

ESSAYS IN BEHAVIORAL ECONOMICS
IN THE CONTEXT OF
STRATEGIC INTERACTION

DISSERTATION

Presented in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy in the Graduate
School of The Ohio State University

By
Asen Ivanov

* * * * *

The Ohio State University

2007

Dissertation Committee:
Professor Dan Levin, Adviser
Professor James Peck
Professor John Kagel
Professor Stephen Cosslett

Approved by

Adviser
Graduate Program in Economics

ABSTRACT

The traditional theoretical concept in game theory, Nash equilibrium, makes strong assumptions about people's rationality and the accuracy of their expectations about others' behavior. As a result, it often provides a poor description of actual behavior. Behavioral Economics seeks to improve the descriptive power of Economics by identifying and studying, often through experiments, actual patterns of behavior and reasoning.

In the first chapter of my dissertation, I study experimentally behavior in one-shot normal-form games. These games allow us to minimize learning and cultural context and to study behavior based mostly on reasoning. In this way, they could provide useful insights into real-life interactions in which people engage without prior experience or clear cultural norms, such as the first spectrum rights auctions or school-matching schemes.

I use a new approach to investigating behavior in one-shot normal-form games. Using subjects' play as well as their stated beliefs about their opponent's play, I study two fundamental dimensions of behavior. The first dimension is whether subjects are naive (do not consider what their opponent might do) or strategic (consider what their opponent might do). The second dimension is whether subjects' behavior is better captured by risk neutrality or by risk aversion.

In treatment A, subjects (graduate students at OSU) play the games without interference from belief elicitation (beliefs are elicited after all games have been played).

I find that (i) only a small minority of subjects is naive, and (ii) the majority of subjects is risk averse. However, these results are not robust to changing the games or the subject population (from graduate to undergraduate students).

Some interesting comparative statics emerge by manipulating treatment A (keeping the games and the subject population fixed). Most notably, when subjects are explicitly prompted to form (and state) beliefs while playing the games (treatment B), then (iii) naive subjects all but disappear, and (iv) the proportion of risk averse subjects decreases dramatically relative to treatment A. A possible explanation for the latter is that seemingly risk averse behavior is actually driven by ambiguity aversion (i.e. by a lack of confidence in one's beliefs rather than by curvature in the utility function). In this case, giving subjects a structured way to think about the games in treatment B may be reducing ambiguity, thus increasing subjects' willingness to take risks. If simply having a structured way to think about a decision situation reduces ambiguity, this has far-reaching implications for behavior under uncertainty.

The second chapter of my dissertation, which is based on joint work with Dan Levin and James Peck, investigates experimentally behavior in a dynamic investment game in which players receive two-dimensional signals (a common-value signal about the market return and a private cost of investing) and timing of investment is endogenous. This game involves two key forces: on the one hand, there is an opportunity to wait and observe investment activity by others; on the other hand, there is a cost to waiting. How these forces play out may have implications for important real world situations. For example, at the end of a recession firms may invest straight away, thus putting an abrupt end to the recession; alternatively, they may wait to observe investment by other firms, thus prolonging the recession.

In an experiment with small (two-player) markets, investment is higher and profits are lower than in Nash equilibrium. The study separately considers whether

a subject draws inferences from the other subject's investment, in hindsight, and whether a subject has the foresight to delay profitable investment and learn from market activity. In contrast to Nash equilibrium, cursed equilibrium, and level-k model predictions, behavior remains the same across the experimental treatments. Maximum likelihood estimates are inconsistent with belief-based theories, but are consistent with the notion that subjects use simple rules of thumb, based on insights about the game.

Dedicated to my mother, father, and sister

ACKNOWLEDGMENTS

There are a number of people who played a crucial role in my graduate studies.

I am deeply indebted to my Adviser, Dan Levin, for his intellectual, moral and financial support. He always acted with my best interest at heart and his advice and encouragement helped me overcome many obstacles along the way. In addition, having him as an Adviser was, frankly, a lot of fun.

I am also very grateful to James Peck who was a major pillar of support throughout my graduate studies. I greatly appreciate his help and scholarly example.

I would also like to express my gratitude to John Kagel who was always ready to give me advice when I needed it. I would also like to thank him for his financial support.

Many thanks to Stephen Cosslett whose help regarding econometric issues was extremely useful.

I would also like to thank Hajime Miyazaki who is responsible for me being accepted to the PhD program at OSU in the first place.

Finally, I would like to thank the NSF for their financial support in the form of a Doctoral Dissertation Research Grant.

VITA

September 08, 1977	Born - Sofia, Bulgaria
2000	B.A., Economics with minor in German Language and Literature, Sofia University, Bulgaria
2003	M.A., Economics, The Ohio State Univer- sity
2003-present	Graduate Teaching and Research Asso- ciate, The Ohio State University

FIELDS OF STUDY

Major Field: Economics

Specialization: Behavioral and Experimental Economics, Microeconomic Theory,
Econometrics

TABLE OF CONTENTS

Abstract	ii
Dedication	v
Acknowledgments	vi
Vita	vii
List of Tables	xi
List of Figures	xiv
1 Strategic Play and Risk Aversion in One-Shot Normal-Form Games:	
An Experimental Study	1
1.1 Introduction	1
1.1.1 Literature	2
1.1.2 Outline of Approach in Current Study	3
1.2 Experimental Design	9
1.2.1 Treatment A	10
1.2.2 Treatment B	11
1.2.3 Treatment C	11
1.2.4 Games	12

1.3	Types and Formal Statistical Model	13
1.4	Results	18
1.4.1	Aggregate-Level Analysis	18
1.4.2	Estimation of Formal Statistical Model	20
1.5	Discussion and Concluding Remarks	26
1.5.1	Type specification	26
1.5.2	Stated Beliefs vs. True Beliefs	27
1.5.3	Robustness Across Games and Subject Population	30
1.5.4	Ambiguity Aversion rather than Risk Aversion?	31
1.5.5	Concluding Remarks	34
2	Hindsight, Foresight, and Insight: An Experimental Study of a Small-Market Investment Game with Common and Private Values	36
2.1	Introduction	36
2.2	Theoretical Framework	41
2.3	Behavioral Considerations	45
2.4	Experimental Design	51
2.5	Results	53
2.5.1	Aggregate-Level Analysis	53
2.5.2	What is Driving Behavior?	63
2.5.3	Treatment 3	71
2.5.4	Personal Characteristics as Determinants of Behavior	76
2.6	Concluding Remarks	77
	Bibliography	79
	Appendix A: Figures and Tables from Ch. 1	84

Appendix B: Instructions for Treatment A from Ch. 1	105
Appendix C: Instructions for Treatment C from Ch. 1	113
Appendix D: Derivation of Symmetric Cursed Equilibrium from Ch. 2	117
Appendix E: Derivation of the Likelihood Function from Ch. 2 . . .	120
Appendix F: Instructions for the Two-Cost Treatment from Ch. 2 .	122
Appendix G: Screen Printout from Alternating One-Cost Treatment from Ch. 2	129

LIST OF TABLES

1.1	Types	3
1.2	Percent of Actions Consistent with each Type.	19
1.3	Percent of Actions Consistent with Naive Types, but not with Strategic Types and vice versa.	19
1.4	Percent of Actions Consistent with Risk Neutral Types, but not with Risk Averse Types and vice versa.	19
1.5	Formal Model Estimation.	21
1.6	Example: Precision.	23
1.7	Hypotheses Tests between Treatments.	25
2.1	Equilibrium Characterization for the Two-Cost Game	43
2.2	Equilibrium Characterization for the Alternating One-Cost Treatment	45
2.3	Nash, Level-k, and Cursed Equilibrium	48
2.4	Aggregate Actions and Frequency of Investment at each History	54
2.5	Compliance with Nash Equilibrium	57
2.6	Investment	59
2.7	Average Profits per Period (in ECU)	59
2.8	Probability of an F Subject's Behavior	66
2.9	Maximum Likelihood Estimates	69
2.10	Maximum Likelihood Estimates - All Treatments	72

2.11 Hypotheses Tests - p-values.	73
2.12 Regression Results for Earnings	76
A.13 Responses of Subjects in Treatment A	96
A.14 Responses of Subjects in Treatment B	104

LIST OF FIGURES

A.1 Games	85
A.2 Marginal Posteriors I	86
A.3 Marginal Posteriors II	87

CHAPTER 1

STRATEGIC PLAY AND RISK AVERSION IN ONE-SHOT NORMAL-FORM GAMES: AN EXPERIMENTAL STUDY

1.1 Introduction

Behavior in a game depends on a combination of reasoning, learning and cultural context. One-shot normal-form games allow us to minimize the effects of learning and cultural context and to study behavior based mostly on reasoning. This approach of isolating reasoning could offer general insights into decision-making in games. On a more practical level, it could provide a useful benchmark for real-life interactions in which individuals engage without prior experience or clear cultural norms, such as the first spectrum rights auctions or school-matching schemes.

Experimental investigation of behavior in one-shot normal-form games is necessary since the theoretical concept, Nash equilibrium, often provides a poor description of behavior in the absence of learning and cultural context.

In the current paper, we focus on two general dimensions of behavior in one-shot normal-form games. The first dimension is whether a player ignores what the opponent might do (i.e. behaves naively) or whether she considers what the opponent

might do (i.e. behaves strategically). The second dimension is whether a player's behavior is better captured by risk neutrality or by risk aversion. Before we outline our approach in more detail, let us briefly review the literature.

1.1.1 Literature

The existing literature has taken two approaches to experimental investigation of behavior in one-shot normal-form games.

The first approach is to specify types of players and to estimate which types describe subjects' behavior best. Each type is characterized by a fixed rule which she uses for making a decision. The most prominent types include *L0* who plays randomly; *L1* who best-responds to *L0*; *L2* who best-responds to *L1*; *Nash* who plays the Nash equilibrium; and *Worldly* who best-responds to a mixture of *L1* and *Nash* players.

Unfortunately, this approach has so far not lead to a clear picture regarding which types are most common in the population. Stahl and Wilson (1995) (SW hereafter) estimate that the most common type is *Worldly* whereas *L1*, *L2* and *Nash* are relatively rarer.¹ On the other hand, in a comprehensive study which uses both players' decisions as well as their patterns of looking up payoffs, Costa-Gomes, Crawford and Broseta (2001) (CGCB hereafter) estimate that 45% of the population are *L1* and 44% are *L2*. The strong presence of *L1* seems to be confirmed by Costa-Gomes and Weizsäcker (2005) (CGW hereafter) who find that subjects choose *L1*'s preferred action most frequently (60% of the time). However, in another twist, Rey Biel (2005) finds that subjects play the Nash equilibrium most frequently (80% of the time), whereas they choose *L1*'s preferred action much less frequently (50% of the time).

¹SW estimate that *Worldly*, *L1*, *L2* and *Nash* comprise 43%, 21%, 2% and 17% of the population, respectively. The remaining 17% are estimated to be *L0*.

	Risk Neutral	Risk Averse
Naive	<i>NRN</i>	<i>NRA</i>
Strategic	<i>SRN</i>	<i>SRA</i>

Table 1.1: Types

The second approach in the literature is to elicit players’ beliefs regarding the opponent’s play and to investigate average best-response rates (assuming risk-neutrality) to (stated) beliefs. CGW find a best-response rate of only 54% in 3×3 games, while Rey Biel (2005) finds a much higher best-response rate of 73% (again in 3×3 games).

1.1.2 Outline of Approach in Current Study

Given the importance of one-shot normal-form games and given that no clear picture of behavior in these games has emerged so far, we take a different approach to studying behavior in these games.

In our experiment (similar to what has already been done in the literature) we let subjects play ten 3×3 one-shot normal-form games and we also elicit beliefs regarding the opponent’s play.

Our approach differs from the existing literature in how we specify types of players. In particular, we specify four types, each of which is characterized by where she falls along the two general dimensions “naive vs. strategic” and “risk neutral vs. risk averse”. Thus, we have a naive risk neutral type (*NRN*), a strategic risk neutral type (*SRN*), a naive risk averse type (*NRA*) and a strategic risk averse type (*SRA*) (see table 1.1). The two risk neutral (risk averse) types are assumed to have linear (logarithmic) utility. The two naive types are modeled as preferring the action which gives them the highest average utility. The two strategic types are modeled as forming a belief over the opponent’s actions and preferring the action which maximizes expected utility given that belief.

A player of type $t \in T = \{NRN, SRN, NRA, SRA\}$ is modeled as choosing t 's preferred action with a probability which depends on an individual-specific precision parameter λ . We treat λ as a random effect which is generated from a distribution with mean μ and standard deviation σ .

This setup allows us to write down the probability of players' chosen actions, conditional on their beliefs, as a function of the proportion of each type in the population ($\{p_t\}_{t \in T}$) and (μ, σ) . Using subjects' stated beliefs as a proxy for their true beliefs enables us to estimate $\{p_t\}_{t \in T}$ and (μ, σ) in each treatment via maximum likelihood. We also compute Bayesian posteriors over the parameters (starting from a uniform prior).

Our approach has several advantages. First, it focuses on two very general dimensions of behavior while imposing as little additional structure as possible. This is in contrast to the existing literature, in which types act according to narrowly fixed rules. In fact, almost all types in this literature either coincide with or are a special case of one of our types. For example, *L1* coincides with *NRN* and the above mentioned *L2*, *Nash* and *Worldly* are special cases of *SRN*. If people do not behave according to narrowly fixed rules which can readily be included in the specification, then more general types are desirable as they reduce the danger of misspecification.

Second, in contrast to previous studies, we allow for risk aversion.² Despite the fact that payoffs are relatively small, it is quite plausible that many subjects' behavior is better captured by risk aversion than by risk neutrality.³ In fact, risk aversion could explain why *L1* does well in predicting behavior in some studies (CGCB and CGW) and not so well in other studies (SW and Rey Biel (2005)). In

²SW use Roth and Malouf's (1979) binary lottery procedure in which a subject's payoff determines the probability of winning a given monetary prize. Although this procedure should, theoretically, eliminate any effects of risk aversion, there is evidence that it often does not work well in practice. See Camerer (2003), p.41 for a brief discussion as well as for further references.

³For example, modeling subjects as having CRRA utility, Holt and Laury (2002) find that 66% of subjects exhibit risk aversion even when payoffs are between \$0.1 and \$3.85.

CGCB and CGW, $L1$ happens to choose the maximin action⁴ in 15 out of 18 and in 12 out of 14 games, respectively; on the other hand, in SW and Rey Biel (2005) this occurs in only 3 out of 12 and in 4.5⁵ out of 10 games, respectively. Given that risk averse subjects have a tendency to guarantee a certain level of payoff and hence may often choose the maximin action, it could be that in studies in which the maximin action and the $L1$ action often coincide, $L1$ is simply masking the presence of risk averse subjects. Actually, risk aversion could also explain why subjects are playing the Nash equilibrium so frequently in Rey Biel (2005): the games in this study are constant-sum so that the Nash action always coincides with the maximin action.

Third, we avoid a possible bias in favor of $L1$ which exists in previous studies. In particular, let S be the simplex which represents all possible beliefs over the opponent's actions in a game and let $S_{L1} \in S$ represent the beliefs for which the $L1$ action is a best response. Then the ratio $\frac{\text{Area of } S_{L1}}{\text{Area of } S}$ has a tendency to be rather large: it is approx. 0.66, 0.75, 0.63 and 0.71 (averaged over games) in SW, CGCB, CGW and Rey Biel (2005), respectively.⁶ This means that in previous studies $L1$ may be masking the presence of subjects who are best-responding to beliefs which are not consistent with any of the specified types. Our approach avoids this bias by explicitly allowing for types which best-respond to their beliefs (whatever these beliefs may be).

Fourth, although CGW and Rey Biel (2005) investigate average best-response rates to (stated) beliefs (assuming risk-neutrality), their approach has two limitations. First, it does not tell us whether non-best-response decisions are simply due to errors or whether they are due to behavior which deviates in a systematic way

⁴The action that guarantees the highest payoff regardless of what one's opponent does.

⁵Averaged over row players and column players.

⁶These ratios are very high given that each game in SW, CGW and Rey Biel (2005) involves a choice between 3 actions and each game in CGCB involves a choice between 2-4 actions.

from expected payoff maximization. Second, because of the focus on average best-response rates, it does not address subject heterogeneity. We address both of these issues.

Finally, our approach is very parsimonious and we need to estimate only five parameters: three parameters for the proportions of the four types in the population as well as (μ, σ) .⁷

However, our approach also relies on two main assumptions. The first assumption is that types are correctly specified. Given the generality of our types this assumption is weaker than in previous studies. However, it is still nontrivial. We discuss this assumption further in section 1.5.

The second assumption is specific to our approach. In particular, even if types are correctly specified, the likelihood function depends on subjects' true beliefs whereas we use stated beliefs in the estimation. This means that the estimation implicitly relies on stated beliefs being a good proxy for true beliefs. It is this assumption which makes the generality and parsimony of our approach possible. As discussed in section 1.5, we believe that it is a reasonable assumption. However, it too is nontrivial.

In addition, although the generality of our types is an advantage, it is also a limitation in that it does not allow us to address more concrete questions about behavior. For example, even though we can estimate the proportion of strategic types, we cannot say much about the kinds of strategic reasoning they employ. Because of this limitation, as well as because of the second assumption above, we view our approach as complimentary to rather than as a substitute for the approaches taken in the literature.

⁷SW estimate 13 independent parameters (11 when they omit one of their types). CGCB estimate 15 independent parameters in the model which looks only at decisions and 67 independent parameters in the model which also incorporates search patterns.

Before proceeding, let us briefly sketch the design of the experiment as well as the main findings. The main part of the experiment consists of three treatments (A, B & C) in which we use graduate students as subjects.

In treatment A, subjects first play all games and only after that beliefs are elicited. This treatment allows us to estimate what proportion of players are each type when the games are played in a natural way without interference from belief elicitation.

Treatments B and C allow us to investigate how behavior along the “naive vs. strategic” and the “risk neutral vs. risk averse” dimension is affected by two manipulations of treatment A.⁸ In treatment B, beliefs are elicited at the same time that each game is played, i.e. players are exogenously prompted to form beliefs while playing the games. In treatment C, we eliminate the belief formation process altogether by having subjects choose between lottery tickets instead of between actions in a game.

If our estimates of naive behavior diminish in B and C relative to A, this would suggest that these estimates are indeed driven by a failure of some subjects to take into account what their opponent might do.

What might be more puzzling is why we should be interested in how behavior along the “risk neutral vs. risk averse” dimension varies across treatments. After all, there is no reason for subjects’ utility function for money to change its shape across treatments. However, we do not take risk aversion too literally. We merely view it as a formal way to capture cautious behavior. Such cautious behavior may actually be driven by something different from curvature in the utility function. For example, if seemingly risk averse behavior is in fact driven by ambiguity aversion,⁹ we can very well expect variation in the estimated proportions of risk neutral subjects given that the ambiguity of decision tasks may differ across treatments.

⁸Keeping the games and the subject population (graduate students) fixed.

⁹A situation is ambiguous if the decision maker is not confident in her belief. As explained in section 1.5, ambiguity aversion and risk aversion have similar implications for behavior.

Regarding our findings in A, we estimate that only a small minority of subjects (12%) is naive and only a minority (39%) is risk neutral. However, these estimates do not seem robust across games with similar formal structure or across subject populations. In particular, in a pilot session using CGW's games¹⁰ and in a follow-up session with our games, but with undergraduate subjects, we obtain quite different estimates. In a sense, this is a negative result because it suggests that it may be difficult to draw general conclusions about behavior in one-shot normal-form games. Perhaps this is the reason why no clear picture has emerged from the existing literature. On a more positive note, variations in behavior across games and subject populations present us with the new challenge of explaining these variations.

Perhaps the more generalizable conclusions come from looking at changes in our estimates across treatments A, B and C. In this regard we find that, as expected, the estimate of the proportion of naive types falls from 12% to 4% and then to 3% in A, B and C, respectively. The estimate of the proportion of risk neutral types increases from A to B almost twofold (from 39% to 74%) and then decreases again in C (to 42%). The increase from A to B is consistent with ambiguity aversion - it is plausible that the decision tasks are perceived as less ambiguous in B than in A because in B subjects are provided with a way to think about the games. The decrease from B to C is consistent with ambiguity aversion only if lottery tickets are perceived as ambiguous. We discuss this pattern at length in section 1.5. At any rate, ambiguity aversion or no ambiguity aversion, the variation in the estimated proportions of risk neutral subjects across treatments suggests that there is something more going on than mere curvature in the utility function.

We proceed as follows: section 1.2 explains the experimental design; section 1.3 presents the formal model; section 1.4 presents the results; section 1.5 discusses some relevant issues and concludes.

¹⁰I thank Costa-Gomes and Weizsäcker for letting me use their games.

1.2 Experimental Design

The main part of the experiment consists of three treatments - A, B and C.¹¹ We conducted 2 sessions of A (24 and 29 participants, respectively), two sessions of B (19 and 24 participants, respectively) and three sessions of C (13, 12 and 13 participants, respectively). Subjects in C were participants from A and B, who accepted the invitation to attend one more session. Subjects were Ohio State University PhD students from a wide range of programs who had never taken Economics courses.

All subjects were paid a \$5 show-up fee in A and B and \$7 in C. In addition, subjects could earn Experimental Currency Units (ECU) which were converted into dollars at the rate 0.1\$ per ECU. Average earnings (including the show-up fee) were \$20.68, \$20.83 and \$11.85 in A, B and C, respectively.

We also conducted 1 pilot session for A (28 participants) and 1 pilot session for B (26 participants). Both pilots were different from the main sessions in that we used CGW's games and there were also slight differences in the design and the instructions.¹²

In addition, we conducted one follow-up session for A (25 participants) in which we used undergraduate students in order to check if the results from A are robust to the subject population.

The experiment was programmed and conducted with the software z-Tree (Fischbacher (1999)). The sessions were held in the Experimental Economics Lab at The Ohio State University.

¹¹The instructions for treatments A and C can be found in the appendix. The instructions for treatment B are similar to those for treatment A.

¹²In the pilot for B, we also used undergraduate students instead of PhD students.

1.2.1 Treatment A

Treatment A consists of three parts. The instructions for each part were handed out and read out immediately before that part. There was no feedback whatsoever until the end of the experiment. Subjects were divided into two groups - row players and column players.

In part I, subjects simply played ten 3×3 one-shot normal-form games. Each subject i 's earnings were determined according to her action and the action of a random player $-i$ from the other group in one randomly determined game.

In part II, each subject i was asked (for each game) to state what, in her opinion, was the probability that $-i$ chose each action in part I. For part II, each subject was paid a lump sum of \$6. We discuss this payment scheme in section 1.5.

