w \pi_{2,w}^n\right],$$

with $\pi_{2,w}^n + I > \pi_{2,w}^a > \pi_{2,d}^a$ given the assumptions above. With Knightian uncertainty, the condition for delayed adaptation to be optimal is

(6) $$\beta \left[\left(\min_{p_w \in \Delta} p_w\right)\left(I + \pi_{2,w}^n - \pi_{2,w}^a\right)\right] > \pi_1^a - \pi_1^n - I,$$

which will be lower than the case without Knightian uncertainty.

When the situation becomes more ambiguous, we assume the possible range of priors increases. In other words, the lowest possible p_w becomes even smaller and the highest potential p_w becomes larger. Since the ambiguity averse firm focuses on the lowest p_w, the increase in ambiguity causes the left-hand side of equation (6) to decrease. Thus, the agent will speed up adaptation. Intuitively, the increase in ambiguity makes the possibility of mal-adaptation deceases while the difference in profits from adopting is not affected. As a result, the expected cost of maladaptation decreases and makes adopting in period 1 more favorable.

While a two-period, two-state model is an over-simplification of the complex dynamics in play under climate change, this example does clearly highlight the countervailing effects of objective risk and ambiguity. We now extend the two-period model into a continuous-time adaptation problem. At any point in time, the firm produces a product using water $N(t)$ whose

supply is exogenously determined by weather at time *t*. The firm has an option to install a more efficient technology with a sunk cost of M. With the available water input, the firm earns a profit

(7) $\pi(t) = a_i(N(t))^{b_i}$.

with $a_i > 0$ and $b_i < 0$ to signal decreasing marginal returns from the water input.[7] The subscript i is an indicator that equals 1 if the firm has adopted the technology and otherwise equals 0. The new technology increases profits, $a_1 > a_0$, and lessens the sensitivity of profit to changes in water supply, $b_1 < b_0$. To maximize long-run profit, the firm must decide the optimal time t^* to switch from the existing technology, a_0 and b_0, to the conservation technology, a_1 and b_1, to decrease the potential loss from decreasing water availability.

To determine the optimal adoption time, the firm uses available climate projections to form expectations of future water availability and the benefits of adaptation. The firm's expectation of future water availability follows geometric Brownian motion

(8) $dN = \alpha N(t)dt + \sigma N(t)dB$

where α is the instantaneous drift rate of the water supply process, σ^2 is the instantaneous variance, and dB is the increment of a Brownian motion with respect to the probability measure $\mathcal{P}$. However, due to the variation across climate projects, the firm has a set of beliefs Q about future water supply $N(t)$ instead of a single probability measure. The set of beliefs is defined as a set of probability distributions, which assigns each possible state (amount of water supply) with a probability. To model Knightian uncertainty in continuous time, we follow Chen and Epstein (2002) and make two assumptions about the beliefs of the firm. The first assumption is that the firm only considers a set of probability measures that have a perfect agreement with $\mathcal{P}$ with respect

[1] It is possible this assumption may be violated when the firm holds a large water right and this water right can be leased to another firm via a water market. In this case, increased water availability lowers the value of the water right. For exposition, we assume that any water rights held by the firm are sufficiently small so as to make the firm a net water buyer.

to zero probability events. This assumption allows for the probability measures to be continuous with respect to one another and allows for a set of Ito processes to be generated via a density generator.[8] We define the set of priors as follows:

(9) $Q^{\theta} = \{\mathcal{P}^{\theta} | \theta \in \Theta\}$,

where θ stands for a single density generator, Θ is a set of density generators, $\mathcal{P}^{\theta}$ is the probability measure generated by θ, and Q^{θ} is the set of all the alternative probability measures. According to the Giranov's Theorem, for each θ we have $dB^{\theta} = dB + \theta dt$. Combining this result with equation (8), we have for any $\theta \in \Theta$: $dN = [\alpha N(t) - \sigma N(t)\theta]dt + \sigma N(t) dB^{\theta}$. Thus, when the firm is following a certain prior, it is equivalent to him following a single stochastic process. Under this continuous time concept of Knightian uncertainty, the firm considers all stochastic differential equations with $\theta \in \Theta$ varying. In other words, there exists only a single stochastic process for $N(t)$ but this process has many interpretations characterized by a set of stochastic differential equations.

The second assumption about climate change beliefs is that the range of priors are bounded by the climate projections. Define θ_w and θ_b as the density generator that brings the worst and best expected profit: $\theta \in [\theta_w, \theta_b]$[9]. With the boundary condition and a hypothetical θ^* that generates the true stochastic process, in the rest of this section we follow Miao and Wang (2011) and use a weighted average between the worst (or best) boundary and θ^* to illustrate how the perceived ambiguity influences decision making.

[2] This continuous time version of Knightian uncertainty is a weaker version of the discrete time concept. According to Girsanov's Theorem, an equivalent probability measure can be generated by an original probability measure via a density generator.

[3] The specification of Knight uncertainty in Nishimura and Ozaki (2007) is a special case based on our setting, where $\theta_w = -k, \theta_b = k$, together with $\theta^* = 0$.

These two assumptions allow us to define the set of priors directly over the drift and volatility terms and the geometric Brownian motion process:

(8') $dN = \alpha_q N(t)dt + \sigma_q N(t)dB^q,$

where $q \in Q$ indexes a single belief and α_q and σ_q are the drift and volatility terms based on this belief. As shown in the appendix, firm profit also follows a geometric Brownian motion

(9) $d\pi = \alpha_{i,q}^p \pi(t)dt + \sigma_{i,q}^p \pi(t)dB^q$

where $\alpha_{i,q}^p = \left[b_i\alpha_q + \frac{1}{2}b_i(b_i-1){\sigma_q}^2\right]$ and $\sigma_{i,q}^p = b_i\sigma_q$. Because $b_i < 0$, increased variability in water supply lowers α_{iq}^p and the expected rate of change in firm profits through Jensen's Inequality. Since $b_1 < b_0$, the conservation technology lowers the variability in firm profits: $\sigma_{1,q}^p < \sigma_{0,q}^p$. However, the conservation technology has an ambiguous effect on the drift rate in profits.

Equation (9) implies that the firm has multiple beliefs about future profit due to multiple beliefs about future water availability. When the firm is ambiguity averse, it cares more about the worst possible future, which we denote q_w. To maximize the present value of profits, the firm needs to choose the optimal time t^* to switch to more efficient technology. At any point in time, the firm assesses the tradeoffs between buying and installing the technology now or waiting until $t + dt$, and the optimal stopping time is when the firm is indifferent between the two options. Mathematically, the firm solves the optimal stopping problem defined as

(10) $V(\pi_\tau) = \max_{t^*} E\left[\int_\tau^{t^*} \pi_0(t)e^{-\rho(t-\tau)}\,dt + \left\{\left[\int_{t^*}^{\infty} \pi_1(t)e^{-\rho(t-t^*)}\,dt - M\right]e^{-\rho(t^*-\tau)}\right\}|\wp\right],$

$\wp = \{(1-\varepsilon)\mu + \varepsilon q_w : q_w = \underset{q\epsilon Q}{argmin}\, V(\pi_t)\},$

subjects to equation (9) and where μ is any distribution over $\pi(t)$ (i.e., any of the set of stochastic differential equations in equation (9), and ε is a probability between zero and one. Here we define μ as the expected distribution used by an ambiguity neutral agent. Following the behavioral

literature (Halvey 2007, Ahn et al. 2014), we assume an ambiguity neutral firm places equal weight on each climate projection.

Following Miao and Wang (2011) and Epstein and Wang (1994), the ambiguity specification can be interpreted as follows: in a setting characterized by ambiguity, the decision maker has a probability of ε of taking the worst possible scenario as the realized pathway and 1-ε of believing the ambiguity-neutral distribution. Thus ε stands as a measure of impact of the ambiguity, and an increase in ε is caused by either an increase in the perceived ambiguity or an increase in the ambiguity aversion level, or both. When $\varepsilon = 0$, the decision-maker acts as if he is making decisions under μ and using expected utility, when $\varepsilon = 1$, the decision maker is extremely ambiguity averse and bases the optimal stopping decision solely on the worst-case scenario. For any $\varepsilon \in (0,1)$, the decision maker will generate a $\wp$ as the weighted average of the ambiguity-neutral future and the worst-case future, and act as if $\wp$ is his single belief. Based on this kinked specification of ambiguity, we have $\alpha_{i,q}^{p} = \varepsilon\alpha_{0,\omega}^{p} + (1-\varepsilon)\overline{\alpha_0^{p}} = \varepsilon\left[b_0\alpha_\omega + \frac{1}{2}b_0(b_0-1){\sigma_\omega}^2\right] + (1-\varepsilon)\left[b_0\bar{\alpha} + \frac{1}{2}b_0(b_0-1)\bar{\sigma}^2\right]$. For convenience, here we make one additional assumption, $\sigma_{0,q}^{p} = b_0\sqrt{\varepsilon\,{\sigma_\omega}^2 + (1-\varepsilon)\,\overline{\sigma}^2}$. With the additional assumption, it is equivalent to assume either the firm is ambiguous about future water supply or future profit.

Prior to adopting the conservation technology, the optimal value function, V_0, must balance the required return from delaying adoption, ρV_0, against the expected return which is the sum of today's dividend, $\pi(t)$, and the expected appreciation/depreciation of the value function. Taking advantage of the geometric Brownian motion, mathematically this relationship can be expressed as:

$$(11) \qquad \rho V_0 = \pi(t) + \frac{E[dV_0]}{dt} = \pi(t) + \alpha_{0,q}^{p}\pi(t)\frac{\partial V_0}{\partial \pi} + \frac{(\sigma_{0,q}^{p}\pi(t))^2}{2}\frac{\partial^2 V_0}{\partial \pi^2}.$$

where $\alpha_{0,q}^{p} = \varepsilon\left[b_0\alpha_\omega + \frac{1}{2}b_0(b_0-1){\sigma_\omega}^2\right] + (1-\varepsilon)\left[b_0\bar{\alpha} + \frac{1}{2}b_0(b_0-1)\bar{\sigma}^2\right]$ and $\sigma_{0,q}^{p} = b_0\sqrt{\varepsilon\,{\sigma_\omega}^2 + (1-\varepsilon)\,\overline{\sigma}^2}$. To guarantee the existence of solution, we need two transversality (boundary) conditions between the region of the state space where the old technology is optimal (the continuation region) and the region of the state space where the new technology is preferred (the termination region) (Dixit and Pindyck 1994). The first condition is known as the value matching condition

(12) $\frac{\pi^*}{\rho-\alpha_{0,q}^{p}} + \eta{\pi^*}^{\kappa} = \frac{\pi^*}{\rho-\alpha_{1,q}^{p}} - M.$

The value matching condition basically states that at the switch point the firm is indifferent between adapting the technology or not. The left hand side of equation (12) is the sum of the expected net present value of firm profits with the old technology, $\frac{\pi^*}{\rho-\alpha_{0,q}^{p}}$, and the option value association with the investment in the conservation technology, $\eta{\pi^*}^{\kappa}$. Following Dixit and Pindyck (1994) η is a unknown constant that needs to be solved later, and κ is the positive solution of the fundamental quadratic function $\frac{1}{2}{\sigma_{0,q}^{p}}^2\kappa(\kappa-1) + \alpha_{0,q}^{p}\kappa - \rho = 0$. The term associated with the negative root of the fundamental quadratic is eliminated by the boundary condition $V_0(0) = 0$, implying the value of investing in a new technology in an industry that generates zero profit is 0. The right hand side of equation (12) is the sum of the expected net present value of firm profits with the conservation technology, $\frac{\pi^*}{\rho-\alpha_{1,q}^{p}}$ net of the sunk investment cost. If $\alpha_{1,q}^{p} > \alpha_{0,q}^{p}$, the expected benefit of adopting the conservation technology is positive and the technology will eventually be adopted. If $\alpha_{1,q}^{p} < \alpha_{0,q}^{p}$, the expected benefit of the technology is negative, and a firm will never adopt it. We also make the standard assumption that $\rho > \alpha_{i,q}^{p}$ to unsure the expected net present values converge to finite values.

The second boundary condition is the smooth pasting condition which ensures the marginal value of firm profits is equal before and after adopting the technology:

(13) $\frac{1}{\rho-\alpha_{0,q}^{p}}+\varphi\eta\pi^{*(\kappa-1)}=\frac{1}{\rho-\alpha_{1,q}^{p}}$.

Solving equation (12) and (13), we obtain the closed form solution of adaptation threshold, π^{*} as

(14) $\frac{\kappa}{(\kappa-1)}\frac{\left(\rho-\alpha_{0,q}^{p}\right)\left(\rho-\alpha_{1,q}^{p}\right)}{\alpha_{0,q}^{p}-\alpha_{1,q}^{p}}M=\pi^{*}$.

When the current profit is smaller than the adaptation threshold, $\pi(t) \quad <\pi^{*}$, the firm installs the technology immediately, while if the current profit is higher than the threshold then the firm will wait. This profit threshold can also be rewritten as the critical level of water supply that triggers adaptation:

(15) $N^{*}=\left(\frac{\pi^{*}}{a_0}\right)^{1/b_0}$.

To understand how the ambiguity attitude influences the adaptation decision, here we follow Nishimura and Ozaki (2007) and focus on the value of technology adaptation. Taking advantage of geometric Brownian motion, the value of the conservation technology is expressed as:

(16) $\int_{t^{*}}^{\infty}\pi^{*}\,e^{\left(-\rho+\varepsilon\alpha_{1,\omega}^{p}+(1-\varepsilon)\overline{\alpha_{1}^{p}}\right)(t-t^{*})}dt-\int_{t^{*}}^{\infty}\pi^{*}e^{\left(-\rho+\varepsilon\alpha_{0,\omega}^{p}+(1-\varepsilon)\overline{\alpha_{0}^{p}}\right)(t-t^{*})}dt$.

Here $\alpha_{i,\omega}^{p}$ and $\overline{\alpha_{i}^{p}}$ corresponding to the drift term of geometric Brownian motion for worst and ambiguity neutral prior, $\overline{\alpha_{i}^{p}}>\alpha_{i,\omega}^{p}$. Equation (16) simplifies to

(17) $\frac{\pi^{*}}{\lambda_1}-\frac{\pi^{*}}{\lambda_0}$, where

$\lambda_i=\rho-\varepsilon\alpha_{i,\omega}^{p}-(1-\varepsilon)\overline{\alpha_{i}^{p}}$.

Equation (17) tells us that when the firm calculates the value of new technology, it will separately consider the worst changing rate and the average changing rate given by different climate

projections, and calculate the expected changing rate based on the value of ε. Here we have $\frac{d\lambda_i}{d\varepsilon} = \bar{\alpha}_i - \alpha_{i,\omega}$, which is always positive since by definition $\alpha_{i,\omega}$ is the lowest changing rate of future profit. We also have $\frac{d\lambda_0}{d\varepsilon} > \frac{d\lambda_1}{d\varepsilon}$ since the new technology reduces the impact of water supply and limits the decrease in the drift gap between the worst prior and the average[10]. An increase in ε leads to a decrease in profit both before and after technology adaptation, however, the decrease is larger under the old technology. Thus, the expected profit from technology adaptation, which is measured by the differences in the net present value before and after switching technology, increases as ε increases which favors adapting earlier. The intuition is that when people put a higher weight on the drier future, the value of the conservation technology will increase, leading to a faster adaptation.

3.3 An Illustrative Example: Irrigated Agriculture in California's Central Valley

In this section, we apply our theoretical model to the decision to invest in more efficient irrigation technology in the northern portion of California's Central Valley - the most productive agricultural region in the country. All parameter values are presented in Table 1. Agricultural production in the Sacramento Valley relies on irrigation and over 70 percent of the annual irrigation supply comes from surface water. Surface water is annually replenished by melting snowpack in the Sierra Nevada Mountains and brought to agricultural producers in the valley via several rivers. The predicted decrease in water availability in recent climate change scenarios downscaled to California is expected to decrease farmland values (Schlenker et al. 2007). One potential response by farmers is to invest in more efficient irrigation technologies. Between 1970

[4] Previously we have assumed $b_0 > b_1$. Since $\alpha_{i,q}^p = b_i\alpha_q + \frac{1}{2}b_i(b_i - 1){\sigma_q}^2$, when $0.5 > b_0 > b_1 > 0$, the old technology brings a higher range of the drift rate of profit.

and 2010, the irrigated area in California that uses inefficient surface irrigation methods such as furrow or flood irrigation has fallen by 44 percent (Tindula, et al. 2013). Whether this transition continues as farmers update their expectations of climate change depends on how expectations account for risk and ambiguity. If climate change is expected to increase the variability in snowfall, investing in more efficient irrigation will become riskier and the shift to more efficient irrigation technologies will slow. However, since climate change introduces ambiguity, it may also create an incentive to speed up the rate of transition to more efficient irrigation.

3.3.1 Irrigated production and water markets

Consider a risk-neutral farmer that uses physical capital, k, and irrigation water, n, to produce a single agricultural commodity $y = f_0(k, n)$ with diminishing returns to scale technology f_0. Let i = 0 denote production using the farmer's current surface irrigation method (e.g, furrow or flood irrigation) and i = 1 correspond to production using a less water-intensive irrigation technology (e.g., center pivot or drip irrigation). Following Berck and Helfand (1990), Letey (1991), and Carey and Zilberman (2002), the farmer's production follows a Von-Liebig type production function

$$(18) \qquad y_i = f_i(k, n_i) = \begin{Bmatrix} \gamma_i k n_i \; when \; n_i < n_i^* \\ y^* \; when \; n_i \geq n_i^* \end{Bmatrix},$$

where y^* is the farmer's optimal level of production per acre, n_i^* is the farmer's optimal water demand under technology i (what he would require for optimal production if unconstrained by water supply), and n_i is the amount of water employed in production. The Von-Liebig (or plateau) model assumes that the production is first responding to the addition of water, and then limited by the capital input (Letey 1991). When $n_i < n_i^*$, water is the limiting input, and at the margin

increases production by $\gamma_i k$, where γ_i is a scalar that reflects the efficiency of technology i with $\gamma_1 > \gamma_0$. Another production input (e.g., soil quality or canal size) becomes limiting at $n_i \geq n_i^*$.

In each period, the farmer is endowed with some water rights $A(N(t))$, which allows him to access a certain level of water with zero cost:

$$(19) \quad A(N) = \begin{Bmatrix} \theta\widetilde{N} \; if \; N(t) \geq \widetilde{N} \\ \theta N(t) \; if \; N(t) < \widetilde{N} \end{Bmatrix}.$$

Here, $N(t)$ is the exogenous supply of water at time t proxied by the flow of a local river, $\widetilde{N}$ is a scarcity threshold which is exogenously determined, and θ is a ratio which implies the amount of cost-free water given the current stock level. When water is relatively abundant, the farmer is guaranteed an amount of water $\theta\widetilde{N}$. During periods of water scarcity, the farmer will be restricted to a fixed proportion of the total stock $\theta N(t)$.

The relationship between available water $A(N)$ and applied water n_i is as follows;

$$(20) \quad n_i = \begin{Bmatrix} n_i^* \; if \; A(N) \geq n_i^* \\ n_i < \; n_i^* \; if \; A(N) < \; n_i^* \end{Bmatrix}.$$

Thus, $A(N) - \; n_i$ is the difference between water the farmer has available for production and how much he actually employs in production. If positive ($A(N) > \; n_i$), this difference represents the total amount of water left over after production and available for sale by the farmer. If negative ($A(N) < \; n_i$), this difference indicates how much water the farmer buys to supplement his available water for production. Investment in efficient irrigation technology helps tilt the difference of the two terms towards a surplus by reducing n_i.

The farmer may buy and lease out water in a formal water market. The demand for water is isoelastic:

$$(21) \quad p(N) = \left(\frac{N(t)}{\varphi}\right)^{-\frac{1}{e}},$$

where $e > 0$ is the constant price elasticity of demand and φ is a positive demand parameter. To predict the future water price, we use data from the Water Transfer Level Dataset compiled by researchers at UC-Santa Barbara to recover the demand parameter φ and the elasticity of water demand in the water market. The Water Transfer Data Base is funded by the National Science Foundation and the California Water Resources Research Center and presented by the Bren School. The dataset includes the year of a water transfer, the acquirer of the water, the supplier, the amount of water transferred, the proposed use of the water, and the real price of the trade (in 1987 dollars). To generate a representative sample for the study area, we limit the transaction data to northern California and at least one of the engagement parties needs to be an agriculture user. This leaves us with 196 transactions from the year 1989 to 2009. Using an ordinary least squares regression, we estimate the elasticity of demand $e = 1.94$ and φ = 899.[11]

We assume the farmer knows the price P for the final product and rent r for capital and solves the following profit-maximization problem:

$$\pi_i(N(t)) = \max_{k,n} P f_i(k,n) - p\big(N(t)\big)\big(n_i - A(N(t))\big) - rk. \tag{22}$$

where π_i stands for the farmer's profit, f_i stands for the production technology. Because farm profits are linear in the amount of water applied, the first-order condition arising from profit maximization problem suggests a bang-bang solution. When $p(N) \leq P\gamma_i k$ the farmer optimally chooses to apply n_i^*. Whether a farmer is a buyer or seller of water depends on the relative magnitude of the farmer's water right and optimal water demand. If $A(N) < n_i^*$, the farmer will purchase water to produce y^* given the von Liebig production function. The need to purchase water will arise when the famer's water rights are lower than the optimal water demand ($\theta\widetilde{N} <$

[5] We first take the log of equation 21 and estimate φ with a restricted OLS.

n_i^*), or aggregate water supplies are sufficiently low. If the farmer's water right exceeds his optimal water demand $(\theta\widetilde{N} > n_i^*)$, $N(t)$ must fall further than $\widetilde{N}$ to trigger water purchases. The von Liebig production function also implies that the farmer will sell excess water when water supplies are relatively abundant $A(N) > n_i^*$. Water sales arise when the farmer's water right exceeds his optimal water demand $(\theta\widetilde{N} > n_i^*)$. Otherwise, an investment in the more efficient technology is needed before the farmer will ever choose to sell water.

But when $p(N) > P\gamma_i k$, the farmer will terminate production and lease all of his available water since the value marginal product of water in irrigated agriculture is lower than the price of water on the spot market. Since it would no longer be in use, the water conservation technology provides no value to the farmer. This suggests that the benefit of water conservation is greatest at intermediate water supply levels. In times of ample aggregate water supplies, the price of water and the benefits of the water conserved are low. In times of severe water scarcity, the price of water will be high, the farmer will lease all water rights and has no need for new irrigation technology.

The profit-maximization problem detailed above implies a power function relationship between profit and water. To estimate the parameters a_i and b_i, we fit the profit function in equation (7) to the observed relationship between water supply and profits implied by the profit maximization problem. We separately calculate the profit from 0.1 million acre-feet to 60 million acre-feet before and after the technology adaptation. The estimation results are included in Table 1. Compared with the traditional technology, the new irrigation technology brings a larger scale coefficient together with a lower power term, which indicates an overall larger profit in the study region along with a lower sensitivity to changes in the water supply. In addition, the gap between

the two profit functions decreases as the level of water increase. This result corresponds to our assumption in the theoretical section.