Part III was included since (as mentioned in the introduction) we suspect that ambiguity aversion, rather than risk aversion, may be at work. Given that our formal framework does not incorporate ambiguity aversion (our types are within the expected utility framework), we will not include the results from part III in the main analysis. However, we will discuss them in section 1.5 when we consider ambiguity aversion as a possible driving force behind seemingly risk averse behavior.

Before we describe part III, let us explain what we mean by a lottery ticket which, for subject i , corresponds to an action in one of the games.¹³ Such a lottery ticket has the same possible payoffs as the action and the probabilities of the payoffs are matched to the belief i stated over $-i$'s play. For example, let us say that the action pays 30, 40 or 50 ECU if $-i$ chooses action 1, 2 or 3, respectively and that i 's stated belief places probabilities of 0.4, 0.5 and 0.1 on $-i$ choosing 1, 2 and 3, respectively.

¹³This kind of lottery ticket will also play a part in treatment C.

Then the lottery ticket pays 30, 40 and 50 ECU with probabilities 0.4, 0.5 and 0.1, respectively. The lottery ticket's final payoff is determined by the computer which randomizes using the respective probabilities.

In part III, we let each subject i choose, for each game, between being paid according to (i) the combination of one of her actions (which is exogenously fixed) and the action $-i$ chose in part I (so that i 's fixed action represents a bet on $-i$'s decision) or (ii) a lottery ticket corresponding to i 's fixed action. Note that an ambiguity averse subject would prefer (i) if she perceives the lottery tickets as more ambiguous than the games and would prefer (ii) if she perceives the games as more ambiguous than the lottery tickets.

In order to determine earnings for part III, one of the ten decisions in that part was taken at random and each subject was paid according to (i) or (ii), depending on which one she chose.

1.2.2 Treatment B

Treatment B was analogous to treatment A with the only difference that subjects stated beliefs and chose actions at the same time so that parts I and II from A were collapsed into one part.¹⁴

1.2.3 Treatment C

In treatment C, each subject i chose not between actions in a game but between lottery tickets. In particular, i chose one lottery ticket from each of 10 triplets of lottery tickets. Each lottery ticket in a triplet corresponded to an action in one of the games which i played in A or B, i.e. the lottery ticket had the same payoffs as the action and the probabilities of the payoffs were matched to i 's stated belief. The

¹⁴Subjects also received a lump sum payment of \$4 rather than \$6 for stating their beliefs since stating beliefs while playing each game is supposedly less additional effort than stating beliefs after all games have been played.

three lottery tickets in a triplet corresponded to different actions in the same game. In this way the decision situations in C were formally identical (within the expected utility framework) to the games played in A and B.

In order to determine earnings in C, one of the ten triplets was taken at random and each subject was paid according to the lottery ticket which she chose from that triplet.

1.2.4 Games

We designed the ten games with the aim of distinguishing between our four types as much as possible.¹⁵ This task was complicated by the fact that the preferred actions of the two strategic types depend on their beliefs. Therefore, we tried to control for the beliefs subjects were likely to form. In order to do this, we exploited a finding from CGW about subjects' stated beliefs: subjects in that study placed, on average, the largest probability on their opponent choosing *L1*'s (or in our terminology *NRN*'s) preferred action. Therefore, our games are designed to give the best identification of types if (i) subjects place a large probability on actions of the opponent which give a high average payoff and (ii) subjects place roughly equal probabilities on actions of the opponent which give roughly equal average payoffs.

Since it is difficult to ensure optimal identification of types for both row and column players and at the same time control for row and column players' beliefs, we designed our games so that row players' actions separate between types while column players' actions are used solely to control for row players' beliefs. Since column players' actions are not designed to separate between types we let a small minority of subjects (2 subjects per session) be column players and we disregard their behavior in the analysis.¹⁶

¹⁵CGW's games, which we used in the pilots, were not designed to separate between our types and did not do so well (especially along the "risk neutral vs. risk averse" dimension).

¹⁶The fact that there were different numbers of row and column players means that subjects were

The ten games are presented in figure A.1. The fourth column to the right of each game indicates the preferred action of each type (for strategic types this is only tentative since it is assuming subjects form beliefs satisfying (i) and (ii) above). As can be seen from the figure, each type’s preferred action will, ideally, differ from that of each other type in at least five games.

For comparability with the literature, all games have a unique pure Nash equilibrium and three of them (1, 2 and 10) are dominance-solvable.

1.3 Types and Formal Statistical Model

In this section, we formally specify the four types of players as well as the statistical framework within which we will estimate the proportion of each type in the population.¹⁷ The four types are: a naive risk neutral type (*NRN*), a strategic risk neutral type (*SRN*), a naive risk averse type (*NRA*) and a strategic risk averse type (*SRA*).

Formally, each type is characterized by the way she evaluates each action in a game.¹⁸ Let $a = (a(1), a(2), a(3))$ be an action which pays $a(1)$, $a(2)$ or $a(3)$ if one’s opponent chooses action 1, 2 or 3, respectively. Let $\bar{b} = (\bar{b}(1), \bar{b}(2), \bar{b}(3))$ be a subject’s belief over her opponent’s actions.¹⁹ Let the respective utility functions for the risk neutral and risk averse types be $u_{RN}(x) = x$ and $u_{RA}(x) = \frac{89}{\ln(9.9)}\ln(x) + 10 - \frac{89}{\ln(9.9)}\ln(10)$.²⁰

not paired, i.e. the fact that $-i$ ’s action was used to determine i ’s payoff does not imply that i ’s action was used to determine $-i$ ’s payoff. This is irrelevant from the point of view of each subject’s own payoff which is determined (just like when subjects are paired) according to the combination of that subject’s action and the action of some other random subject.

¹⁷This statistical framework is similar to that in many experimental papers, including Camerer and Harless (1994), SW, CGCB, CGW. The main difference in our paper is that we will have individual-specific precision parameters which are treated as random effects.

¹⁸Of course, in C subjects are evaluating not actions in a game (which pay differently depending on what the opponent does) but lottery tickets (which pay differently depending on the computer’s randomization). We will generically talk about actions in a game with the implicit understanding that in the case of C we actually mean lottery tickets.

¹⁹In the case of C, \bar{b} represents the exogenously given probabilities with which the computer randomizes.

²⁰The constants in $u_{RA}(\cdot)$ ensure that risk averse and risk neutral types’ utility functions are

Finally let $T = \{NRN, SRN, NRA, SRA\}$ be the set of types and $V_t(a; \bar{b})$ be the value that type $t \in T$ with belief \bar{b} attaches to a . Then $V_t(a; \bar{b})$ for each type is specified as follows:

$$\begin{aligned} V_{NRN}(a; \bar{b}) &= \frac{1}{3}u_{RN}(a(1)) + \frac{1}{3}u_{RN}(a(2)) + \frac{1}{3}u_{RN}(a(3)) \\ V_{SRN}(a; \bar{b}) &= \bar{b}(1)u_{RN}(a(1)) + \bar{b}(2)u_{RN}(a(2)) + \bar{b}(3)u_{RN}(a(3)) \\ V_{NRA}(a; \bar{b}) &= \frac{1}{3}u_{RA}(a(1)) + \frac{1}{3}u_{RA}(a(2)) + \frac{1}{3}u_{RA}(a(3)) \\ V_{SRA}(a; \bar{b}) &= \bar{b}(1)u_{RA}(a(1)) + \bar{b}(2)u_{RA}(a(2)) + \bar{b}(3)u_{RA}(a(3)) \end{aligned}$$

Thus, the two naive types evaluate each action according to the average utility of its payoffs. The two strategic types form a belief over the opponent's actions and evaluate each of their own actions according to its expected utility given that belief.

We interpret the naive types as focusing on their own payoffs and ignoring what the opponent might do. Of course, one could alternatively interpret them as thinking about what the opponent might do and always coming up with a uniform belief, but this hardly seems plausible. Moreover, such an interpretation is at odds with subjects' stated beliefs which are rarely uniform. Note that with our interpretation there is nothing to stop naive types from forming and stating non-uniform beliefs when they are explicitly asked to state beliefs. In fact, if naive types always stated uniform beliefs they would be indistinguishable from strategic types. The whole idea

on a similar absolute scale ($u_{RN}(10) = u_{RA}(10)$ and $u_{RN}(99) = u_{RA}(99)$ where 10 and 99 are the minimum and maximum possible ECU earnings in a game). This is in anticipation of the fact that, within the multinomial logit model which we will introduce shortly, a player's precision parameter is related to the scale of her utility function. Ensuring that all types have utility functions on a similar scale will allow us to reject the hypothesis that precision parameters are generated from distributions with *type-specific* means and standard deviations and hence will allow us to reduce the number of parameters.

behind our ability to distinguish between naive and strategic types relies precisely on all types' ability to form and state non-uniform beliefs when beliefs are explicitly elicited. Then subjects who are picking actions with high average utility rather than best-responding to stated beliefs will be evidence in favor of naive types. Analogously, subjects who are best-responding to stated beliefs rather than picking actions with high average utility will be evidence in favor of strategic types.

The two risk averse types are modeled as having logarithmic utility. Logarithmic utility leads to behavior which is both sufficiently different from risk neutrality so as to allow us to distinguish between risk neutral and risk averse types and at the same time is not so extreme as to seem implausible.

Now that we have specified our types, we need to specify how they make choices. If we simply say that each $t \in T$ chooses the action with highest $V_t(a; \bar{b})$, many subjects' behavior will not fit any type. Therefore, we model subjects' choices within the logit multinomial model, i.e. the probability that a type t player with belief \bar{b} chooses action a is²¹:

$$Pr(a|t, \bar{b}; \lambda) = \frac{e^{\lambda V_t(a; \bar{b})}}{\sum_{a'} e^{\lambda V_t(a'; \bar{b})}}$$

where the summation in the denominator is over all actions a' in the game.

$Pr(a|t, \bar{b}; \lambda)$ depends on λ which plays the role of a precision parameter. If $\lambda = 0$, then for any action a , $Pr(a|t, \bar{b}; \lambda) = \frac{1}{3}$. As $\lambda \rightarrow \infty$, the probability of the action a with highest $V_t(a; \bar{b})$ being chosen goes to 1.

We assume that for a subject of type t , λ is an individual-specific random effect which is generated (independently across subjects) from a gamma distribution with type-specific mean $\mu_t > 0$, standard deviation $\sigma_t \geq 0$, and cumulative den-

²¹We will generically denote by $Pr(\cdot|\cdot; \cdot)$ the probability of the first term conditional on the second term given the parameter(s) after the semi-colon.

sity (reparameterized in terms of its mean and standard deviation) $G(\lambda; \mu_t, \sigma_t)$. We chose the gamma distribution because it has non-negative support, because it can be characterized in terms of its mean and standard deviation, and because it has a thin (exponentially decreasing) tail. The latter property is desirable since it prevents implausibly high values of λ from driving up the mean.

Let us introduce some additional notation which we will need to write down the likelihood function. Let $a_j^g = (a_j^g(1), a_j^g(2), a_j^g(3))$ be action j in game g . Let $x_i^g \in \{a_1^g, a_2^g, a_3^g\}$ be subject i 's chosen action in game g ; let $x_i = (x_i^1, \dots, x_i^{10})$ be i 's choices in all games; let $x = (x_1, \dots, x_N)$ be all N subjects' choices. Analogously, let $b_i^g = (b_i^g(1), b_i^g(2), b_i^g(3))$ be subject i 's belief for game g ; let $b_i = (b_i^1, \dots, b_i^{10})$ be i 's beliefs in all games; let $b = (b_1, \dots, b_N)$ be all N subjects' beliefs. Let p_t be the proportion of type t in the population and θ be the vector of parameters $\{p_t, \mu_t, \sigma_t\}_{t \in T}$. We will also denote by $G(\lambda|b_i; \mu_t, \sigma_t)$ the cumulative density of λ conditional on b_i given μ_t and σ_t .²² Given all this we can write the probability of i 's choices, conditional on i 's beliefs, by summing (integrating) over t and λ :

$$Pr(x_i|b_i; \theta) = \sum_{t \in T} Pr(t|b_i; \theta) \int_0^\infty Pr(x_i|t, b_i, \lambda) dG(\lambda|b_i; \mu_t, \sigma_t) \quad (1.1)$$

In order to obtain an explicit form for $Pr(t|b_i; \theta)$ and $G(\lambda|b_i; \mu_t, \sigma_t)$, we assume that t and λ are jointly independent of b_i .²³ In order to check if this is a realistic assumption, we perform the following test of independence. First, we categorize each subject according to the type whose preferred action she chose most often²⁴ and according to whether she chose this type's preferred action in at least eight games (high precision) or in strictly less than eight games (low precision). Thus, each

²²Note that $G(\lambda|b_i; \mu_t, \sigma_t)$ is not necessarily a gamma cumulative density.

²³For now it would have been enough to make the weaker assumption that t is independent of b_i and that λ is independent of b_i . However, we assume joint independence as this will play a role when we write down the Bayesian posterior over θ .

²⁴In case a subject chose the preferred action of more than one type the same number of times, she is randomly assigned to one of these types.

subject is assigned to one of eight categories, each category being a combination of a type and a precision. Next, we assign each subject's belief in each game to one of four categories as follows: a belief falls in category $j = \{1, 2, 3\}$ if it assigns strictly more than 0.5 weight to action j of the opponent; otherwise it falls in category 4.²⁵ Given this, we can test for each game the hypothesis that a subject's type-precision category is independent of the category her belief falls into. Performing 30 Fisher's exact tests (10 tests each for A, B and C), we can reject the null hypothesis of independence at the 5% level in 3 cases (game 6 in A, game 2 in B and the lottery ticket triplet corresponding to game 10 in C). For 30 comparisons this seems well within the limits of chance.²⁶

Although this test of independence is admittedly crude, it does make it unlikely that b_i provides much information about t and λ . Thus, replacing $Pr(t|b_i; \theta)$ by $Pr(t|\theta) = p_t$ and $G(\lambda|b_i; \mu_t, \sigma_t)$ by $G(\lambda; \mu_t, \sigma_t)$ seems reasonable. Using this to rewrite (1.1) and taking the product over subjects, we obtain the probability of all subjects' choices conditional on their beliefs:

$$Pr(x|b; \theta) = \prod_{i=1}^N Pr(x_i|b_i; \theta) = \prod_{i=1}^N \sum_{t \in T} p_t \int_0^\infty Pr(x_i|t, b_i, \lambda) dG(\lambda; \mu_t, \sigma_t) \quad (1.2)$$

Plugging in subjects' stated beliefs for b will allow us to estimate θ by maximizing the above conditional (on b) maximum likelihood function.²⁷ Given that for A, B

²⁵This way of discretizing beliefs in order to test hypotheses is used in CGW and Rey-Biel (2005).

²⁶Of course, only the distribution of the test statistic for each separate comparison is known, but not the joint distribution of the test statistic in all 30 comparisons. Therefore the probability of getting three or more rejections at the 5% significance level is unknown. If the test statistic is independent across the 30 comparisons, then this probability is 0.188.

²⁷This function is continuous in θ for every x and b . If beliefs are not always uniform, it has a strict maximum at the true θ . Assuming that $\forall t, \epsilon \leq \mu_t \leq \bar{\mu}$ and $\sigma_t \leq \bar{\sigma}$ (where ϵ is some small number and $\bar{\mu}$ and $\bar{\sigma}$ are some large numbers), the parameter space is compact. All other technical requirements (as given in theorems 13.1 and theorem 13.2 in Wooldridge (2001)) hold so that the ML estimator is consistent and asymptotically normal (for asymptotic normality the true θ also needs to be interior).

and C we cannot reject the hypothesis that (μ_t, σ_t) are the same for all types²⁸, we assume that $(\mu_t, \sigma_t) = (\mu, \sigma)$ for all t , so that the parameter vector to be estimated is $\theta = (p_{NRN}, p_{SRN}, p_{NRA}, p_{SRA}, \mu, \sigma)$ which has five independent parameters (since $\sum_{t \in T} p_t = 1$).

Based on (1.2) we can also compute, starting from a uniform prior $f(\cdot)$ ²⁹, the Bayesian posterior over the parameters:

$$f(\theta|x, b) = \frac{Pr(x|b; \theta)f(\theta)}{\int Pr(x|b; \psi)f(\psi)d\psi} = \frac{Pr(x|b; \theta)}{\int Pr(x|b; \psi)d\psi} \quad (1.3)$$

The second equality follows since $f(\cdot)$ is uniform.³⁰

1.4 Results

1.4.1 Aggregate-Level Analysis

Figure A.1 shows³¹, for each game, row subjects' mean beliefs in A and B as well as aggregate actions in all three treatments. Mean beliefs are very similar between A and B. Aggregate actions are however quite different between A (C) and B. Comparison of aggregate actions between A and B for each game via Fisher's exact tests leads to significant differences (at the 5% level) in three of the ten games (games 5, 7 and 8). Aggregate actions in B and C are significantly different in two games (games 5 and 7). Aggregate actions in A and C are not significantly different in any game.

²⁸The likelihood-ratio test yields p-values of 0.64, 0.95 and 0.18 for A, B and C.

²⁹We set $\epsilon = 0.01$, $\bar{\mu} = \bar{\sigma} = 0.6$ (see footnote 27) in defining the support of $f(\cdot)$.

³⁰Technically, in order to write $f(\cdot)$ in (1.3) without conditioning on b , we need to assume that b and θ are independent. This is hardly a strong assumption given that $\forall i$, b_i and θ are independent. The latter holds since θ is merely used to generate the individual types and precisions which, by assumption, are independent of individuals' beliefs.

³¹Apart from excluding column players (as explained earlier), in all of the subsequent analysis we also exclude one subject from A (who said after the experiment that she was familiar with game theory), one subject from B (who said that she had mistaken column players' payoffs for row players' payoffs and vice versa) and one subject from C (who said that he thought that "the lottery ticket pays" meant that he had to pay so that he chose tickets with low payoffs). Excluding these subjects has a negligible effect on the analysis.

It is worth noting that the games in which behavior is significantly different between A (C) and B are all meant to distinguish between risk neutral and risk averse types. In particular, subjects in A (C) are choosing more the actions which we (tentatively) expect to be the preferred actions of the risk averse types, whereas subjects in B are doing the opposite.³²

Table 1.2 shows how many percent of subjects' actions in each treatment coincide with the preferred actions of each type. In A and C, *SRA* predicts the largest percentage of actions, followed by *SRN*. In B, *SRN* predicts the largest percentage of actions, followed by *SRA*.

Type	Treatment A	Treatment B	Treatment C
<i>NRN</i>	33%	42%	33%
<i>SRN</i>	55%	72%	68%
<i>NRA</i>	34%	21%	27%
<i>SRA</i>	59%	60%	70%

Table 1.2: Percent of Actions Consistent with each Type.

Type	Treat. A	Treat. B	Treat. C
Naive but not Strategic	17%	9%	8%
Strategic but not naive	34%	37%	43%

Table 1.3: Percent of Actions Consistent with Naive Types, but not with Strategic Types and vice versa.

Type	Treat. A	Treat. B	Treat. C
Risk neutral but not risk averse	15%	25%	15%
Risk averse but not risk neutral	26%	16%	24%

Table 1.4: Percent of Actions Consistent with Risk Neutral Types, but not with Risk Averse Types and vice versa.

³²CGW and Rey Biel (2005) do not find significant differences between aggregate actions in treatments in which beliefs are elicited before the games and treatments in which beliefs are elicited after all games have been played. Perhaps this is the case because their games are not designed to distinguish between risk neutral and risk averse behavior.

Table 1.2 has the drawback that actions which are consistent with a particular type may simply have been chosen because they are also consistent with other types. By concentrating on “naive vs. strategic” and “risk neutral vs. risk averse” behavior at a time we can easily eliminate such overlap.

Table 1.3 shows the percentage of actions which are consistent with naive, but not with strategic types and vice versa. The table shows that:

Result 1.4.1.1 *Subjects in all three treatments chose considerably more actions which are unequivocally strategic rather than unequivocally naive. The difference is larger in B than in A and is largest in C.*

Table 1.4 shows the percentage of actions which are consistent with risk neutral, but not with risk averse types and vice versa. The table shows that:

Result 1.4.1.2 *Subjects in A and C chose more actions which are unequivocally risk averse than unequivocally risk neutral (the difference is slightly larger in A). The opposite is true for subjects in B.*

Of course, the above analysis is not particularly illuminating regarding subject heterogeneity. It also does not take into account whether actions which are not consistent with a given type are costly mistakes for that type.

1.4.2 Estimation of Formal Statistical Model

In this section, we present the results based on the framework from section 1.3. We perform the analysis separately for A, B and C, pooling the data from the sessions within a treatment.³³

Table 1.5 presents the main results. The first column corresponding to each treatment shows the ML estimates of θ , the log-likelihood as well as the estimate of

³³A likelihood ratio test of the hypothesis that the true θ is the same in all sessions within a treatment yields p-values of 0.101, 0.767 and 0.413 for A, B and C, respectively.

	Treatment A		Treatment B		Treatment C	
	MLE	Posterior	MLE	Posterior	MLE	Posterior
p_{NRN}	0.009 (0.031)	0.043 [0.01,0.11]	0.039 (0.038)	0.068 [0.01,0.15]	0 N/A	0.023 [0.01,0.07]
p_{SRN}	0.383 (0.094)	0.368 [0.23,0.51]	0.703 (0.10)	0.648 [0.49,0.79]	0.416 (0.101)	0.403 [0.25,0.55]
p_{NRA}	0.109 (0.061)	0.114 [0.03,0.21]	0 N/A	0.03 [0.01,0.09]	0.026 (0.029)	0.052 [0.01,0.13]
p_{SRA}	0.50 (0.099)	0.485 [0.33,0.63]	0.259 (0.095)	0.254 [0.11,0.41]	0.558 (0.102)	0.528 [0.37,0.67]
μ	0.108 (0.014)	0.115 [0.085,0.145]	0.148 (0.022)	0.158 [0.125,0.205]	0.251 (0.054)	0.298 [0.195,0.445]
σ	0.05 (0.02)	0.057 [0.025,0.105]	0.073 (0.026)	0.089 [0.045,0.145]	0.143 (0.066)	0.213 [0.075,0.415]
$\log\text{-lik.}$	-383.6033	N/A	-251.6819	N/A	-201.5477	N/A
$p_{NRN} + p_{NRA}$	0.118 (0.055)	0.153 [0.07,0.25]	0.039 (0.038)	0.097 [0.03,0.19]	0.026 (0.029)	0.069 [0.01,0.15]
$p_{NRN} + p_{SRN}$	0.391 (0.096)	0.401 [0.25,0.55]	0.742 (0.095)	0.711 [0.55,0.85]	0.416 (0.101)	0.428 [0.29,0.59]

Table 1.5: Formal Model Estimation.

the proportion of naive types (simply the sum of the relevant previous lines) and the estimate of the proportion of risk neutral types (again, simply the sum of the relevant previous lines). Below each ML estimate we show the estimated standard error.³⁴ The second column corresponding to each treatment shows summary information (means and 90% confidence intervals) about marginal posteriors over elements (or combinations of elements) of θ .

³⁴In B (C) the estimate of p_{NRA} (p_{NRN}) is on the boundary of the parameter space. We do not compute the standard error for this estimate since the standard error does not have the usual interpretation in terms of confidence intervals. The standard errors for the elements of θ which are not on the boundary are computed by estimating a restricted model in which p_{NRA} (p_{NRN}) is set equal to 0. Note that in all treatments estimated standard errors should be treated with caution given that at least one element of θ is less than 1.96 estimated standard errors from the boundary of the parameter space.

ML Estimates

First, let us discuss the ML estimates given in table 1.5. The first thing to notice is that the estimate of the proportion of naive types in A is rather small (0.118). This is lower than the estimate in SW (0.21) and is much lower than the estimate in CGCB (0.45). The other noticeable fact is that only a minority of the population in A is estimated to be risk neutral (0.391). This estimate increases almost twofold in B (to 0.742) and then in C, drops almost all the way back down to the level in A (to 0.416).

Several features of the estimates in table 1.5 make good sense and are encouraging news about the appropriateness of our specification.

The estimate of the proportion of naive types drops from 0.118 in A to 0.039 in B and then to 0.026 in C. This is precisely as expected. Explicitly telling subjects to think about what their opponent will do is likely to reduce naive behavior in B relative to A and giving subjects clear probabilities in C is likely to reduce naive behavior even further.

The estimate of μ increases from A to B. This is to be expected as subjects whose attention is focused on forming beliefs are likely to best-respond with less noise. The estimate of μ increases even further in C. This is again as expected given that choosing between lottery tickets is clearly simpler than choosing between actions in a game.

The absolute values of the estimates of μ are important. If subjects' mean precision is low, then types' behavior is erratic and this undermines the whole idea of types who behave in a systematic way. If μ is reasonably high this is encouraging news for the adequacy of our specification of types.

In order to interpret the absolute values of the estimates of μ in each treatment, we consider an example of how a risk neutral type with precision parameter equal to

the estimate of μ would make choices in each treatment. In particular, let's say that she has to choose between three actions, each of which has a certainty equivalent of $\$y$, $\$(y + 1)$ and $\$(y + 2)$, respectively.³⁵ Table 1.6 shows the probability with which each action will be chosen. As can be seen, the probabilities in the table are reasonably high (certainly much better than random play).³⁶

Certainty Equivalent:	Treatment A	Treatment B	Treatment C
$\$y$	0.08	0.04	0.01
$\$(y + 1)$	0.23	0.18	0.07
$\$(y + 2)$	0.69	0.78	0.92

Table 1.6: Example: Precision.

The ML estimates of σ seem to suggest that there is non-negligible heterogeneity in subjects' individual precision parameters.³⁷

Based on the above let us summarize:

Result 1.4.2.0.1 *Only a small proportion of the population in A is estimated to be naive. This estimate drops further in B and C.*

Result 1.4.2.0.2 *Only a minority of the population in A and C is estimated to be risk neutral. The estimate in B is much higher.*

Result 1.4.2.0.3 *The estimates of μ in each treatment are reasonably high, suggesting that subjects' behavior has a strong systematic component which is captured by our types.*

³⁵i.e. each of the three actions is valued by the decision-maker as if it paid the constant amount $\$y$, $\$(y + 1)$ or $\$(y + 2)$, respectively.

³⁶The corresponding table for a risk averse type will be different (since $u_{RN}(\cdot)$ and $u_{RA}(\cdot)$ do not coincide between 10 and 99 ECU so that the same value of the precision parameter has different implications for choices; see footnote 20) and will also depend on the value of y . For example if y equals $\$2$, $\$3$, $\$4$, or $\$5$, then a risk averse type with precision parameter equal to the estimate in A will choose the $\$(y + 2)$ action with probability 0.74, 0.66, 0.61, or 0.57, respectively.

³⁷Heterogeneity in individual λ 's may in part be driven by the fact that $u_{RN}(\cdot)$ and $u_{RA}(\cdot)$ do not coincide between 10 and 99 ECU. See footnotes 20 and 36.

Results 1.4.2.0.1 and 1.4.2.0.2 (which are based on our formal statistical framework) are in accord with results 1.4.1.1 and 1.4.1.2 (which were based on crude aggregate-level data).

Posteriors and Hypothesis Tests

In this section, we would first like to draw conclusions (beyond point estimates) about the true θ in each treatment. Second, we would like to check if the differences across treatments suggested by the ML estimates are significant.

Regarding the first issue, the first thing that comes to mind is to test hypotheses about whether the proportions of different types, μ or σ are statistically different from zero.³⁸ Unfortunately, it is difficult to test hypotheses on the boundary of the parameter space. Instead, we look at the marginal posteriors over elements (or combinations of elements) of θ .

The marginal posteriors over $p_{NRN} + p_{NRA}$ (the proportion of naive types) and $p_{NRN} + p_{SRN}$ (the proportion of risk neutral types) in each treatment are depicted in figure A.2. These posteriors are relatively tight (especially relative to the uniform prior we start with) which suggests that the design has accomplished reasonable identification of types.

In A, the marginal posterior over $p_{NRN} + p_{NRA}$ places most of the weight well away from zero so that naive types probably do exist. This effect is less pronounced in B and C.

The marginal posteriors over $p_{NRN} + p_{SRN}$ in all treatments suggest that it is very likely that both risk neutral and risk averse types are present. There is also a pronounced shift to the right in the marginal posterior in B which is in accord with the higher ML estimate of $p_{NRN} + p_{SRN}$ in B.

³⁸To be precise, ϵ , rather than 0, in the case of μ . See footnote 27.

The marginal posteriors over μ and σ in each treatment are depicted in figure A.3. As can be seen they are quite tight (except in C) and place most of the weight well away from zero. This suggests that it is both very likely that players exhibit systematic behavior which is captured by our types ($\mu \gg 0$) and that there is heterogeneity in terms of individual λ 's ($\sigma \gg 0$).³⁹

Now we turn to the issue of whether behavior across treatments is statistically different. In particular, using a likelihood-ratio test, we test the hypothesis whether the proportion of naive types is the same in different pairs of treatments as well as the hypothesis whether the proportion of risk neutral types is the same in different pairs of treatments. Table 2.11 shows the p-values for each hypothesis test for each pair of treatments.⁴⁰

	A and B	A and C	B and C
$H_0: p_{NRN} + p_{NRA}$ is same	0.072	0.139	0.791
$H_0: p_{NRN} + p_{SRN}$ is same	0.013	0.86	0.024

Table 1.7: Hypotheses Tests between Treatments.

The p-values for the comparison of $p_{NRN} + p_{NRA}$ between A and B as well as between A and C are rather low, but not significant at the 5% level. This is probably the case since naive types are already rare in A so that any reductions in B and C fail to be significant. Perhaps more importantly we can state:

Result 1.4.2.0.4 *The proportions of risk neutral types are significantly different between A and B as well as between B and C.*

³⁹Although, see footnote 36.

⁴⁰These p-values are computed assuming that subjects in each treatment are drawn randomly from the population: the computations ignore the fact that each subject in C also participated in A or B. Given that a subject who is naive/risk neutral in A or B is more likely to be naive/risk neutral in C, the p-values in the last two columns are probably larger than the true p-values.

1.5 Discussion and Concluding Remarks

1.5.1 Type specification

As noted in the introduction, one of the main assumptions in our approach is that types are correctly specified, i.e. that each subject acts according to one of our types. Although our types are quite general, there are behaviors which do not fit into any of them.

An obvious behavior not captured by our types is risk loving behavior. This is not a serious omission in our experiment. If some subjects are risk loving, they would tend to choose the action with the highest possible payoff which in all games coincides with NRN's preferred action and we may therefore mistakenly interpret them as being naive. However, this cannot be occurring too much given that the estimates of naive types are very low anyway. Moreover, if these estimates were driven by risk loving behavior there would be no reason to see them drop in treatments B and C.

Another issue is that specifying naive behavior as picking the action with the highest average utility may seem ad hoc. If a player is self-interested and satisfies the minimum requirement on rationality that she evaluates each action based on its possible consequences and then makes her choice based on this evaluation, we can think of only two other plausible kinds of naive play. The first kind is that a naive player may simply pick the action with the highest possible payoff.⁴¹ However, in all our games, this action happens to coincide with NRN's preferred action so that a player who picks the action with the highest possible payoff will correctly be interpreted as naive.

⁴¹CGCB call such a player *optimistic*.

The second kind is that a naive player may simply pick the maximin action.⁴² However, in all our games, the maximin action happens to coincide with NRA’s preferred action so that a player who picks the maximin action will correctly be interpreted as naive.

However, the assumption that players are self-interested and that they satisfy the mentioned minimum requirement on rationality is not trivial. In fact in an informal questionnaire conducted at the end of the sessions, we find that a large proportion of subjects express considerations which seem to violate this assumption. These subjects seem either to be altruistic or to be trying to “choose” a box in the payoff table which looks good for both players, perhaps guided by a home-grown rule of thumb that one should compromise with others in order to be well-off oneself. Trying to “choose” a box is clearly irrational since (regardless of whether one is self-interested or not) one should evaluate and compare actions and not boxes.⁴³

Fortunately, excluding the subjects for which we have a strong suspicion based on the questionnaire that they were either being altruistic or that they were irrationally trying to “choose” a box in the payoff table (instead of an action), does not change any of our main results.⁴⁴

1.5.2 Stated Beliefs vs. True Beliefs

As mentioned in the introduction, the validity of our estimation relies on stated beliefs being a good proxy for true beliefs.⁴⁵ There are three key issues here.

The first is whether subjects are aware of their true beliefs. In the context of tasks with which people have a lot of experience, it is quite plausible that they

⁴²CGCB call such a player *pessimistic*.

⁴³Subjects’ responses to the main question in the questionnaire can be found in the appendix in tables A.13 and A.14.

⁴⁴The excluded subjects were approximately one-third of all subjects. Their ID’s are placed in brackets in tables A.13 and A.14.

⁴⁵Here, we are maintaining the assumption that subjects have “true” beliefs, i.e. beliefs according to which they act.

are not aware of their true beliefs since decision processes are to a large degree automated. However, in our context subjects are confronted with a new unfamiliar task. When they are asked to state their beliefs they have to deliberately think about their opponent's play. In this context, it is more difficult to imagine that subjects have beliefs which are in any sense more true than the beliefs they are deliberately forming.

The second issue is whether subjects are not deliberately misrepresenting beliefs. There are several studies which suggest that subjects state beliefs truthfully under a variety of payment schemes (including a lump sum payment). In particular, Friedman and Massaro (1998) find no significant differences in performance between subjects paid according to the quadratic scoring rule⁴⁶ and unpaid subjects. In a study by Sonnemans and Offerman (2001), there are no significant differences either in performance or in exerted effort for subjects paid according to the quadratic scoring rule and subjects who are paid a lump sum. Nelson and Bessler (1989) find that, even with feedback, stated beliefs in the first five periods are indistinguishable between a group rewarded according to the quadratic scoring rule and a group rewarded according to the blatantly non-incentive-compatible linear scoring rule.⁴⁷ Given the very different incentive structure of a linear scoring rule, a quadratic scoring rule and schemes which give no performance-based monetary reward (no payment and lump sum payment), it is difficult to imagine that subjects are deliberately misrepresenting their beliefs in the same way under all these schemes.⁴⁸ In addition, CGW and Rey-Biel (2005) find that subjects' stated beliefs predict actual frequencies rather well.

⁴⁶This scoring rule is incentive compatible for risk neutral expected utility maximizers.

⁴⁷The linear scoring rule pays proportionately to the probability placed on the actual outcome. A risk neutral expected utility maximizer should state a degenerate belief with weight 1 on the mode of her true belief.

⁴⁸Some early psychological studies (Beach and Philips (1967) and Jensen and Peterson (1973)) also suggest that there is no difference between paying subjects according to a scoring rule and not paying them based on performance.

All this evidence is reassuring that deliberate belief misrepresentation is not a serious practical problem in the laboratory. We opt for a lump sum payment since it is easy to explain to subjects and since with a lump sum payment subjects do not even have a theoretical incentive to misrepresent beliefs regardless of risk preferences. The latter is not true for scoring rules since they are incentive compatible only for particular risk preferences (usually risk neutrality).

The third issue, which pertains to A, is a more serious concern. In particular, it is possible that beliefs shift between the time strategic subjects choose actions (part I) and the time they state beliefs (part II).⁴⁹ There are a few points which are reassuring in this respect.

First, if beliefs shifted frequently in A, then one would expect the difference between subjects' estimated mean precision in A and B (in which beliefs and actions are entered simultaneously) to be larger than what it is (see table 1.6).

Second, even if beliefs shift, this is unlikely to invalidate the conclusion about the low proportion of naive types in A - if anything shifting beliefs are likely to favor naive types. The reason for this is that in our games (similar to the games in other studies as explained in the introduction) naive types pick actions which maximize expected utility given any belief in a disproportionately large area of the simplex.⁵⁰ So if a subject best responds to a certain belief and then states a completely different belief, chances are that her behavior would look more like being naive than like being a best response to the stated belief.

Third, the conclusion that more subjects behave in a risk neutral way in B than in C is unaffected by any possibility of shifting beliefs in A.

⁴⁹In the case of naive types, shifting beliefs is not a concern since they, by definition, do not form beliefs in part I.

⁵⁰In our games *NRN* (*SRN*) picks an action which maximizes expected utility given any belief in 61% (46%) of the simplex (averaged over games).

1.5.3 Robustness Across Games and Subject Population

As mentioned in the introduction, behavior in A does not seem to be robust across games with similar formal structure or across subject populations.

In particular, in the pilot for A we used CGW's games. These games are similar in formal structure to our games: they are 3×3 , each one has a unique pure Nash equilibrium and 10 of the 14 games are dominance solvable. However, the ML estimate for the proportion of naive players in the pilot is 45% which is significantly different from the estimate in A ($p = 0.008$).

One plausible explanation of this difference is that the extent of naive play may depend on the cost of being naive (in the sense of the opportunity cost of not being strategic). In particular, in 9 out of the 14 games in CGW, the $L1$ action guarantees a payoff of at least 45 ECU out of a maximum of 99 ECU and can potentially earn much more. In 12 out of the 14 games, the $L1$ action coincides with the maximin action. Given this, it seems there is little incentive for subjects to engage in much strategic thinking so they simply go with the $L1$ action.

Another, related conjecture is that the extent of naive play depends on how costly it is to engage in strategic thinking. For example, players may be briefly scanning the opponent's actions to see if some of these actions can be ruled out based on their low average payoffs. If this cheap decision-making procedure yields no insights subjects simply play naively. This may result in more naive play in CGW's games since in these games the average payoffs of the opponent's actions are closer than in our games.⁵¹ In other words, in CGW's games it may be more costly to rule out one (or more) of the opponents' actions - hence, there is more naive play.

⁵¹In particular, in our games the ratio of the average payoffs of the opponent's action with the lowest average payoff and the opponent's action with the highest average payoff is 0.61 (averaged over games). The corresponding ratio in CGW's games is 0.72.

The issue of how naive play depends on the cost of being naive and on the costliness of strategic thinking seems like a promising topic for further research.

We also ran a follow-up session for A using undergraduate subjects in order to check the robustness of our results across subject populations. The ML estimate for the proportion of naive players in the follow-up session is 47% which is significantly different from the estimate in A ($p = 0.01$). The direction of the difference (even if not necessarily its size) is not particularly surprising. In addition, the estimate for the proportion of risk neutral players in the follow-up session is 81% which is also significantly different from the estimate in A ($p = 0.007$). A tentative conjecture for the difference is simply that graduate students are more risk averse.

Overall, the pilot and the follow-up session for A seem to suggest that it is difficult to draw general conclusions about behavior in one-shot normal-form games as it may vary across games (even if they have a similar formal structure) and across subject populations. The question remains open whether the variations can be explained convincingly in a parsimonious and intuitive way.

1.5.4 Ambiguity Aversion rather than Risk Aversion?

Risk aversion, taken literally, cannot explain the variation across treatments in our estimates of the proportion of risk neutral types because there is no reason for subjects' utility function to change its curvature across treatments. This suggests that there is something more at work than mere curvature of the utility function.

An obvious candidate explanation is ambiguity aversion. Risk aversion and ambiguity aversion have similar implications for behavior - intuitively, both lead to cautious behavior.⁵² Therefore, what we are estimating as risk averse behavior within

⁵²More formally, risk aversion diminishes the importance of high payoffs through the concavity of the utility function; ambiguity aversion has the same effect through pessimistic priors (within the Gilboa and Schmeidler (1989) multiple prior framework) or through subadditive beliefs (within the Schmeidler (1989) framework of nonadditive beliefs).

each treatment may actually be driven by ambiguity aversion. In addition, unlike risk aversion, ambiguity aversion could potentially provide an explanation for any variation across treatments in the estimated proportions of risk neutral subjects given that the ambiguity of decision tasks may differ across treatments.

In particular, ambiguity aversion is a plausible explanation for the increase in the estimated proportion of risk neutral subjects from A to B: by prompting subjects in B to form beliefs about what the opponent might do, we are giving them a way to think about the games and this may reduce ambiguity.⁵³

The finding that the estimated proportion of risk neutral subjects in C is much lower than in B and is only slightly higher than in A is quite surprising. It can be explained with ambiguity aversion only if the lottery tickets in C are perceived by subjects as ambiguous. The latter is possible if, for example, subjects are not very comfortable with lottery tickets because of their abstract nature or even if subjects simply have doubts about whether the computer (as programmed by the experimenter) randomizes with the given probabilities. Although not implausible, this explanation is, at this point, merely a speculation. It is also likely that something different from both risk aversion and ambiguity aversion is behind the smaller estimates of the proportion of risk neutral subjects in C compared to B. At this point, we do not have a clear explanation.

Note, however, that our finding that the estimated proportion of risk neutral subjects is higher in B than in C and is (slightly) higher in C than in A is consistent with what has previously been found in the literature. In particular, Heath and Tversky (1991) find that (i) subjects prefer to bet on ambiguous gambles in areas in which they are competent than to bet on lottery tickets and (ii) prefer to bet on lottery tickets than to bet on ambiguous gambles in areas in which they are not

⁵³An alternative interpretation is that (i) subjects perceive the decision situations as *less* ambiguous in A than in B and (ii) they are *ambiguity loving*. This interpretation seems implausible since the games could hardly be perceived as less ambiguous in A than in B.

competent.⁵⁴ ⁵⁵ If one draws an analogy between gambles in areas of competence and our treatment B (subjects are given a way to think about the games) and between gambles in areas of no competence and our treatment A (subjects are not given a way to think about the games), then there is an obvious similarity between our finding and that of Heath and Tversky: our (Heath and Tversky's) finding suggests that subjects are more willing to take risks in B (with gambles in areas of competence) than in C (with lottery tickets) and are more willing to take risks in C (with lottery tickets) than in A (with gambles in areas of no competence).⁵⁶

If the interpretation is correct that the differences in the estimates of the proportion of risk neutral types between A and B are driven by ambiguity aversion, this suggests the following hypothesis: When people have a clear, structured way to think about a decision situation, this reduces ambiguity and hence ambiguity averse behavior. If correct, this hypothesis has far-reaching implications. First, it sheds light on ambiguity aversion as being, to a large extent, the result of people's inability to think clearly about a decision situation. In this way, it could explain, for example, findings in the literature that people are more willing to take risks in areas of competence (Heath and Tversky (1991)) and that people with higher cognitive abilities are more willing to take risks (Benjamin and Shapiro (2006)). Second, this hypothesis implies that there may be strong framing effects, depending on whether a problem is formulated in a way which suggests a clear way to think about it. Third, if one is of the opinion that ambiguity aversion is normatively incorrect (e.g. because it violates

⁵⁴See especially their Experiment 4.

⁵⁵The probabilities of the different outcomes of the lottery ticket are matched to the subject's belief over the outcomes of the ambiguous gamble.

⁵⁶Note that Heath and Tversky's finding cannot be explained with any attitude to ambiguity if lottery tickets are unambiguous. If a lottery ticket is unambiguous an ambiguity averse subject would prefer it to any ambiguous gamble (regardless of the subject's competence in the area of the ambiguous gamble); an ambiguity loving subject would prefer any ambiguous gamble to the lottery ticket; and an ambiguity neutral subject would be indifferent. Therefore, if a lottery ticket is perceived as unambiguous, no attitude to ambiguity could explain the lottery ticket being ranked strictly between two ambiguous gambles.

the independence axiom), then this hypothesis has normative implications: if one wishes to reduce ambiguity averse behavior, one needs to induce decision-makers to think about the situation at hand in a clear, structured way. Given the significance of this hypothesis, its validity needs to be checked in further studies.⁵⁷

Note that when subjects choose between being paid according to a bet on the opponent's action in a game or according to a corresponding lottery ticket (part III in A and part II in B), there is little evidence for ambiguity aversion. In particular, subjects chose the bet on the opponent's action over the lottery ticket 54% of the time in A and 51% in B.⁵⁸ Two offsetting effects could be responsible for these numbers. On the one hand, subjects have already stated beliefs in both A and B by the time they have to choose between betting on the opponent's action or on a lottery ticket: this may reduce the ambiguity of the bet on the opponent's action and may make subjects more inclined to choose it. On the other hand, the direct comparison between the bet on the opponent's action and a lottery ticket may make it salient that the probabilities are given for the lottery tickets but not for the games and this may bias subjects towards the lottery tickets.^{59 60}

1.5.5 Concluding Remarks

In this study, we introduce a new approach to specifying types of players in one-shot normal-form games. The reasonableness of estimated parameters (e.g. precision parameters) and the plausible ways in which parameters vary across treatments (especially the estimated proportion of naive subjects from A to B to C and the estimated proportion of risk neutral subjects from A to B) are reassuring regarding the appro-

⁵⁷I am currently devising an individual-choice experiment to test this hypothesis.

⁵⁸These numbers are not statistically different from 50% ($p=0.112$ and $p=0.616$, respectively).

⁵⁹See Fox and Tversky (1995).

⁶⁰In the pilots rather than have subjects *choose* between betting on the opponent's action in a game and betting on a lottery ticket, we had subjects *price* these two options (using the Becker, DeGroot, Marschak (1964) procedure). Stated prices were not significantly different.

priateness of our specification. The width of Bayesian 90% confidence intervals and our ability to reject interesting hypotheses are reassuring regarding the usefulness of our approach.