3.3.2 Climate change beliefs

We use historical data and climate projections of the Tuolumne River streamflow, a single basin in central California, to estimate streamflow trend, variability, and climate uncertainty. Our data includes the annual streamflow in the Tuolumne River at LaGrange from 1951 to 2000 to estimate historical streamflow trends and variability. Our climate projections come from Coupled Model Intercomparison Project Phase 5 (CMIP5) (U.S. Department of the Interior & Bureau of Reclamation 2013, U.S. Department of Interior & Bureau of Reclamation 2014). CMIP5 projections are the most recently published data, and provide the latest set of global greenhouse emissions scenarios, called representative concentration pathways (RCPs). The data contains four pathways, namely RCP2.6, RCP4.5, RCP6, and RCP8.5, which are labeled based on radiative forcing values in the year 2100 (2.6, 4.5, 6.0, and 8.5 W/m2, respectively). Our climate data (air temperature, precipitation, and relative humidity) are from 16 general circulation models (GCMs). Together with 4 RCPs, our data contains 64 different climate projections of $N(t)$, which capture regional meteorology and a range of emissions and economic conditions from the year 2000 to the year 2100. We treat these 64 climate projections as the set of possible climate beliefs possibly held by the farmer: $Q = 64$.

Corresponding to the theoretical assumption, we assume water supply follows a geometric Brownian motion process. We estimate the parameters in equation (8) for each of the 64 climate projections. Each estimated stochastic process succinctly characterizes a single belief about future climate. Figure 1 compares the estimated parameter values for the 64 climate projections (blue color) and the parameter values for the same stochastic process estimates from historical

streamflow data (green color). As can be balanced from the Figure, first, compared with the historical streamflow, most projections convey a lower trend of water flow. In addition, almost all the projections predict lower volatility compared with the historical record. The results suggest that extreme weather events (droughts and floods) will be less likely in the future due to climate change, which makes this type of adaptation investment more favorable.

Based on our theory specification (Appendix), we also calculate the drift rate and volatility of profit. Figure 2 compares the estimated parameter values estimated by the 64 climate projections and the historical data. Compared with the historical result, about half of the projections suggest a decrease in further profit while others suggest an increase. With the adaptation of technology, we notice a reduction in both the range of the drift term and volatility, which corresponds to our theoretical specification as the new technology reduces the impact of water supply. In addition, we find that a higher drift term or lower volatility in the water supply implies a smaller increase in profit. Based on our result, miroc5.1.rcp26 is selected as the worst projection since it provides the lowest drift rate of profit. Following previous experimental research on ambiguity, we define the 64 projections as the set of priors, and each of them encompasses all the possible future states. We take the average of the 64 projections and treat it as the ambiguity neutral prior in the approximation. Figure 3 represents the comparison between the worst projections, the average across projections, and the projections based on historical data. From both the parameter values and the graph, the average projection brings a higher increase in profit, together with higher volatility.

3.3.3 Results

To investigate how the farmer's ambiguity attitude might impact the adaptation decision, we estimate how the optimal adoption time changes as the framer becomes ambiguity averse. We

take advantage of the closed form solution and directly calculate the adaptation threshold for different scenarios. Based on our parameter values, our profit function has a relatively large constant term together with a relatively smaller value of power term. Together with our mathematical expression for the drift term of profit, our parameters of profit function make the adaptation time response slowly to the changes in the drift rate of water supply while responding sharply to the changes in the volatility of streamflow.

We choose the flow level in the year 2002 as our initial condition to calculate the optimal switching time. Our empirical analysis provides the adaptation threshold and the trend of water supply. Those two results together provide us with the expected hitting time. Although our data based on the 100-years projections, with a lower threshold and slow changing speed of water supply, the optimal switching time could exceed 100 years. Since we chose the year 2002 as the initial point, our range of optimal adaption timing is from 2002 to infinity. As an example, if the optimal adoption year is 150 years later, it means switch technology at the year 2152. Table 2 summarizes the optimal conservation technology adoption time when farmers do and do not update expectations of water availability to account for climate change. Compared with most projections, historical data leads to the highest average supply of water and the highest changing rate of profit. Based on our result, the framers will never adopt new technology if they only consider the historical data as the expected water supply keeps increasing. When an ambiguity neutral framer considers the climate projections and puts equal weight on each prior, he recognizes that the future is drier relative to the historical record and switches the technology in 241 years. On the other hand, when the framer is extremely ambiguity averse and focuses only on the worst projection, he will

become conservative and install the technology 193 years. Regardless of ambiguity preferences, updating expectations to account for climate change will hasten adaptation.[12]

Table 3 considers how changes in risk and ambiguity impact the adoption time for different combinations for the ambiguity and risk level. In the table, we have two levels of ambiguity, extremely averse ($\varepsilon = 1$) and ambiguity neutral ($\varepsilon = 0$), together with three levels of risk, low (95% of original volatility), benchmark, and high (105% of original volatility). Across different risk levels, an increase in the risk level delays the optimal adoption time, which means the farmer has less incentive to adapt, and prefers to wait until the water supply becomes even lower. This result fits with the prediction of Pindyck (1991), where the author claims an increase in risk leads to longer waiting. Across different volatility levels, an increase in ambiguity aversion leads to an earlier adaptation, as the framer becomes more concerned about the drier future and has a higher incentive to switch technology. This result fits with our theoretical prediction.

In Table 4 we present results from a sensitivity analysis. In particular, we calculate how changes in the underlying parameter values alter the optimal timing decision. For brevity we focus on the adaptation time for the extremely ambiguity-averse farmer. In general, our results show that any parameter change that increases the additional profits generated by the conservation technology hastens adaptation, while any change that reduces the additional profits generated by the conservation technology delays the investment. Adoption times are incredibly sensitive to changes in a_i, which is a multiplicative parameter in the production function. The results suggest that seemingly subtle differences in this term between the adoption/no adoption states, i.e., $a_1 - a_0$, can dramatically alter the optimal decision from one of adopting immediately to not adopting in the foreseeable future. In contrast, the optimal timing of adaptation is relatively insensitive to

[6] Our general results hold even if the water supply follows a mean-reverting process. See Appendix B for additional details.

the curvature of the profit function b_i. To understand this result, note that b_i controls the relative influence of the expected rate and volatility in streamflow on the expected rate of change in profits: $\alpha_{i,q}^{p} = \left[b_i \alpha_q + \frac{1}{2} b_i (b_i - 1) {\sigma_q}^2\right]$. When the drift rate and volatility rate of the water supply process are of similar magnitude, shifting the weight between these terms will have a minimal effect on the expected change in profit and the optimal timing of adaptation all else constant. The optimal timing of adaptation is also relatively insensitive to the expected rate of change in streamflow. The extreme curvature of the profit function (low value of b_i) consistent by the profit maximization problem facing the farmer implies the expected rate of change in profits is relatively insensitive to changes in the expected rate of change in the water supply. The greater the curvature in the profit function, the more the timing of adaptation is driven by the volatility in streamflow as opposed to the expected change in streamflow.

3.4 Conclusion

This chapter investigates the impact of changing climate uncertainty on a firm's adaptation decision. We separately consider two types of uncertainty. One is the objective risk as characterized by the volatility in a single climate projection, and the other one is the ambiguity, which depends on the number of climate projections considered. Using both theoretical analysis and numerical simulation, we find that when the volatility in one projection increases, the decision-maker has a higher incentive to wait since the expected cost of maladaptation increases. On the other hand, when the firm's degree of ambiguity aversion, or the number of projections increases, the decision-maker will accelerate the installment of technology as he believes that the natural resources are more likely to be scarce.

A successful climate policy requires both mitigation and adaptation. Besides understanding the relationship between the two kinds of policies, a behavioral model that helps researchers better

predict an individual's adaptation behavior is also essential. By considering ambiguity aversion, we emphasize how the information provided by the government is potentially important to the firm's optimal response to climate change. Besides the information itself, how the information is provided also matters. When too much conflicting information is present at the same time, along with the information within each single projection, the complicated environment also affects the adaptation behavior. Thus, when the government has a clear individual adaptation target, our results provide another channel to help influence the adoption decision.

Our research provides an illustration of how ambiguity theory can be applied to predict adaptation behavior in the climate change context. Further refinements to the modeling are of course possible, and may lead to additional insight. First, the ambiguity specification we consider, while typical to the closely related literature, does not distinguish the effects of ambiguity attitude (i.e., how adverse to ambiguity the decision-maker is) from the effects of a change in ambiguity (Billion et al. 2018). Second, we assumed the demand parameter in the model is non-stochastic. However, it may be more realistic to assume this follows another Ito process. Third, we do not consider both adaptation and mitigation. As stated by several previous papers, a systemic welfare analysis requires the consideration of both (Bosello, Carraro et al. 2010, Hashida and Lewis 2019, Li and Rus 2019, De Bruin, Dellink et al. 2009, Wang and McCarl 2013). Thus, our model would need to be combined with possible mitigation mechanisms to better understand the implications of ambiguity for optimal policy. The design of optimal policy would further need to specify parameters of the ambiguity model, such as the degree of ambiguity aversion.

References

Ahn, D., Choi, S., Gale, D., & Kariv, S. (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5(2), 195-223.

Baillon, A., Huang, Z., Selim, A., & Wakker, P. P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86(5), 1839-1858.

Berck, P., & Helfand, G. (1990). Reconciling the von Liebig and differentiable crop production functions. *American Journal of Agricultural Economics*, 72(4), 985-996.

Bosello, F., Carraro, C., & De Cian, E. (2010). Climate policy and the optimal balance between mitigation, adaptation and unavoided damage. *Climate Change Economics*, 1(2), 71-92.

Bryan, G. (2019). Ambiguity aversion decreases the impact of partial insurance: Evidence from African farmerfirms. *Journal of the European Economic Association*, 17(5), 1428-1469.

Carey, J. M., & Zilberman, D. (2002). A model of investment under uncertainty: modern irrigation technology and emerging markets in water. *American Journal of Agricultural Economics*, 84(1), 171-183.

Chambers, R. G., & Melkonyan, T. (2017). Ambiguity, reasoned determination, and climate-change policy. *Journal of Environmental Economics and Management*, 81, 74-92.

Chambwera, M., Heal, G., Dubeux, C., Hallegatte, S., Leclerc, L., Markandya, A., ... & Neumann, J. E. (2014). *Economics of adaptation*.

Chen, Z., & Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70(4), 1403-1443.

De Bruin, K. C., Dellink, R. B., & Tol, R. S. (2009). AD-DICE: an implementation of adaptation in the DICE model. *Climatic Change*, 95(1-2), 63-81.

Dinar, A., & Letey, J. (1991). Agricultural water marketing, allocative efficiency, and drainage reduction. *Journal of Environmental Economics and Management*, 20(3), 210-223.

Dixit, A. K., Dixit, R. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton university press.

Easley, D., & O'Hara, M. (2009). Ambiguity and nonparticipation: The role of regulation. *The Review of Financial Studies*, 22(5), 1817-1843.

Epstein, L. G., & Ji, S. (2013). Ambiguous volatility and asset pricing in continuous time. *The Review of Financial Studies*, 26(7), 1740-1786.

Epstein, L. G., & Wang, T. (1994). Intertemporal asset pricing under Knightian uncertainty. *Econometrica*, vol. 62, 283-322.

Gilboa, I., & Schmeidler, D. (2004). Maxmin expected utility with non-unique prior. In *Uncertainty in economic theory* (pp. 141-151). Routledge.

Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica*, 75(2), 503-536.

Hashida, Y., & Lewis, D. J. (2019). The intersection between climate adaptation, mitigation, and natural resources: An empirical analysis of forest management. *Journal of the Association of Environmental and Resource Economists*, 6(5), 893-926.

Heal, G., & Millner, A. (2013). *Uncertainty and decision in climate change economics* (No. w18929). National Bureau of Economic Research.

Jensen, S., & Traeger, C. P. (2014). Optimal climate change mitigation under long-term growth uncertainty: Stochastic integrated assessment and analytic findings. *European Economic Review*, 69, 104-125.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3), 930-976.

Le Roux, S. (2018). Climate change catastrophes and insuring decisions: A study in the presence of ambiguity. *Journal of Economic Behavior & Organization*.

Lemoine, D. M., & Traeger, C. P. (2012). *Tipping points and ambiguity in the economics of climate change* (No. w18230). National Bureau of Economic Research.

Li, H., & Rus, H. A. (2019). Climate Change Adaptation and International Mitigation Agreements with Heterogeneous Countries. *Journal of the Association of Environmental and Resource Economists*, 6(3), 503-530.

Margulis, S., Hughes, G., Schneider, R., Pandey, K., Narain, U., & Kemeny, T. (2010). *Economics of adaptation to climate change: Synthesis report.*

Miao, J., & Wang, N. (2011). Risk, uncertainty, and option exercise. *Journal of Economic Dynamics and Control*, 35(4), 442-461.

Millner, A., Dietz, S., & Heal, G. (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics*, 55(1), 21-46.

Nishimura, K. G., & Ozaki, H. (2004). Search and Knightian uncertainty. *Journal of Economic Theory*, 119(2), 299-333.

Nishimura, K. G., & Ozaki, H. (2007). Irreversible investment and Knightian uncertainty. *Journal of Economic Theory*, 136(1), 668-694.

Pachauri, R. K., Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., ... & Dubash, N. K. (2014). *Climate change 2014: synthesis report. Contribution of Working Groups I, II and III to the fifth assessment report of the Intergovernmental Panel on Climate Change* (p. 151). Ipcc.

Pindyck, R. S. (1990). Irreversibility, uncertainty, and investment (No. w3307). National Bureau of Economic Research.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. "Water availability, degree days, and the potential impact of climate change on irrigated agriculture in California." *Climatic Change* 81.1 (2007): 19-38.

Tebaldi, C., & Knutti, R. (2007). The use of the multi-model ensemble in probabilistic climate projections. *Philosophical transactions of the royal society A: mathematical, physical and engineering sciences*, 365(1857), 2053-2075.

Tindula, G. N., Orang, M. N., & Snyder, R. L. (2013). Survey of irrigation methods in California in 2010. *Journal of irrigation and drainage engineering*, 139(3), 233-238.

Tindula, G.N., M.N. Orang, and R.L. Snyder. 2013. "Survey of irrigation methods in California in 2010." *Journal of Irrigation and Drainage Engineering*, 139:233-238.

Treasury, H. M. S. (2009). Accounting for the effects of climate change: Supplementary green book guidance. *UK Government, HM Treasury, London*.

Wang, W., & McCARL, B. A. (2013). Temporal investment in climate change adaptation and mitigation. *Climate Change Economics*, *4*(02), 1350009.

APPENDICES

Appendix A: Certification in lemons market

Appendix A1: Theory Derivation

Proof of proposition 3:

Given the cost function, under the seller certification, the seller solves the profit-maximizing problem

(1) $\max_{e\in[0,1]} e(v_h - P_c) + (1-e)E(v) - e^{\lambda}$,

and choses effort based on

(2) $FOC: \Delta v - P_c = \lambda e^{\lambda-1}$,

Which yields

(3) $e_s(P_c, \Delta v) = (\frac{\Delta v - P_c}{\lambda})^{\frac{1}{\lambda-1}}$,

where Δq stands for the quality gap between high- and low-quality item. Using backward induction, the certifier solves the profit-maximization problem

(4) $\max_{P_c} e_s(P_c, \Delta q)P_c$.

Solving the system, we find the optimal price for certification equals to

(5) $P_c = \frac{(\lambda-1)\Delta v}{\lambda}$,

which means the marginal benefit of additional effort in the seller certification equals to

(6) $MB_b = \tilde{P} - v_l$.

For the buyer certification (weakly), the certifier's problem is

(7) $\max_{P_c}[e_b(P_c) + (1-e_b(P_c))\sigma_l^*(\tilde{P}(P_c))]\sigma^*(s_h(P_c))P_c$.

By substituting in the expression for the two probabilities, we reframe the problem as

(8) $\max_{P_c} P_c \frac{e_b(P_c)(\tilde{P}(P_c) - v_l)}{\tilde{P}(P_c)\tilde{\mu}(P_c)}$.

Define $v_l = n\Delta v$ and solve the system, we find the optimal price for certification as:

(9) $$P_c^* = -\frac{\Delta q(-2n\lambda-\lambda+n+1)\sqrt{(n+1)(\frac{1}{4}+n(\lambda-\frac{1}{2})^2)}+(n+1)(\lambda^2-(4n+1)\lambda+n+1)}{2\lambda^2}.$$

Substituting P_c^* back to the FOC condition for seller's profit maximizing problem and compare with the case of seller-certification, we have

(10) $$\frac{\Delta v}{2}+\frac{\sqrt{\Delta v(\Delta v-4P_c^*)}}{2}-\frac{\Delta v}{\lambda}.$$

Setting equation (10) equal to zero we find that when $\lambda_1 = 1, \lambda_2 = \frac{2}{\frac{2n+1+\sqrt{5n^2+6n+1}}{n+2}+1}$ and $\lambda_3 = \frac{2}{\frac{-2n-1+\sqrt{5n^2+6n+1}}{n+2}+1}$, the two certification mechanism cases will provide the same level of incentive. Clearly, since n is a non-negative number, λ_2 is less than 1, which is outside of the range of λ, while λ_3 is larger than 2. Since the above equation is continuous when $n > 0\ and\ \lambda > 1$, thus for any $\lambda \in [1, \lambda_3]$ conveys the same sign. When $\lambda = 2$, the equation above becomes

(11) $$\frac{\sqrt{1+(-3n-1)\sqrt{(n+1)(\frac{1}{4}+\frac{9}{4}n)}+2(n+1)(\frac{9}{4}n-\frac{1}{4})}}{2},$$

which is larger than

(12) $$\frac{\sqrt{1+(-3n-1)\sqrt{(n+1)(\frac{1}{4}+\frac{1}{4}n)}+2(n+1)(\frac{9}{4}n-\frac{1}{4})}}{2} = \frac{\sqrt{3n^2+2n}}{2} \geq 0.$$

Thus, buyer-certification will provide more incentive to the seller for any $\lambda \in (1, +\infty]$ for any value of the low-quality item. Also, we observe that the difference will become larger as value of low-quality item increases.

Proof for Extension 2

Under seller certification, the seller's profit maximization problem is

(13) $$\max_{e\in[0,1]} \tau e(P_h - P_c) + e(1-\tau)(P_L - P_c) + (1-e)P_L - c(e),$$

(14) $FOC\colon\ \tau(P_h - v_L) - P_c = c'(e)$,

and optimal $P_c = \frac{(\lambda-1)(P_h-v_L)}{\lambda}$. At the same time, the problem needs to satisfy one participation constrain that makes sure the seller has no incentive to cheat:

(15) $\pi^* \geq (1-\tau)P_h + \tau v_l - P_c$.

If the constrain is satisfied, $P_h = \tau\Delta v + v_l$ and $P_L = e(1-\tau)\Delta v + v_l$

On the other hand, for the buyer-certification, we keep all our assumptions in case 3 except changing the accuracy of certification. Thus, the seller still needs to send out a price signal to make buyer indifferent among s_b, s_h and s_n. Receiving the price signal, the buyer will be indifferent among the three actions only if the following three equations hold:

(16) $P_c = \tau\mu(P)(v_h - P) + (1-\tau)(1-\mu)(v_l - P)$,

(17) $P_c = \tau(1-\mu)(P - v_l) + (1-\tau)\mu(P - v_h)$,

(18) $P = \mu(P)v_h + \big(1-\mu(P)\big)v_l$.

Where τ means the probability that the certification indicates high-quality item while $1-\tau$ represents the chance of mistake. Solving the equations, we have

(19) $\tilde{P}' \equiv \frac{\left(v_h+v_l+\sqrt{\Delta v\left(\Delta v-\frac{4P_c}{2\tau-1}\right)}\right)}{2}$ & $\tilde{\mu}' \equiv \frac{1+\sqrt{1-\frac{4P_c}{\Delta v(2\tau-1)}}}{2}$,

which indicates a lower equilibrium price for the high-quality commodity and a smaller price range for certification. Following the same logic in case 3, the buyer chose the chance of using certification service $\sigma'^*(s_h)$,to make the low type seller being indifferent between ask for $\tilde{P}'$ or q_l, and

(20) $\sigma'^*(s_h) = \frac{\tilde{P}'-v_l}{\tilde{P}'\tau}$.

Also, based on the consistency of belief, we would have

(21) $\sigma_l'^{*}(\tilde{P}) = \frac{\tilde{e}(1-\tilde{\mu}')}{\tilde{\mu}'(1-\tilde{e})}$.

Substituting all the terms back to seller's problem, we have

(22) $\max_{e\in[0,1]} \tau e\tilde{P} + (1-e)v_l - c(e)$,

(23) $FOC\colon\ \tau(\tilde{P} - v_l) - (1-\ \tau)v_l = c'(e)$.

From the profit maximizing condition, we notice that when the accuracy rate decrease together value of low-quality item increase, the effort level in the buyer certification case will decrease even further compared with the seller certification case. This is because under the buyer certification, high type seller will be punished with an additional opportunity cost, which is not able to sell the item at any price, if the item is certified as low quality. From the comparisons, the seller certification could dominate buyer certification as the additional punishment increase. Based on our simulation result, when $v_h = 2v_l$ and $\tau = 0.7$, the seller certification will lead to a higher effort input.

Appendix A2: Experiment Instructions and Screenshots

Instructions for seller imperfect certification treatment

Thank you for participating in today's study. Please follow the instructions carefully. At any time, please feel free to raise your hand if you have a question.

You have been randomly assigned an ID number for this session. You will make decisions using a computer. You will never be asked to reveal your identity to anyone. Your name will never be associated with any of your decisions. In order to keep your decisions private, please do not reveal your choices or otherwise communicate with any other participant. Importantly, please refrain from verbally reacting to events that occur.

Today's session has three parts: Experiment 1, Experiment 2, and a short questionnaire. You will have the opportunity to earn money in both experiments based on your decisions. You will be paid your earnings privately, and in cash, at the end of the experiment session. **We will proceed through the written materials together. Please do not enter any decisions on the computer until instructed to do so.**

Instructions for Experiment 1

Please refer to your computer screen while we read the instructions.

We would like you to make a decision for each of 10 scenarios. Each scenario involves a choice between playing a lottery that pays either $4 or $0 according to specified chances (Option A) or receiving $2 for sure (Option B).

You will notice that the only differences across scenarios are the chances of receiving the high or low prize for the lottery. At the end of the today's session, ONE of the 10 scenarios will be selected at random and you will be paid according to your decision for this selected scenario ONLY. Each scenario has an equal chance of being selected.

Please consider your choice for each scenario carefully. Since you do not know which scenario will be played out, it is in your best interest to treat each scenario as if it will be the one used to determine your earnings.

Before making decisions, are there any questions?

Please proceed to entering decisions on your computer. Once you are ready to submit your decisions, please click the "Submit" button.

Instructions for Experiment 2

In this experiment, you will be randomly placed into a two-person group to form a trading market. One member of your group will be a seller and the other will be a buyer.

There will be many decision rounds in the experiment. You will not know the number of rounds until the experiment has been completed. Each decision round is separate from the other rounds, in the sense that the decisions you make in one round will not affect the outcome or earnings of any other round.

In this experiment, all money amounts are denominated in lab dollars, and will be exchanged at a rate of 200 lab dollars to 1 US dollar at the end of the experiment. There are five parts to each decision period:

- The seller decides how much effort to put into "product" quality. Increasing effort improves the chance that the product is of "high" rather than "low" quality.
- Product quality (high or low) is revealed to the seller. The seller then has the option to purchase a certification service. If a product is certified to be of high quality, this information is made known to the buyer. Otherwise, the buyer does not know product quality when making a purchase decision.
- The seller makes a take-it-or-leave-it price offer to the buyer.
- The buyer chooses whether to accept the offer.
- The computer calculates earnings.

We will now go through the details of the seller and buyer tasks in each decision round.

Information for Sellers:

The seller produces and offers to sell one unit of a "product". The product may be of "low" or "high" quality. Buyers are willing to pay more for a high quality product, but producing a high quality product can be more costly.

The seller first makes an effort choice. Production cost increases with effort, but higher effort increases the chance that the product is high quality. The relationship between effort, product quality, and production cost is presented in the following table.

Effort choice	Chance of high product quality	Production cost
0	0%	0
1	10%	1
2	20%	4
3	30%	9
4	40%	16
5	50%	25
6	60%	36
7	70%	49
8	80%	64
9	90%	81
10	100%	100

For example, if the seller selects an effort of 0, production cost will be 0 lab dollars and there will be a 0% chance that the product is high quality. As another example, if the seller selects an effort of 10, production cost will be 100 lab dollars and there will be a 100% chance the product is high quality.

After the effort choice is made, product quality will be revealed to the seller, but it will not be known to the buyer. The seller can purchase a certification service at a known cost. However, the certification is error prone. If the service is purchased, and the actual quality is High, the service has a 70% chance of revealing to the buyer that the product quality is High; If the actual quality is Low, the service has a 30% chance of revealing to the buyer that the product is High. If the service is purchased, and the service determines the product is low quality, this information will not be revealed to the buyer.

The seller next makes a take-it-or-leave-it price offer to the buyer. This is the only opportunity to sell the product.

In each round, the seller will receive 100 lab dollars in fixed income. This amount does not depend on any decisions made. If the buyer accepts the offer, seller earnings are equal to the fixed income plus price received, minus any production cost. If the buyer rejects the offer, seller earnings are equal to fixed income minus any production cost. To summarize,

If the buyer accepts the offer:

Seller earnings = Fixed income + Price received – Production cost – Certification Cost (if any)

If the buyer rejects the offer:

Seller earnings = Fixed income – Production cost – Certification Cost (if any)

Information for Buyers:

The value of the product to the buyer depends on whether it is high or low quality. If the buyer purchases a high-quality product, they will receive 500 lab dollars. If the buyer purchases a low-quality product, the buyer will instead receive 250 lab dollars. If the seller purchases the certification service and the product is certified to be high quality, this information will be revealed to the buyer. Otherwise, the quality of product is not known to the buyer when making a purchase decision.

The buyer can accept any price offer they choose. There is no budget constraining this purchase.

In each round, the buyer will receive 100 lab dollars in fixed income. This amount does not depend on any decisions made. If the buyer accepts the seller's offer, the buyer's earnings are equal to the fixed income plus the value of the product (which depends on actual quality), minus the price paid. If the buyer rejects the offer, buyer earnings are equal to the fixed income. To summarize,

If the buyer accepts the offer:

Buyer earnings = Fixed Income + Value of product – Price Paid

If the buyer rejects the offer:

Buyer earnings = Fixed Income

Proceeding through the experiment:

Prior to each new decision round, you will be randomly matched with a different person in the room. The computer is programmed such that, when you are a seller, you will not be matched with the buyer from the previous round. Similarly, when you are a buyer, you will not be matched with the seller from the previous round. The decisions you made in prior rounds will not be known to your trading partner.

Your role in the experiment (buyer or seller) and the cost of the certification service may change from one round to the next. Please look carefully at the information on your computer screen before making any decisions.

We will begin with two training rounds to help you understand the procedures. In one training round you will play the role of a buyer and in the other you will be a seller.

Aside from decisions in the training rounds, you will be paid based on the outcome of each decision round. This means that it is very important to consider each decision prior to making it. Before we continue, do you have any questions?