The main findings of the paper are twofold. First, we estimate that in A only a small minority of subjects is naive and only a minority is risk neutral. However, these results may not be robust to the games used or to the subject population. The next challenge is to explain the variation.

Second, we find that behavior is much closer to risk neutrality when subjects are prompted to state beliefs over the opponent's actions (treatment B) compared to when subjects simply play the games (treatment A). This result may be driven by ambiguity aversion. The hypothesis that simply giving subjects a clear way to think about a decision situation may reduce ambiguity has far-reaching implications and needs to be studied further.

It is not clear why behavior is much closer to risk neutrality when subjects are prompted to state beliefs over the opponent's actions (treatment B) compared to when subjects choose between lottery tickets (treatment C).

CHAPTER 2

HINDSIGHT, FORESIGHT, AND INSIGHT: AN EXPERIMENTAL STUDY OF A SMALL-MARKET INVESTMENT GAME WITH COMMON AND PRIVATE VALUES

2.1 Introduction

The theoretical literature on herding with endogenous timing, pioneered by Chamley and Gale (1994), explores the important issue of how market activity aggregates and transmits private information. Will firms with favorable information about the investment climate act on that information, thereby providing benefits to others, or will they postpone investment, to acquire information by observing other firms' investment activity? In Chamley and Gale (1994), firms receive a signal correlated with the unobserved investment return, which is common to all firms, and then face a sequence of decisions about whether or not to invest.¹ They find that the incentive

¹The previous generation of herding models assumes that agents have one opportunity to invest, and must decide in an exogenously specified sequence. See Banerjee (1992) or Bikhchandani, Hirshleifer, and Welch (1992). Also, Anderson and Holt (1997) provided the first experimental tests of herding models with exogenous timing. For a nice overview of the theoretical literature on herding, see Chamley (2003).

to delay leads to inefficiency and the possibility of investment collapse. Indeed, firms are no better off than in a static game, in which firms must invest without learning anything about other firms' information. Levin and Peck (2006) introduce a second signal, about the cost of undertaking the investment, which is firm specific and independent of the costs faced by other firms. Observing the investment decisions of other firms could be used to improve inference about the aggregate state, but firms must disentangle whether another firm invests because it receives a favorable signal about investment returns or simply has a low cost.

Experimentally testing the theoretical implications of the endogenous timing herding literature is important for several reasons. First, we can compare the theoretical implications for aggregate investment activity on markets with what actually occurs in the lab. One intriguing possibility is that subjects with favorable information underestimate the option value of waiting and become more likely to invest immediately, thereby providing more information to the market and a more efficient outcome than the theoretical analysis would predict.^{2 3} Second, interactions are purely informational in our investment market, so decision making is not complicated by payoff externalities, as opposed to an auction environment where subjects must decide how much to bid, how others will respond, and so on. We can test separately (i) whether a subject understands the expected profits of investing, (ii) whether a subject draws inferences from the other's behavior (hindsight), and (iii) whether a subject looks ahead by considering the option value of delaying investment (foresight). On this score, we add to the work of SgROI (2003) and Ziegelmeyer et al (2006).⁴ Third, because interactions are purely informational, the setting offers a

²The other experimental work on endogenous timing herding models, by SgROI (2003) and Ziegelmeyer et al (2005), are not well suited to address issues of aggregate investment, because in those games it is always optimal to "invest" eventually.

³There are examples in the experimental literature on exogenous timing herding in which departures from the theoretical predictions lead to increased social learning. See Kübler and Weizsäcker (2004) and Goeree et al. (2006)

⁴SgROI (2003) uses an "urn" setting, in which signals either strongly point to the (mostly) red

sharp test of various recently advanced behavioral theories, such as *level-k beliefs* (see Crawford and Iriberri (2005)) and *cursed equilibrium* (see Eyster and Rabin (2005)).

Let us sketch our experimental treatments. Subjects are matched in two-player trials, and each subject observes her cost of investment and another signal that is correlated with the common investment return. Ranging from least favorable to most favorable, possible types are $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$, where the first component is the common-value signal (0 is low and 1 is high) and the second component is the investment cost (L is low and H is high). At the beginning of each trial, the subjects observe their signals and simultaneously decide whether to invest in round 1. Then, subjects that do not invest in round 1 observe whether the other subject invested in round 1, and decide whether to invest in round 2, and so on. Investment is irreversible, and can be done at most once per trial per subject. In our Two-Cost Treatment, each subject's investment cost is high or low with probability one half, independent of the other subject's cost and the investment return. In our Alternating One-Cost Treatment, trials alternate between common knowledge of high cost and common knowledge of low cost.

We find that the typical subject is quite good at determining whether investment is profitable, based on her type. In particular, a type $(0, L)$ subject is far less likely to invest in round 1 than a type $(1, H)$ subject, even though it may not be obvious that a bad signal and low cost is slightly unprofitable and a good signal and high cost is slightly profitable. Thus, this basic aspect of rationality is satisfied. Subjects usually treat investment by the other subject in the trial as good news about the investment return; they are more likely to invest in round 2 following investment in round 1

urn, strongly point to the (mostly) white urn, or are completely neutral. The most profitable urn to pick and the proper inference from the other subjects' behavior are somewhat obvious. In Ziegelmeyer et al (2006), the most profitable investment choice is obvious for the first mover, and the only inference required is an understanding that the first mover has a stronger signal than someone who waits. Both papers exhibit an interesting tradeoff between the cost of delay and the information gained.

than following no investment in round 1. Thus, subjects show that they can draw inferences from the other subject's behavior, in hindsight. The evidence regarding foresight is nuanced. On the one hand, a significant fraction of type $(1, H)$ subjects, who would receive positive expected profits by investing in round 1, prefer costly waiting and invest in round 2 if and only if the other subject invested in round 1. It seems that these subjects have the foresight to understand that waiting will provide useful information. On the other hand, we argue below that it is unlikely that these subjects are computing the option value of waiting, which is then compared to the profits of investing immediately.

Here are our main results about aggregate investment and information flows. We find that the frequency of investment is higher and overall profits are lower than those predicted by the Nash equilibrium. Investment by the other subject generates an informational externality, which could be either positive or negative. There is a positive informational externality due to investment by subjects with a high common-value signal, and a negative informational externality due to investment by subjects with a low common-value signal. Not surprisingly, we find that the overall externality is positive in all our treatments. The more interesting comparison is to the Nash prediction. In our Two-Cost Treatment, the incremental positive informational externality (over Nash), due to overinvestment by subjects with a high common-value signal, is balanced by the negative externality created by unprofitable investment by subjects with a low common-value signal. The overall externality is as large as in the Nash equilibrium, in the sense that a subject best responding to the empirical distribution of strategies receives as high a profit as she would receive if everyone were playing Nash. In our Alternating One-Cost Treatment, the incremental positive externality due to overinvestment disappears, but it is the theoretical yardstick and not investment behavior that is different.

For our two main treatments, there are essentially only three strategies that are consistent with Nash, level-k beliefs, or cursed equilibrium. The *self-contained* strategy, S , responds optimally to one's own information, but ignores the behavior of the other subject. The *myopic* strategy, M , is to invest in round 1 whenever investment is profitable (even if waiting is more profitable), but to properly infer (in hindsight) that investment by the other subject in round 1 is good news, and act accordingly. The *foresight* strategy, F , is to delay investment whenever valuable information can be revealed, but otherwise to invest when profitable. We perform maximum likelihood estimation on the proportion of S subjects, M subjects, and F subjects, allowing for errors. Even though the theory differs across our two main treatments, our maximum likelihood estimates are nearly identical. The proportion of F subjects is greater than one half, and the proportion of S subjects is greater than one third. These findings are inconsistent with any symmetric cursed equilibrium (see Eyster and Rabin (2005)), which allows for either F or S , depending on the level of cursedness, but not the coexistence of F and S . As far as level-k beliefs are concerned,⁵ a level-1 subject plays S , and a level-2 subject plays F , so behavior across our two main treatments is broadly consistent with level-k beliefs.

Cursed equilibrium and level-k beliefs are belief-based theories, in which players form beliefs as specified in the theory and choose best responses accordingly. Our alternative interpretation of behavior is that subjects choose particular rules of thumb, based on various sorts of *insights* that subjects may experience. First, there is the insight that investment in round 1 is profitable for types $(1, H)$ and $(1, L)$, and not for the other types. Second, there is the insight that investment by the other subject signals the high common-value signal, calling for an updating of beliefs. Third, there is the insight that delaying investment may provide useful information. Thus, S , M , and F are interpreted as rules of thumb, reflecting the degree of insight acquired. To

⁵See Stahl and Wilson (1995), Nagel (1995), or Crawford and Iriberri (2005).

better separate our insight story from the belief-based explanations, we introduce a new treatment. In Treatment 3, a type $(1, H)$ subject receives higher expected profits by investing in round 1 than waiting, even if waiting revealed the other subject's type with probability one. Playing F is strictly dominated, so it is inconsistent with any theory of behavior involving a best response to beliefs about the other subject's strategy. On the other hand, it is plausible that our subject has the insight that waiting will yield useful information, but does not perform a calculation to see that F is dominated. Maximum likelihood estimates for Treatment 3 indicate that the proportion of F subjects remains above one half, providing support for the insight story and evidence against the belief-based theories.

2.2 Theoretical Framework

Our theoretical framework is based on the general model studied in Levin and Peck (2006). There are $n \geq 2$ risk-neutral players or potential investors. Let $Z \in \{0, 10\}$ denote the true investment return, common to all investors, with $pr(Z = 0) = pr(Z = 10) = \frac{1}{2}$. Each player $i = 1, \dots, n$ observes a signal correlated with the investment return, $X_i \in \{0, 1\}$, which we call the “common-value” signal of player i . We assume that signals are independent, conditional on Z . The accuracy of the signal is given by the parameter, $\alpha \in [\frac{1}{2}, 1]$:

$$pr(Z = 0 \mid X_i = 0) = pr(Z = 10 \mid X_i = 1) = \alpha.$$

When we have $\alpha = \frac{1}{2}$, common-value signals have no information content at all, and when we have $\alpha = 1$, a common-value signal fully reveals the aggregate state. Thus, α effectively captures the informativeness of the common-value signal, X_i .

Each player i also privately observes a second signal, representing the idiosyncratic cost of undertaking the investment, c_i . We assume that c_i is independent of all other variables, and distributed according to a distribution function defined over the support, $[\underline{c}, \bar{c}]$. The structure of signals is common knowledge.

Impatience is measured by the discount factor, $\delta < 1$. If player i has cost c_i and the state is Z , her profits are zero if she does not invest, and $\delta^{t-1}(Z - c_i)$ if she invests in round t . We now describe the game. First, each player observes her signals, (X_i, c_i) . For $t = 1, 2, \dots$, each player observes the history of play through round $t - 1$, and players not yet invested simultaneously decide whether to invest in round t . A strategy for player i is a mapping from signal realizations and histories (including the null history) into a decision of whether to invest following that history. We require that a player can invest at most once.

Although Levin and Peck (2006) consider continuous cost distributions, our experimental design considers a discrete distribution containing either one or two points. This simplifies the decision making required of subjects and simplifies the data analysis. At the same time, it maintains the essential tradeoff between the incentive to delay and gain information by observing investment activity, versus the associated shrinkage of the (expected) pie due to discounting. We now define the three games relevant to our experiment, and solve for the Bayesian Nash equilibria. For the remainder of this section, we restrict attention to the parameter values, $n = 2$, $\delta = 0.9$, and $\alpha = 0.7$.

Game 1 (two costs):

There are two equally likely cost realizations, $L = 3.5$ and $H = 6.5$. Thus, we have four possible types of players, based on the common-value signal and the cost: $(0, H)$, $(0, L)$, $(1, H)$, and $(1, L)$. Equilibrium path play is uniquely determined, and involves pure strategies. A type $(0, H)$ player will never invest under any circumstances,

	Equilibrium Strategy	Expected Profits for each type
type $(0, H)$	N	0
type $(0, L)$	W	0.2835
type $(1, H)$	W	0.5085
type $(1, L)$	1	3.5

Table 2.1: Equilibrium Characterization for the Two-Cost Game

because her expected profits from investing are negative even if she knows that the other player has the high common-value signal.⁶ Similarly, a type $(1, L)$ player finds it profitable to invest even if she knows that the other player has the low common-value signal, and therefore invests in round 1. A type $(0, L)$ player will not invest in round 1, because the expected return is 3, while her cost is 3.5. However, since we have established that investment in round 1 must come from a player with the high common-value signal, a type $(0, L)$ player will invest in round 2 if the other player invests in round 1, because her conditional expected return becomes 5.

The remaining type, $(1, H)$, is the most interesting. Investment in round 1 yields positive expected profits of 0.5, but profits are slightly higher by taking advantage of the option value of waiting, investing in round 2 if and only if the other player invests in round 1. If all other type $(1, H)$ players wait, profits from waiting are 0.5085. If some other type $(1, H)$ players would invest in round 1, the advantage of waiting is even greater. Thus, we have characterized the equilibrium path, which is given in Table 1.⁷ To simplify the discussion, denote the choice to invest in round 1 as “1”, denote the choice never to invest as “ N ”, and denote the choice to wait and invest immediately following investment by the other player as “ W ”.

⁶In such a case, her conditional expected return would be 5, but her cost is 6.5.

⁷Several specifications of behavior and beliefs off the equilibrium path are consistent with sequential equilibrium, all yielding the same equilibrium path. After no one invests in round 1 and one player invests in round 2, beliefs about the investor’s common-value signal can affect the remaining player’s decision. However, the beliefs and subsequent decision do not affect the original investor’s profit, so the equilibrium path is unaffected.

Game 2 (low cost):

There is only one possible cost realization, 3.5. Thus, we have two possible types of players, $(0, L)$ and $(1, L)$. It is easy to see that a type $(1, L)$ player would want to invest, no matter what it believes about the other player's type, so she invests in round 1. A type $(0, L)$ player finds it unprofitable to invest in round 1, but will invest in round 2 if the other player invests in round 1 (implying a high common-value signal).

Game 3 (high cost):

There is only one possible cost realization, 6.5. Thus, we have two possible types of players, $(0, H)$ and $(1, H)$. It is easy to see that a type $(0, H)$ player would not want to invest, no matter what it believes about the other player's type, so she never invests. A type $(1, H)$ player compares the profits of choices "1" and "W". If all other type $(1, H)$ players choose "1", then she can act with full information in round 2, and "W" is a best response. If all other type $(1, H)$ players choose "W", then nothing is learned from waiting, and "1" is a best response. Therefore, type $(1, H)$ players mix, choosing "1" with probability 0.4916 and choosing "W" with probability 0.5084.

In our Alternating One-Cost Treatment, the subjects alternate between Game 2 and Game 3. Because matching is random and anonymous, folk theorem issues do not arise, so the equilibrium characterization combines the equilibria of Game 2 and Game 3, as given in Table 2.

The Bayesian Nash equilibria characterized in Table 1 and Table 2 allow for inefficient delay. A type $(1, H)$ player does not take into account the positive informational externality that investing in round 1 provides to the other player. Thus,

	Equilibrium Strategy	Expected Profits for each type
type (0, H)	N	0
type (0, L)	W	0.567
type (1, H)	1 with probability 0.4916	0.5
	W with probability 0.5084	
type (1, L)	1	3.5

Table 2.2: Equilibrium Characterization for the Alternating One-Cost Treatment

if our experimental subjects are more likely to invest than the theory predicts when they are type (1, H), it is possible that profits are higher than the theory predicts. Of course, it is also possible that subjects invest in round 1 with the low common-value signal, leading to a negative informational externality and lower profits.

2.3 Behavioral Considerations

In the investment games we study here, interactions between players are purely informational, with no direct payoff consequences. This simple structure allows us to test separately whether subjects invest when investment is unprofitable; whether subjects draw inferences, from round 1 investment by the other subject (*hindsight*); and whether they delay profitable investment in order to observe the other subject's decision (*foresight*). These different aspects of rationality would be more difficult to disentangle in auctions or other games with payoff externalities.⁸ In particular, we focus on three strategies, which will be relevant for the three behavioral theories we consider. A subject who disregards the behavior of the other subject will invest if and only if investment is profitable based upon her own signals. Such a subject chooses the type-dependent strategy $S \equiv (N, N, 1, 1)$, where S stands for *self-contained*.⁹

⁸For example, Garratt and Keister (2006) experimentally test a model of bank runs, where a player's decision to withdraw deposits simultaneously involves Bayesian updating, considering the option of withdrawing in the future, and anticipating the strategies of the other players.

⁹The type-dependent strategy, (A, B, C, D) , means that type (0, H) players choose A , type (0, L) players choose B , type (1, H) players choose C , and type (1, L) players choose D .

A subject who invests whenever investment is profitable, incorporating information from the history of play, chooses the type-dependent strategy $M \equiv (N, W, 1, 1)$, where M stands for *myopic*. Thus, an M subject is able to draw inferences from market activity, in *hindsight*, but does not look with *foresight* to the informational benefits of waiting. Note that it is possible that an M subject of type $(1, H)$ might understand that there is a benefit to waiting, but feel that investment in round 1 is the better choice.¹⁰

Consider a subject who updates beliefs, in hindsight, and looks with *foresight* to the useful information that can be gained by waiting when her type is $(0, L)$ or $(1, H)$.¹¹ Such a subject chooses the type-dependent strategy $F \equiv (N, W, W, 1)$. Note that this *foresight* does not necessarily mean that she actually calculates the value of waiting, or even that such a calculation would justify waiting.

The strategies, S , M , and F , have interpretations that depend on the theory used to justify behavior. We focus on two belief-based theories, cursed equilibrium and level-k beliefs, and a theory of boundedly rational rules of thumb, based on insights a subject might acquire about understanding the game. As it turns out, we can characterize the Nash equilibrium, level-k behavior, and the cursed equilibrium, all in terms of these three strategies, so it is natural to focus on them.¹²

There is considerable evidence that experimental subjects fail to pick actions in accordance with the relevant Nash equilibrium. Such discrepancies are particularly acute in tasks that require players to make inferences and update their priors based

¹⁰Another interpretation is that the subject is somehow more impatient than that implied by the discount factor, 0.9.

¹¹A type $(0, H)$ subject does not receive useful information by waiting, because the expected profits from investment are always negative. A type $(1, L)$ subject also does not receive useful information by waiting, because the expected profits from investment are always positive.

¹²There is one other strategy in which a subject never makes an unprofitable investment, and invests immediately when nothing can be learned to make investment unprofitable. This strategy is $(N, N, W, 1)$. Maximum likelihood estimation, allowing for all four strategies, does not change our conclusions. Indeed, allowing this strategy would stack the deck against cursed equilibrium and level-k beliefs, in favor of our explanation of insight-based rules of thumb.

on other players' actions in games with incomplete information.¹³ There are also several studies claiming such failure in real markets.¹⁴ Faced with such overwhelming evidence, it is not surprising that there were several attempts to explain these discrepancies. Stahl and Wilson (1995) and Nagel (1995) use a non-equilibrium model of "level-k" beliefs, where L_0 players behave in some pre-determined way (usually randomly), and for $k = 1, 2, \dots$, the L_k players choose a best response to the strategy chosen by the L_{k-1} players. See Crawford and Iriberri (2005) for a survey and an explanation for the winner's curse in auctions, based on level-k beliefs. Eyster and Rabin (2005) propose an alternative theory, which they call "cursed equilibrium." Players are assumed to best respond to the other players' strategies in a certain sense. In a χ -cursed equilibrium, players believe that with probability χ , each other player j chooses an action that is type-independent, and whose distribution is given by the ex ante distribution of player j 's actions. Also, players believe that with probability $1 - \chi$, each other player j chooses an action according to player j 's type-dependent strategy. Thus, if $\chi = 0$, we have a standard Bayesian Nash equilibrium, and if $\chi = 1$, players draw no inferences about other players' types. Both level-k beliefs and cursed equilibrium weaken the "usual" requirements of correct beliefs regarding other players' strategies, while insisting on players rationally choosing best responses to the more flexible belief structures that are allowed.

For our Two-Cost Treatment (Game 1) and our Alternating One-Cost Treatment (Games 2 and 3), the Nash equilibrium, level-k behavior, and the cursed equilibrium

¹³Leading examples from laboratory studies include failures in the *takeover game* (see Ball and Bazerman (1991) and Charness and Levin (2005)), and systematic overbidding and the winner's curse in common-value auctions (see Bazerman and Samuelson (1983), Kagel and Levin (1986), Levin, Kagel, and Richard (1996), Holt and Sherman (1994), and Kagel and Levin (2002)).

¹⁴Leading examples include oil and gas lease auctions (see Capen, Clapp, and Campbell (1971), Mead, Moseidjord, and Sorensen (1983) and (1984), and the opposite conclusions reached in Hendricks, Porter, and Boudreau (1987)), baseball's free agent market (see Cassing and Douglas (1980) and Blecherman and Camerer (1998)), book publishing (see Dessauer (1981)), initial public offerings (see Levis (1990) and Rock (1986)), and corporate takeovers (see Roll (1986)).

	Two-Cost Game	Alternating One-Cost Game
Nash Equilibrium	F	M with probability 0.4916
		F with probability 0.5084
Level-k		
L_1	S	S
L_2	F	F
L_3	F	M
L_4	F	F
Cursed Equilibrium		
$0 < \chi < \frac{17}{756}$	F	M with probability q^{alt}
		F with probability $1 - q^{alt}$
$\frac{17}{756} < \chi < \frac{517}{756}$	M with probability q^{2cost}	M with probability q^{alt}
	F with probability $1 - q^{2cost}$	F with probability $1 - q^{alt}$
$\frac{517}{756} < \chi < \frac{3}{4}$	M	M
$\frac{3}{4} < \chi < 1$	S	S

Table 2.3: Nash, Level-k, and Cursed Equilibrium

are characterized in Table 3, all in terms of the three strategies, S , M , and F .¹⁵

The probabilities given in Table 3, q^{2cost} and q^{alt} , are the following.

$$q^{2cost} = \frac{756\chi - 17}{9(113 - 84\chi)} \quad \text{and} \quad q^{alt} = \frac{500}{9(113 - 84\chi)}.$$

For our Two-Cost Treatment (Game 1), the Nash equilibrium strategy is F ,¹⁶ and for our Alternating One-Cost Treatment (Games 2 and 3), the symmetric Nash equilibrium involves mixing over M and F according to the probabilities given. Now

¹⁵We are being a little loose when we refer to F , M , and S as strategies, because behavior is not specified after unexpected contingencies that do not affect play in the NE, the cursed equilibrium, or under level-k beliefs. In our maximum likelihood estimation (Section 5), we are careful to specify the decisions following investment by the other subject in round 2 (after not investing in round 1), and following a tremble that is inconsistent with the strategy itself.