Figure 1.1. Subject screen for risk elicitation task

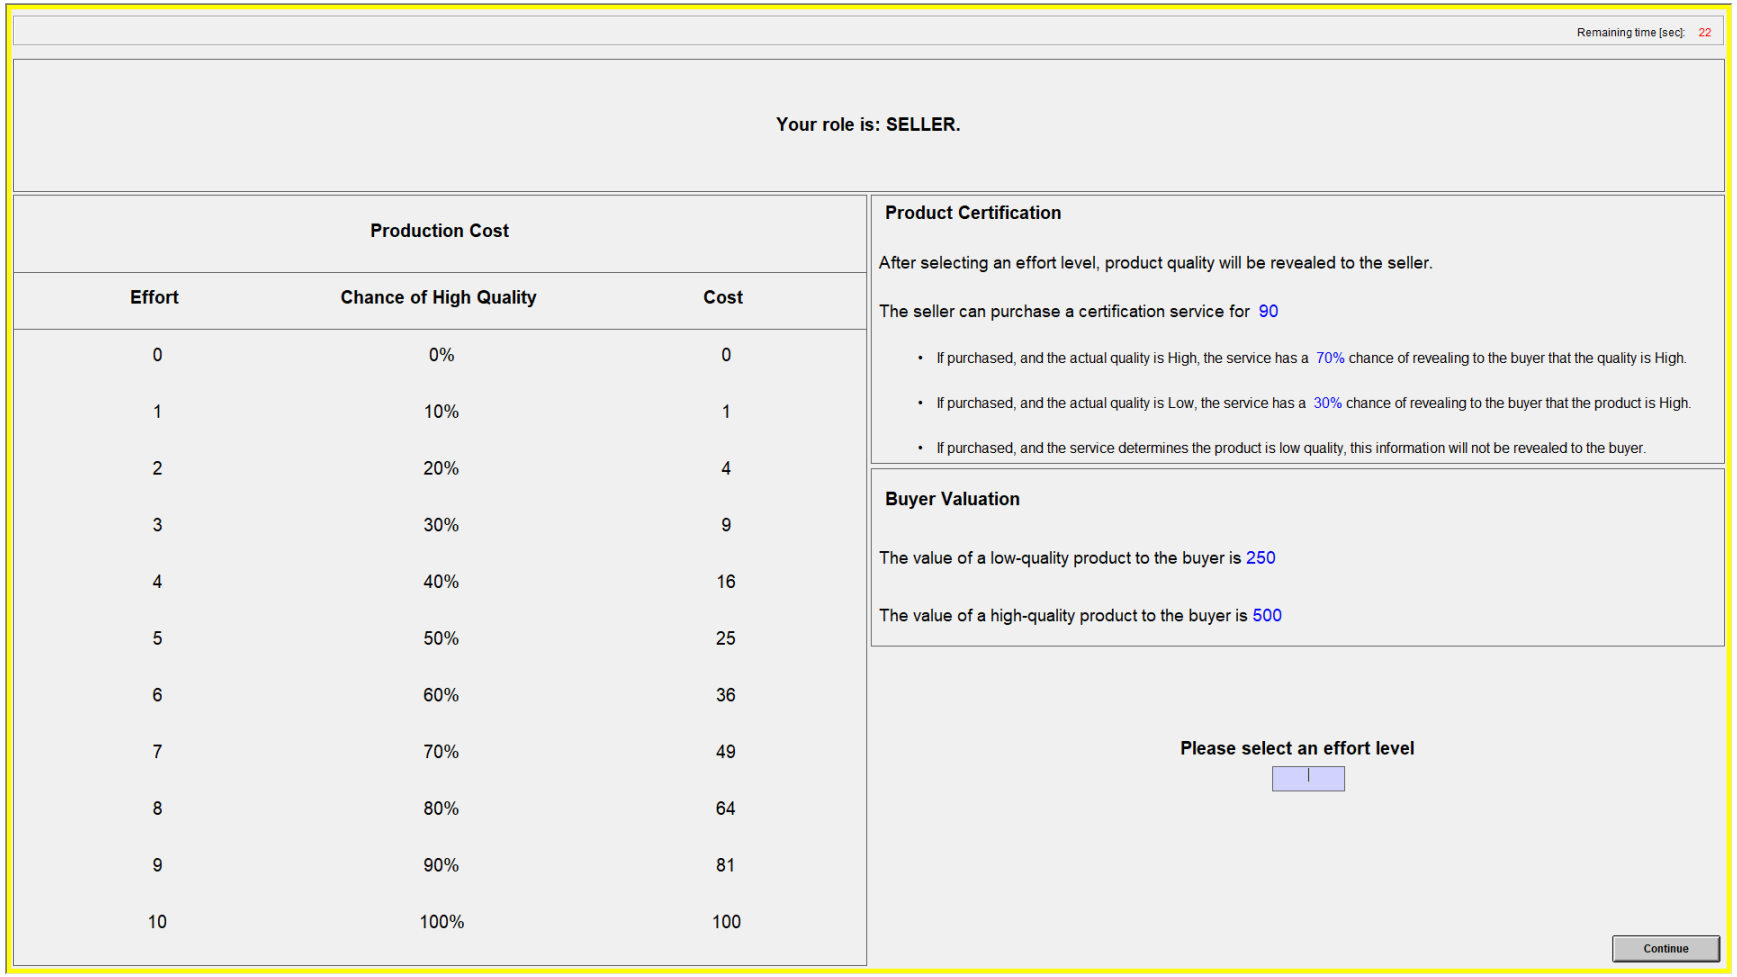

Figure 1.2. Subject screen for seller's effort choice

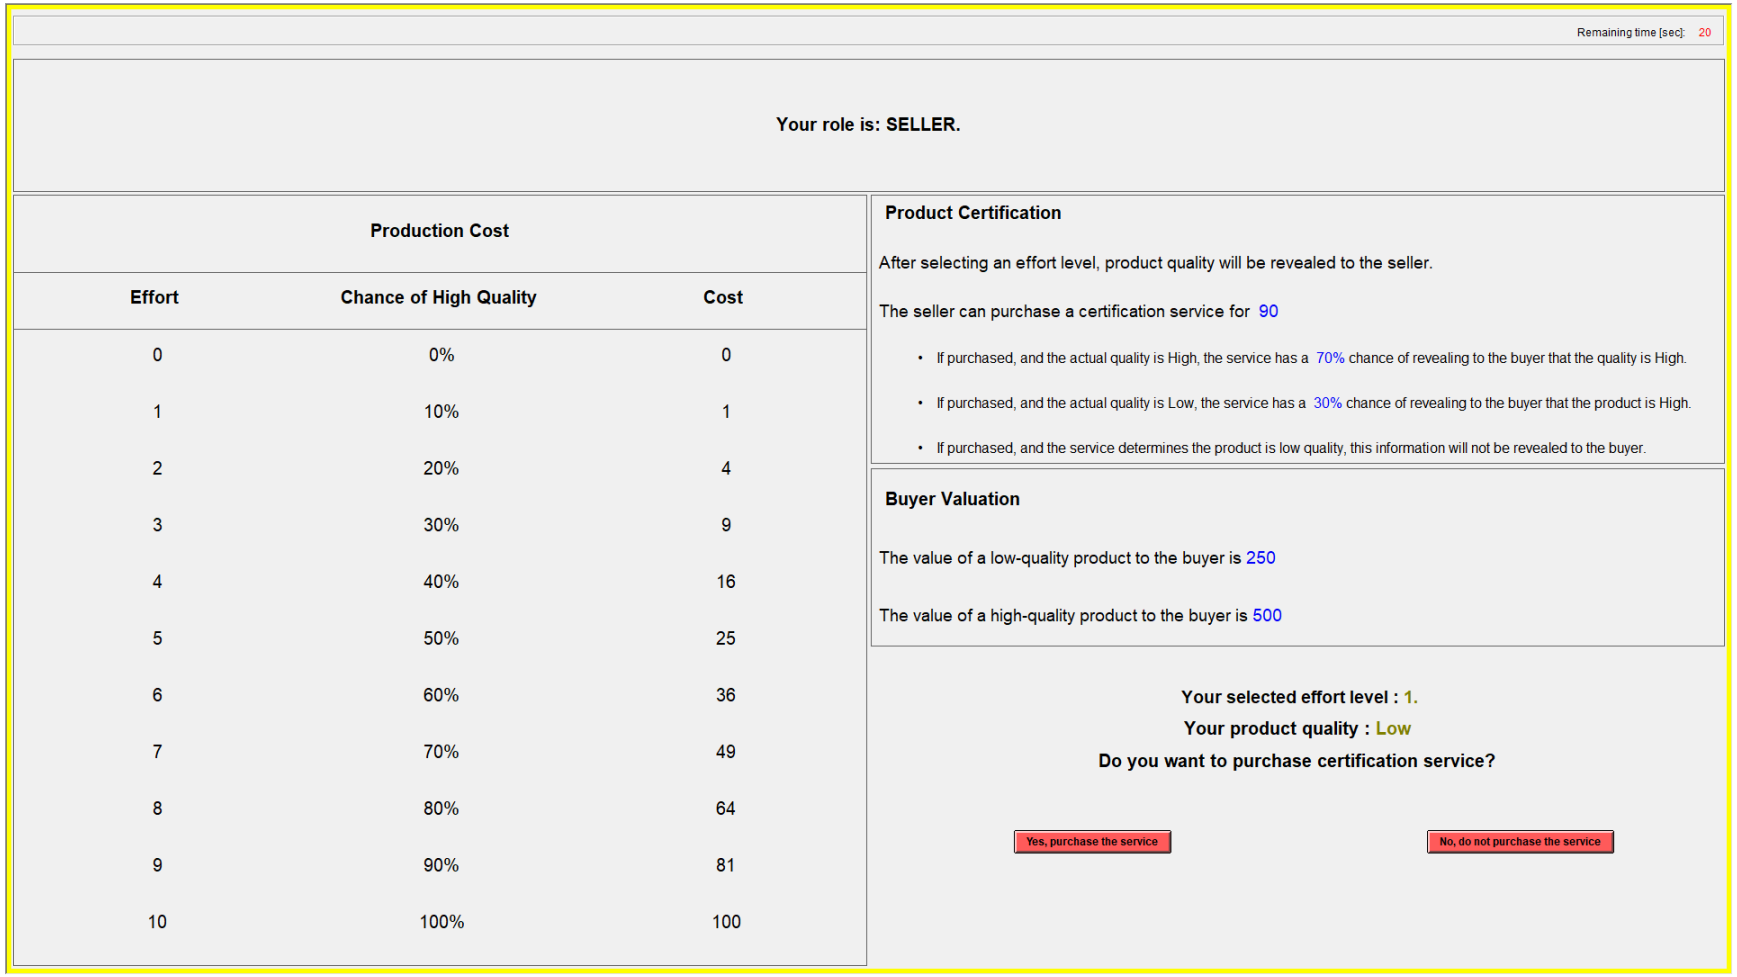

Figure 1.3. Subject screen for seller's certification decision

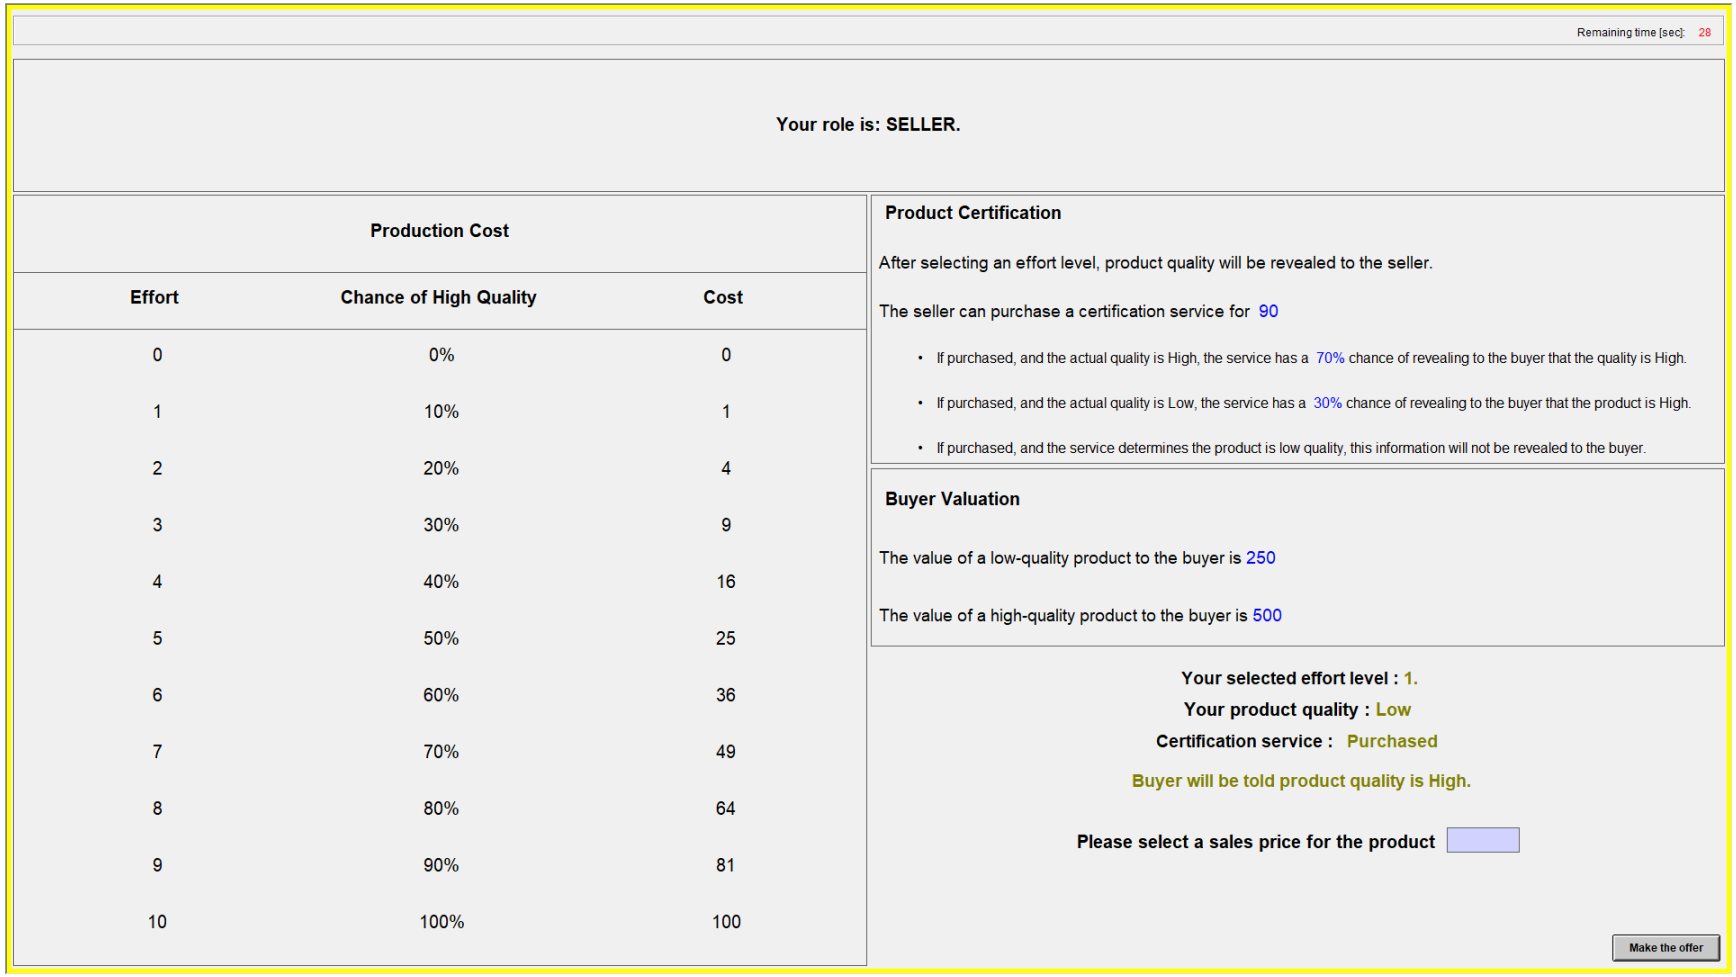

Figure 1.4. Subject screen for seller's price offer

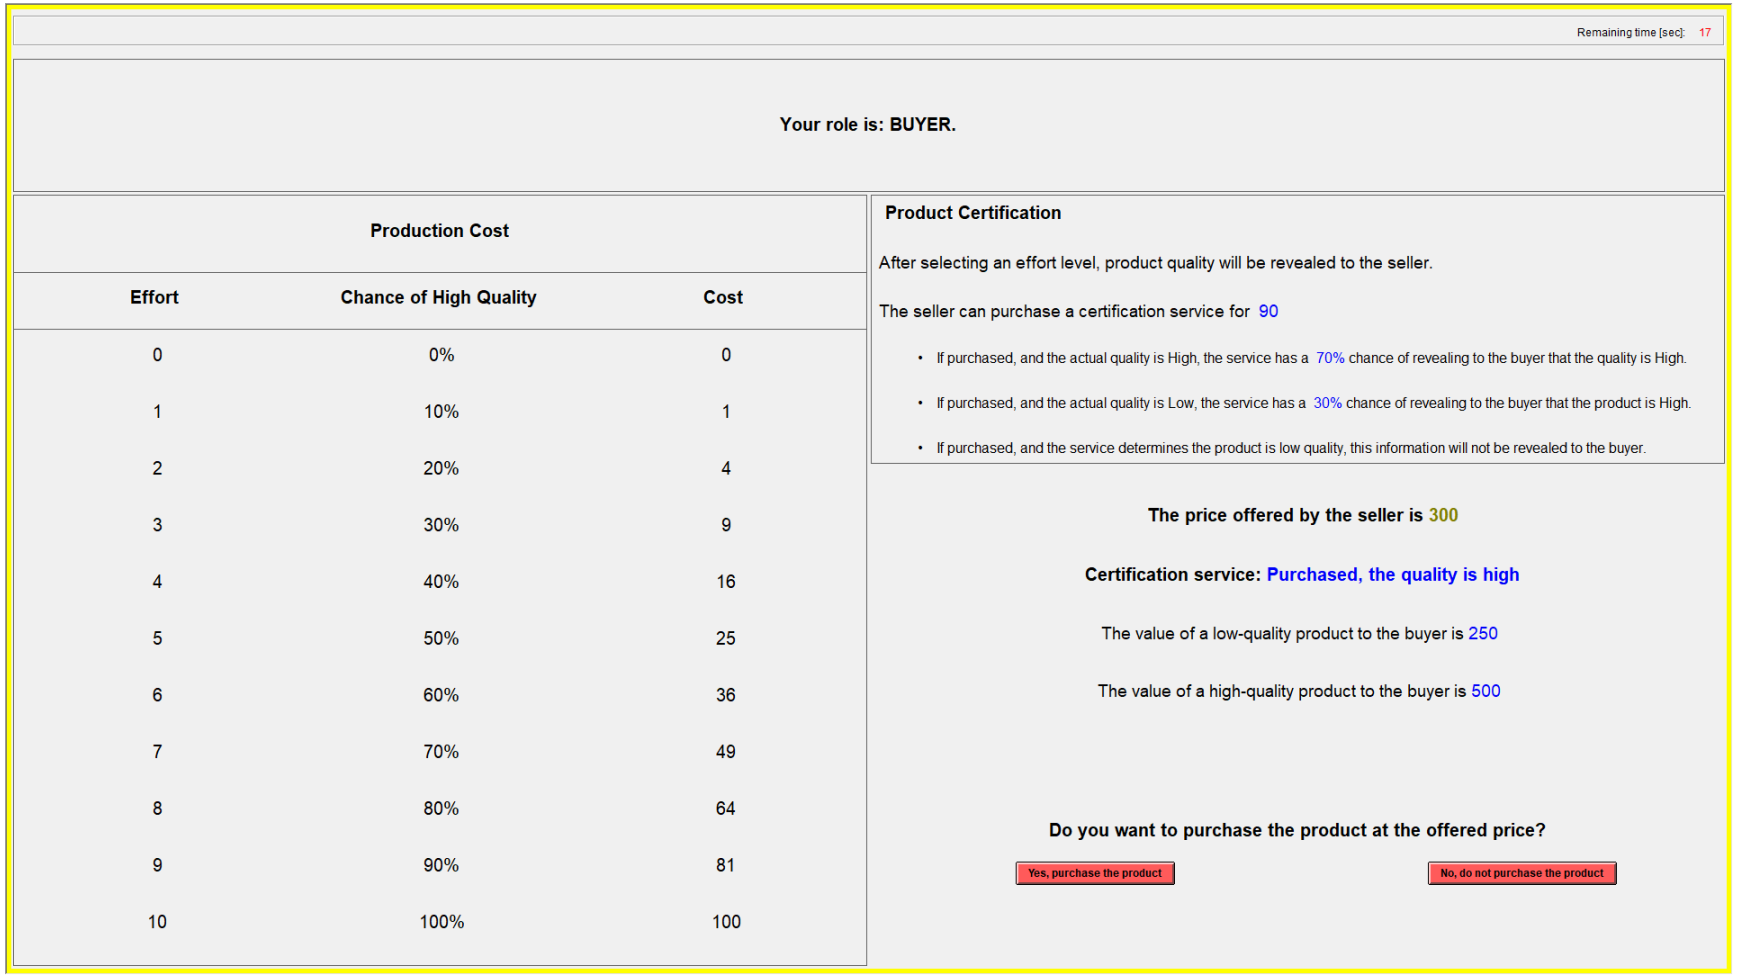

Figure 1.5. Subject screen for buyer's purchase decision

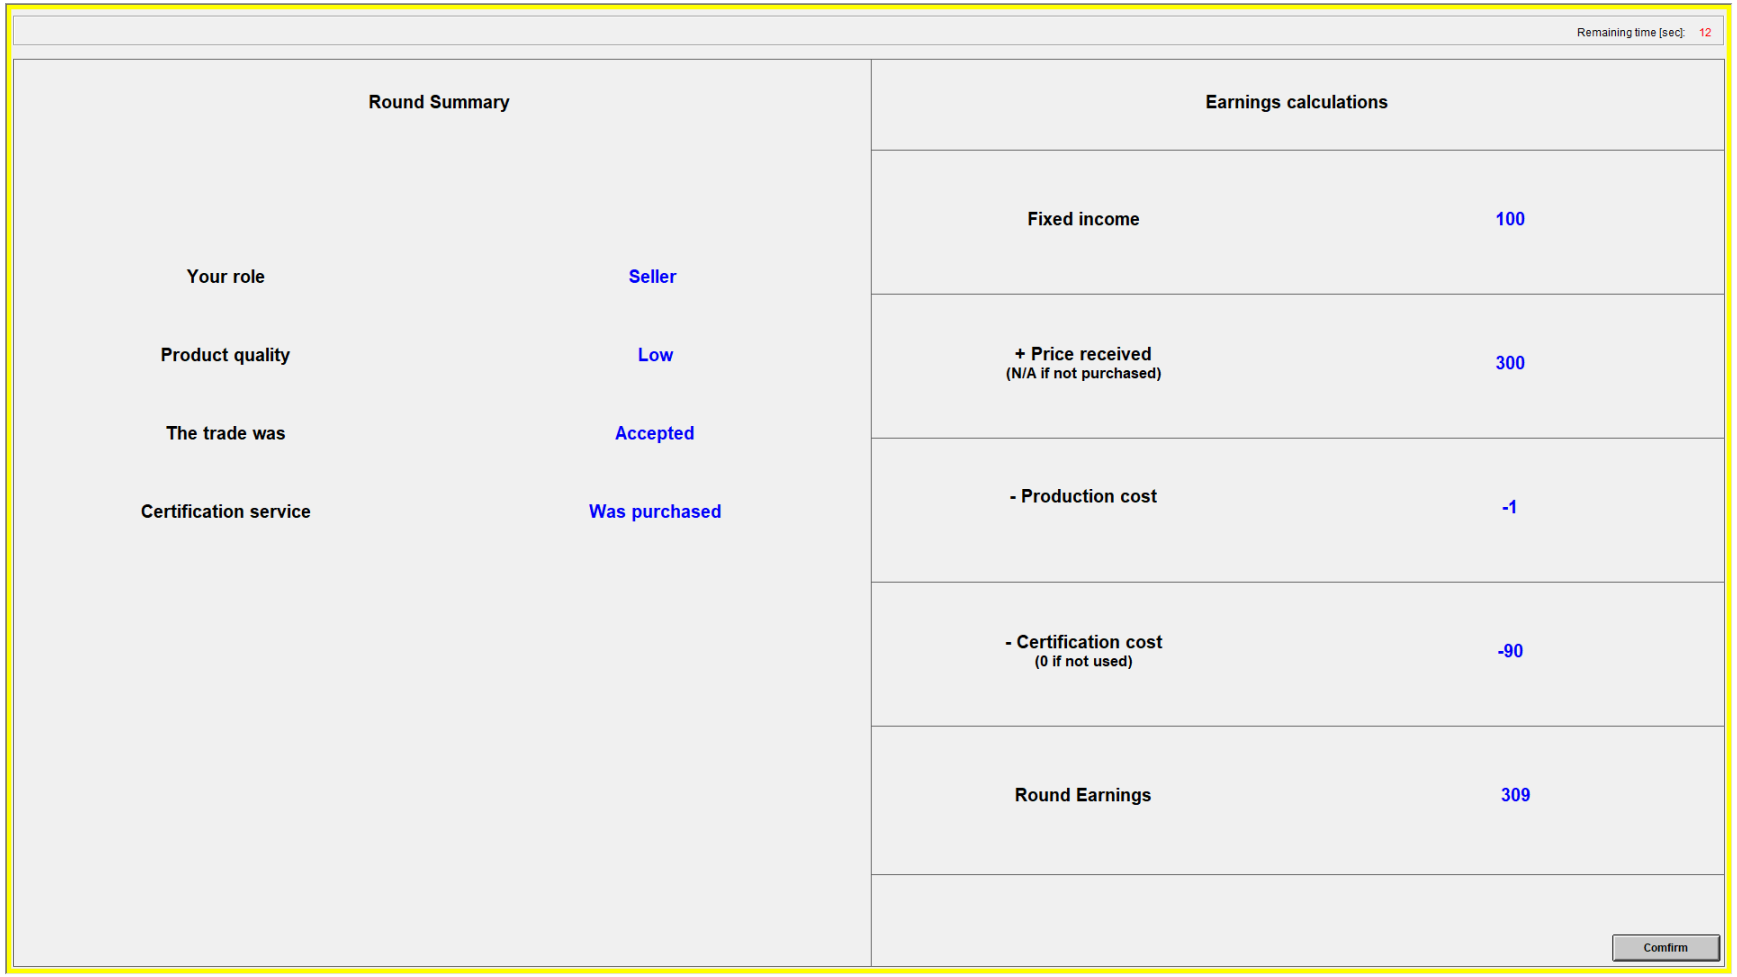

Figure 1.6. Subject screen for seller's round summary

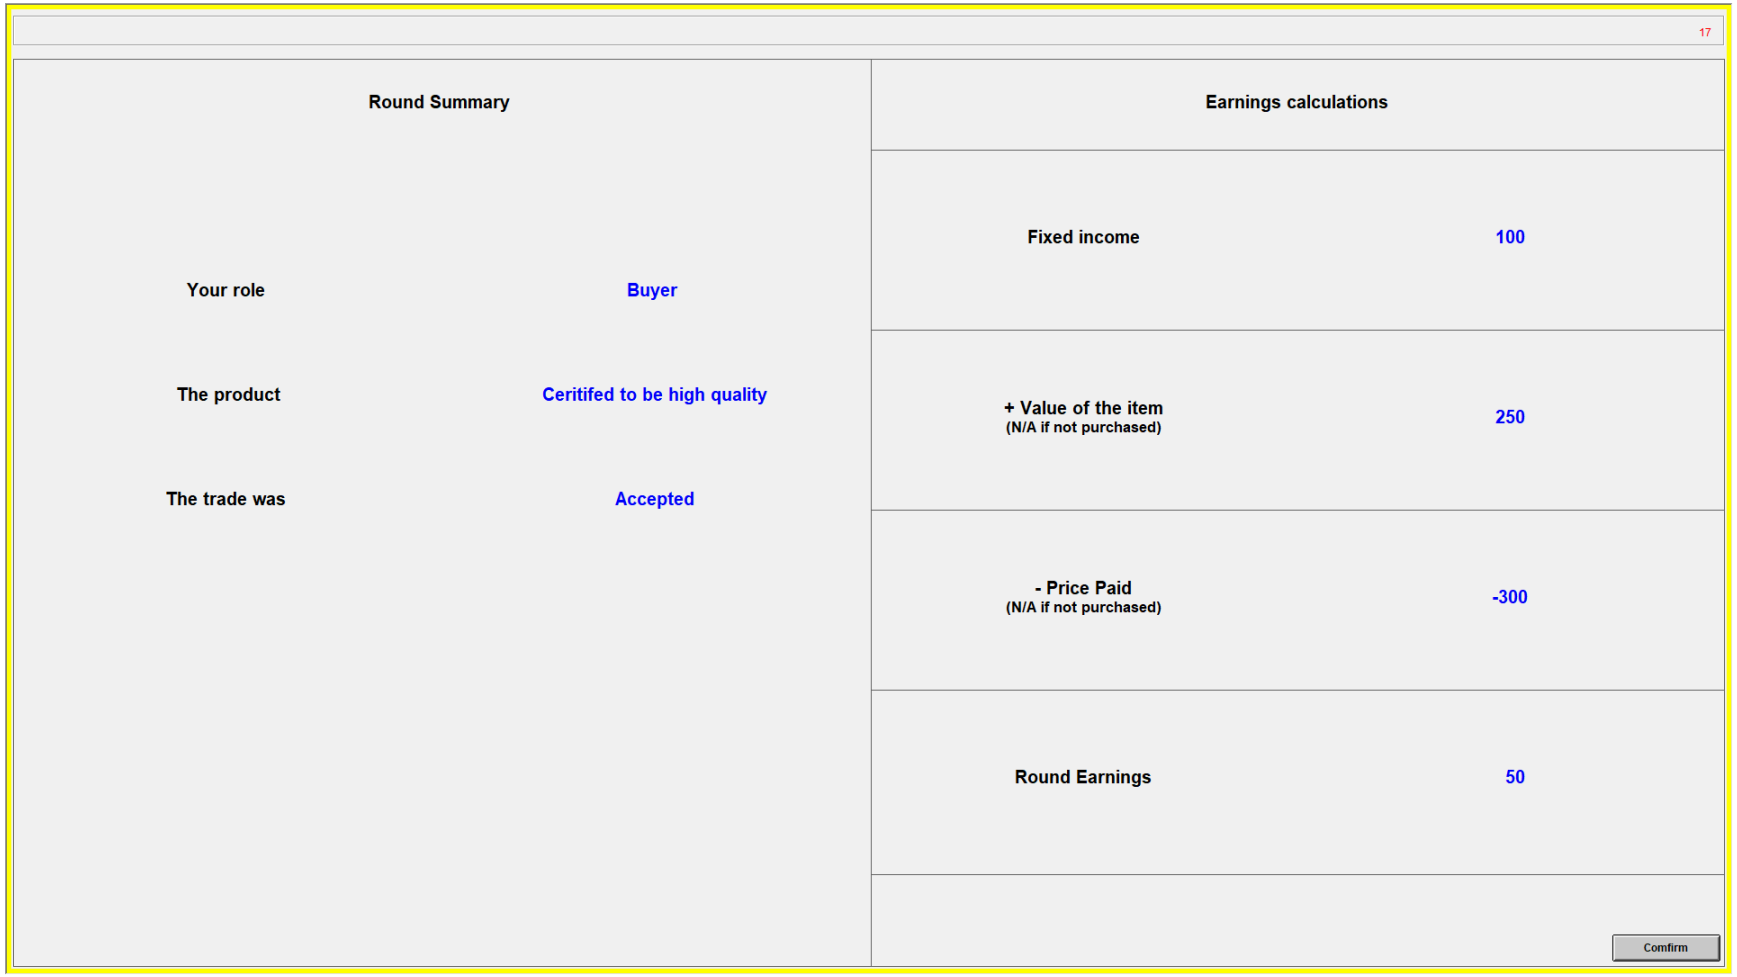

Figure 1.7. Subject screen for buyer's round summary

Appendix A3: Tables

Table 1.1. Experiment treatments

Treatment	Available Certification	Certification process	Certification price
Control	N/A	N/A	N/A
SPA(60)	Seller	Perfect	60 (alternative)
SPM(125)	Seller	Perfect	125 (monopolistic)
BPA(125)	Buyer	Perfect	125 (alternative)
BPM(60)	Buyer	Perfect	60 (monopolistic)
SIA(20)	Seller	Imperfect	20 (alternative)
SIM(90)	Seller	Imperfect	90 (monopolistic)
BIA(90)	Buyer	Imperfect	90 (alternative)
BIM(20)	Buyer	Imperfect	20 (monopolistic)

Note: Under an imperfect certification process, there is a 70% chance that a high-quality item will be certified as high quality and a 30% chance that a low-quality item will be certified as high quality. Monopolistic prices are those derived directly from the theory. The alternative price for seller (buyer) certification is the monopolistic price from the comparable buyer (seller) certification treatment.

Table 1.2. Theoretical predictions

Treatment	Seller effort (probability of high-quality product)	Probability certification purchased	Social welfare (efficiency)
Control	0 (0%)	N/A	62.5%
SPA(60)	10 (100%)	100%	100%
SPM(125)	6 (60%)	60%	91%
BPA(125)	0 (0%)	0%	62.5%
BPM(60)	8 (80%)	33.86%	89.73%
SIA(20)	6 (40%)	60%	91%
SIM(90)	4 (40%)	40%	83.5%
BIA(90)	0 (0%)	0%	62.5%
BIM(20)	2 (20%)	18.70%	70.94%

Table 1.3. Data Description

Variable Name	Description	Mean	S.D.
Seller Effort	Seller's effort choice, 0 to 10	3.70	2.90
Product Sold	=1 if buyer purchased product	0.663	0.472
Certification Purchased	=1 if certification service purchased	0.393	0.488
Certification Price	price of certification, in lab dollars	70.62	34.45
Social Welfare	actual social welfare divided by maximum possible social welfare, multiplied by 100%	50.6	46.62
Control	=1 if control condition	0.102	0.303
SPA(60)	=1 if seller perfect certification with alternative price treatment	0.096	0.295
SPM(125)	=1 if seller perfect certification with monopolistic price treatment	0.097	0.296
BPA(125)	=1 if buyer perfect certification with alternative price treatment	0.097	0.296
BPM(60)	=1 if buyer perfect certification with monopolistic price treatment	0.096	0.295
SIA(20)	=1 if seller imperfect certification with alternative price treatment	0.097	0.296
SIM(90)	=1 if seller imperfect certification with monopolistic price treatment	0.096	0.295
BIA(90)	=1 if buyer imperfect certification with alternative price treatment	0.097	0.296
BIM(20)	=1 if buyer imperfect certification with monopolistic price treatment	0.096	0.295
Risk Averse	=1 if participant selected the safe option at least six times in the risk elicitation task	0.675	0.469
Employed	=1 if participant employed full or part-time	0.602	0.489
Female	=1 if participant is female	0.419	0.494
Preexp	=1 if participated other experiment before	0.623	0.486
Saleexp	=1 if participant has sales experience before	0.390	0.489
Extraversion	measure of personality trait "extraversion", 1 to 7	4.03	1.44
Agreeableness	measure of personality trait "agreeableness", 1 to 7	4.13	1.24
Conscientiousness	measure of personality trait "extraversion", 1 to 7	5.11	1.23
Emotional Stability	measure of personality trait "emotional stability", 1 to 7	4.17	1.31
Openness	measure of personality trait "openness to experience", 1 to 7	4.78	1.08

Table 1.4. Treatment effects regressions

	(1) Seller Effort	(2) Certification Purchased	(3) Product Sold	(4) Social Welfare	(5) Product market efficiency
SPA(60)	2.969***	0.582***	0.196***	27.60***	18.87***
	(0.505)	(0.0538)	(0.0529)	(4.341)	(3.985)
SPM(125)	2.027***	0.476***	0.237***	29.56***	14.67***
	(0.515)	(0.0573)	(0.0437)	(3.890)	(3.419)
BPA(125)	0.563	0.147***	-0.0077	2.062	-2.534
	(0.489)	(0.0293)	(0.0574)	(4.199)	(4.079)
BPM(60)	2.180***	0.447***	0.0821	15.55***	8.648*
	(0.567)	(0.0365)	(0.0509)	(4.972)	(5.070)
SIA(20)	1.016**	0.688***	0.0840*	8.025**	4.584
	(0.472)	(0.0562)	(0.0496)	(3.808)	(3.796)
SIM(90)	0.569	0.371***	0.0604	7.354	-0.984
	(0.534)	(0.0549)	(0.046)	(4.834)	(4.712)
BIA(90)	0.539	0.153***	0.031	3.723	0.282
	(0.540)	(0.0290)	(0.0563)	(4.751)	(4.725)
BIM(20)	0.798	0.282***	0.0546	7.035*	5.623
	(0.561)	(0.0346)	(0.0543)	(3.884)	(3.854)
Constant	2.673***		0.592***	41.11***	41.11***
	(0.340)		(0.0336)	(2.172)	(2.172)
Observations	1,760	1,360	1,760	1,760	1,760
R^2	0.107	0.477	0.027	0.049	0.023

Notes: *,**, and *** denote estimates that are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively. Standard errors are clustered at the seller-level.

Table 1.5. Selected hypothesis tests

Hypothesis	Treatment 1	Treatment 2	Difference of means
	Seller versus buyer certification		
H2	Buyer Perfect Monopolistic	Seller Perfect Monopolistic	0.152
H3	Seller Imperfect Monopolistic	Buyer Imperfect Monopolistic	-0.229
H4	Seller Perfect Monopolistic	Buyer Perfect Alternative	1.465***
H4	Seller Perfect Alternative	Buyer Perfect Monopolistic	0.758
	Perfect versus imperfect certification		
H5	Seller Perfect Monopolistic	Seller Imperfect Monopolistic	1.459***
H5	Buyer Perfect Monopolistic	Buyer Imperfect Monopolistic	1.382**
	Changes in certification prices		
	Seller Perfect Alternative	Seller Perfect Monopolistic	0.941**
	Buyer Perfect Monopolistic	Buyer Perfect Alternative	1.617***
	Seller Imperfect Alternative	Seller Imperfect Monopolistic	0.447
	Buyer Imperfect Monopolistic	Buyer Imperfect Alternative	0.261

Notes: *,**, and *** denote estimates that are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively. Test results based on regressions presented in Table 4.

Table 1.6. Conditional certification purchases

Seller Certification	Item Quality	Certification Purchased (%)
Perfect Monopolistic	High	91.9
Perfect Monopolistic	Low	1.2
Perfect Alternative	High	97.9
Perfect Alternative	Low	6.7
Imperfect Monopolistic	High	74.1
Imperfect Monopolistic	Low	17.8
Imperfect Alternative	High	89.5
Imperfect Alternative	Low	58.4
Buyer Certification	**Price**	**Certification Purchased (%)**
Perfect Monopolistic	>250	62.6
Perfect Monopolistic	≤250	9.6
Perfect Alternative	>250	22.1
Perfect Alternative	≤250	6.0
Imperfect Monopolistic	>250	36.7
Imperfect Monopolistic	≤250	20.5
Imperfect Alternative	>250	18.5
Imperfect Alternative	≤250	11.9

Table 1.7. Dynamic specification

	(1) Seller Effort	(2) Certification Purchased	(3) Product Sold
SPA(60)	0.900	0.495***	-0.0114
	(1.398)	(0.152)	(0.117)
SPM(125)	0.855	0.359**	0.117
	(1.413)	(0.167)	(0.105)
BPA(125)	1.127	0.296**	-0.256**
	(1.388)	(0.116)	(0.128)
BPM(60)	0.586	0.645***	-0.143
	(1.390)	(0.105)	(0.124)
SIA(20)	0.975	0.873***	-0.232**
	(1.264)	(0.147)	(0.0998)
SIM(90)	-0.864	0.194	-0.220*
	(1.361)	(0.120)	(0.113)
BIA(90)	0.902	0.162*	-0.245*
	(1.361)	(0.0971)	(0.137)
BIM(20)	1.303	0.617***	-0.154
	(1.498)	(0.116)	(0.120)
Lasthalf×Control	-0.151		0.175*
	(1.150)		(0.0928)
Lasthalf×SPA	-0.622	-0.0328	0.0412
	(1.459)	(0.187)	(0.121)
Lasthalf×SPM	-0.289	-0.0361	-0.124
	(1.469)	(0.211)	(0.123)
Lasthalf×BPA	2.581**	0.0815	-0.159
	(1.017)	(0.0935)	(0.171)
Lasthalf×BPM	-2.721*	0.0737	-0.327*
	(1.481)	(0.127)	(0.177)
Lasthalf×SPA	1.051	0.200	-0.00746
	(1.240)	(0.200)	(0.150)
Lasthalf×SPM	-0.237	0.146	0.0651
	(1.443)	(0.204)	(0.191)
Lasthalf×BPA	2.118	0.0960	0.00511
	(1.456)	(0.104)	(0.188)
Lasthalf×BPM	-0.508	0.198*	-0.290*
	(1.210)	(0.106)	(0.162)
Period×Control	-0.133		-0.0159**
	(0.108)		(0.00782)
Period×SPA	0.0757	0.00848	0.00515
	(0.124)	(0.0167)	(0.0116)
Period×SPM	-0.00988	0.0113	0.00480
	(0.120)	(0.0174)	(0.0111)
Period×BPA	-0.291***	-0.0144	0.0189

Table 1.7. Continued

	(0.101)	(0.0109)	(0.0146)
Period×BPM	0.0984	-0.0176	0.0228
	(0.134)	(0.0131)	(0.0140)
Period×SIA	-0.177*	-0.0237	0.0158
	(0.101)	(0.0174)	(0.0113)
Period×SIM	-0.0120	0.00697	0.00939
	(0.117)	(0.0164)	(0.0137)
Period×BIA	-0.230*	-0.00405	0.0121
	(0.121)	(0.0106)	(0.0149)
Period×BIM	-0.132	-0.0340***	0.0185
	(0.119)	(0.0112)	(0.0128)
Risk Averse	-0.115	0.00196	0.0127
	(0.0906)	(0.00817)	(0.00823)
Female	0.136	0.0547*	0.0263
	(0.320)	(0.0293)	(0.0307)
Preexp	-0.860***	-0.0336	0.0194
	(0.312)	(0.0274)	(0.0296)
Saleexp	-0.605*	0.0160	0.00149
	(0.326)	(0.0298)	(0.0280)
Extraversion	-0.0612	0.00199	-0.00385
	(0.108)	(0.0103)	(0.00970)
Agreeableness	-0.132	0.0127	-0.00601
	(0.131)	(0.0131)	(0.0111)
Conscientiousness	-0.111	-0.00844	0.0130
	(0.137)	(0.0109)	(0.0129)
Emotional stability	-0.0718	-0.0153	-0.0140
	(0.123)	(0.0135)	(0.0110)
Openness	0.254	0.0239	-0.00392
	(0.162)	(0.0175)	(0.0126)
Constant	6.148***	-0.0977	0.650***
	(1.672)	(0.0880)	(0.116)
Observations	1,740	1,740	1,740
R-squared	0.206	0.274	0.053

Notes: *,**, and *** denote estimates that are statistically different from zero at the 10%, 5%, and 1% significance levels, respectively. Standard errors are clustered at the seller-level.

Table 1.8. Quality misrepresentation

	Quality misrepresentation (%)	P_H	P_L
Control	43.0	320.20	283.95
Seller Perfect Monopolistic	3.5	420.05	211.83
Seller Perfect Alternative	5.3	418.45	224.08
Buyer Perfect Monopolistic	22.3	382.35	276.16
Buyer Perfect Alternative	24.7	389.09	268.63
Seller Imperfect Monopolistic	22.3	365.29	251.07
Seller Imperfect Alternative	31.1	353.56	271.70
Buyer Imperfect Monopolistic	15.0	369.04	246.78
Buyer Imperfect Alternative	27.0	345.67	278.03

Notes: product quality misrepresentation is defined as a case where the actual product quality is low, and the seller offers a price that is higher than the buyer's reservation price for a low-quality item. The last two columns provide the average prices offered for all high and low-quality products, respectively.