¹⁶In fact, for the Two-Cost Treatment, suppose a player knows that everyone is playing one of the strategies, S , M , or F . Then the best response is F , no matter how many of the other players are choosing each of the three strategies.

let us consider level- k beliefs, starting with the Two-Cost Treatment. A player choosing a best response to an L_0 player who plays randomly can learn nothing from the other player's behavior, so an L_1 player chooses the strategy S . An L_2 player chooses the best response to S , which is F . For $k = 3, 4, \dots$, an L_k player chooses the best response to F , which is F . For the Alternating One-Cost Treatment, an L_1 player chooses the strategy S , and an L_2 player chooses the best response to S , which is F . However, the situation is different from the Two-Cost Treatment, where a type $(1, H)$ is better off waiting if others are playing F , based on the hope that the other player is type $(1, L)$. In the Alternating One-Cost Treatment, a type $(1, H)$ is playing the high-cost game (Game 3), and knows that the other player also has high cost. Therefore, if the other player chooses F , this implies that the other player never invests in round 1, in which case the best response is to invest. It follows that the best response to F is M . On the other hand, the best response to M is F , because if the other player invests in round 1 when type $(1, H)$, a type $(1, H)$ is better off waiting.

Computing the symmetric cursed equilibria for our games is a bit more involved, and details are given in the Appendix. In the Two-Cost Treatment, when the cursedness parameter, χ , is small, the Nash equilibrium strategy, F , continues to be played. For $\frac{17}{756} < \chi < \frac{517}{756}$, if type $(1, H)$ players wait, the informativeness of round 1 investment is sufficiently weakened that it does not pay them to wait. However, if all type $(1, H)$ players invest in round 1, now it pays to wait. Therefore, the cursed equilibrium involves mixing by type $(1, H)$ players, but a type $(0, L)$ player will still want to invest if she sees the other player invest. Thus, players mix between M and F , with the probability of choosing the myopic strategy, $q^{2\text{cost}}$, increasing in χ . For $\frac{517}{756} < \chi < \frac{3}{4}$, the level of cursedness is high enough so that a type $(1, H)$ player is better off investing in round 1 than waiting, even if all other $(1, H)$ players are also

investing in round 1. A type $(0, L)$ player will still want to invest if she sees the other player invest, so in the χ -cursed equilibrium the players choose the pure strategy, M . Finally, for $\chi > \frac{3}{4}$, the level of cursedness is so high that no useful inferences can be made, and the equilibrium strategy is S .

For the Alternating One-Cost Treatment, when $\chi < \frac{517}{756}$, players mix between M and F in the cursed equilibrium, with the probability, q^{alt} , of choosing the myopic strategy M increasing in χ . The probability of choosing M ranges from the Nash equilibrium value, $q^{alt} = 0.4916$ for $\chi = 0$, to unity, $q^{alt} = 1$ for $\chi = \frac{517}{756}$. For $\frac{517}{756} < \chi < \frac{3}{4}$, the informativeness of round 1 investment is too low to justify waiting for a type $(1, H)$, but it is high enough for a type $(0, L)$ to be willing to invest after she sees the other player invest. Thus, the cursed equilibrium strategy is M . For $\chi > \frac{3}{4}$, even a type $(0, L)$ is unwilling to invest after she sees the other player invest, and the cursed equilibrium strategy is S .

Of course, there are many possibilities for bounded rationality that go beyond misspecification of beliefs.¹⁷ For our investment games, a small but nonzero proportion of type $(0, H)$ subjects invests in round 1, which is a mistake indicating a lack of understanding of the game. If mistakes like this occur, perhaps less extreme departures from Nash behavior reflect a lack of insight about some of the fine points of the game.

Let us elaborate on the insights we have in mind. Investment in round 1 is profitable for types $(1, H)$ and $(1, L)$, and not for the other types. We start with the notion that subjects have the basic insight that the expected revenue is 7 when receiving the high common-value signal and 3 when receiving the low common value signal. The second level of sophistication occurs when a subject, upon seeing the other subject invest, revises upwards the expected investment return, because she realizes that the other subject is very likely to have the high common-value signal.

¹⁷See, for example, Simon (1972) and Rubinstein (1998), including references therein.

Based on our parameters, any upward revision should be enough to induce a type $(0, L)$ subject to invest in round 2 after the other subject invests in round 1 (a numerical computation is not required). The third level of sophistication occurs when a subject has the insight that there is a tradeoff, between potentially higher profits of investing early versus the information gained by waiting. Since a type $(1, H)$ subject with this third level of sophistication does not compute the expected profits from waiting, she might invest in round 1 or wait, depending on how she evaluates this tradeoff.

We call these strategies rules of thumb because they are not based on computation of expected profits, by applying the discount factor and using Bayes' rule explicitly. Not only would such a computation be time consuming, but it would require deeper insights about the game that few subjects are capable of acquiring during the session.¹⁸ Think of the subjects, either as the instructions are read or early in the session, as having a "Eureka moment" giving them a partial understanding of the game. Under this view, it is reasonable to suppose that the exact parameter values, or the difference between the Two-Cost Treatment and the Alternating One-Cost Treatment in how costs are determined, should not affect the generation of these insights. Fewer than 10% of our subjects make choices that are inconsistent with F , M , or S in more than half of their trials.¹⁹

2.4 Experimental Design

The experiment consisted of the Two-Cost Treatment, the Alternating One-Cost Treatment, and Treatment 3, which was adopted to clarify some of the behavioral is-

¹⁸A subject with total command of the game would need to think about the other subject's type-contingent probability of investing in round 1, in order to compute the probability that the other subject has the high common-value signal, conditional on behavior and one's own type.

¹⁹Excluding these subjects from the sample has almost no effect on our estimates for the proportion of F , M , and S subjects, although the error term is reduced.

sues. We conducted two sessions of the Two-Cost Treatment (18 participants in each session), two sessions of the Alternating One-Cost Treatment (26 and 20 participants, respectively), and one session of Treatment 3 (28 participants).²⁰

In the Two-Cost Treatment, subjects played Game 1 (see Section 2) in each trial, so that each subject's private cost of investment was randomly selected to be either 3.5 or 6.5. In the Alternating One-Cost Treatment, subjects played Game 2 in odd numbered trials and Game 3 in even numbered trials, so that each subject's cost alternated between 3.5 and 6.5 from trial to trial (but was the same for all subjects within a trial). Treatment 3 was the same as our Alternating One-Cost Treatment, except that the discount factor was given by $\delta = 0.8$ and the costs alternated between 3.5 and 5.7. To guarantee that the trials ended, without changing the equilibria, subjects were told that the trial ended after either both subjects had invested or there were two consecutive rounds with no investment.

Each session in all treatments consisted of two practice trials and 24 trials in which subjects played for real money. At the start of each trial, subjects were randomly and anonymously matched in pairs to form separate two-player markets which bore no relation to each other. There was a new random matching from trial to trial. Subjects were given an initial cash balance of 20 experimental currency units (ECU). In addition, they could gain (lose) ECU in each trial, which were added to (subtracted from) their cash balances. At the end of the session, ECU were converted into dollars at a rate of 0.6 \$/ECU in our two main treatments, and 0.5 \$/ECU in Treatment 3. Subjects were paid the resulting dollar amount or \$5, whichever was greater. If a subject's cash balances fell below 0 at any point during the session, that subject was paid \$5 and was asked to leave.²¹ Average earnings for the Two-

²⁰We also conducted a pilot session for the Two-Cost Treatment (14 participants).

²¹This occurred for two subjects in the Two-Cost Treatment, for one subject in the Alternating One-Cost Treatment, and for one subject in Treatment 3. After a subject goes bankrupt, if the number of subjects in a session is no longer even, one subject was randomly selected to sit out

Cost Treatment, the Alternating One-Cost Treatment, and Treatment 3 were \$22.82, \$23.53, and \$20.27 respectively. Including the reading of instructions, sessions lasted between 1 hour 45 minutes and 2 hours.

Subjects in the experiment were students at The Ohio State University who were enrolled in undergraduate classes in Economics. The sessions were held at the Experimental Economics Lab at OSU. Before starting the trials, the experimenter read the instructions aloud as subjects read along, seated at their computer terminals. Subjects were invited to ask questions, including after the practice trials. Once the real trials began, no more questions were allowed. See the Appendix for our Instructions and a printout of the screen seen by a player with cost $c_i = 6.5$ and signal $X_i = 1$, who is deciding whether to invest in round 2 after the other player has invested in round 1.

The experiment was programmed and conducted with the software z-Tree (Fischbacher (1999)).

2.5 Results

2.5.1 Aggregate-Level Analysis

Table 2.4 presents aggregate-level decisions in the Two-Cost and the Alternating One-Cost Treatments. There are only six possible histories after which a subject can invest: the null history, $\{\}$; the history following no investment in round 1, $\{0\}$; the history following one subject investing in round 1, $\{1\}$; the history following no investment in round 1 and one subject investing in round 2, $\{0, 1\}$; and so on. The first six rows of the table show, for each treatment, type, and history, how many times a subject facing a decision in that situation invested. There are only three possible histories after which no investment would end the game: $\{0\}$, $(1, 0)$,

during each trial.

History	Two-Cost				Alternating One-Cost			
	(0,H)	(0,L)	(1,H)	(1,L)	(0,H)	(0,L)	(1,H)	(1,L)
{}	20 (9%)	27 (14%)	67 (35%)	148 (78%)	16 (6%)	35 (13%)	109 (42%)	206 (74%)
{0}	12 (9%)	12 (11%)	22 (27%)	11 (46%)	19 (9%)	14 (10%)	30 (27%)	13 (46%)
{1}	9 (15%)	22 (37%)	26 (59%)	18 (100%)	9 (15%)	54 (53%)	24 (62%)	34 (79%)
{0,1}	2 (9%)	5 (45%)	7 (58%)	4 (100%)	5 (14%)	5 (28%)	7 (88%)	1 (100%)
{1,0}	6 (12%)	17 (46%)	6 (33%)	0 N/A	8 (16%)	16 (33%)	9 (60%)	4 (44%)
{0,1,0}	0 (0%)	1 (17%)	0 (0%)	0 N/A	3 (10%)	4 (31%)	1 (100%)	0 N/A
no {0}	101	83	49	9	160	106	74	14
no {1,0}	44	20	12	0	43	32	6	5
no {0,1,0}	20	5	5	0	27	9	0	0
Total	214	192	194	190	290	275	260	277

Table 2.4: Aggregate Actions and Frequency of Investment at each History

and $\{0,1,0\}$. Rows 7-9 show, for each treatment, type, and history for which no investment would end the game, how many times a subject in that situation made the final decision *not* to invest. The numbers in parenthesis show what percentage of decisions made at a particular history were decisions to invest. For example, in the Two-Cost Treatment, the 22 times a type $(0, L)$ subject invested after history $\{1\}$ represent 37% of all decisions taken at that history.²² The last line of the table shows the total number of realizations of each type. Note that a subject's play in a given trial is counted in one, and only one, cell of the table. Hence the last line is simply the sum of all previous lines.

²²To see how the 37% figure is computed, note that the number of times a type $(0, L)$ subject passed through history $\{1\}$ is the sum of the 22 times a type $(0, L)$ subject invested after $\{1\}$, plus the 17 times a type $(0, L)$ subject invested after $\{1,0\}$, plus the 20 times a type $(0, L)$ subject who experienced history $\{1,0\}$ ended up not investing. Thus, $\frac{22}{22+17+20} = 0.37$.

We are interested in performing statistical tests, based on investment decisions and realized profits in our experiment. However, given that there is dependence between the investment decisions and profits of the two subjects in any trial, we will arbitrarily call one subject in each market an A subject (the one with the lower subject ID) and the other subject a B subject. Then we will perform the relevant test separately for A subjects and B subjects. Note that now each test will be based on observations, no two of which were from the same market trial.²³

Let us start by checking if subjects' aggregate behavior satisfies some basic requirements on rationality. First note that, in the aggregate, subjects in both the Two-Cost and Alternating One-Cost Treatments respond to their own information (common value signal and investment cost). This can be seen when we consider investment in round 1, and when we consider investment in any round.

The higher the expected profit from investment, given a subject's type, the more likely she is to invest in round 1. In particular, the frequencies with which subjects of type $(0, H)/(0, L)/(1, H)/(1, L)$ invest in round 1 are 9%/14%/35%/78% in the Two-Cost and 6%/13%/42%/74% in the Alternating One-Cost Treatment. The increase in the probability of investing in round 1, per ECU increase in the expected profit of investing in round 1, is 0.134 ($se = 0.017; p < 0.001$) in the Two-Cost Treatment and 0.125 ($se = 0.01; p < 0.001$) in the Alternating One-Cost Treatment. The estimating procedure used was a random effects probit.

The higher the expected payoff from investment, given a subject's type, the more likely she is to invest during some round. In particular, the frequencies with which subjects of type $(0, H)/(0, L)/(1, H)/(1, L)$ invest are 23%/44%/66%/95% in the Two-Cost and 21%/47%/69%/93% in the Alternating One-Cost Treatment.²⁴ The

²³There is no dependence between both subjects' decisions to invest conditional on any particular history, so we do not have to split the sample when we test behavior after any particular history.

²⁴These numbers are obtained from Table 2.4 by dividing the sum of the first six numbers in a column (the number of times that type invested) by the last number in the column (the total

increase in the probability of investing in some round, per ECU increase in the expected profit of investing in round 1, is 0.158 ($se = 0.016; p < 0.001$) for A subjects and 0.132 ($se = 0.014; p < 0.001$) for B subjects in the Two-Cost Treatment. The corresponding effect is 0.131 ($se = 0.011; p < 0.001$) for A subjects and 0.133 ($se = 0.011; p < 0.001$) for B subjects in the Alternating One-Cost Treatment (estimated via random effects probit).

The data suggest that subjects understand when investment is profitable and when it is not profitable, based on their signals. The fact that slightly over one in five type $(0, H)$ subjects eventually invests is somewhat high, because one should realize that even when the other subject has the high signal, the two signals would cancel and expected revenue is 5. However, nearly four in five type $(0, H)$ subjects get it right and never invest. It is quite impressive that less than 15% of type $(0, L)$ subjects invest in round 1, even though expected profits are only slightly negative.²⁵ Also impressive is that a type $(1, H)$ subject is far more likely to invest than a type $(0, L)$, even though the expected profits are only slightly positive.²⁶

We now move on to the question of whether subjects respond to the behavior of the other subject in their trial. We are interested in determining whether subjects were more likely to invest after seeing the other subject invest in round 1, as compared to seeing the other subject not invest in round 1. In the Two-Cost Treatment, the frequency with which subjects invest immediately after the history $\{1\}$ is 42%, and the frequency with which subjects invest immediately after the history $\{0\}$ is 16%. In the Alternating One-Cost Treatment, the corresponding frequencies are 50% and

number of times that type occurred).

²⁵Of course, the percentage of type $(0, L)$ subjects who eventually invest is higher, but much of that investment is following an investment by the other subject in the trial, which is consistent with equilibrium.

²⁶It is not a problem that the percentages of type $(1, H)$ subjects who invest in round 1 are only 35% in the Two-Cost Treatment and 42% in the Alternating One-Cost Treatment. Many subjects who do not invest understand that investment is profitable, but are waiting to obtain more information.

15%. Using a random effects probit, the marginal effect of the other subject investing in round 1, on the probability that a subject invests in round 2 (controlling for subjects' types) is 0.292 ($se = 0.051; p < 0.001$) in the Two-Cost Treatment and 0.306 ($se = 0.039; p < 0.001$) in the Alternating One-Cost Treatment.

Let us summarize our results about the ability of subjects to invest only when profitable and to draw inferences from the investment of the other subject.

Result 1 *In the aggregate, for both treatments, (i) types with higher expected profits are more likely to invest in round 1 and are more likely to invest eventually, and (ii) subjects are more likely to invest in round 2 after the other subject invests in round 1 than after the other subject does not invest in round 1. Both results are statistically significant.*

Now that we have established that behavior satisfies some basic requirements on rationality, let us turn to the question of how well it complies with Nash equilibrium (NE). Table 2.5 shows the percentage of decisions in each treatment which comply with NE, broken down by type and overall.²⁷

	(0,H)	(0,L)	(1,H)	(1,L)	Overall
Two-Cost	77%	60%	45%	78%	65%
Alternating One-Cost	79%	63%	82%	74%	75%

Table 2.5: Compliance with Nash Equilibrium

Compliance with NE for $(0, H)$ and $(1, L)$ types is high in both treatments. Note that these are the types for which there is no need for strategic interaction with the other subject. A type $(0, H)$ subject should not invest, even if she knew that the other subject's common-value signal were 1, and a type $(1, L)$ subject should invest, even if she knew that the other subject's common-value signal were 0. The

²⁷In computing these numbers, we treat any decisions as consistent with NE, following a history of no investment in round 1 and the other subject investing in round 2. See footnote 7.

rate of compliance with NE for $(0, L)$ and $(1, H)$ types is more interesting, since decisions are less clear cut, and these types must draw inferences from the other subject's behavior in the NE. The table shows that for type $(0, L)$ subjects in both treatments, and for type $(1, H)$ subjects in the Two-Cost Treatment, a substantial percentage of decisions deviates from the NE. Compliance with NE for type $(1, H)$ subjects in the Alternating One-Cost Treatment is high, but this is due to the fact that the NE involves mixing, so it is consistent with NE to invest in round 1, and also to wait and invest in round 2 if the other subject invests in round 1.

One might conjecture that behavior moves closer to NE in later trials. In the Two-Cost Treatment, this is indeed the case. Using random effects probit estimation, the marginal effect of the trial number on the probability that a decision is consistent with NE (controlling for subjects' types) is 0.0057 ($se = 0.0026$; $p = 0.03$). Thus, the probability of a subject playing her NE strategy increases by $24 \times 0.0057 = 0.137$ over the 24 trials, which is not a small effect. However, the same marginal effect in the Alternating One-Cost Treatment is actually negative (although insignificantly so), suggesting that there is no movement towards NE in that treatment.

Result 2 *When only one strategy is consistent with NE and the choice is not clear cut (i.e., types $(0, L)$ and $(1, H)$ in the Two-Cost Treatment and type $(0, L)$ in the Alternating One-Cost Treatment), then a substantial percentage of decisions (37%-55%) deviates from the NE. The probability that a decision is consistent with NE increases with time in the Two-Cost Treatment, but not in the Alternating One-Cost Treatment.*

Let us now turn to the question of how actual investment and profit outcomes compare with those in the NE.

Table 2.6 shows, for the Two-Cost and Alternating One-Cost Treatments, the actual frequency of investment as well as the ex ante expected frequency of investment

	Two-Cost					Alternating One-Cost				
	(0,H)	(0,L)	(1,H)	(1,L)	Overall	(0,H)	(0,L)	(1,H)	(1,L)	Overall
Actual	23%	44%	66%	95%	56%	21%	47%	69%	93%	57%
NE	0%	21%	29%	100%	38%	0%	42%	64%	100%	51%

Table 2.6: Investment

in the NE (broken down by type as well as overall). In both treatments, the actual frequency of investment exceeds the NE frequency for all types, except type $(1, L)$, where the actual frequency comes close to the NE frequency of 100%. In both treatments, the actual overall frequency of investment is higher than the NE frequency. This overinvestment is both economically (especially in the Two-Cost Treatment) and statistically significant ($p < 0.001$ for both A and B subjects in the Two-Cost Treatment and $p = 0.011/p = 0.003$ for A/B subjects in the Alternating One-Cost Treatment²⁸).

Result 3 *In both treatments, the actual overall frequency of investment is significantly higher than the NE frequency of investment. In the Two-Cost Treatment, overinvestment is especially pronounced (actual frequency of investment is 56% vs. 38% in the NE).*

	Two-Cost					Alternating One-Cost				
	(0,H)	(0,L)	(1,H)	(1,L)	Overall	(0,H)	(0,L)	(1,H)	(1,L)	Overall
Actual	-0.56	0.07	0.21	3.61	0.79	-0.34	0.03	0.44	3.05	0.79
NE	0	0.28	0.51	3.50	1.07	0	0.57	0.50	3.50	1.14

Table 2.7: Average Profits per Period (in ECU)

Table 2.7 shows, for the Two-Cost and Alternating One-Cost Treatments, the average actual profits per period as well as the NE expected profits per period (broken

²⁸These p-values are obtained in the following way. In order to account for possible dependence between investment decisions made by the same player, we perform a random effects probit estimation with only a constant as a right-hand-side variable. Then we test the hypothesis that the estimated constant is the same as the constant which leads to a predicted probability of investment equal to the NE ex ante probability of investment.

down by type as well as overall). In both treatments, actual average profits are lower than the NE expected profits. Due to the high variance of realized profits, this difference is not statistically significant at the 5% level ($p = 0.215/p = 0.092$ for A/B subjects in the Two-Cost Treatment and $p = 0.078/p = 0.064$ for A/B subjects in the Alternating One-Cost Treatment).²⁹ The reason that profits are lower than the NE prediction is primarily due to unprofitable investment by subjects with the low common-value signal, and the ensuing unprofitable investment by subjects drawing the wrong inference. Result 4 summarizes our findings about aggregate profits. Determinants of profits at the individual level will be considered later.

Result 4 *Actual average profits per trial are 73% of expected NE profits in the Two-Cost Treatment, and 69% of expected NE profits in the Alternating One-Cost Treatment.*

Central to our study is the informational interaction between subjects. A positive informational externality exists when subjects are likely to invest in round 1 with type $(1, H)$ or $(1, L)$, but unlikely to invest in round 1 with type $(0, H)$ or $(0, L)$.³⁰ Not surprisingly, this is indeed the case: in the Two-Cost Treatment, subjects with common-value signal 1 invest in round 1 55% of the time, and subjects with common-value signal 0 invest in round 1 only 12% of the time. In the Alternating Two-Cost Treatment, the corresponding percentages are 59% and 9%. In the NE, these percentages would be 50% and 0% in the Two-Cost Treatment, and 75% and 0% in the Alternating Two-Cost Treatment.³¹

²⁹These p-values are obtained in the following way. In order to account for possible dependence between profits earned by the same player in different periods, we perform a random effects regression with only a constant as a right-hand-side variable. Then we test the hypothesis that the estimated constant is the same as NE ex ante profits.

³⁰The informational externality created by investment after the history $\{0\}$ is more tricky, because this behavior is off the equilibrium path, and it is unclear that one should infer that the investor has the high common-value signal. This behavior is relatively rare in our experiment, and we ignore it in our analysis of the informational externality.

³¹Even in the NE, the informational externality is inefficiently low, because a type $(1, H)$ player

The difference between the actual informational externality and the NE prediction involves three effects. First, the difference is increased if type $(1, H)$ subjects invest in round 1 more often than in the NE. Second, the difference is decreased if type $(1, L)$ subjects fail to invest in round 1. Third, the difference is decreased if subjects with common-value signal 0 invest in round 1. The first effect is present in the Two-Cost Treatment, but not in the Alternating One-Cost Treatment. The second and third effects are present in both treatments (see tables 2.1, 2.2 and 2.4).

How do we combine these three effects to compare the actual informational externality with the NE prediction? To answer this question we compare the ex ante profits that a subject would receive, in the NE, with the ex ante profits that a subject would receive, based on best responding to the empirical frequencies of strategies chosen in each treatment. We will refer to this hypothetical best responder as a BR subject.³² Thus, we compare the long run profits earned by choosing the optimal strategy in a market where others play the NE strategy, and the long run profits earned by choosing the optimal strategy in a market where others play according to the empirical distribution of strategies chosen in our experiment. The ex ante profits in the NE and for a BR subject are 1.07 ECU in the Two-Cost Treatment. In the Alternating One-Cost Treatment, the ex ante profits in the NE are 1.14 ECU, and the profits for a BR subject are 1.08 ECU. The difference in profits for the Alternating One-Cost Treatment is not negligible, given that a subject who decides whether to invest based solely on her own information (and ignores any information provided by the other subject) earns expected profits of 1 ECU. Our findings about information flows are summarized in Result 5.

ignores the benefit that investing provides to the other player. Thus, it is possible that the actual play could be more informationally efficient than the NE play. We are currently running large market experiments ($n = 10$) to explore this possibility.

³²A BR player plays F (and invests after $\{0, 1\}$ when she is type $(1, H)$) in the Two-Cost Treatment and plays M (and invests after $\{0, 1\}$ when she is type $(0, L)$ or $(1, H)$) in the Alternating One-Cost Treatment.

Result 5 *Behavior in both the Two-Cost and Alternating One-Cost Treatments creates (in the aggregate) a positive informational externality. This externality is as large as the NE externality in the Two-Cost Treatment, and is smaller than the NE externality in the Alternating One-Cost Treatment (by 5%).*

Inspection of Tables 2.4, 2.6 and 2.7 indicates that behavior and outcomes in the Two-Cost and Alternating One-Cost Treatments are remarkably similar.³³ To quantify this similarity, we employ random effects probit estimation with treatment dummies as right hand side variables, to test the hypothesis that there is no treatment effect on (i) key history and type-contingent investment choices, (ii) aggregate investment, and (iii) aggregate profits. First, let us compare behavior across treatments for some key types, after the history, $\{\}$, and after the history, $\{1\}$. In the Two-Cost/Alternating One-Cost Treatment, type $(1, L)$ subjects invest in round 1 78%/74% of the time ($p = 0.373$). In the Two-Cost/Alternating One-Cost Treatment, type $(1, H)$ subjects invest in round 1 35%/42% of the time ($p = 0.397$). In the Two-Cost/Alternating One-Cost Treatment, type $(1, H)$ subjects invest after the history, $\{1\}$, 59%/62% of the time ($p = 1$). In the Two-Cost/Alternating One-Cost Treatment, type $(0, L)$ subjects invest after the history, $\{1\}$, 37%/53% of the time ($p = 0.305$).

As can be seen from table 2.6, frequency of investment is very similar between treatments: 56% vs. 57%. Testing the hypothesis that there is no treatment effect on the probability of investing (controlling for subjects' types) yields a p-value of 0.885/0.554 for A/B subjects (within a random effects probit model). From table 2.7, average profits per period are very similar between treatments: 0.79 ECU in both treatments. Testing the hypothesis that there is no treatment effect on profits (by regressing profits on a treatment dummy within a random effects model, controlling

³³Table 5 shows the compliance with NE, and differs across the two treatments. This is because the NE themselves differ across the treatments, not the actual behavior.

for subjects' types) yields a p-value of 0.813/0.982 for A/B subjects. Summarizing these comparisons across the two treatments, we can state the following result:

Result 6 *We cannot reject the hypothesis that there is no treatment effect on (i) key history and type-contingent investment choices, (ii) aggregate investment, and (iii) profits.*

2.5.2 What is Driving Behavior?

Given the substantial departure from NE behavior and the finding that there is no clear tendency towards NE behavior as the trials progress (Result 2), it remains to study what drives the actual behavior. Our task is simplified by the fact that there are essentially only three strategies that are consistent with level- k beliefs or cursed equilibrium: F , M and S . We say “essentially” because there are variations on these strategies, based on how decisions are made following a mistake, or when the other subject invests in round 2 after no one invests in round 1. Furthermore, these same three strategies can also be interpreted as rules of thumb, based on insights a subject might have about the value of observing the other subject’s behavior. We will proceed as follows. First, we consider a model for each treatment, in which subjects are drawn from a population of subjects who play one of the strategies F , M or S , fully specified off the equilibrium path. For $j \in \{F, M, S\}$, the probability of a subject being drawn from class j is denoted by p_j . At each decision node a subject faces, we assume that the subject chooses as dictated by her strategy class with probability $(1 - \varepsilon)$, and makes the other decision with probability ε . We then estimate, via maximum likelihood, the parameters $(p_F, p_M, p_S, \varepsilon)$.³⁴ Next, armed with our estimates of the proportions of the population in each strategy class for

³⁴This statistical framework is similar to that in many experimental papers, especially Camerer and Harless (1994) and Costa-Gomes, Crawford and Broseta (2001). The main difference is that, in our context, subjects are making a sequence of decisions in each game.

the Two-Cost Treatment, the Alternating One-Cost Treatment, and Treatment 3 (described below), we evaluate three behavioral theories: cursed equilibrium, level-k beliefs, and the insight-based rules of thumb described above.

Before proceeding to the maximum likelihood estimation, we must fully specify the strategies, F , M and S . For now, we specify the behavior of a subject playing according to her strategy class, and we will introduce errors later. The basic principle we use is that a subject corrects her own departures, and chooses each action with probability $\frac{1}{2}$ following an unexpected choice by the other subject (i.e. following histories $\{0, 1\}$ and $\{0, 1, 0\}$).³⁵ Consider first a subject of class F . When type $(1, L)$, she invests with probability 1 following all histories. When type $(1, H)$, she invests with probability 1 following the histories $\{1\}$ and $\{1, 0\}$, invests with probability $\frac{1}{2}$ following the histories $\{0, 1\}$ and $\{0, 1, 0\}$, and does not invest following the histories $\{\}$ and $\{0\}$. When type $(0, L)$, she invests with probability 1 following the histories $\{1\}$ and $\{1, 0\}$, invests with probability $\frac{1}{2}$ following the histories $\{0, 1\}$ and $\{0, 1, 0\}$, and does not invest following the histories $\{\}$ and $\{0\}$. When type $(0, H)$ she does not invest following all histories.

Consider next a subject of class M . When type $(1, L)$, she invests with probability 1 following all histories. When type $(1, H)$, she invests with probability 1 following the histories $\{\}$, $\{1\}$, and $\{1, 0\}$, invests with probability $\frac{1}{2}$ following the histories $\{0, 1\}$ and $\{0, 1, 0\}$, and does not invest following the history $\{0\}$. When type $(0, L)$, she invests with probability 1 following the histories $\{1\}$ and $\{1, 0\}$, invests with probability $\frac{1}{2}$ following the histories $\{0, 1\}$ and $\{0, 1, 0\}$, and does not invest following the histories $\{\}$ and $\{0\}$. When type $(0, H)$ she does not invest following all histories.

Finally, consider a subject of class S . When type $(1, L)$ or type $(1, H)$, she invests with probability 1 following all histories. When type $(0, L)$ or type $(0, H)$, she does not invest following all histories.

³⁵Other specifications yield similar results, because these departures are relatively rare.

Our model of behavior is that, each time a subject of class $j \in \{F, M, S\}$ has to make a decision to invest or not invest, given her type and the observed history of the other subject's behavior, she makes the decision prescribed by strategy j with probability $1 - \varepsilon$ and makes an "error" with probability $\varepsilon \in [0, 0.5]$. Errors are assumed to be i.i.d. across types and histories, trials, and subjects. Table 2.8 shows the probability that a class F subject invests, or makes a final decision never to invest, conditional on her type, conditional on the history, and conditional on the fact that the other subject's behavior allows the history to occur. To understand how the table is constructed, consider the following examples. The probability that an F subject of type $(0, L)$ invests after the history $\{1\}$ is $(1 - \varepsilon)^2$, since she acted according to her strategy class twice: by not investing in round 1 and then by investing after the other subject invests in round 1. The probability that an F subject of type $(0, L)$ invests after the history $\{1, 0\}$ is $\varepsilon(1 - \varepsilon)^2$, since she acted according to her class by not investing in round 1, then she made an "error" by not investing after $\{1\}$, and finally she acted according to her class by recovering from her error and investing after $\{1, 0\}$. The probability that an F subject of type $(0, L)$ ends up not investing after experiencing history $\{0\}$ is $(1 - \varepsilon)^2$, since she made two type-consistent decisions: she did not invest either after $\{\}$ or after $\{0\}$.

For M subjects and S subjects, we can construct tables analogous to Table 2.8. An M subject behaves in the same way as an F subject, except when her type is $(1, H)$. Thus, columns 1, 2, and 4 are as in Table 2.8, but the entries in column 3 should be: $(1 - \varepsilon), \varepsilon^2, \varepsilon(1 - \varepsilon), \frac{1}{2}\varepsilon(1 - \varepsilon), \varepsilon^2(1 - \varepsilon), \frac{1}{4}\varepsilon(1 - \varepsilon), \varepsilon(1 - \varepsilon), \varepsilon^3$, and $\frac{1}{4}\varepsilon(1 - \varepsilon)$. An S subject behaves in the same way as an F subject when her type is $(0, H)$ or $(1, L)$. Furthermore, her behavior when her type is $(0, L)/(1, H)$ is identical to her behavior when her type is $(0, H)/(1, L)$. Therefore, columns 1 and 2 are identical to column 1 in Table 2.8, and columns 3 and 4 are identical to column 4 in Table 2.8.

History	(0,H)	(0,L)	(1,H)	(1,L)
{}	ε	ε	ε	$(1 - \varepsilon)$
{0}	$\varepsilon(1 - \varepsilon)$	$\varepsilon(1 - \varepsilon)$	$\varepsilon(1 - \varepsilon)$	$\varepsilon(1 - \varepsilon)$
{1}	$\varepsilon(1 - \varepsilon)$	$(1 - \varepsilon)^2$	$(1 - \varepsilon)^2$	$\varepsilon(1 - \varepsilon)$
{0,1}	$\varepsilon(1 - \varepsilon)^2$	$\frac{1}{2}(1 - \varepsilon)^2$	$\frac{1}{2}(1 - \varepsilon)^2$	$\varepsilon^2(1 - \varepsilon)$
{1,0}	$\varepsilon(1 - \varepsilon)^2$	$\varepsilon(1 - \varepsilon)^2$	$\varepsilon(1 - \varepsilon)^2$	$\varepsilon^2(1 - \varepsilon)$
{0,1,0}	$\varepsilon(1 - \varepsilon)^3$	$\frac{1}{4}(1 - \varepsilon)^2$	$\frac{1}{4}(1 - \varepsilon)^2$	$\varepsilon^3(1 - \varepsilon)$
no {0}	$(1 - \varepsilon)^2$	$(1 - \varepsilon)^2$	$(1 - \varepsilon)^2$	ε^2
no {1,0}	$(1 - \varepsilon)^3$	$\varepsilon^2(1 - \varepsilon)$	$\varepsilon^2(1 - \varepsilon)$	ε^3
no {0,1,0}	$(1 - \varepsilon)^4$	$\frac{1}{4}(1 - \varepsilon)^2$	$\frac{1}{4}(1 - \varepsilon)^2$	ε^4

Table 2.8: Probability of an F Subject's Behavior

Before constructing the likelihood function, we need some more notation. We number all of a subject's trials by $t = 1, 2, \dots, 24$. Let B_i^t denote the full behavior of subject i during (i 's) trial t . By full behavior, we mean the round in which she invests, if at all. We formalize B_i^t as a four dimensional vector of zeros and ones. $B_i^t = (0, 0, 0, 0)$ signifies that the subject did not invest, the vector $B_i^t = (0, 0, 1, 0)$ signifies that she invested in round 3, and so on. Let $-i_t$ denote the subject matched with subject i during trial t , and let B_{-i}^t denote the full behavior of subject $-i_t$ during trial t .³⁶ Denote the behavior of subject i over all trials as B_i , where we have $B_i = (B_i^1, \dots, B_i^t, \dots, B_i^{24})$, and denote the behavior of all the subjects matched with subject i (during the trials they are matched with i) as B_{-i} , where we have $B_{-i} = (B_{-i}^1, \dots, B_{-i}^t, \dots, B_{-i}^{24})$. Finally, let $T_i^t \in \{(0, H), (0, L), (1, H), (1, L), -1\}$ denote subject i 's type during trial t , where type -1 means that subject i was sitting out or bankrupt, let $T_i = (T_i^1, \dots, T_i^t, \dots, T_i^{24})$, and let $T = (T_1, \dots, T_i, \dots, T_n)$.

³⁶Define $B_i^t = B_{-i}^t = (-1, -1, -1, -1)$ if i did not participate in trial t (either because she sat out or because she went bankrupt).

From Table 2.8, or the analogous tables corresponding to M subjects and S subjects, the probability of B_i^t is determined, given that the subject is of strategy class j , type T_i^t , and given that the other subject's behavior is B_{-i}^t .³⁷ We denote this probability, which also depends on the parameter ε , as

$$\Pr(B_i^t|j, T_i^t, B_{-i}^t; \varepsilon).$$

For example, suppose that in trial t , subject i is type $(1, H)$ and the other subject invests in round 1, $T_i^t = (1, H)$ and $B_{-i}^t = (1, 0, 0, 0)$. If subject i is an F subject, then the probability of $B_i^t = (1, 0, 0, 0)$ is ε (row 1, column 3 of Table 2.8), the probability of $B_i^t = (0, 1, 0, 0)$ is $(1 - \varepsilon)^2$ (row 3, column 3 of Table 2.8), the probability of $B_i^t = (0, 0, 1, 0)$ is $\varepsilon(1 - \varepsilon)^2$ (row 5, column 3 of Table 2.8), and the probability of $B_i^t = (0, 0, 0, 0)$ is $(1 - \varepsilon)\varepsilon^2$ (row 8, column 3 of Table 2.8). Given that the other subject invests in round 1, investing in round 4 is impossible, so the probability of $B_i^t = (0, 0, 0, 1)$ is zero.

The probability that subject i chooses behavior B_i , given that her strategy class is j , given her type realizations T_i , and given the behavior of the other subjects she faces is B_{-i} , is given by

$$\Pr(B_i|j, T_i, B_{-i}; \varepsilon) = \prod_{t=1}^{24} \Pr(B_i^t|j, T_i^t, B_{-i}^t; \varepsilon).$$

³⁷If a subject is sitting out trial t or has gone bankrupt, then $B_i^t = B_{-i}^t = (-1, -1, -1, -1)$ with probability one.

Thus, we can compute the probability that subject i chooses behavior B_i , given her type realizations T_i , and given that the behavior of the other subjects she faces is B_{-i} ,³⁸

$$\Pr(B_i|T_i, B_{-i}; p_F, p_M, p_S, \varepsilon) = \sum_{j \in \{F, M, S\}} p_j \Pr(B_i|j, T_i, B_{-i}; \varepsilon).$$

In the Appendix, we show that the likelihood function is given by³⁹

$$\Pr(B|T; p_F, p_M, p_S, \varepsilon) = \prod_{i=1}^n \Pr(B_i|T_i, B_{-i}; p_F, p_M, p_S, \varepsilon). \quad (2.1)$$

For each treatment, we estimate the vector of parameters $\theta = (p_F, p_M, p_S, \varepsilon)$, by maximizing the likelihood function (2.1).⁴⁰ Table 2.9 shows our estimates, along with estimated standard errors, for the Two-Cost and the Alternating One-Cost Treatments.⁴¹

As can be seen from Table 2.9, in both of our main treatments, the population frequency of class F is estimated to be more than one half; the population frequency of class S is estimated to be more than one third; and the population frequency of

³⁸ T_i and B_{-i} do not affect the probability of subject i being in class F , M , or S , except, conceivably, when $T_i^t = -1$ and $B_{-i}^t = (-1, -1, -1, -1)$ because a subject has gone bankrupt. The latter is not a concern in practice because bankruptcies were exceedingly rare, and because the probabilities of each strategy class (conditional on bankruptcy) would not change very much. The stronger inference is that the subject made many “errors.” We ignore this complication.

³⁹The likelihood function is also implicitly conditional on the realized matching of subjects.

⁴⁰This function is continuous and differentiable in θ for every B and T and has a strict maximum at the true θ . The parameter space is obviously compact. All other technical requirements (as given in theorems 13.1 and theorem 13.2 in Wooldridge (2001)) hold so that the ML estimator is consistent and asymptotically normal (for asymptotic normality the true θ also needs to be interior).

⁴¹In the Two-Cost Treatment, the estimate of p_M is on the boundary of the parameter space. We do not compute the standard error for this estimate since the standard error does not have the usual interpretation in terms of confidence intervals. The standard errors for the elements of θ which are not on the boundary are approximate, since they are computed by estimating a restricted model in which p_M is set equal to 0. Estimated standard errors in the Alternating One-Cost Treatment should be treated with caution, since the estimate of p_M is only 1.58 (rather than at least 1.96) estimated standard errors from the boundary of the parameter space.

	Two-Cost	Alternating One-Cost
p_F	0.607 (0.102)	0.567 (0.083)
p_M	0 N/A	0.092 (0.058)
p_S	0.393 (0.102)	0.341 (0.078)
ε	0.192 (0.011)	0.176 (0.01)
<i>log-likelihood</i>	-757.5402	-1002.5052

Table 2.9: Maximum Likelihood Estimates

class M is estimated to be very small (0% in the Two-Cost Treatment).⁴² Error rates are not very high, and parameters are nearly the same across the two treatments.

We wish to demonstrate that the population contains both class F and class S subjects, but formal testing is complicated by boundary issues and the possibility that test statistics are not asymptotically normal. However, we are able to test, using a likelihood-ratio test, the hypotheses that (i) $p_F = 0.25$ in the Two-Cost Treatment, (ii) $p_S = 0.2$ in the Two-Cost Treatment, (iii) $p_F = 0.25$ in the Alternating One-Cost Treatment, and (iv) $p_S = 0.16$ in the Alternating One-Cost Treatment.⁴³ The corresponding p-values are 0.001 or less for (i), (iii) and (iv); the p-value for (ii) is 0.071. Note that the p-value for (ii) would be lower if we could test the hypothesis that $p_S = 0$ in the Two-Cost Treatment. The estimates of the parameter vector θ are very similar in both treatments. We cannot reject (using a likelihood-ratio test) any of the hypotheses that $p_F/p_M/p_S/\varepsilon/p_F \& p_M \& p_S/p_F \& p_M \& p_S \& \varepsilon$ are equal in both treatments (p=0.754/0.341/0.675/0.235/0.635/0.532). Let us summarize:

⁴²Our estimates are robust to the specification of errors following an error. We estimated a model in which, after a subject makes a mistake, she chooses to invest or not invest in each subsequent round with probability one half. The estimate of θ in the Alternating One-Cost Treatment changed only negligibly. In the Two-Cost Treatment, the estimate of p_F decreased by 0.085, and the estimate of p_M increased by 0.073.

⁴³These values are chosen so that they are at least 1.96 (estimated) standard errors from the boundary of the parameter space.

Result 7 (i) *More than half the population in each treatment is estimated to be F ; slightly more than one third is estimated to be S ; and only a small minority is estimated to be M . The estimates of p_F and p_S are statistically different from 0.*⁴⁴

(ii) *The estimates of θ are very similar across treatments. Any differences are statistically insignificant.*

How well do the various behavioral theories explain our estimation results? First consider the cursed equilibrium framework. There are several reasons why symmetric cursed equilibrium (or NE, which is a special case) is probably not driving our maximum likelihood estimates.⁴⁵ Referring back to Table 2.3, we see that no level of the cursedness parameter, χ , can explain the simultaneous presence of F and S subjects. Thus, symmetric cursed equilibrium is inconsistent with result 7. Also, in the Alternating One-Cost Treatment, the strategy, M , is played with probability at least 0.4916, for any level of the cursedness parameter in which F is also played. This restriction is inconsistent with our estimate of p_M , which is significantly different from 0.496 ($p < .01$). Furthermore, the low level of χ required to explain the high frequency of F requires different behavior across treatments. In particular, for fixed small χ we would expect a higher fraction of M subjects in the Alternating One-Cost Treatment.⁴⁶

Next, consider the level-k framework. Recall that, in both treatments, L_1 plays S and L_2 plays F . According to our estimates, the majority of the population is indeed F or S . The estimates of p_F and p_S are nearly the same across treatments, which is

⁴⁴At the 10% level in the case of the estimate of p_S in the Two-Cost Treatment.

⁴⁵We consider here a common cursedness parameter for all subjects. When we introduce our estimates for Treatment 3, we will discuss the possibility of heterogeneity.

⁴⁶We performed separate maximum likelihood estimations of χ , for a model in which subjects play a symmetric cursed equilibrium with errors. In particular, each time a subject has to make a decision, she chooses the strategy prescribed by the cursed equilibrium with probability $1 - \varepsilon$ and makes an error with probability $\varepsilon \in [0, 0.5]$. The estimate of χ is 0 in the Alternating One-Cost Treatment, and 0.282 in the Two-Cost Treatment. The difference across treatments is troubling, and the low levels suggests that cursed equilibrium adds little explanatory power over NE.

what one would expect if the proportions of L_1 and L_2 players in the population are stable across the treatments. Therefore, our estimates based on the Two-Cost and Alternating One-Cost Treatments are consistent with the level-k framework.

Finally, consider the framework in which each subject uses a rule of thumb prescribing either F , M or S . Because the Two-Cost and Alternating One-Cost Treatments are essentially the same, in terms of the nature and difficulty of the insights discussed above, one would expect to see nearly identical behavior across the two treatments. This is indeed the case. Therefore, our estimates based on the Two-Cost and Alternating One-Cost Treatments are consistent with the framework in which subjects use a rule of thumb, based on the various insights discussed above.