Appendix B: Updating under ambiguity

Appendix B1: Theory Derivations

Support for Hypothesis 1:

We first start with the α-MEU model. For task 1A, a player's utility can be described as

$$V_{1A} = \alpha u(5) + (1-\alpha)u(15),$$

as the best (worst) prior is drawing a white (black) ball for sure. After trimming out the black ball, utility is:

$$V_{1A}{}' = \alpha u(10) + (1-\alpha)u(15),$$

as the worst prior becomes getting a red ball for sure. The value of self-protection in this case is:

$$V_{1A}^{P} = V_{1A'} - V_{1A} = \alpha[u(10) - u(5)],$$

which is independent of the utility of the best prior. The value of self-protection based on tasks 1B and 1B′ is:

$$V_{1B}^{P} = V_{1B'} - V_{1B} = V_{1A}^{P} = \alpha[u(10) - u(5)].$$

Thus, changing the payoff from the white ball has no effect on the value of self-protection.

On the other hand, using the re-weighting Bayesian rule, the value for lottery 1A and 1A′ are, respectively,

$$V_{1A,REU} = w_1\varphi[u(5)] + w_2\varphi[u(10)] + w_3\varphi[u(15)] \text{ and}$$

$$V_{1A',REU} = \frac{w_2}{1-w_1}\varphi[u(10)] + \frac{w_3}{1-w_1}\varphi[u(15)].$$

The value of self-protection is then

$$V_{1A,REU}^{P} = V_{1A,REU} - V_{1A',REU} = \frac{w_2 w_1}{1-w_1}\varphi[u(10)] + \frac{w_3 w_1}{1-w_1}\varphi[u(15)] - w_1\varphi[u(5)].$$

The value of self-protection based on comparing 1B and 1B′ is

$$V_{1B,REU}^{P} = \frac{w_2 w_1}{1-w_1}\varphi[u(10)] + \frac{w_3 w_1}{1-w_1}\varphi[u(30)] - w_1\varphi[u(5)],$$

and the difference in the values of self-protection is positive, i.e.,

$$V^{P}_{1B,REU} - V^{P}_{1A,REU} = \frac{w_3 w_1}{1-w_1}\{\varphi[u(30)] - \varphi[u(15)]\} > 0,$$

given both φ and u are monotonically increasing. This result implies that regardless the ambiguity averse level, for a subject with REU, the value self-protection is higher when the value of highest prize increases.

Support for Hypothesis 2:

Since the best and worst priors have not changed, the α-MEU model predicts the same valuations for both lotteries, assuming α is held fixed. For REU, an agent's utility for lottery 2A is

$$V_{2A} = w_1\varphi[u(5)] + w_3\varphi[u(15)],$$

where w_1 is the weight on the prior of all black balls and w_3 is the weight on the prior of all white balls. With Bayesian updating on the weights, when adding a third prior, w_2 (all red balls), we have

$$V_{2B} = w_1(1-w_2)\varphi[u(5)] + w_2\varphi[u(10)] + w_3(1-w_2)\varphi[u(15)].$$

Given the concavity of $\varphi(.)$, it follows that

$$2\varphi[u(10)] > \varphi[u(5)] + \varphi[u(15)].$$

If $w_1 = w_3, w_1 + w_3 = 1$ then we have

$$V_{2A} - V_{2B} = w_2\{(w_1\varphi[u(5)] + w_3\varphi[u(15)]) - \varphi[u(10)]\} < 0.$$

If $w_1 > w_3$, it follows that

$$\begin{aligned} V_{2A} - V_{2B} &= w_2\{(w_1\varphi[u(5)] + w_3\varphi[u(15)]) - \varphi[u(10)]\} \\ &< w_2\{(0.5\varphi[u(5)] + 0.5\varphi[u(15)]) - \varphi[u(10)]\} < 0 \end{aligned}$$

We do not consider the case where $w_1 < w_3$, as putting a higher weight on the best prior potentially contradicts the assumption that agents are ambiguity averse.

Support for Hypothesis 3:

For α-MEU with a compound lottery as a prior, we have

$$V_{3A}^{r} = \alpha u(10) + (1-\alpha)\left[0.5\left(\tfrac{2}{3}u(10) + \tfrac{1}{3}u(15)\right) + 0.5\left(\tfrac{1}{3}u(10) + \tfrac{2}{3}u(15)\right)\right], \text{ and}$$

$$V_{3B}^{r} = \alpha u(10) + (1-\alpha)\left[\tfrac{1}{2}u(10) + \tfrac{1}{2}u(15)\right].$$

It follows that $V_{3A}^{r} = V_{3B}^{r}$. For α-MEU with a single lottery as a prior, we have

$$V_{3A}^{u} = \alpha u(10) + (1-\alpha)\left[\tfrac{1}{3}u(10) + \tfrac{2}{3}u(15)\right],$$

$$V_{3B}^{u} = \alpha u(10) + (1-\alpha)\left[\tfrac{1}{2}u(10) + \tfrac{1}{2}u(15)\right], \text{ and}$$

$$V_{3A}^{r} > V_{3B}^{r}.$$

Under REU we have:

$$V_{3A} = w_1\varphi[u(10)] + w_2\varphi[\tfrac{2}{3}u(10) + \tfrac{1}{3}u(15)] + w_3\varphi[\tfrac{1}{3}u(10) + \tfrac{2}{3}u(15)], \text{ and}$$

$$V_{3B} = w_1\varphi[u(10)] + 2w_3\varphi\left[\tfrac{1}{2}u(10) + \tfrac{1}{2}u(15)\right].$$

It follows that $V_{3A} < V_{3B}$ as $w_2 = w_3$ and $\varphi(.)$ is a concave function.

Support for Hypothesis 4:

For sophisticated α-MEU it follows that

$$V_{5\mathrm{A}}^{s} = \alpha' u(10) + (1-\alpha')u(15), \text{ and}$$

$$V_{5\mathrm{B}}^{s} = \alpha u(10) + (1-\alpha)u(15).$$

As a "sophisticated" bidder will have $\alpha' < \alpha$, since the prior with only white balls is more likely to happen for task 5A, it follows that $V_{5\mathrm{A}}^{s} > V_{5\mathrm{B}}^{s}$.

For naïve α-MEU preferences we have

$$V_{5\mathrm{A}}^{n} = V_{5\mathrm{B}}^{n} = \alpha u(10) + (1-\alpha)u(15).$$

For the REU model,

$$V_{5\mathrm{A}} = w_1\varphi[u(10)] + w_2\varphi[\tfrac{2}{3}u(10) + \tfrac{1}{3}u(15)] + (w_3 + w_4)\varphi[u(15)], \text{ and}$$

$$V_{5\mathrm{B}} = w_1\varphi[u(10)] + 2w_2\varphi\left[\tfrac{1}{3}u(10) + \tfrac{2}{3}u(15)\right] + w_4\varphi[u(15)].$$

It follows that $V_{5\mathrm{A}} < V_{5\mathrm{B}}$ since $w_2 = w_3$ and $\varphi(.)$ is a concave function.

Appendix B2: Experiment Instructions

(Note: Instructions are unaltered, with the exception of changing the task labels to reflect those used in the manuscript)

Thank you for participating in today's study. If you have a question at any time, please raise your hand. We ask that you do not communicate with other study participants, unless instructed to do so. Your decisions will not be associated with your name or other identifying information. Your name will not be linked in any way to the results of the study.

The session is divided into three experiments. You will have the opportunity to earn money in each experiment based on your decisions. You will also receive $10 for completing all three experiments, along with a post-experiment questionnaire. You will be paid your earnings privately, and in cash, at the end of the experiment session. **We will proceed through the written materials together. Please do not make any decisions until instructed.**

Instructions for Experiment 1

Please refer to the Decision Sheet for Experiment 1 as we read the instructions.

We would like you to make a decision for each of 10 scenarios. Each scenario involves a choice between receiving $2 for sure (Option A) or playing a lottery that pays $3.85 or $0.10 with the stated chances (Option B).

You will notice that the only differences across scenarios are the chances of receiving the high or low prize for the lottery.

At the end of the session, ONE of the 10 scenarios will be selected at random and you will be paid according to your decision for this selected scenario ONLY. Each scenario has an equal chance of being selected.

Please consider your choice for each scenario carefully. Since you do not know which scenario will be played out, it is in your best interest to treat each scenario as if it will be the one used to determine your earnings.

Before making decisions, are there any questions?

Decision Sheet for Experiment 1

This is your Decision Sheet. Please indicate which option you prefer by circling the letter in the "choice" column.

No.	Option A	Option B	Choice
1	receive $2 for sure	10% chance of $3.85 90% chance of $0.10	A B
2	receive $2 for sure	20% chance of $3.85 80% chance of $0.10	A B
3	receive $2 for sure	30% chance of $3.85 70% chance of $0.10	A B
4	receive $2 for sure	40% chance of $3.85 60% chance of $0.10	A B
5	receive $2 for sure	50% chance of $3.85 50% chance of $0.10	A B
6	receive $2 for sure	60% chance of $3.85 40% chance of $0.10	A B
7	receive $2 for sure	70% chance of $3.85 30% chance of $0.10	A B
8	receive $2 for sure	80% chance of $3.85 20% chance of $0.10	A B
9	receive $2 for sure	90% chance of $3.85 10% chance of $0.10	A B
10	receive $2 for sure	100% chance of $3.85 0% chance of $0.10	A B

At the end of the session, after a scenario is played out, please fill-in the information below:

The selected scenario: _____.

The option you selected for this scenario: _____.

Your earnings from the selected scenario: _____.

Please also record your earnings on the record sheet

Instructions for Experiment 2

Please refer to the Decision Sheet for Experiment 2 as we read the instructions.

The Decision Sheet contains 20 separate Decisions numbering 1 through 20. Each of these Decisions is a choice between drawing a ball from "Urn A" or "Urn B". You will select a color, White or Black, and this will be your **Success Color**. Your earnings will be determined by whether the ball drawn from the Urn matches your Success Color.

In each of the 20 decisions, Urn A has 50 White balls and 50 Black balls, and pays 4 dollars if the ball drawn from Urn A matches your Success Color, and 0 if it does not match. Since each color has a 50% chance of being drawn, this means that drawing from Urn A pays 2 dollars with a chance of 50%, and pays 0 with a chance of 50%.

Urn B, on the other hand, has an unknown number of white and Black balls (with a total of 100 balls). It pays a positive amount if the ball drawn from Urn B matches your Success Color, and 0 if it does not match. Since the chance of each color being drawn is unknown, the chance of Urn B paying a positive amount is unknown as well.

You will notice that the only differences across scenarios is the amount paid when a ball matching your Success Color is drawn from Urn B.

At the end of the session, ONE of the 20 scenarios will be selected at random and you will be paid according to your decision for this selected scenario ONLY. Each scenario has an equal chance of being selected.

Please consider your choice for each scenario carefully. Since you do not know which scenario will be played out, it is in your best interest to treat each scenario as if it will be the one used to determine your earnings.

Before making decisions, are there any questions?

Decision Sheet for Experiment 2

My Success Color is (please circle one): **White** **Black**

.

	Urn A	Urn B	
No.	50 White balls, 50 Black balls	? White balls, ? Black balls	Choice
1	$4 if Success Color, 0 if not	$3.28 if Success Color,0 if not	A B
2	$4 if Success Color, 0 if not	$3.44 if Success Color,0 if not	A B
3	$4 if Success Color, 0 if not	$3.60 if Success Color,0 if not	A B
4	$4 if Success Color, 0 if not	$3.76 if Success Color,0 if not	A B
5	$4 if Success Color, 0 if not	$3.92 if Success Color,0 if not	A B
6	$4 if Success Color, 0 if not	$4.08 if Success Color,0 if not	A B
7	$4 if Success Color, 0 if not	$4.24 if Success Color,0 if not	A B
8	$4 if Success Color, 0 if not	$4.40 if Success Color,0 if not	A B
9	$4 if Success Color, 0 if not	$4.56 if Success Color,0 if not	A B
10	$4 if Success Color, 0 if not	$4.72 if Success Color,0 if not	A B
11	$4 if Success Color, 0 if not	$4.88 if Success Color,0 if not	A B
12	$4 if Success Color, 0 if not	$5.04 if Success Color,0 if not	A B
13	$4 if Success Color, 0 if not	$5.20 if Success Color,0 if not	A B
14	$4 if Success Color, 0 if not	$5.36 if Success Color,0 if not	A B
15	$4 if Success Color, 0 if not	$5.52 if Success Color,0 if not	A B
16	$4 if Success Color, 0 if not	$5.68 if Success Color,0 if not	A B
17	$4 if Success Color, 0 if not	$5.84 if Success Color,0 if not	A B
18	$4 if Success Color, 0 if not	$6.00 if Success Color,0 if not	A B
19	$4 if Success Color, 0 if not	$6.16 if Success Color,0 if not	A B
20	$4 if Success Color, 0 if not	$6.32 if Success Color,0 if not	A B

At the end of the session, after a scenario is played out, please fill-in the information below:

The selected scenario: _____.

The urn you selected for this scenario: _____.

Your earnings from the selected scenario: _____.

Please also record your earnings on the record sheet

Instructions for Experiment 3

This experiment involves many decision "rounds". In each, the moderator will offer an item for sale. Your task will be to place a bid to buy the item for sale. As the purchase procedure will be new to you, we will first go through training rounds.

We will use the following purchase procedure in all rounds:

1. You will place a bid on the item. You will not know the price prior to bidding.
2. The price of the item will be randomly drawn. A volunteer will be asked to roll dice to determine this price. The random price will be the same for all participants.
3. **If your bid is equal to or higher than the random price, you buy the item and pay the random price (<u>not</u> your bid!). If your bid is lower than the random price, you do not buy the item.**

Here are some possible scenarios based on the purchase procedure:

— You bid $2. The random price is drawn to be $1.50. Since your bid is equal to or higher than the random price, you buy the item at a price of $1.50.

— You bid $5. The random price is drawn to be $5. Since your bid is equal to or higher than the random price, you buy the item at a price of $5.

— You bid $3. The random price is drawn to be $3.50. Since your bid is lower than the random price, you do not buy the item.

It is important to point out some aspects of the procedure. First, different from auctions, you are <u>not</u> bidding against other players. The bids of other players do <u>not</u> impact whether you buy an item. If, for example, everyone bids an amount higher than the random price, each person will pay the random price and each person will receive the item. Second, different from some auctions, if you buy something, the price is <u>not</u> equal to your bid. Instead, you pay the randomly selected price.

Third, your bid sets the highest price for which you agree to buy the good. For example, if you bid $6.25, this means that you agree to buy the item as long as the price is something less than or equal to $6.25. Your bid of $6.25 guarantees that you do <u>not</u> buy the item at prices above $6.25.

Before bidding you should ask yourself "what is the highest price I am willing to pay for the item?" It is in your best interest to place a bid equal to this highest price.

"What If" Scenarios

To help you understand the procedures, we ask that you consider a number of "what if" scenarios. Here is the good news: you will be paid 25 cents for each scenario you answer correctly. There is a bonus question, and you will be paid 50 cents for a correct answer to this.

The item for sale in these scenarios is a $5 bill. Remember: If your bid is equal to or higher than the random price, you buy the item and pay the random price (not your bid!). If your bid is lower than the random price, you do not buy.

1. Suppose you **bid $2.50**. Then, a volunteer draws a **random price** of **$4.00**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

2. Suppose you **bid $3.12**. Then, a volunteer draws a **random price** of **$6.37**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

3. Suppose you **bid $5.00**. Then, a volunteer draws a **random price** of **$4.25**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

4. Suppose you **bid $5.00**. Then, a volunteer draws a **random price** of **$6.56**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

5. Suppose you **bid $7.16**. Then, a volunteer draws a **random price** of **$4.12**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

6. Suppose you **bid $8.00**. Then, a volunteer draws a **random price** of **$6.50**.
   Based on the procedure we described would you purchase the $5 bill? Yes No
   If you answered "Yes", what price would you pay? $ _____

Bonus question. Given the purchase procedure, how much should you bid for the $5 bill? Keep in mind that it is in your best interest to place a bid equal to the highest price you are willing to pay.

You should bid: $ _____ . _____

Please raise your hand when you are ready to have your calculations checked.

Training Round 1

In this training round you will have the opportunity to earn money.

Your task in this round is to place a bid to buy a ticket to draw ONE ball from an urn containing three balls. If you successfully buy a ticket, you will receive an amount of money based on the color of the ONE ball drawn. All three balls are black. Drawing a black ball pays $2.

After everyone has placed a bid, a volunteer will roll dice to determine the random price. Although you will not know the price range before you bid, know that three dice will be rolled. The first will determine the dollars and the other two will determine the cents.

If your bid is equal to or higher than the random price, you will receive the amount of money for the ball drawn and pay the random price. If you make a purchase at a price that is higher than the value of the ball drawn you will in fact have negative earnings (lose money).

If your bid is less than the random price, you will not buy a ticket. You will not pay the random price. You will earn $0.

For training purposes, we will play out the procedures several times. However, you will only bid once. **You will not be able to change your bid after the random prices are determined**.

Your bid (in dollars and cents): $ _____ . _____

Trial 1

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Training Round 1—Continued

Trial 2

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Trial 3

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Trial 4

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Training Round 2

In this training round you will have the opportunity to earn money.

Your task in this round is to place a bid to buy a ticket to draw ONE ball from an urn containing three balls: one black ball, one red ball and one white ball. If you successfully buy a ticket, you will receive an amount of money based on the color of the ONE ball drawn. The black ball pays 2 dollars, the red ball pays 3 dollars and the white ball pays 4 dollars if drawn.

After everyone has indicated their bid, a volunteer will roll dice to determine the random price. Although you will not know the price range before you bid, know that three dice will be rolled. The first will determine the dollars and the other two will determine the cents.

If your bid is equal to or higher than the random price, you will receive the amount of money for the ball drawn and pay the random price. If you make a purchase at a price that is higher than the value of the ball drawn you will in fact have negative earnings (lose money).

If your bid is less than the random price, you will not buy a ticket. You will not pay the random price. You will earn $0.

We will play out the procedures several times. However, you will only bid once. **You will not be able to change your bid after the random prices are determined**.

Your bid (in dollars and cents): $ _____ . _____

Trial 1

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Training Round 2—Continued

Trial 2

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Trial 3

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Trial 4

Random price: $_____ . _____

Is your bid equal to or higher than the random price? (check the box below)

☐ **Yes**. You bought a ticket.

Your earnings are equal to: $_____ – _____ = $ _____ . _____

(value of ball) (random price)

☐ **No**. You did not buy a ticket. Your earnings for this trial are $0.

Record your earnings on your Record Sheet.

Experiment 3 Decision Rounds – Part A

The decision rounds in Experiment 3 are divided into two parts – Part A and Part B. There are six decision rounds in this part.

Your task in each round is to place a bid to buy a ticket to draw ONE ball from an urn. You will be told the possible colors of the balls in the urn, and the payoffs for each color.

In this part, may not know exactly how many balls of a particular color there are. As an example, you might know there are three balls in the urn, and that the possible ball colors are black, red and white. This could mean that all three balls are the same color, or that there are two of one color and one of a second color, or one of each color. Any combination is possible.

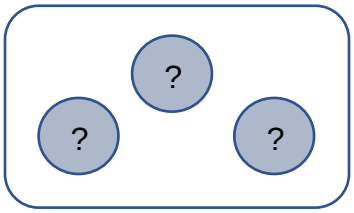

The purchase procedure is the same as before. Your bid will be compared to a random price. You will purchase the ticket only if your bid is equal to or higher than the random price. The random price will be a randomly drawn number between $0.00 and a maximum price. The maximum price will an amount equal to or lower than the payoff associated with the highest-valued ball in the urn. For example, if the highest-valued ball is $15, the maximum possible price may be as high as $15. Any price within this range is equally likely to be chosen.

As before, it is in your best interest to place a bid equal to the highest price you are willing to pay for the ticket. By doing so, you will only purchase the ticket at prices you are willing to pay. You will not purchase the ticket at prices you are not willing to pay.

If you instead bid lower than the highest price you are willing to pay, you risk not purchasing the ticket at prices favorable to you. If you instead bid more than the highest price you are willing to pay, you risk purchasing the ticket at prices that are not favorable to you.

Only one decision round, which may be from Part A or Part B, will be implemented for real. After all decision rounds are completed, we will have a volunteer roll dice to determine which round this is. Since you will not know which round will be selected prior to making any decisions, it is in your best interest to take each decision seriously as if it will determine an actual purchase.

After this paid round is selected, a volunteer will roll dice to determine the random price. Another volunteer will draw a ball. Before we proceed to the decision rounds, are there any questions?

In the next FOUR rounds, the same urn will be used. The urn contains three balls of unknown colors.

Please, proceed to the next decision round.

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. The possible ball colors are black, red and white.

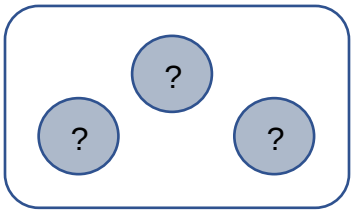

Payoff if drawn	Black - $5	Red - $10	White - $15
Total number of balls in an urn		3	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. The possible ball colors are black, red and white.

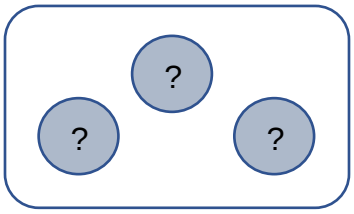

Payoff if drawn	Black - $5	Red - $10	White - $30
Total number of balls in an urn		3	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. The possible ball colors are black, red and white. IF this urn contains one or more black balls, these balls will be replaced with red or white balls before a ball is drawn.

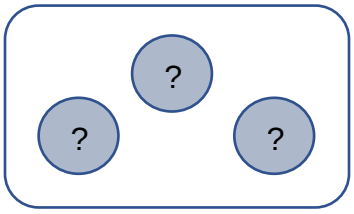

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	3	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. The possible ball colors are black, red and white. IF this urn contains one or more black balls, these balls will be replaced with red or white balls before a ball is drawn.

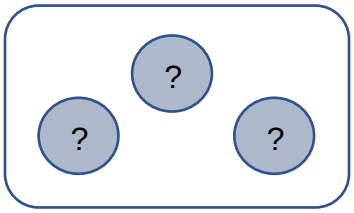

Payoff if drawn	Red - $10	White - $30
Total number of balls in an urn	3	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. Either all balls in the urn are black or all balls in the urn are white. There are no other possibilities.

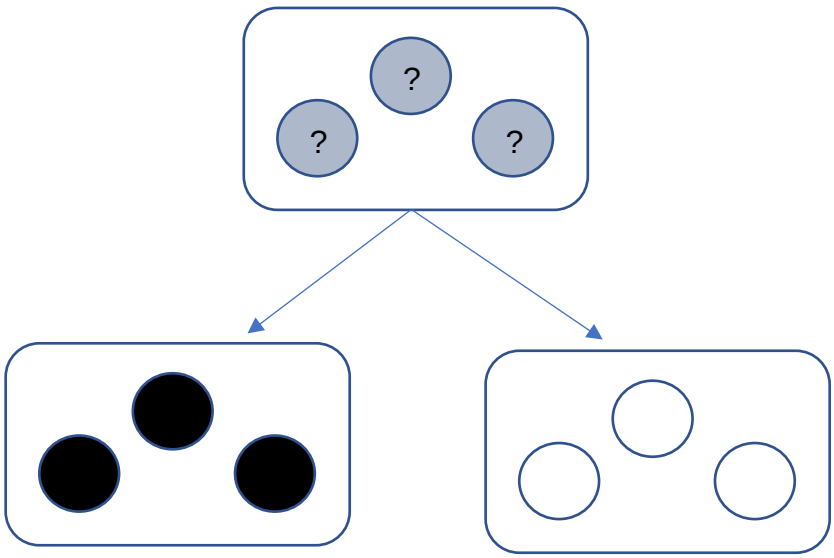

Payoff if drawn	Black - $5	White - $15
Total number of balls in an urn	3	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the urn, and pay the random price.