2.5.3 Treatment 3

To distinguish better between the possible explanations of behavior, and especially between the belief-based theories and insight-based rules of thumb, we conducted an additional treatment (one session, 28 participants). This treatment, which we call Treatment 3, is almost the same as the Alternating One-Cost Treatment, with the sole differences being that the high investment cost is 5.7, rather than 6.5, and the discount factor is 0.8, rather than 0.9.⁴⁷ With the new parameters, a type $(1, H)$ subject has a dominant strategy to invest in round 1. The expected profits from investing in round 1 are greater than the expected profits of waiting, even if waiting would reveal the other subject's type. Therefore, the strategy F is never the optimal strategy for a risk-neutral subject within the expected utility framework, regardless of her beliefs. If behavior in our Two-Cost and Alternating One-Cost Treatments is driven by cursed equilibrium (even allowing each subject to have her own separate χ), level-k beliefs, or some other belief-based behavioral theory such as Quantal

⁴⁷To maintain roughly the same expected earnings as before, the exchange rate was changed to \$0.50/EUCU.

	Two-Cost	Alternating One-Cost	Treatment 3
p_F	0.607 (0.102)	0.567 (0.083)	0.574 (0.115)
p_M	0 N/A	0.092 (0.058)	0.287 (0.118)
p_S	0.393 (0.102)	0.341 (0.078)	0.139 (0.093)
ε	0.192 (0.011)	0.176 (0.01)	0.198 (0.013)
<i>log-likelihood</i>	-757.5402	-1002.5052	-600.378

Table 2.10: Maximum Likelihood Estimates - All Treatments

Response Equilibrium (see McKelvey and Palfrey (1995) and (1998)), the resulting maximum likelihood estimation for Treatment 3 should show a collapse of p_F .⁴⁸

Suppose behavior is driven by rules of thumb. Some subjects play S because they do not acquire the insight that investment by the other subject is good news for them. Some subjects play M because either (i) they acquire the insight that investment by the other subject is good news, but not the insight that there is a tradeoff between the cost of waiting and the information gained by waiting, or (ii) they acquire both of the above insights, but simply resolve the tradeoff in favor of investing in round 1. Some subjects play F , because they acquire both of the above insights, but resolve the tradeoff in favor of gathering information by waiting. The new parameters in Treatment 3 should not affect the difficulty of acquiring these insights. Therefore, one would expect to see a significant proportion of subjects who continue to play F , because they do not explicitly perform the computation to determine that F is dominated.

Table 2.10 shows the estimates of θ for all three treatments along with estimated standard errors.⁴⁹ Amazingly, the estimate of p_F in Treatment 3 is very similar to

⁴⁸The symmetric cursed equilibrium is to play M for $\chi < 0.75$, and S for $\chi > 0.75$. In any cursed equilibrium, symmetric or not, F is never played. For all $k > 0$, a subject with level- k beliefs plays M .

⁴⁹Estimated standard errors in Treatment 3 should be treated with caution, since the estimate of

	Two-Cost & Alternating One-Cost	Two-Cost & Treatment 3	Alternating One-Cost & Treatment 3	All three treatments
p_F	0.754	0.824	0.959	0.95
p_M	0.341	0.04	0.098	0.086
p_S	0.675	0.074	0.089	0.143
ε	0.235	0.714	0.139	0.268
p_F, p_M, p_S	0.635	0.082	0.124	0.173
$p_F, p_M, p_S, \varepsilon$	0.532	0.159	0.091	0.17

Table 2.11: Hypotheses Tests - p-values.

that in the initial two treatments. The point estimate is actually slightly higher than in the Alternating One-Cost Treatment. The estimate of p_M is higher than in the initial two treatments and the estimate of p_S is correspondingly lower.⁵⁰ Using likelihood-ratio tests, we can test hypotheses about whether certain elements of θ are the same in different pairs of treatments, as well as across all three treatments. Entry (i, j) in Table 2.11 shows the resulting p-value, under the null hypothesis that the parameters in row i are restricted to be the same across all treatments in column j .

Any belief-based theory would imply large differences across treatments, and in particular, no F subjects in Treatment 3. Table 2.10 shows that the estimates of p_F are virtually indistinguishable across treatments (high p-values), and we can reject the hypothesis that $p_F = 0.25$. This supports the view that behavior in our experiment is, to a large extent, driven by boundedly rational rules of thumb, rather than explicit belief formation about the behavior of the other subject. Moreover, among type $(1, H)$ subjects who perceive a tradeoff between the costs and benefits of waiting, the proportion deciding to wait does not change as we vary the high-cost parameter.

p_M is 1.49 (rather than at least 1.96) estimated standard errors from the boundary of the parameter space. See footnote 41.

⁵⁰We performed the same robustness check that we did for the initial two treatments (see footnote 42), where we assume that after a subject makes a mistake, she chooses to invest or not invest in each subsequent round with probability one half. The estimate of θ changed only negligibly.

Result 8 *Any differences in the estimates of p_F in all three treatments are both economically and statistically insignificant. In Treatment 3, we can reject the hypothesis that p_F is as low as 0.25.*

Overall, there is no strong evidence of statistically significant differences in behavior across treatments. The p-values below 0.10 are due to higher estimates of p_M and lower estimates of p_S in Treatment 3. These higher estimates for p_M and lower estimates for p_S are primarily driven by the fact that type $(0, L)$ subjects in Treatment 3 invested 70% of the time after history $\{1\}$, versus 37% in the Two-Cost Treatment and 53% in the Alternating One-Cost Treatment. These differences seem anomalous to us. It is particularly hard to see why changing the high investment cost would affect behavior when the investment cost is low (Alternating One-Cost Treatment vs. Treatment 3), so we do not attach much significance to the different estimates of p_M and p_S .

Our maximum likelihood estimation assumes that the strategy classes F , M , and S are drawn from the population at the beginning of the experiment, and do not evolve as the trials progress. This specification was made for simplicity, and because learning issues are not our main focus. However, there seems to be some learning going on, which sheds light on our interpretation of behavior as rules of thumb. Random effects probit estimation is used to study the effect of the trial number on the probability that a type $(1, H)$ subject invests in round 1. In the Two-Cost Treatment, the marginal effect is -0.018 (standard error = 0.006, $p = 0.002$). In the Alternating One-Cost Treatment, the marginal effect is -0.018 (standard error = 0.005, $p = 0.001$). In Treatment 3, the marginal effect is 0.001 (standard error = 0.007, $p = 0.89$). The negative marginal effect in the Two-Cost and Alternating One-Cost Treatments indicates that subjects are “learning” to wait and observe the

behavior of the other subject.⁵¹ This learning moves behavior towards the NE in the Two-Cost Treatment, but moves behavior beyond and away from NE in the Alternating One-Cost Treatment. This learning could be due to a “Eureka” effect, where some subjects suddenly acquire the insight that waiting gives them useful information about the other subject. Why is learning absent in Treatment 3 (the estimated marginal effect is insignificant and of the wrong sign)? Perhaps while some subjects acquire the insight that waiting gives them useful information, other subjects acquire the additional insight that the benefits are not adequate to compensate for the discounting.⁵² These two opposing effects may offset each other.

Treatment 3 addresses the potential criticism that the incentives of a type $(1, H)$ subject are weak, so that drawing conclusions about behavior is problematic. In the Two-Cost and Alternating One-Cost Treatments, the differences in the ex ante expected payoff of playing F , M , and S are quite small (both in the NE and given the empirical frequencies of play).⁵³ In Treatment 3, the expected payoff of playing $F/M/S$, given NE beliefs, is 1.32/1.326/1 ECU. However, the expected payoff of playing $F/M/S$, given the empirical frequencies is 1.141/1.29/1 ECU. Therefore, the advantage of M over F in Treatment 3 is quite substantial. Over 24 trials, the expected profit gain of playing M rather than F is 3.576 ECU or \$1.79.

⁵¹This effect is large in the Two-Cost and Alternating One-Cost Treatments. The predicted probability of investment in round 1 by a type $(1, H)$ player is higher in trial 1 than trial 24, by about 0.43.

⁵²After all, even if the other subject is revealed to have the low common-value signal, expected losses are only 0.7 ECU in Treatment 3, while it is 1.5 ECU in the other treatments. Also the other subject may choose to wait with the high common-value signal, thereby weakening the inference.

⁵³The expected payoff of playing $F/M/S$, given the empirical frequencies of play, is 1.071/1.05/1 ECU in the Two-Cost Treatment and 1.084/1.085/1 ECU in the Alternating One-Cost Treatment. The expected payoff of playing $F/M/S$, given NE beliefs, is 1.073/1.071/1 ECU in the Two-Cost Treatment and 1.142/1.142/1 ECU in the Alternating One-Cost Treatment.

GPA	SAT	π_{BR}	$D_{\text{Two-Cost}}$	D_{Alt}	$D_{\text{Treatment 3}}$
0.1032** (0.0438)	- -	0.8337*** (0.0113)	-0.4011*** (0.1408)	-0.3336** (0.141)	-0.4176*** (0.1452)
-0.03 (0.072)	0.0008139*** (0.0002856)	0.8402*** (0.0148)	-0.9679*** (0.2888)	-0.8884*** (0.2969)	-0.9311*** (0.3003)

Table 2.12: Regression Results for Earnings (*/**/*** indicates significance at the 10%/5%/1% level.)

2.5.4 Personal Characteristics as Determinants of Behavior

In this section we investigate whether subject's personal characteristics (GPA and SATs in particular), affect their earnings in the experiment.

Table 2.12 shows the results of two random effects regressions. The first one regresses earnings in each period on a subject's GPA, the earnings a subject would have made had she played the best response to the empirical frequencies (π_{BR}), and treatment dummies. The second regression is the same as the first with the sole difference that it also includes subjects' SAT scores as an explanatory variable. The earnings a subject would have made had she played the best response to the empirical frequencies are included in order to eliminate any noise in earnings due to a favorable combination of a subject's cost signal and behavior of the other subject.⁵⁴

The first regression suggests that a subject's GPA is positively correlated with her earnings. However, this effect is wiped out when one includes SATs in the regression. The effect of a higher SATs score on a subject's earnings comes out strongly significant. This effect is also economically significant: it implies a difference in earnings per period of \$0.244 between a subject who is one standard deviation above the mean SAT score and a subject who is one standard deviation below it. This translates into \$5.849 over 24 periods.

⁵⁴We also used another, more direct, method to control for these factors. In particular, we included a dummy variable for each possible combination of a player's cost, signal, whether the other person invested in round 1 or not and the treatment she is in (48 dummies in total). This yields very similar results regarding the effect of GPA and SATs.

2.6 Concluding Remarks

To summarize our main results, we find that subjects are more likely to invest as their signals become more favorable, even for the subtle comparison between type $(0, L)$ and type $(1, H)$. Subjects overinvest relative to the Nash benchmark. To the extent that subjects with the low common-value signal invest in round 1, a negative informational externality is created, and to the extent that subjects with the high common-value signal invest in round 1, a positive informational externality is created. When we compare behavior with the theoretical predictions in the Two-Cost Treatment, the negative externality is balanced by a “theoretically excessive” positive externality, so a subject best responding to actual play receives the same profit that would be received if everyone were playing Nash. In the Alternating One-Cost Treatment, the positive externality is no longer excessive, so best responding to actual play yields lower profits than what would be received if everyone were playing Nash. This difference across treatments is due entirely to difference in the theoretical predictions, because we cannot reject the hypothesis that there is no treatment effect on behavior or profits.

Maximum likelihood estimates for our two main treatments indicate strong evidence of both S and F subjects in the population, which is inconsistent with symmetric cursed equilibrium. Level- k beliefs can account for these estimates, due to the flexibility to allow for subject heterogeneity.⁵⁵ We can also account for these estimates if subjects choose rules of thumb, based on insights about how to understand the game (hindsight and foresight).

⁵⁵We feel that a better comparison would be to some notion of asymmetric cursed equilibrium that allows for heterogeneous subjects. We did not go to the considerable trouble of considering such a concept, because Treatment 3 would rule it out in any case.

To separate these explanations, we introduce Treatment 3, in which F is strictly dominated and is inconsistent with any theory of best responding to beliefs. We find that the proportion of F subjects does not decline significantly, and remains above 50%.

In conclusion, we do not want to discredit the belief-based theories. Cursed equilibrium formalizes the notion that subjects do not fully draw inferences about others' types from their behavior. Level- k beliefs allow for heterogeneous beliefs about how sophisticated the other players are, but requires best responding to those beliefs. These theories provide important generalizations to Bayesian Nash equilibrium, and can account for many behavioral anomalies. In our context, other forms of bounded rationality seem to do a better job of explaining the data.

BIBLIOGRAPHY

- [1] Anderson, L. R. and C. A. Holt (1997), "Information Cascades in the Laboratory," *American Economic Review*, 87, 847-862.
- [2] Ball, S. B., M.H. Bazerman, and J.S. Carroll (1991), "An Evaluation of Learning in the Bilateral Winner's Curse," *Organizational Behavior and Human Decision Processes*, 48, 1-22
- [3] Banerjee, A. V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107 (3), August 1992, 797-817.
- [4] Bazerman, M.H. and W.F. Samuelson (1983), "I Won the Auction But Don't Want the Prize," *Journal of Conflict Resolution*, 27, 618-634.
- [5] Beach, L.R. and Philips, L.D. (1967). "Subjective Probabilities Inferred from Estimates and Bets," *Journal of Experimental Psychology* 75, 354-359
- [6] Becker, G.M.; DeGroot, M.H.; Marschak, J. (1964). "Measuring Utility by a Single-Response Sequential Method," *Behav. Sci.* 9, 226-32
- [7] Benjamin, Daniel J. and Shapiro, Jesse M. (2006) "Who is "Behavioral"? Cognitive Ability and Anomalous Preferences," working paper, available at <http://home.uchicago.edu/jmshapir/>
- [8] Bikhchandani, S., D. Hirshleifer, and I. Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, Vol. 100 (1992), pp. 992-1026.
- [9] Blecherman, B. and C.F. Camerer (1998), "Is There a Winner's Curse in the Market for Baseball Players?" mimeograph, Brooklyn Polytechnic University, Brooklyn, NY.
- [10] Camerer, Colin F. (2003). "Behavioral Game Theory: Experiments in Strategic Interaction," Princeton University Press
- [11] Camerer, C. F. and D. W. Harless (1994), "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, Vol. 62, 1251-1289.

- [12] Capen, E.C., R.V. Clapp, and W.M. Campbell (1971), "Competitive Bidding in High-Risk Situations," *Journal of Petroleum Technology*, 23, 641-653.
- [13] Cassing, James and Richard W. Douglas (1980), "Implications of the Auction Mechanism in Baseball's Free agent Draft," *Southern Economic Review*, 47, 110-121.
- [14] Chamley, C. (2003), "Rational Herds: Economic Models of Social Learning," Cambridge University Press
- [15] Chamley, C. and D. Gale (1994), "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica*, Volume 62, Issue 5, 1065-1085.
- [16] Charness, Gary and dan Levin, 2005, "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Psychological Affect", *American Economic Review*, Vol. 95, No. 4 (September 2005), 1300-1309.
- [17] Costa-Gomes, M., V. Crawford, and B. Brosetta (2001), "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, Vol. 69, 1193-1235.
- [18] Costa-Gomes, Miguel and Weizsäcker, Georg (2005). "Stated Beliefs and Play in Normal-Form Games," working paper, available at <http://personal.lse.ac.uk/weizsack/>
- [19] Crawford, V. P. and N. Iriberri (2005), "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" UCSD Working Paper.
- [20] Dessauer, John P., (1981) Book Publishing, New York: Bowker.
- [21] Eyster, E. and M. Rabin (2005), "Cursed Equilibrium," *Econometrica*, Vol. 73, No. 5, 1623-1672.
- [22] Fischbacher, Urs (1999). "z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual," Working Paper Nr. 21, Institute for Empirical Research in Economics, University of Zurich.
- [23] Fox, Craig R. and Tversky, Amos (1995). "Ambiguity Aversion and Comparative Ignorance," *The Quarterly Journal of Economics* 110, 585-603
- [24] Friedman, Daniel and Massaro, Dominic (1998). "Understanding Variability in Binary and Continuous Choice," *Psychonomic Bulletin and Review* 5, 370-389
- [25] Garratt, R. and T. Keister (2006), "Bank Runs as Coordination Failures: An Experimental Study," UCSB and ITAM Working Paper.

- [26] Gilboa, Itzhak and Schmeidler, David (1989). "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics* 18, 141-153
- [27] Goeree, Jacob K.; Palfrey, Thomas R.; Rogers, Brian W. and McKelvey, Richard D. (2006) "Self-correcting Information Cascades," Working paper
- [28] Heath, Chip and Tversky, Amos (1991). "Preference and Belief: Ambiguity Aversion and Competence in Choice under Uncertainty," *Journal of Risk and Uncertainty* 4, 5-28
- [29] Hendricks, K., R. Porter, and B. Boudreau (1987), "Information and Returns in OCS Auctions, 1954-1969," *Journal of Industrial Economics*, 35, 517-42.
- [30] Holt, Charles A. and Laury, Susan K. (2002). "Risk Aversion and Incentive Effects," *American Economic Review* 92, 1644-1655
- [31] Holt, C.A. Jr., and R. Sherman (1994), "The Loser's Curse and Bidder's Bias," *American Economic Review*, 84, 642-652.
- [32] Jensen, F.A. and Peterson, C.R. (1973). "Psychological Effects of Proper Scoring Rules," *Organizational Behavior and Human Performances* 9, 307-317
- [33] Kagel, J. H. and D. Levin (1986), "The Winner's Curse and Public Information in Common Value Auctions," *American Economic Review*, 76, 894-920.
- [34] Kagel, John H. and Levin, Dan (1991), "The Winner's Curse and Public Information in Common Value Auctions: Reply," *American Economic Review*, 81, 362-69.
- [35] Kagel, John H. and Levin, Dan (1999), "Common Value Auctions with Insider Information," *Econometrica*, 67, 1219-1238.
- [36] Kagel, John H. and Levin, Dan (2002), *Common Value Auctions and the Winner's Curse*, Princeton: Princeton University Press.
- [37] Kübler, Dorothea and Weizsäcker, Georg (2004), "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, v. 71, iss. 2, pp. 425-41
- [38] Levin D., Kagel J. H. and J. F. Richard (1996), "Revenue Effects and Information Processing in English Common Value Auctions" *American Economic Review*, 86, 442-460.
- [39] Levin, D. and J. Peck (2006), "Investment Dynamics with Common and Private Values," Ohio State University Working Paper.

- [40] Levis, Mario (1990), "The Winner's Curse Problem, Interest Costs and the Underpricing of Initial Public Offerings," *Economic Journal*, 100, 76-89.
- [41] McKelvey, R. D. and T. R. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10, 6-38.
- [42] McKelvey, R. D. and T. R. Palfrey (1998), "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 9-41.
- [43] Mead, W. J., A. Moseidjord, and P.E. Sorensen (1983), "The Rate of Return Earned by Leases Under Cash Bonus Bidding in OCS Oil and Gas Leases," *Energy Journal*, 4, 37-52.
- [44] Mead, Walter J.; Moseidjord, Asbjorn and Sorensen, Philip E. (1984), "Competitive Bidding Under Asymmetrical Information: Behavior and Performance in Gulf of Mexico Drainage Lease Sales, 1954-1969," *Review of Economics and Statistics*, 66, 505-08.
- [45] Nagel, R. (1995), "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313-1326.
- [46] Nelson, Robert .G and Bessler, David A. (1989). "Subjective Probabilities and Scoring Rules: Experimental Evidence," *American Journal of Agricultural Economics* 71, 363-369
- [47] Rey Biel, Pedro (2005). "Equilibrium Play and Best Response to (Stated) Beliefs in Constant Sum Games," working paper, available at <http://pareto.uab.es/prey/Research.htm>
- [48] Rock, Kevin (1986), "Why New Issues are Underpriced," *Journal of Financial Economics*, 15, 187-212.
- [49] Roll, Richard (1986), "The Hubris Hypothesis of Corporate Takeovers," *Journal of Business*, 59, 197-216.
- [50] Roth, Alvin and Malouf, Michael (1979). "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review* 86, 574-594
- [51] Rubinstein, A. (1998), *Modeling Bounded Rationality*, The MIT Press (Cambridge, MA and London, England).
- [52] Samuelson, W.F. and M.H. Bazerman (1985), "The winner's curse in bilateral negotiations," In V. Smith (Ed.), *Research in Experimental Economics* (Vol. 3), Greenwich, CT: JAI Press.
- [53] Schmeidler, David (1989), "Subjective Probability and Expected Utility without Additivity," *Econometrica* 57, 571-587

- [54] Sgroi, D. (2003), “The Right Choice at the Right Time: A Herding Experiment in Endogenous Time,” *Experimental Economics* 6, 159-180.
- [55] Simon, H. A. (1972), “Theories of Bounded Rationality,” in C. B. McGuire and R. Radner, eds., *Decisions and Organization*, 161-176, Amsterdam: North Holland.
- [56] Sonnemans, Joep and Offerman, Theo (2001). “Is the Quadratic Scoring Rule Incentive Compatible?”, working paper, available at <http://www1.fee.uva.nl/creed/people/offerman/index.shtml>
- [57] Stahl, Dale and Wilson, Paul (1995). “On Players’ Models of Other Players: Theory and Experimental Evidence,” *Games and Economic Behavior* 10, 218-254
- [58] Wooldridge, Jeffrey (2001). “Econometric Analysis of Cross Section and Panel Data,” The MIT Press; 1st edition
- [59] Ziegelmeyer, A., K. B. My, J.-C. Vergnaud, and M. Willinger (2005), “Strategic Delay and Rational Imitation in the Laboratory,” Discussion Paper on Strategic Interaction 35-2005, Max Planck Institute of Economics, Jena, Germany.