Description:

There is one urn containing three balls. Either all balls in the urn are black, all balls in the urn are red or all balls in the urn are white. There are no other possibilities.

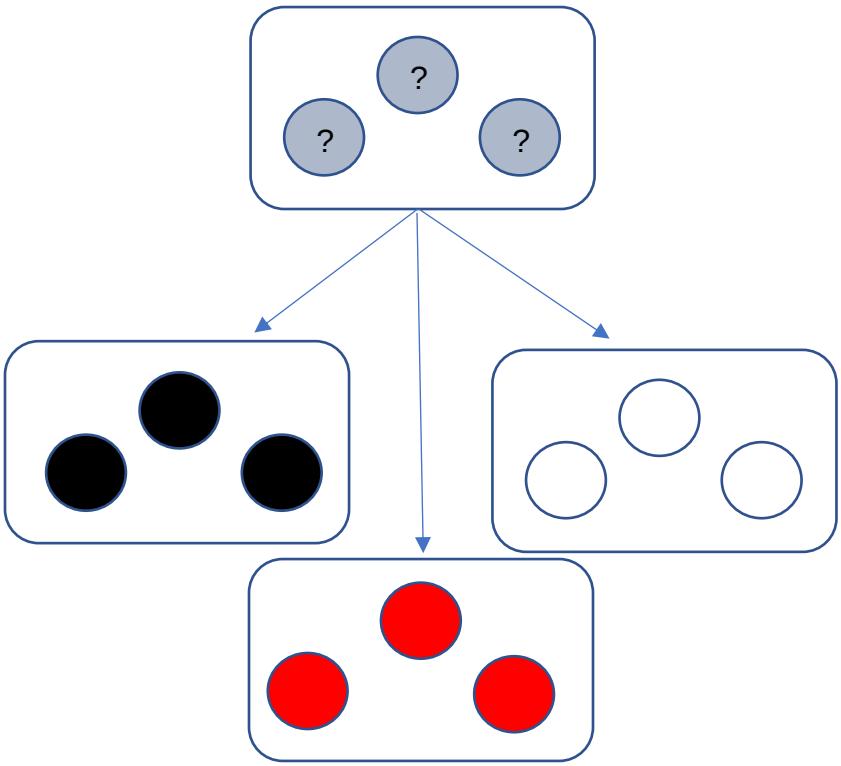

Payoff if drawn	Black - $5	Red - $10	White - $15
Total number of balls in an urn		3	

Your bid (in dollars and cents): $_____ . _____

Experiment 3 Decision Rounds – Part B

The decision rounds in Experiment 3 are divided into two parts – Part A and Part B. There are eight decision rounds in this part.

Your task in each round is to place a bid to buy a ticket to draw ONE ball from an urn. You will be told the possible colors of the balls in the urn, and the payoffs for each color.

In this part, a ball may be drawn from one of several possible urns. If you buy a ticket to draw a ball from an urn, you will first draw a card to determine what urn is in play, and then draw a ball from the selected urn. To illustrate how this works, we will go through an example.

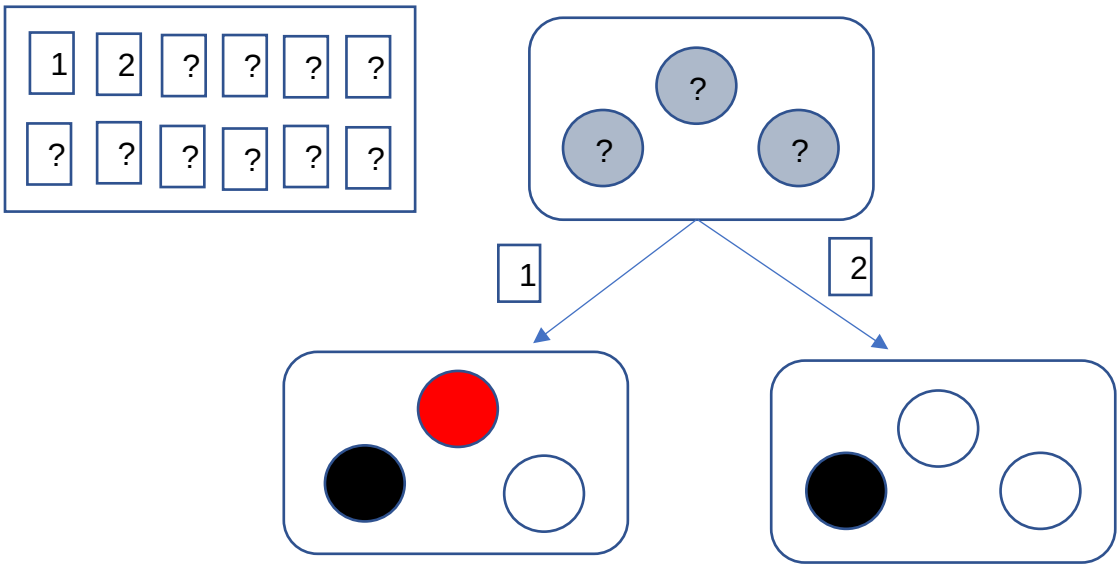

In this example, there are two different urns, numbered 1 and 2. Urn 1 contains one black ball, one red ball and one white ball. Urn 2 contains one black ball and two white balls. If you purchased a ticket in this situation, you would first draw a card from a stack of TWELVE cards. You know for sure that one of the cards is for Urn 1 and one is for Urn 2. You do not know the number on the other cards – they can all be 1s or all be 2s or any combination of 1s and 2s. The number on the card drawn identifies the urn in play, and then a ball is selected from this urn. As before, each ball will be worth a particular amount of money based on its color.

The purchase procedure is the same as before. Your bid will be compared to a random price. You will purchase the ticket only if your bid is equal to or higher than the random price. The random price will be a randomly drawn number between $0.00 and a maximum price. The maximum price will an amount equal to or lower than the payoff associated with the highest-valued ball in the urn. For example, if the highest-valued ball is $15, the maximum possible price may be as high as $15. Any price within this range is equally likely to be chosen.

As before, it is in your best interest to place a bid equal to the highest price you are willing to pay for the ticket. By doing so, you will only purchase the ticket at prices you are willing to pay. You will not purchase the ticket at prices you are not willing to pay.

If you instead bid lower than the highest price you are willing to pay, you risk not purchasing the ticket at prices favorable to you. If you instead bid more than the highest price you are willing to pay, you risk purchasing the ticket at prices that are not favorable to you.

Only one decision round, which may be from Part A or Part B, will be implemented for real. After all decision rounds are completed, we will have a volunteer roll dice to determine which round this is. Since you will not know which round will be selected prior to making any decisions, it is in your best interest to take each decision seriously as if it will determine an actual purchase.

After this paid round is selected, a volunteer will roll dice to determine the random price. Another volunteer will select a ball from the urn.

Before we proceed to the decision rounds, are there any questions?

In the next two rounds, the same stack of cards will be used to determine the number of the urn selected.

Please, proceed to the next decision round.

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are four different urns, numbered 1, 2, 3 and 4. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #3 and #4 cards. IF a black ball is selected from the urn, you would repeat the procedure – draw a different card, and then draw a ball from the selected urn – until a red or white ball is chosen.

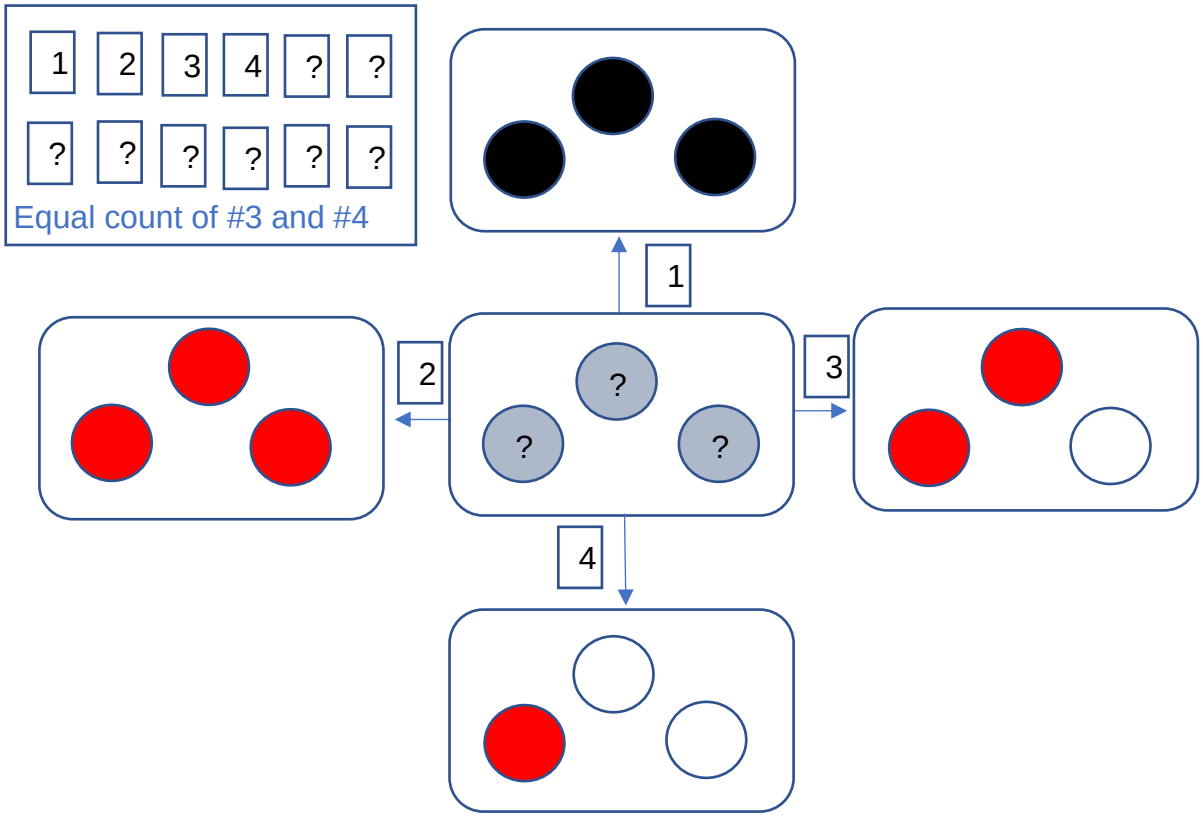

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	3	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are four different urns, numbered 1, 2, 3 and 4. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #3 and #4 cards. IF a black ball is selected from the urn, you would repeat the procedure – draw a different card, and then draw a ball from the selected urn – until a red or white ball is chosen.

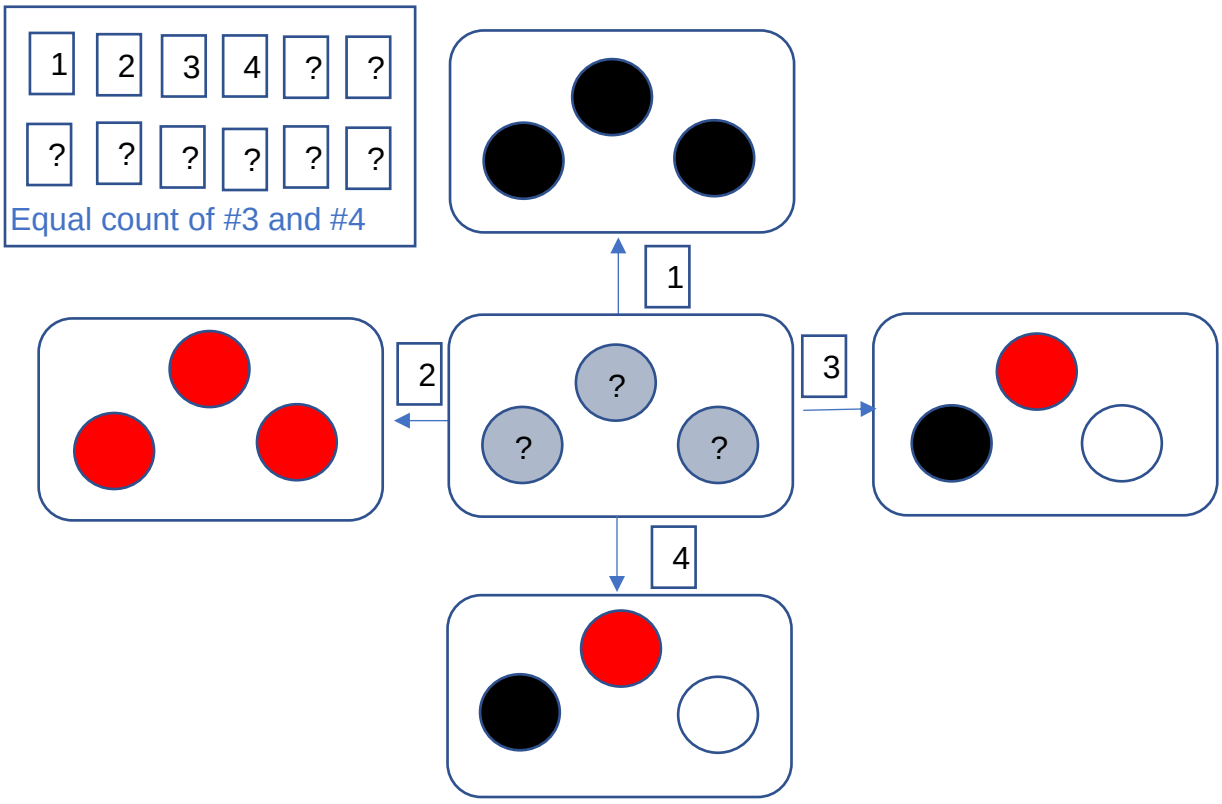

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	3	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

In the next two rounds, the same stack of cards will be used to determine the number of the urn selected.

Please, proceed to the next decision round.

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are three different urns, numbered 1, 2 and 3. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #2 and #3 cards.

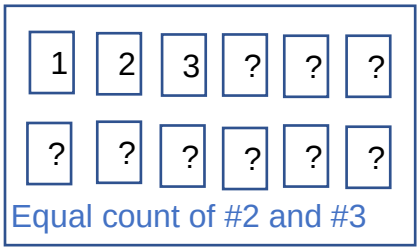

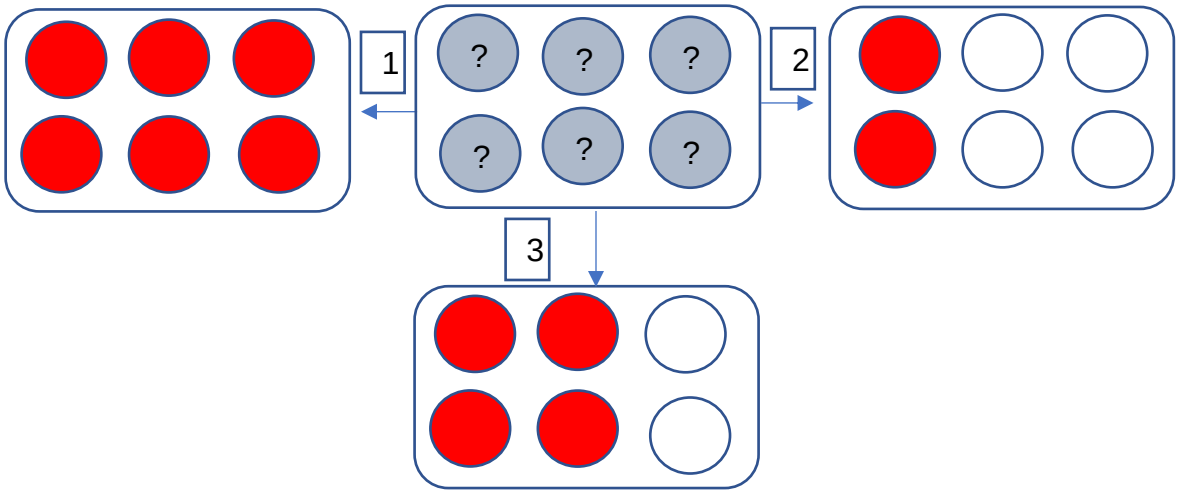

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	6	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are three different urns, numbered 1, 2 and 3. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #2 and #3 cards.

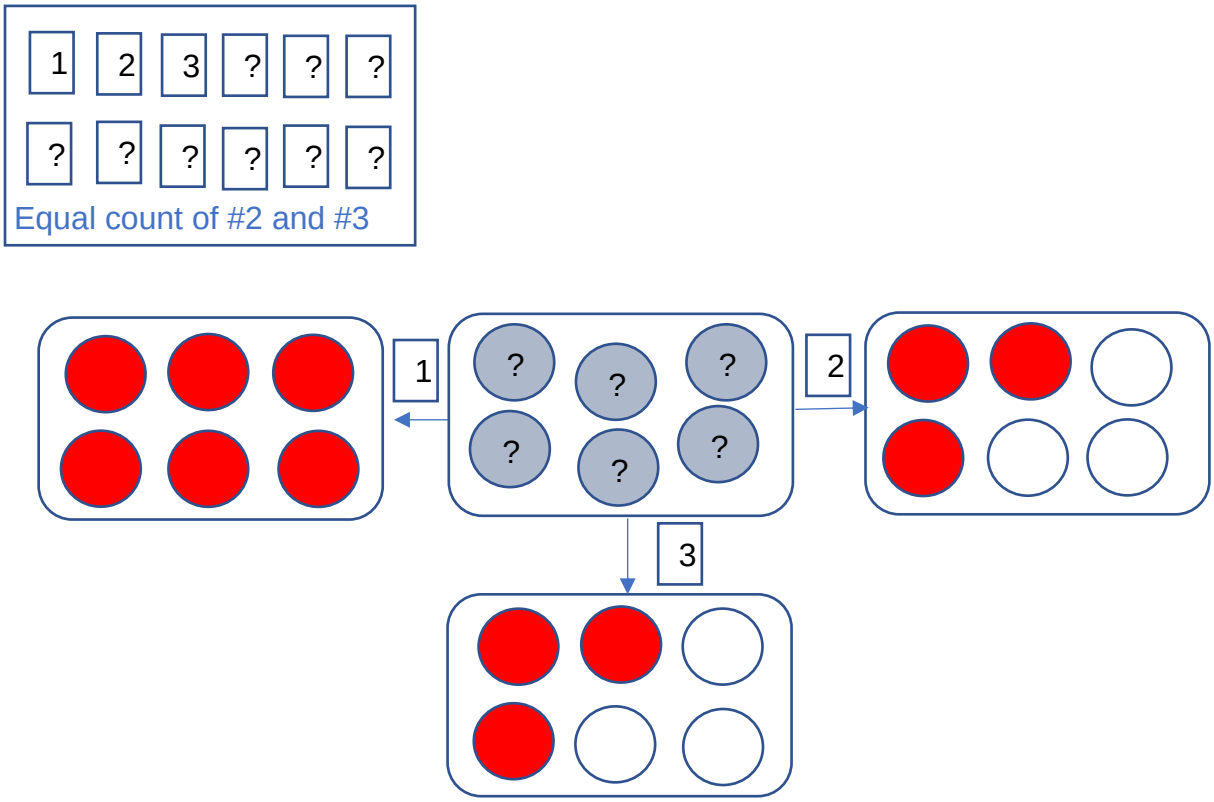

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	6	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

In the next two rounds, the same stack of cards will be used to determine the number of the urn selected.

Please, proceed to the next decision round.

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are five different urns, numbered 1, 2, 3, 4 and 5. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #3 and #4 cards. IF a black ball is selected from the urn, you would repeat the procedure – draw a different card, and then draw a ball from the selected urn – until a red or white ball is chosen.

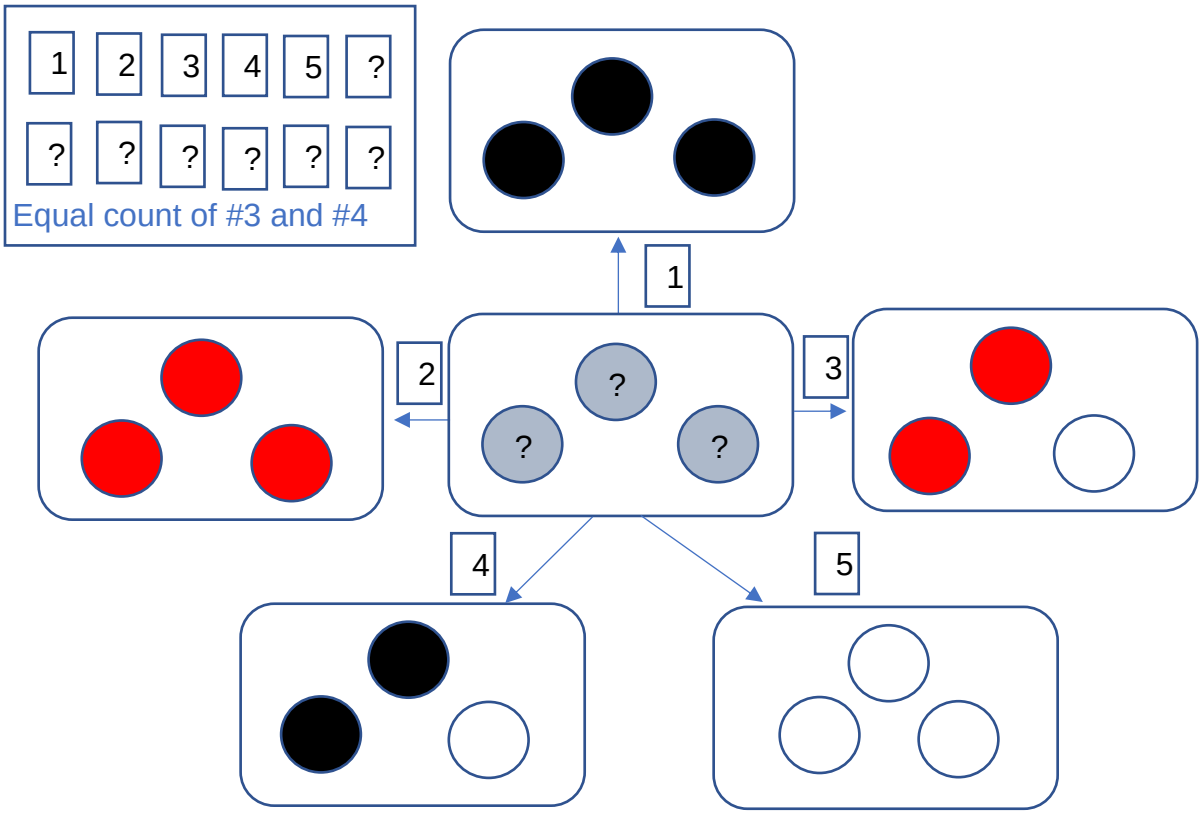

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	3	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are five different urns, numbered 1, 2, 3, 4 and 5. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #3 and #4 cards. IF a black ball is selected from the urn, you would repeat the procedure – draw a different card, and then draw a ball from the selected urn – until a red or white ball is chosen.

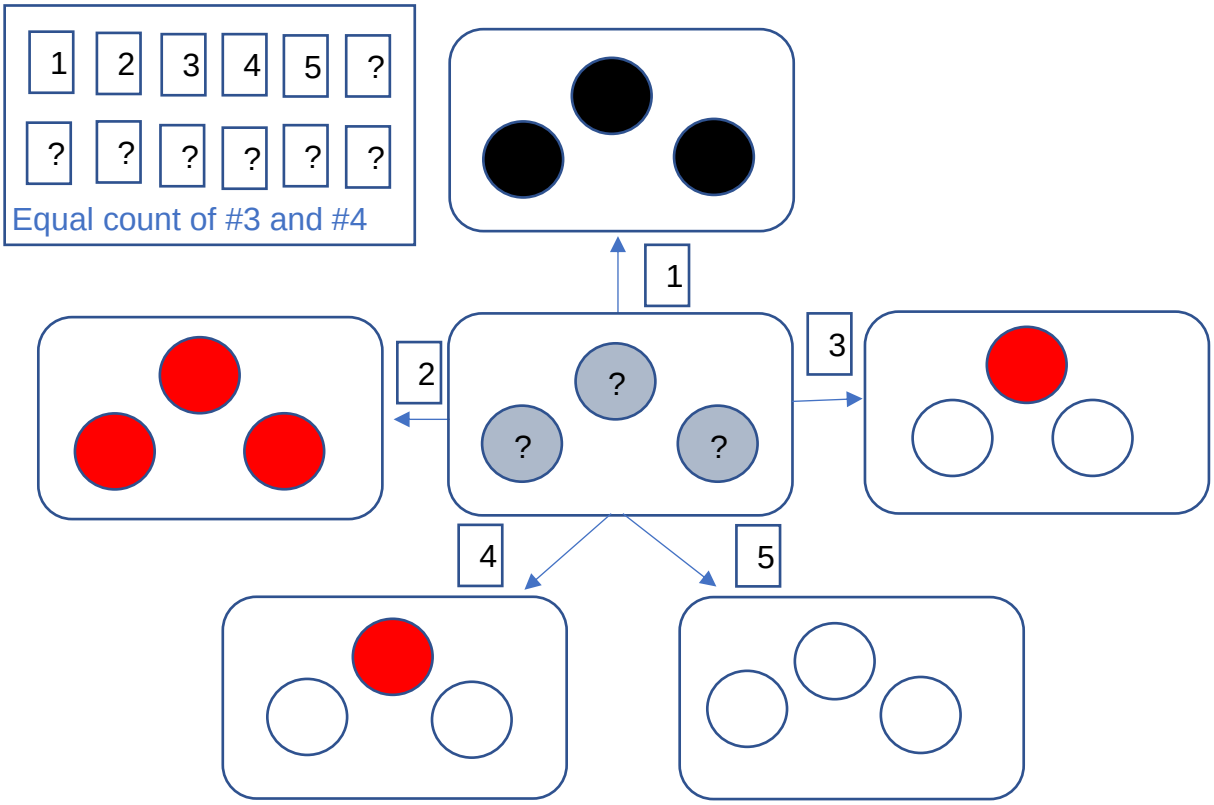

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	3	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

In the next two rounds, the same stack of cards will be used to determine the number of the urn selected.

Please, proceed to the next decision round.

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are four different urns, numbered 1, 2, 3 and 4. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #2 and #3 cards.

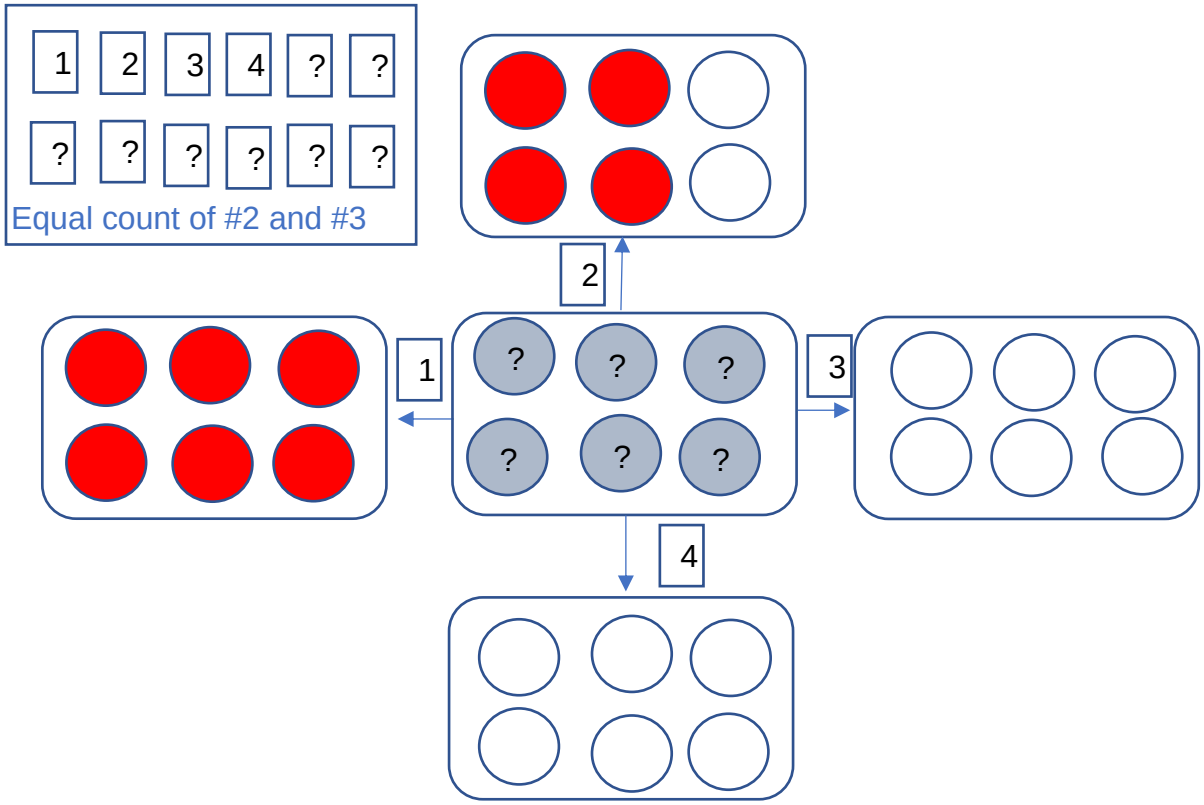

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	6	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

Please read the description carefully. It will be in your best interest to place a bid equal to the highest price you are willing to pay. If you buy a ticket, you will receive the payoff associated with the color of the ONE ball drawn from the selected urn, and pay the random price.

Description:

There are four different urns, numbered 1, 2, 3 and 4. The ball colors in each are illustrated in the figure below. If you purchase a ticket, you would first draw a card from a stack of TWELVE cards to select what urn is used. There are an equal number of #2 and #3 cards.

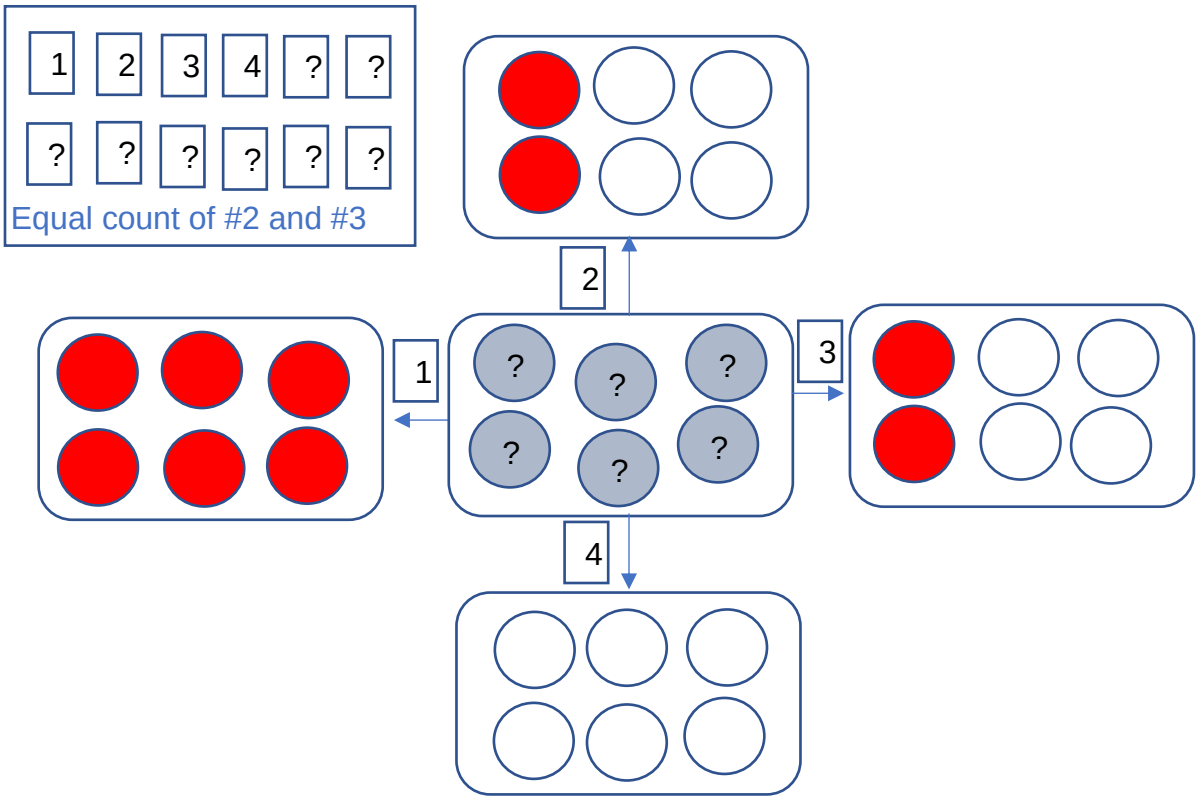

Payoff if drawn	Red - $10	White - $15
Total number of balls in an urn	6	
Total number of cards	12	

Your bid (in dollars and cents): $_____ . _____

Earnings Record Sheet

Payment for completing all experiments and the questionnaire	$10.00
Earnings from Experiment 1	
Earnings from Experiment 2	
Practice calculations	
Training round 1 – trial 1	
Training round 1 – trial 2	
Training round 1 – trial 3	
Training round 1 – trial 4	
Training round 2 – trial 1	
Training round 2 – trial 2	
Training round 2 – trial 3	
Training round 2 – trial 4	
Paid Decision Round **Is your bid equal to or higher than the random price?** □ **Yes**. You bought a ticket. Your earnings are equal to: $_____ – _____ (value of ball) (random price) □ **No**. You did not buy a ticket. Your earnings are $0.	

Total earnings (add up <u>all</u> amounts above): $_____._____

Round up your total earnings to the **next highest quarter**, and record this amount here and on your receipt form:

$_____._____

Questionnaire

We would now like for you to complete a short survey. Please note that all answers are strictly confidential and will be used for statistical purposes only.

1. Have you previously participated in an economics experiment?

(*circle one*) YES NO

2.Did you understand the instructions for the experiment today? Please rate your understanding on a scale from 1 to 5. (Circle one number.)

I understood very poorly				I understood very well
1	2	3	4	5

3.Did you feel that you were well-compensated for your participation in this experiment? Please rate your satisfaction with the compensation on a scale from 1 to 5. (Circle one number.)

I was compensated very poorly				I was compensated very well
1	2	3	4	5

4. What is your age? _____

5. What is your gender? (circle one) Male Female

6. What is your major? (*be specific*)

_____

7. What is your classification for the spring 2017 semester? (*circle one*)

Freshman Sophomore Junior Senior

Master's Student Law Student Doctoral Student

Other _____

8. What is your student status for the current semester? (*circle one*)

Full-time student Part-time student (taking fewer than 12 hours/sem)

Not a student Other (please specify)

9. How many economic courses have you taken at the university level? (include this semester)

_____

10. How would you best describe your current employment situation? (*circle one*)

Full-time employment outside of the university

Part-time employment outside of the university

Student only

Work at the university/research assistantship

11.In what range is your GPA? (*circle one*)

0 to 2.0 / 2.1 to 2.5 / 2.6 to 3.0 / 3.1 to 3.5 / 3.6 to 4.0

Here are a number of personality traits that may or may not apply to you. Please write a number next to each statement to indicate the extent to which *you agree or disagree with that statement*. You should rate the extent to which the pair of traits applies to you, even if one characteristic applies more strongly than the other.

Disagree strongly	Disagree moderately	Disagree a little	Neither agree or disagree	Agree a little	Agree moderately	Agree strongly
1	2	3	4	5	6	7

I see myself as:

1. _____ Extraverted, enthusiastic.
2. _____ Critical, quarrelsome.
3. _____ Dependable, self-disciplined.
4. _____ Anxious, easily upset.
5. _____ Open to new experiences, complex.
6. _____ Reserved, quiet.
7. _____ Sympathetic, warm.
8. _____ Disorganized, careless.
9. _____ Calm, emotionally stable.
10. _____ Conventional, uncreative.

Please use the space below to write any comments you may have about the experiment.

Appendix B3: Tables

Table 2.1. Description of lottery tasks

Set	Task	Payoffs ($B/$R/$W)	Description
1	1A	5/10/15	One urn, containing three balls. The possible ball colors are black, red and white.
1	1A′	trimmed/10/15	Same as 1A, except that IF this urn contains one or more black balls, these balls are replaced with red or white balls before a ball is drawn.
1	1B	5/10/30	Same as 1A, except payoff of white ball increased to 30.
1	1B′	trimmed/10/30	Same as 1A′, except payoff of white ball increased to 30.
2	2A	5/–/15	One urn, containing three balls. Either all are black or all are white.
2	2B	5/10/15	Same as 2A, except that it is now also possible that all balls are red.
3	3A	trimmed/10/15	Four urns, each containing three balls. The possible ball colors are black, red and white. Urn 1: all are black. Urn 2: all are red. Urn 3: two red and one white. Urn 4: one red and two white. An urn is selected by drawing a card from a stack of 12 cards, numbered 1 to 4. All that is known is that there is an equal chance of selecting Urn 3 or Urn 4. IF a black ball is selected from the urn, the urn selection procedure is repeated until a red or white ball is drawn.
3	3B	trimmed/10/15	Same as 3A, except that Urn 3 and Urn 4 now contain one ball of each color.
4	4A	–/10/15	Three urns, each containing six balls. The possible ball colors are red and white. Urn 1: all are red. Urn 2: four white and two red. Urn 3: four red and two white. An urn is selected by drawing a card from a stack of 12 cards, numbered 1 to 3. All that is known is that there is an equal chance of selecting Urn 2 or Urn 3.
4	4B	–/10/15	Same as 4A, except Urn 2 and Urn 3 now each contain three red and three white balls.
5	5A	trimmed/10/15	Five urns, each containing three balls. The possible ball colors are black, red and white. Urn 1: all are black. Urn 2: all are red. Urn 3: two red and one white. Urn 4: two black and one white. Urn 5: all are white. An urn is selected by drawing a card from a stack of 12 cards, numbered 1 to 5. All that is known is that there is an equal chance of selecting Urn 3 or Urn 4.IF a black ball is selected from the urn, the urn selection procedure is repeated until a red or white ball is drawn.
5	5B	trimmed/10/15	Same as 5A, except that Urn 3 and Urn 4 now each contain one red and two white balls.
6	6A	–/10/15	Four earns, each containing six balls. The possible ball colors are black, red and white. Urn 1: all are red. Urn 2: four white and two red. Urn 3: all are white. Urn 4: all are white. An urn is selected by drawing a card from a stack of 12 cards, numbered 1 to 4. All that is known is that there is an equal chance of selecting Urn 2 or Urn 3.
6	6B	–/10/15	Same as 6A, except that Urn 2 and Urn 3 now each contain two red and four white balls.

Table 2.2 . Variables and descriptive statistics

Variable name	Description	Mean	Std. Dev.
Female	=1 if identified gender is female	0.42	0.49
Employed	=1 if participant has a part time or full time job	0.40	0.49
Age	age, in years	20.71	1.74
GPA	cumulative GPA; midpoint of selected range	3.37	0.43
STEM	=1 if majoring in a STEM field	0.42	0.49
Extraversion	measure of personality trait “extraversion”, 1 to 7	4.54	1.51
Agreeableness	measure of personality trait “agreeableness”, 1 to 7	4.75	1.27
Conscientiousness	measure of personality trait “conscientiousness”, 1 to 7	5.49	0.96
Emotional Stability	measure of personality trait “emotional stability”, 1 to 7	4.90	1.35
Openness	measure of personality trait “openness to experience”, 1 to 7	5.31	1.17
Risk	number of Lottery A (safe) choices selected in risk MPL	6.26	1.50
Risk Averse	=1 if number of Lottery A (safe) choices >5 in risk MPL	0.71	0.45
Ambiguity Averse	=1 if number of Urn A choices >5 in ambiguity MPL	0.89	0.31
Comprehension	stated understanding of the experiment; Likert-scale with 1 “very poorly” and 5 “very well”	4.33	0.75
Compensation	stated satisfaction about experiment compensation; Likert-scale with 1 “very poorly” and 5 “very well”	4.46	0.70
Earnings	earnings from experiment session, in $	21.91	4.58

Table 2.3. Lottery valuations, by task (in $)

Task	Mean	Std. Dev.	Min	Max
1A	7.70	2.79	4.00	15.00
1A′	10.65	2.01	5.00	15.00
1B	10.41	5.76	4.99	30.00
1B′	13.98	5.44	5.00	30.00
2A	7.52	3.33	1.37	20.00
2B	8.30	3.29	2.25	25.00
3A	10.26	2.09	4.19	19.99
3B	9.91	1.86	4.00	15.00
4A	10.47	1.90	2.21	15.00
4B	10.34	2.01	3.75	15.00
5A	10.30	2.32	2.00	15.00
5B	11.05	2.33	3.00	15.00
6A	11.26	2.15	5.25	15.00
6B	11.09	2.38	4.25	15.00

Table 2.4. Hypothesis test results, based on lottery valuations

Null hypothesis	Paired *t*-test	Wilcoxon signed-rank test	Negative difference	Zero difference	Positive difference
(a) Tests of kinked versus smooth specifications					
H1: $V_{1A'} - V_{1A} = V_{1B'} - V_{1B}$	-1.79 (0.08)	-1.12 (0.26)	32%	42%	26%
H2: $V_{2A} = V_{2B}$	-3.63 (<0.01)	-4.12 (<0.01)	43%	44%	14%
H3a: $V_{3A} = V_{3B}$	2.08 (0.04)	3.30 (0.01)	10%	59%	31%
H3b: $V_{4A} = V_{4B}$	1.24 (0.21)	0.95 (0.34)	17%	62%	21%
H4a: $V_{5A} = V_{5B}$	-3.91 (<0.01)	-4.29 (<0.01)	40%	48%	12%
H4b: $V_{6A} = V_{6B}$	1.74 (0.09)	1.34 (0.18)	15%	63%	21%
(b) Tests of task complexity					
H5a: $V_{3A} = V_{4A}$	-1.40 (0.16)	-1.77 (0.08)	26%	57%	16%
H5b: $V_{3B} = V_{4B}$	-2.64 (<0.01)	-2.42 (0.02)	34%	48%	18%
H5c: $V_{5A} = V_{6A}$	-5.79 (<0.01)	-5.98 (<0.01)	49%	44%	8%
H5d: $V_{5B} = V_{6B}$	-0.25 (0.80)	0.16 (0.87)	21%	56%	23%

Notes: *p*-values in parentheses. The last three columns indicate the percentage of participants with negative, zero, or positive differences when evaluating the null hypothesis; e.g., for H2, 43% of respondents placed higher bids for lottery 2B, and 14% of participants placed bids higher bids for 2A.

Table 2.5. Linear regression models, bid differences

Variable	Model 1: Theory Specification	Model 2: Task Complexity
Female	0.102 (0.195)	-0.141 (0.226)
Risk	-0.128** (0.052)	-0.032 (0.054)
GPA	-0.141 (0.217)	0.933** (0.330)
STEM	-0.455** (0.208)	-0.007 (0.207)
Extraversion	0.054 (0.065)	0.047 (0.062)
Agreeableness	-0.082 (0.079)	0.022 (0.079)
Conscientiousness	-0.010 (0.099)	0.007 (0.096)
Emotional Stability	-0.020 (0.057)	-0.129 (0.089)
Openness	0.083 (0.074)	0.028 (0.110)
H1	-0.586 (0.356)	
H2	-0.797** (0.223)	
H3a	0.382** (0.176)	
H3b	0.088 (0.107)	
H4a	-0.746** (0.202)	
H4b	0.165 (0.108)	
H5a		-0.135 (0.150)
H5b		-0.428** (0.162)
H5c		-0.954** (0.162)
H5d		-0.043 (0.173)
Controls for task order?	Yes	Yes
R^2	0.079	0.166
F-stat	3.83**	3.66**
Number of observations	666	444

Notes: cluster-robust standard errors in parentheses. * and * indicate estimate is statistically significant at the 10% and 5% significance levels, respectively.

Table 2.6. Multinomial probit models, bid difference categories

Variable	Model 3: Theory Specification		Model 4: Task Complexity	
	Negative difference	Positive difference	Negative difference	Positive difference
Female	-0.030	0.069	-0.251	-0.305
	(0.236)	(0.253)	(0.302)	(0.320)
Risk	-0.109	-0.220**	-0.137	-0.144
	(0.069)	(0.077)	(0.089)	(0.101)
GPA	-0.644**	-0.961**	-0.132**	-0.937**
	(0.289)	(0.305)	(0.352)	(0.367)
STEM	0.263	0.058	-0.113	0.236
	(0.218)	(0.218)	(0.262)	(0.274)
Extraversion	0.077	0.159**	0.056	0.202**
	(0.068)	(0.077)	(0.098)	(0.091)
Agreeableness	-0.058	-0.113	-0.123	-0.181
	(0.092)	(0.090)	(0.104)	(0.115)
Conscientiousness	0.033	-0.026	0.008	0.104
	(0.130)	(0.122)	(0.156)	(0.164)
Emotional Stability	-0.204**	-0.219**	-0.053	-0.244**
	(0.085)	(0.086)	(0.113)	(0.102)
Openness	-0.015	0.011	0.170	0.162
	(0.099)	(0.098)	(0.118)	(0.116)
Hypothesis indicators?	Yes		Yes	
Controls for task order?	Yes		Yes	
Log-likelihood	-604.276		-383.463	
Wald-stat	130.65**		112.17**	
Number of observations	666		444	

Notes: cluster-robust standard errors in parentheses. * and * indicate estimate is statistically significant at the 10% and 5% significance levels, respectively. Reference category is Zero difference.

Table 2.7. Hypothesis test results, excluding participants who poorly understood instructions

Null hypothesis	**Paired *t*-test (*p*-value)**	**Wilcoxon signed-rank test (*p*-value)**
H1: $V_{1A'} - V_{1A} = V_{1B'} - V_{1B}$	-1.52 (0.16)	-0.89 (0.38)
H2: $V_{2A} = V_{2B}$	-3.38 (<0.01)	-3.81 (<0.01)
H3: $V_{3A} = V_{3B}$	2.03 (0.05)	3.25 (<0.01)
H3: $V_{4A} = V_{4B}$	0.77 (0.44)	0.49 (0.62)
H4: $V_{5A} = V_{5B}$	-4.12 (<0.01)	-4.39 (<0.01)
H4: $V_{6A} = V_{6B}$	0.84 (0.40)	0.26 (0.80)
H5: $V_{3A} = V_{4A}$	-1.38 (0.18)	-1.87 (0.06)
H5: $V_{3B} = V_{4B}$	-2.9 (<0.01)	-2.86 (<0.01)
H5: $V_{5A} = V_{6A}$	-5.96 (<0.01)	-5.87 (<0.01)
H5: $V_{5B} = V_{6B}$	-0.85 (0.40)	-0.33 (0.74)

Appendix B4: Figures

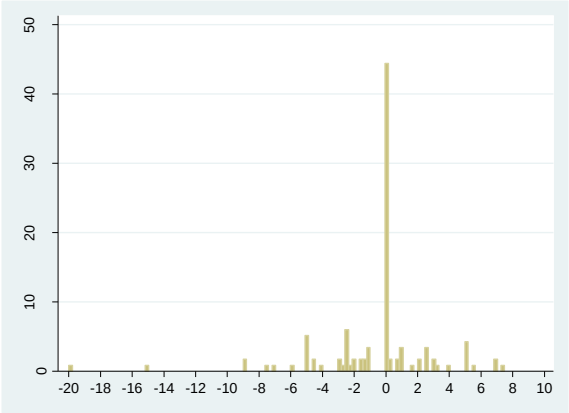

(a) H1, $V_{1A'} - V_{1A} - (V_{1B'} - V_{1B})$

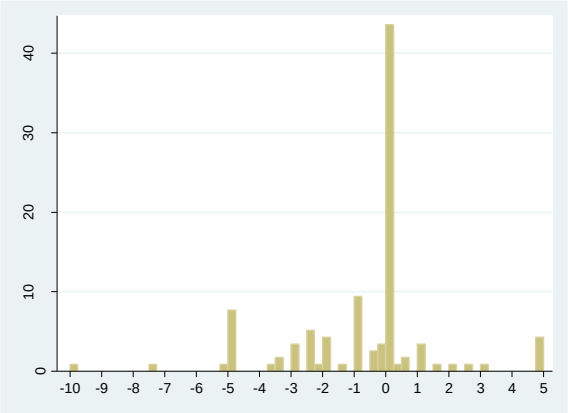

(b) H2, $V_{2A} - V_{2B}$

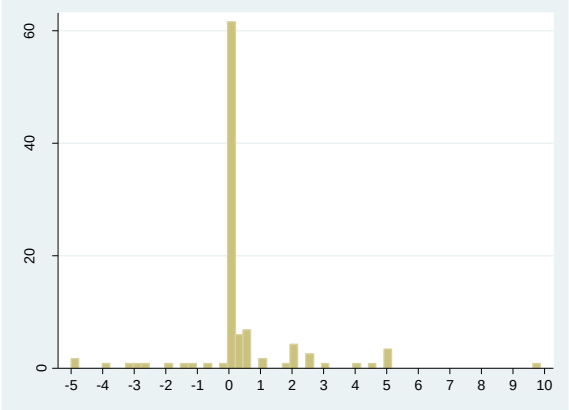

(c) H3a, $V_{3A} - V_{3B}$

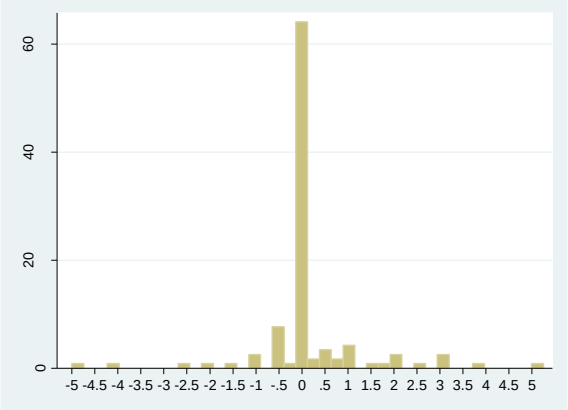

(d) H3b, $V_{4A} - V_{4B}$

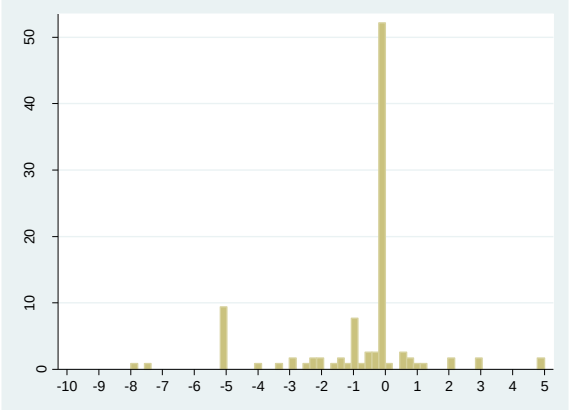

(e) H4a, $V_{5A} - V_{5B}$

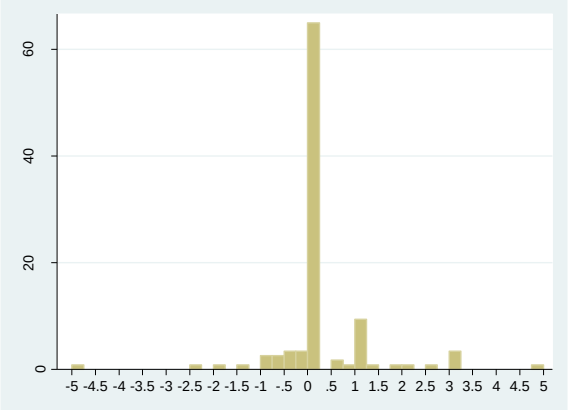

(f) H4b, $V_{6A} - V_{6B}$

Figure 2.1. Individual-level results, kinked versus smooth specifications (Hypotheses 1 to 4)

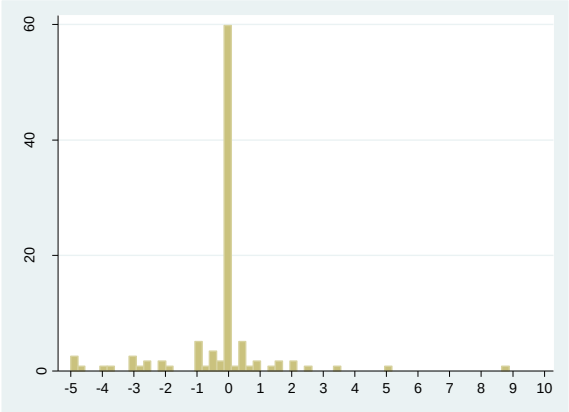

(a) H5a, $V_{3A} - V_{4A}$

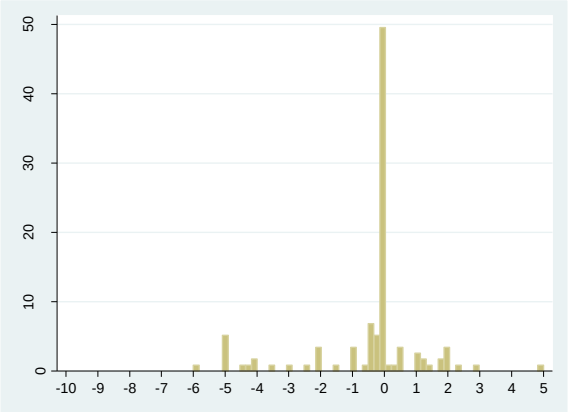

(b) H5b, $V_{3B} - V_{4B}$

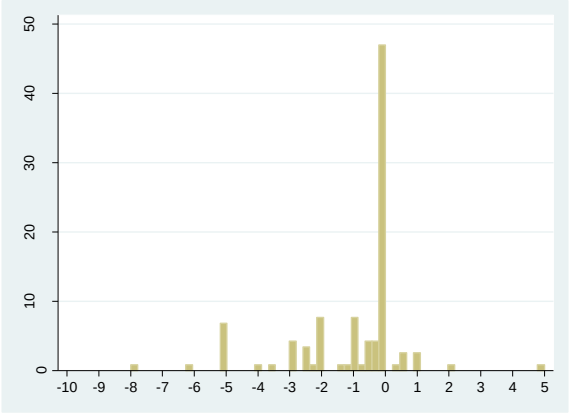

(c) H5c, $V_{5A} - V_{6A}$

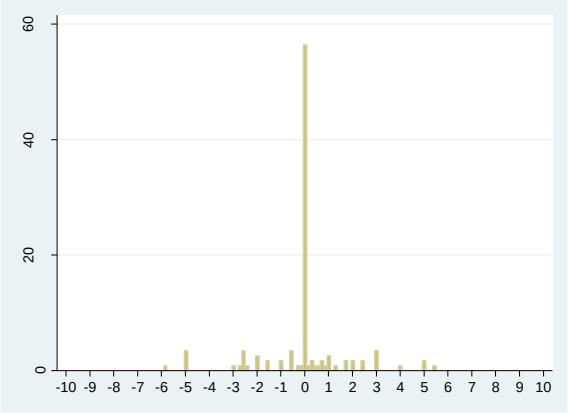

(d) H5d, $V_{5\mathrm{B}} - V_{6B}$

Figure 2.2. Individual-level results, task complexity (Hypothesis 5)

Appendix C: Climate adaptation under deep uncertainty

Appendix C1: Theory Derivations

Assume streamflow, $N(t)$, is a stochastic variable and follows a geometric Brownian motion:

(1) $$dN_t = \alpha N_t dt + \sigma N_t dB.$$

For an initial value of streamflow, N_0, the stochastic differential equation above has the analytic solution

(2) $$N(t) = N_0 e^{\left[\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma B(t)\right]}.$$

Equation (A2) implies that $N(t)$ is a log-normally distributed random variable with expected value $E[N] = N_0 e^{\mu t}$ and variance $Var[N] = N_0^2 e^{2\mu t}\left(e^{\sigma^2 t} - 1\right)$. Now assume profits are a concave function of water supply

(3) $$\pi(t) = \alpha N(t)^{\beta}.$$

Since $\pi(t)$ is a function of $N(t)$, it is also a stochastic function. Applying Ito's lemma, we find

(4) $$d\pi = \frac{\partial \pi}{\partial t} dt + \frac{\partial \pi}{\partial N} dN + \frac{1}{2}\frac{\partial^2 \pi}{\partial N^2}(dN)^2,$$

where $(dN)^2 = (\sigma N)^2 dt$. We can rewrite as

(5) $$d\pi = \frac{\partial \pi}{\partial t} dt + \frac{\partial \pi}{\partial N}[\mu N dt + \sigma N dB] + \frac{1}{2}\frac{\partial^2 \pi}{\partial N^2}(\sigma N)^2 dt$$

$$d\pi = \left[\frac{\partial \pi}{\partial t} + \frac{\partial \pi}{\partial N}\mu N + \frac{1}{2}\frac{\partial^2 \pi}{\partial N^2}(\sigma N)^2\right] dt + \frac{\partial \pi}{\partial N}\sigma N dB.$$

Given the specification of the profit function above, we have

(6) $\frac{\partial \pi}{\partial t} = 0,$

(7) $\frac{\partial \pi}{\partial N} = \alpha\beta N(t)^{\beta-1} = \frac{\beta}{N}\pi,$

(8) $\frac{\partial^2 \pi}{\partial N^2} = \alpha\beta(\beta-1)N(t)^{\beta-2} = \frac{\beta(\beta-1)}{N^2}\pi.$

Substituting (A6), (A7) and (A8) into (A5), we have

(9) $d\pi = \left[\frac{\beta}{N}\pi\mu N + \frac{1}{2}\frac{\beta(\beta-1)}{N^2}\pi(\sigma N)^2\right]dt + \frac{\beta}{N}\pi\sigma N dB$

$$d\pi = \left[\beta\mu + \frac{1}{2}\beta(\beta-1)\sigma^2\right]\pi dt + b\sigma\pi dB$$

$$d\pi = \alpha^p \pi dt + \sigma^p \pi dB.$$

where

$$\alpha^p = \left[\beta\mu + \frac{1}{2}\beta(\beta-1)\sigma^2\right]$$

$$\sigma^p = \beta\sigma$$

Appendix C2: Tables

Table 3.1. Parameter values based on a farm producing alfalfa that is considering switching from flood to center pivot irrigation

Parameter	Value	Source
P	$197/ton	2013 Yuba County Crop Report
y	5 tons per acre annually	USDA 2013 Farm and Ranch Irrigation Survey
γ_0	1.1	USDA 2013 Farm and Ranch Irrigation Survey
γ_1	2	Brown, 2008
M	$1,000/acre	Natural Resources Conservation Service
θ	0.000006	
$\widetilde{N}$	855,150 annual acre feet	
ρ	0.09	
e	1.94	estimated from *Water Strategist* data
φ	899	estimated from *Water Strategist* data
ε	0.8	
a_0	1031	
b_0	0.01448	
a_1	1094	
b_1	0.0008122	

Table 3.2. Optimal adoption time (t*, in years)

Expectations of water availability based on historical weather		Never adopt
Expectations of water availability based on climate projections	Ambiguity averse **($\varepsilon = 1$)**	193.1
	Ambiguity neutral**($\varepsilon = 0$)**	241.1

Table 3.3. Effect of changing risk and ambiguity on the expected adoption time (t*, in years)

		Risk	
	Low	Benchmark	High
Ambiguity averse ($\varepsilon = 1$)	190.3	193.1	257.4
Ambiguity neutral **($\varepsilon = 0$)**	199.3	241.1	330.8

Table 3.4. Sensitivity of results to parameter changes (t*, in years)

		5% Decrease	5% Increase
a_0	Profit scalar under old technology	Immediately	Never
b_0	Sensitivity of profits to water supply under old technology	196.9	189.9
a_1	Profit scalar under conservation technology	938.2	Immediately
b_1	Sensitivity of profits to water supply under conservation technology	192.0	194.3
α_q	Expected rate of change in water supply	168.7	217.7

Table 3.5. Effect of changing risk and ambiguity on the critical adaptation threshold (N*, million acre feet)

		Risk	
	Low	Benchmark	High
Ambiguity averse ($\boldsymbol{\varepsilon = 1}$)	5.387	5.283	4.489
Ambiguity neutral ($\boldsymbol{\varepsilon = 0}$)	4.639	3.491	2.544

Farmer will switch technology once the streamflow went below the threshold. Thus, a higher threshold means a larger incentive for the technology adaptation.

Appendix C3: Figures

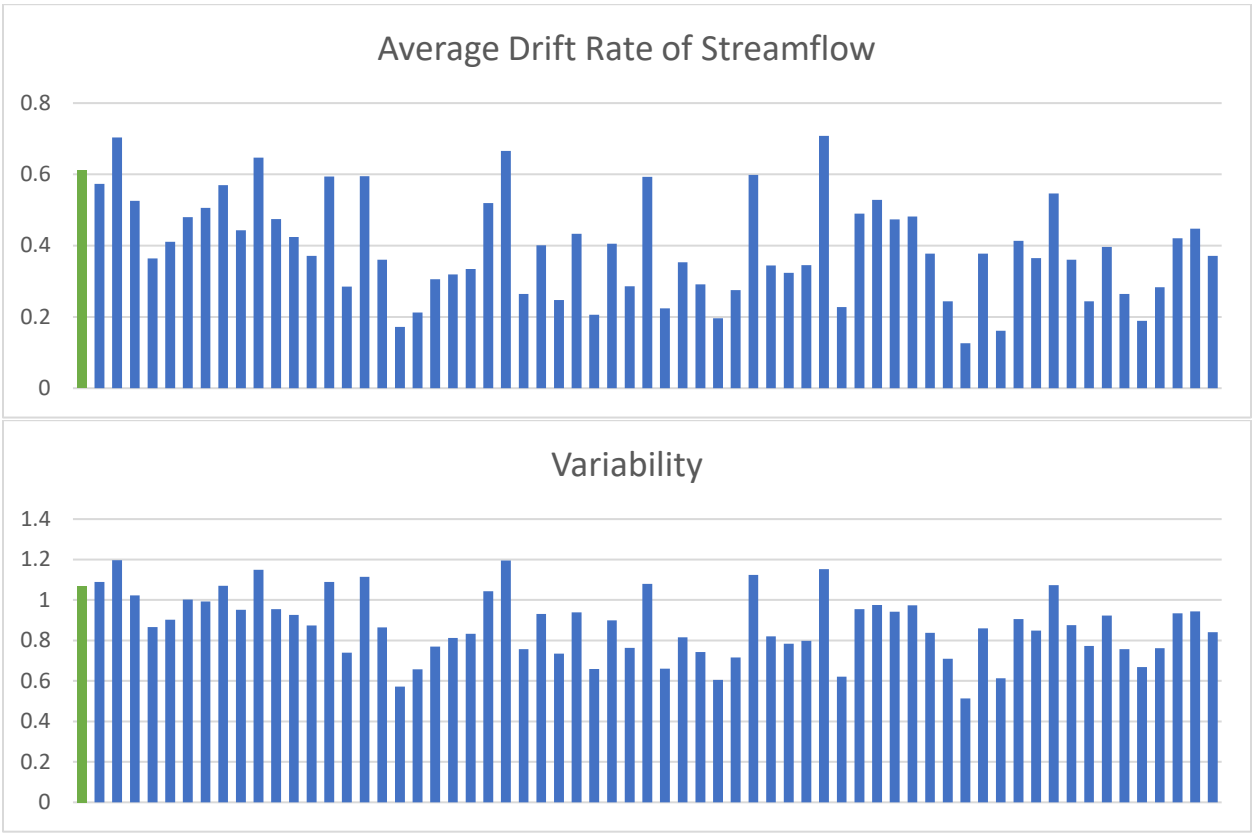

Figure 3.1. Streamflow Parameters (million acre feet)

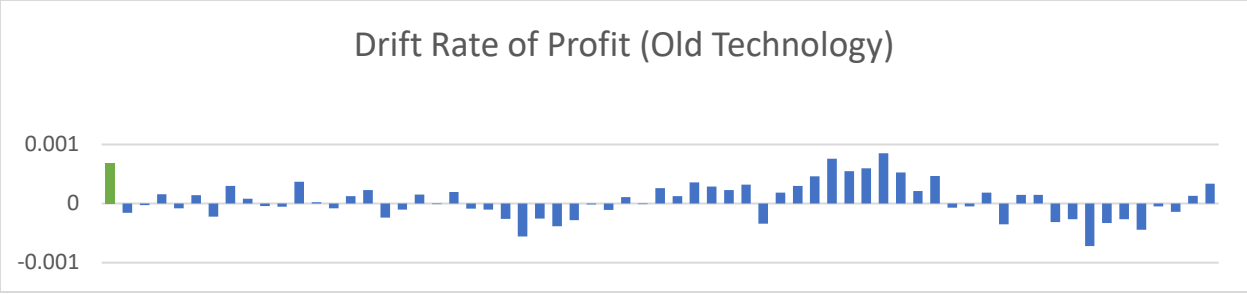



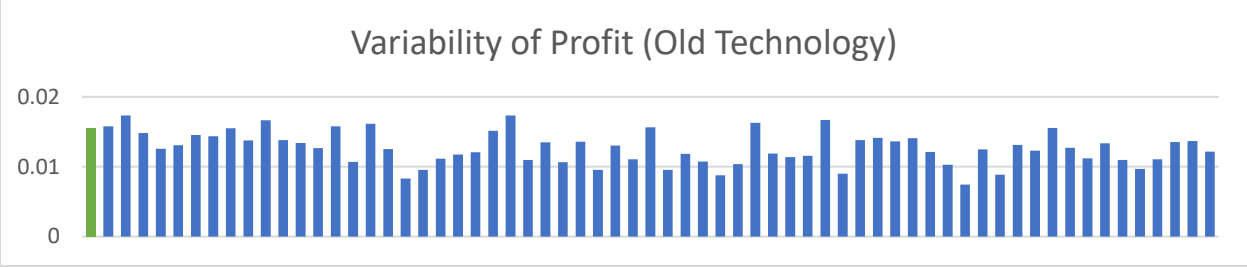

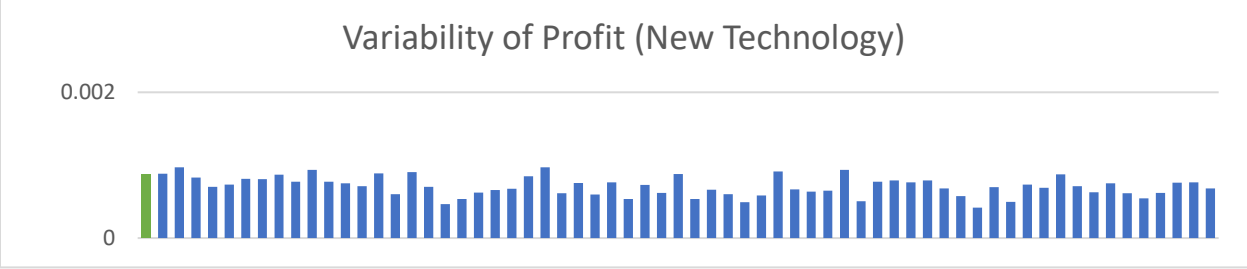

Figure 3.2. Parameters values for stochastic process for profit

		α^p	σ^p
Worst-case prior (miroc5.1.rcp26)		-0.000083	0.011181
Ambiguity neutral prior (average of all projections)		0.000041	0.012665
Backward-looking (historical) prediction		0.00069	0.015479

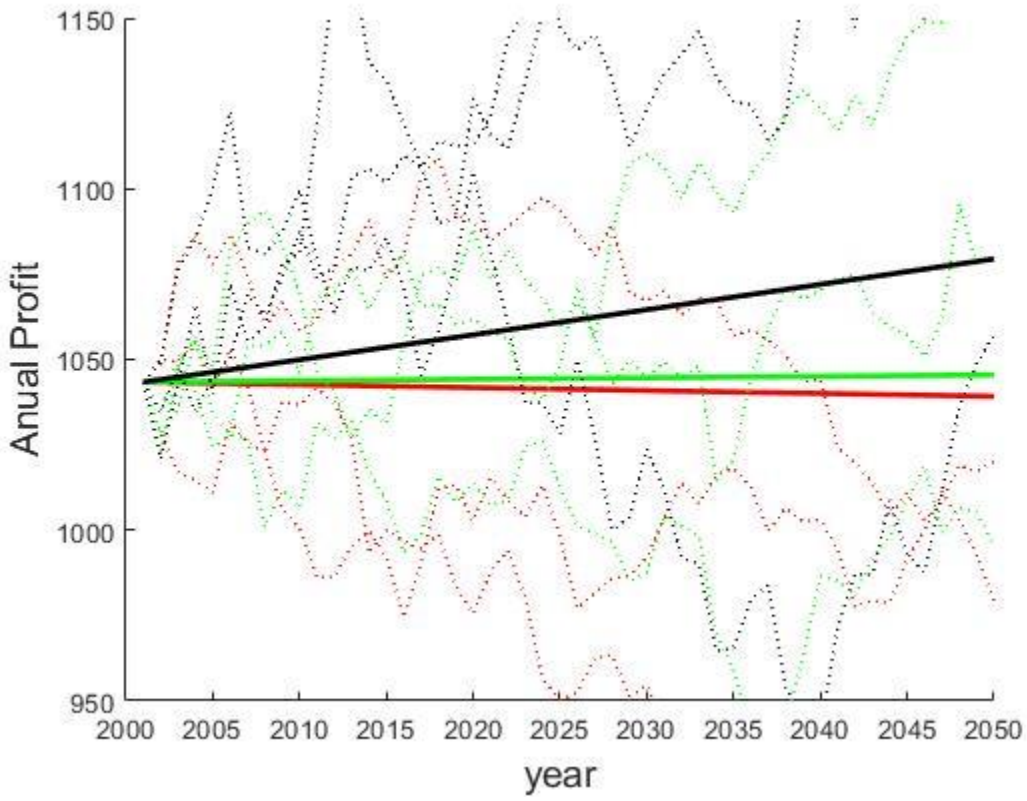

Figure 3.3. Comparison over expected profit trend

VITA

Dong Yan was born in the Northeast part of China. He earned a bachelor's degree in Economics from the Central University of Finance and Economics, and a master's degree in Economics from the University of Illinois at Urbana-Champaign. In 2015, he joined the University of Tennessee and started his life as a Ph.D. student of economics. During his five years of study, Dong independently taught two undergraduate courses, and held research assistant positions in the Boyd Center for Business and Economic Research, Howard H. Baker Jr. Center for Public Policy, and the Department of Economics. These experiences helped him develop important technical skills and establish a research agenda. Dong will begin his academic career in August 2020 as an assistant professor at Renmin University of China, which consistently ranks among the very best universities in China for economics.