APPENDIX A: FIGURES AND TABLES FROM CH. 1

Game 1

0.13	0.45	0.41	
0.19	0.43	0.38	
94;29	34;28	25;69	16 12 13 <i>NRN,NRA</i>
89;35	22;99	36;62	5 4 3
13;39	67;50	50;43	27 22 21 <i>SRN,SRA</i>

-

Game 3

0.08	0.36	0.55	
0.11	0.32	0.56	
50;11	12;95	21;95	0 0 0
96;10	10;61	89;72	23 25 20 <i>NRN,SRN</i>
51;68	48;57	50;48	25 13 17 <i>NRA,SRA</i>

-

Game 5

0.38	0.33	0.29	
0.49	0.26	0.25	
47;25	52;56	46;74	13 7 7 <i>NRA,SRA</i>
62;76	96;26	12;50	23 29 20 <i>NRN,SRN</i>
50;54	46;72	47;29	12 2 10

-

Game 7

0.46	0.28	0.26	
0.47	0.28	0.26	
17;51	10;51	21;50	0 0 0
49;51	48;51	50;51	26 10 21 <i>NRA,SRA</i>
99;50	57;50	10;51	22 28 16 <i>NRN,SRN</i>

-

Game 9

0.35	0.31	0.33	
0.40	0.31	0.30	
62;85	96;35	10;59	19 19 17 <i>NRN,SRN</i>
52;34	48;65	50;83	21 11 13 <i>NRA,SRA</i>
10;63	64;81	91;38	8 8 7

Game 2

0.40	0.51	0.09	
0.43	0.44	0.13	
32;72	30;60	94;12	10 5 5 <i>NRN,NRA</i>
61;59	59;69	11;14	35 33 32 <i>SRN,SRA</i>
40;52	36;55	10;74	3 0 0

Game 4

0.18	0.42	0.41	
0.16	0.42	0.43	
89;56	31;48	25;47	18 9 6 <i>NRN,NRA</i>
14;17	58;75	50;96	26 27 26 <i>SRN,SRA</i>
16;12	44;57	56;35	4 2 5

Game 6

0.35	0.58	0.07	
0.34	0.58	0.08	
12;66	99;60	12;15	23 21 16 <i>SRN</i>
48;72	50;78	15;17	21 15 20 <i>SRA</i>
26;37	30;37	89;45	4 2 1 <i>NRN,NRA</i>

Game 8

0.31	0.55	0.15	
0.32	0.45	0.23	
10;99	14;99	12;10	0 2 0
11;60	99;55	99;67	22 26 24 <i>NRN,SRN</i>
49;55	55;60	51;10	26 10 13 <i>NRA,SRA</i>

Game 10

0.06	0.22	0.71	
0.06	0.18	0.76	
12;14	40;45	78;75	45 31 33 <i>SRN,SRA</i>
98;27	62;28	10;65	1 5 1 <i>NRN</i>
47;19	46;59	44;71	2 2 3 <i>NRA</i>

Figure A.1: Row players' mean beliefs (first (treatment A) and second (treatment B) row above each game) and aggregate choices (first (treatment A), second (treatment B) and third (treatment C) column to the right of each game) along with each type's (tentative) preferred choice.




Figure A.2: Marginal posteriors of $p_{NRN} + p_{NRA}$ (solid line) and $p_{NRN} + p_{SRN}$ (dashed line) in treatments A (top), B (middle) and C (bottom).






Figure A.3: Marginal posteriors of μ (solid line) and σ (dashed line) in treatments A (top), B (middle) and C (bottom).

Table A.13: Responses of Subjects in Treatment A to Questionnaire Question: “What were your key considerations when making decisions in part I?”

Subject	Response
(1)	I alternated between a maximin strategy and a strategy of predicting my opponent's choice and using that to my advantage. When I felt that my opponent had a fairly clear choice, and the risks for me were not too large, I deviated from a maximin strategy. Otherwise I used that strategy.
2	I first consider how my partner would choose, then I choose the answer that would most likely to give me more ECU.
3	the possible ECP and what the other participant will choose
4	i was trying to figure out what option my partner will choose. based on this judgment i choose my own option that will yield the most amount of points.
5	I attempted to determine what column the other person would select prior to making my selection. This also made section II a lot easier to calculate.
6	How much money I could make; blocking other party from making more, when my options were poor
7	The ECU I could get and the decision my partener probably made.
8	HOW THE OTHER PART WILL CHOOSE, AND THEN MAKE MY DECISION TO MAXIMIZE MY PROFIT
(9)	Which combinations of numbers have the highest sum and would be most likely to be selected by the other person.
10	Delete the one that he/she would not choose and select one of two rows I like.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
11	compare my decisions with the other person
12	I chose among two different algorithms. In the first one, I considered the minimax of the table, the option that maximizes his gain given a particular choice. If there was a significant gain by choosing this algorithm, I proceeded with the best gain, which gave me his best gain. If it was not satisfactory, I considered the option which is good for both of us (and not best for each probably).
(13)	i tried to look for the tickets that would pay off the most for both of us. this usually narrowed down the ticket choices to 2 or 3 tickets. then i decided my partner would probably choose the ticket of those three that benefits them the most. at first i chose this ticket if it would also benefit me. then towards the end of the experiment i began choosing the highest total for me that would occur in that row.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
(14)	Trying to be as cooperative as possible with my co-participant. If we could find a win-win situation where the discrepancy between what I would make and s/he would make was small, I would generally choose that option. If the discrepancy between what I could make and what s/he would make was large, I would be somewhat riskier with my decision-making if I was fairly confident that I would make about 40ECU (or close to that figure). Before looking at my own data (I was a row participant), I would always look at the values for the column participant and pinpoint the highest value. If my value was 40 or higher, I would consider it as an option. If not, I would automatically discard it.
15	First, look at the opponent's options and determine if there might be a clear favorable or unfavorable choice. This usually narrowed my probable options. Then I tended to take the less risky of my remaining options. I didn't go for the big prize if there was a chance of a very small one. If all options for the opponent were about equal, I assumed s/he was bright enough to assume I'd go for the biggest payoff possible.
16	the possibility of my partner's selection and also how much I can earn.
17	my partner's choice
18	worst possible outcome, expected actions of s/he, risk.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
(19)	stand on the feet of the other part. Make some compromise. Make a balance between my benefits and those of the other part.
20	I chose to make the safest move financially, rather than going for the highest payout based on my speculation that my partner might be a bigger risk taker.
21	the possibility for me to get the most ECU
22	I consider first about what he/she might choose, standing on his point of view (highest possible ECU for him/her). Then I will pick my choice based on his/her choice I assumed. In the case of similar probability for her, I will assume he/she will also consider about me (stand on my view)
23	Making money. Specifically, making choices that would insulate myself from choices made by s/he. That is, playing conservatively.
24	trying to figure out the likelihood of what he/she would choose and then choosing mine to maximize my income.
(25)	how to maximize profit for both me and the other participant
(26)	I was trying to be nice to s/he. picked what seemed to be most beneficial for both of us.
27	I was trying to play the odds. I tried to make a decision based on me making on average the most \$ while trying to guess which decision was good for the s/he as well.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
28	the row with the least amount of small numbers
29	I assumed my partner would be greedy first, benevolent second, and that my partner would assume I was greedy. Therefore, I chose the option best for me based on first my partner being greedy, but second, recognizing that my partner would expect me to greedy. I chose at first the options that ensured that I didn't get a very low payout, then later in the session, I chose the highest average payout. Over time, this is best, but in a one-shot scenario, I could lose significantly.
(30)	I tried to choose boxes that would be mutually beneficial to both myself and s(he). Though I initially began by choosing based on the thinking that s(he would go for the highest numbers, I then decided that s(he was probably doing the same thing that I was doing.
31	I tried to make the most money I could , but I also tried to make decisions knowing that the other person would be trying to make decisions that were attempting to maximize profit. I tried to find options that gave me the best chance of making a fair amount of money, I uusually did not try to make the larger sums, because I knew others would not try it that way.
32	My partners's possible decision is the key factor for my own decision.
33	maximize the money earned

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
(34)	maximize the gains for both I and my partner
(35)	i mostly chose the one with low risk for both me and my unknown partner
36	I was trying to make money for myself. I attempted to consider which decision had the highest probability of making the most money on average. This involved looking at what the total sum of all the profits in each row were and also what the other person was thinking.
(37)	which choice s/he was most likely to choose and then i chose one based on that, and which would jointly benefit us both.
(38)	I focused mainly on what I thought the other person would choose to make his/her best profit and tried to compromise, basing my decisions off what would be best for both of us. On one question, the numbers for s/he were all 50/51, so for that question I made the decision that would give me the most profit because for s/he it was a difference of only \$.1 ECU.
39	maximize my profit. base on what it's the most likely would be chosen by another participant
(40)	i was trying to make sure both of us would make some kind of money choosing particular column or a row, so i would not choose a row that would not make money for the he/she
41	I tried to avoid my lowest income, when my highest value is not guaranteed.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
42	Most of the time, I was trying to figure out which row had the highest total point value, to hedge his/her choices. Sometimes, though, I looked to see if any of the columns (his/hers) were so low for him/her that s/he could not possibly choose it, and then made my decision based on the two remaining columns' values in my rows.
43	playing safe so that i would be able get a decent ECU regardless of whatever the other person chooses.
(44)	TRYING TO FIND OUT THE BEST WAY TP BENEFIT ME AND MY PARTNER
45	I just tried analytical rather than considering my earnings, which means I tried to predict other people choices as well.
(46)	I am considering both myself and he/she. 1. to make our earnings differ not much. 2. also considering the my risk and his/her risk of getting pretty low earning.
47	I mainly considered how I could make the money. I took some risks, when I thought the benefits were high enough, but mostly just tried to make sure that I earn a decent amount. I only tried to be nice to s/he when it made little difference to my earnings to do so.
48	I considered how he/she would react based on his/her options and I tried to maximize my earnings based on my guesses.

(Continued on next page)

Table A.13 – continued from previous page

Subject	Response
(49)	I first think about what I will get most, then I check what my partner will get if I choose this option. Generally I try to make both of us happy.

Table A.13:

Table A.14: Responses of Subjects in Treatment B to Questionnaire Question: “What were your key considerations when choosing one of your three options in part I?”

Subject	Response
(1)	how much money i would make; but i realized after i did part 1 that i was looking at the wrong numbers for me (i looked at the numbers in the top right of the boxes, instead of the numbers in the bottom left of the boxes
2	I pick the options based on the option i thought was best for my partner, i.e. chose according to what I thought they would pick.
3	what my partner will think, what's the best situation for my partner. and based on that, what will be most beneficial for me.
4	The size of the payoffs in my rows. I also tried to put myself in the position of my partner participant to guess which option they might choose. Where two choices on their part would lead to generally equivalent payoffs for them, I assumed that they might choose the option that would benefit me (as they would anticipate that would be the option I chose).
5	at first, what would be best for me and then what would be best for the other person.
6	seeing the values and then computing the probabilities of its occurrence
7	THE MAXIMUM OF THE SUM OF ROWS AND COLUMNS
8	I looked at the options that my he/she was facing, and if they were higher in one column/row, then I would choose my row/column accordingly.

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
9	I tried to figure out what he/she ould choose and based on that tried to figure out which would be the highest number for me (though I wasn't sure of the arithmetic)...
10	Minimum ECU I can get .
11	BASED ON THE PROBABILITY OF HOW THE OTHER PARTICIPANT WILL CHOOSE HIS/HER OPTION, I WILL CHOOOSE MY OWN TO MAXIMIZE MY EARNINGS
(12)	How the other part will make decision, and how both of us will win the maximam combination.
(13)	I tried to choose a row that would give equal advantage to both parties, i.e. where if the column chooser had used same criteria as I did, we'd both profit most from each table.
14	I thought about what the other he/she might be thinking in making a decision. I saw what rows led to my best options and checked if the corresponding columns were also good for the other guy.
(15)	to earn the most while at the same time allow my partner earn the most too. found a slot where both of us could earn general appropriate, but with my earnings being priority.
16	My benifit and other's benifit
(17)	SOMETHING THAT WOULD ENSURE MAXIMUM ECU FOR ME AND MY S/HE

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
(18)	While I tried to pick the option that would give me the most money, I also took into consideration what the other person might choose. For instance, if there was one square that gave me the highest amount of money but gave the other person the lowest, I would instead choose the square that gave me less but gave the other person more. I thought that the odds of the other person choosing a square that gave them more money were greater than them choosing the square that would give me the most money.
(19)	I was not concerned with myself over my partner. I felt the to increase my odds of being rewarded more \$ I had to also be cognizant of the amount of \$ that my partner would be most interested in receiving. I chose the options that, if chosen would reward us both a reasonable amount of \$ when compared to other options.
(20)	Firstly identifying the rows that would be beneficial to me in order of priority. Then determining which columns would benefit s/he and finding a likely pairing that would be mutually beneficial.
21	I first guess the response from s/he, then bet my decision for favorable outcome.

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
22	My main concern was making the most money for myself possible. When making this choice I assumed that he/she was also trying to maximize thier own money and sometimes adjusted my decision accordingly.
23	Based on the partner's choice with extimated probabilitly, make the best of my money
(24)	Having a higher number for both of usSometimes, a balanced approach so that each one of us benefits.
25	first, consider what a smart guy in another group would choose.second, consider what a normal (neither stupid nor smart) guy in another guy would choose.combine the above two and make my decision
(26)	Initially I wasn't sure how to approach the decison making, but when I realized how my decision may also affect he/she, I began to think clearly prior to choosing the option.
(27)	I choose after considering two options: what would my counter part do to max his/her income and if what he/she do if she/he wanted to max. both our incomes. I choose each one with the one i thought had the highest probability of happening.
28	the possibility to get money as much as possible. I should also consider my partner's chioce.

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
(29)	The option I and s/he can get the optimal amount of ECU is my key considerations. In fact, sometimes I think it didn't work, but it'll work if I and s/he consider all of these tasks as a continuous communicating process.
(30)	Firstly, I was concerned with what I thought the other person might like to choose, i.e. what was the best payoff for the both of us. I assumed the other person might do the same as well. For a few of the items I changed this strategy and looked for the places where if they would the choose there squares based on this strategy I could make the most money—second guessing.
31	The ECU that I could earn and my partener could earn.
32	The average number is large in that row. And the options for the other person are not bad. Or firstly consider the average number for the other person is large, choose the most probable column he/she may choose, then find out my best option from the rows corresponding to the column.
33	more attention was paid on her/his options that s/he will be able to choose.
(34)	If my paired partner would consider the best choice for both of us. I assumed she or he does. Therefore, I also considered what choices will be better for her, not just looking for the best choice for me.

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
35	I was thinking only about how to make the most money possible. I used the probabilities to determine the average pay earned based on the other player's decision.
(36)	Maximize both parties earning where possible. Otherwise try to get something on average for both parties
37	Considering both s/he and I at the same time.
38	<p>1. Trying to figure out first what s/he would choose. I assumed they would want to minimize risk by choosing the column with a high average.</p> <p>2. I also assumed that they would look at my options (rows) and assume that I too would want to minimize risk.</p> <p>3. When s/he could see that I would clearly not choose one of my rows (eg all values were very low), I would narrow my considerations accordingly.</p> <p>4. I assumed they would maximize their own profit based on what I choose.</p> <p>5. I did not try to assure a decent profit to them, and I did not assume that they would do so either.</p> <p>6. Since only 1 of 10 would pay out, I took more risks than if every one paid out</p>

(Continued on next page)

Table A.14 – continued from previous page

Subject	Response
39	<p>First, I think about the probabilities for my parteners to select each options. Then, Consider, what's my best and worst situation. Normally, I don't select for the options that can lead me to my worst situations. So, basically, I can exclude one option. Then, look into my best income, to choose the final answer. I also consider, what are the answers that can bring maximum benefit to my parteners.</p>

Table A.14:

APPENDIX B: INSTRUCTIONS FOR TREATMENT A FROM CH. 1

INSTRUCTIONS

Welcome!

This experiment consists of three parts followed by a brief questionnaire and is expected to last 2 hours.

All participants will receive a \$5 show-up fee. In addition you will have the opportunity to earn Experimental Currency Units (ECU) in each of the three parts of the experiment. ECU will be converted into dollars at a rate of \$0.1 per ECU (i.e. 100 ECU are worth \$10). This means that your total dollar earnings will equal:

$$\$5_{\text{show-up fee}} + 0.1 * (\text{ECU earned in part I} + \text{ECU earned in part II} + \text{ECU earned in part III})$$

Note that your decisions are likely to considerably affect your earnings. You will be paid in cash immediately after the experiment. Please feel free to earn as much money for yourself as you possibly can.

Note that all of you have been recruited using the same procedure and that the same participation requirements apply to all of you. In particular, you are eligible if:

- you are at least 18 years old
- you are an OSU student currently enrolled in an Econ 200 level class
- you have never taken any college-level Econ courses prior to the current quarter.

Caution: This is a serious experiment and talking, looking at others' screens or exclaiming aloud are not allowed. Should you have any questions please raise your hand and an experimenter will come to you.

PART I

The participants in this experiment are randomly divided into two groups - Row participants and Column participants. Each of you is randomly assigned a participant from the other group whom we refer to as s/he. This assignment is anonymous, i.e. no one will find out who s/he was for them.

Each of you will be presented with 10 decision situations. Your earnings in a decision situation will depend not only on your decision but also on the decision s/he made. Each participant will make their decision without being told what decision s/he made.

Each decision situation will be presented in the form of a table like the following one (only the numbers will differ from one decision situation to the next):

	LEFT	CENTER	RIGHT
UP	32 71	93 35	65 55
MIDDLE	69 21	12 63	35 47
DOWN	49 69	88 75	14 58

First we explain the above table from the point of view of a Row participant. If you are a Row participant you will have to choose a row: UP, MIDDLE or DOWN. S/he (who will be a Column participant) will be choosing a column: LEFT, CENTER or RIGHT. The combination of your decision and his/her decision will determine a cell in the table. Your earnings will equal the number in the lower left corner of that cell. E.g.:

- If you choose UP and s/he chooses LEFT you will earn 32 ECU;
- If you choose MIDDLE and s/he chooses RIGHT you will earn 35 ECU;

- If you choose DOWN and s/he chooses CENTER you will earn 88 ECU.

Now we explain the above table from the point of view of a Column participant. If you are a Column participant you will have to choose a column: LEFT, CENTER or RIGHT. S/he (who will be a Row participant) will be choosing a row: UP, MIDDLE or DOWN. The combination of your decision and his/her decision will determine a cell in the table. Your earnings will equal the number in the upper right corner of that cell. E.g.:

- If you choose LEFT and s/he chooses UP you will earn 71 ECU;
- If you choose RIGHT and s/he chooses MIDDLE you will earn 47 ECU;
- If you choose CENTER and s/he chooses DOWN you will earn 75 ECU.

For the purpose of determining your earnings for part I, one randomly selected decision situation from the 10 decision situations will be used. Your earnings will be based on the combination of the decision you made and the decision s/he made in that decision situation.

Please make sure you understand all of the above. If you have any questions please raise your hand. Otherwise proceed to the following quiz which will give you some practice in reading a table like the one above.

PRACTICE QUIZ

Consider the following table:

	LEFT	CENTER	RIGHT
UP	62 78	46 48	86 15
MIDDLE	28 37	64 42	66 69
DOWN	43 25	68 54	31 94

If you are a Row participant, how many ECU will you earn if:

- you choose MIDDLE and s/he chooses CENTER
- you choose UP and s/he chooses RIGHT
- you choose DOWN and s/he chooses LEFT

If you are a Column participant, how many ECU will you earn if:

- you choose CENTER and s/he chooses MIDDLE
- you choose RIGHT and s/he chooses UP
- you choose LEFT and s/he chooses DOWN

Once you are finished please wait for an experimenter to come and check your answers.

Once we start part I of the experiment you will have $2\frac{1}{2}$ min for each decision situation. Please do not take longer than that. Once you have confirmed your choice in a decision situation you cannot go back to change it.

Finally, write down the ID number which you will see on the top of your screen. You will need this number to be paid!

PART II

In part II we will once again present you with the 10 decision situations from part I.

Now, for each decision situation, we are interested in your best estimate of the probability (in percent) that s/he chose each of his/her three options in part I? That is, if you are a Row participant we would like to know: In your opinion, what is the probability that s/he chose LEFT, what is the probability that s/he chose CENTER and what is the probability that s/he chose RIGHT? Similarly, if you are a Column participant we would like to know: In your opinion, what is the probability that s/he chose UP, what is the probability that s/he chose MIDDLE and what is the probability that s/he chose DOWN?

To come up with your three probabilities for a given decision situation you can ask yourself the following question: If 100 different people were put in the same situation as s/he, how many would choose each of his/her three options?

Note that the three probabilities you give for each decision situation need to sum to 100%.

For your efforts in this part of the experiment we will pay each of you 60 ECU.

For each decision situation you will have $1\frac{1}{2}$ min to make your entry. Please do not take longer than that. Once you have confirmed an entry you cannot go back to change it.

PART III

This part of the experiment will involve two types of tickets, type P and type Q (we picked the names arbitrarily). A ticket of type P and a ticket of type Q pay ECU in a different way (as explained below). We will present you with 10 pairs of tickets. Each pair will consist of one ticket of type P and one ticket of type Q. Your task will be, for each pair, to choose the ticket according to which you would prefer to be paid.

How a Ticket of type P pays

A ticket of type P will involve a decision situation from part I. One of your options in that decision situation will be marked with a check on your screen. The ticket pays ECU according to the combination of your checked option and the option s/he chose in part I of the experiment. E.g. say the ticket involves a decision situation like the one below:

	LEFT	CENTER	RIGHT
UP	71 32	35 93	55 65
MIDDLE	21 69	63 12	47 35
DOWN	69 49	75 88	58 14

If you are a Row participant and your checked option is, say, MIDDLE, then the ticket pays as follows:

- 69 ECU if s/he chose LEFT in part I;
- 12 ECU if s/he chose CENTER in part I;
- 35 ECU if s/he chose RIGHT in part I.

If you are a Column participant and your checked option is, say, RIGHT, then the ticket pays as follows:

- 55 ECU if s/he chose UP in part I;
- 47 ECU if s/he chose MIDDLE in part I;
- 58 ECU if s/he chose DOWN in part I.

For your information your screen will also show your estimate (the one you gave in part II) of the probability that s/he chose each of his/her options.

How a Ticket of type Q pays

A ticket of type Q will pay one of three possible amounts of ECU. Which of the three amounts it will pay, will be determined randomly by the computer. The computer will select each of the three amounts with a certain probability.

A ticket of type Q will be presented in the following form (only the numbers will differ):

<i>35%</i>	<i>15%</i>	<i>50%</i>
53	22	91

This is to be understood as follows. The computer selects 53 ECU with probability 35%, 22 ECU with probability 15% and 91 ECU with probability 50%. The ticket pays the amount of ECU that the computer selected.

For the purpose of determining your earnings for part III, one randomly selected pair from the 10 pairs of tickets will be used. Your earnings will be calculated according to the ticket that you chose from that pair.

For each pair of tickets you will have 1 min to choose one of the tickets. Please do not take longer than that. Once you have confirmed an entry you cannot go back to change it.

APPENDIX C: INSTRUCTIONS FOR TREATMENT C FROM CH. 1

INSTRUCTIONS

Today's session consists of one single part and is expected to last less than 1 hour.

All participants will receive a \$7 show-up fee. In addition to that you will have the opportunity to earn Experimental Currency Units (ECU). ECU will be converted into dollars at a rate of \$0.1 per ECU (i.e. 100 ECU are worth \$10).

Note that your decisions are likely to considerably affect your earnings. You will be paid in cash immediately after the experiment. Please feel free to earn as much money for yourself as you possibly can.

Caution: This is a serious experiment and talking, looking at others' screens or exclaiming aloud are not allowed. Should you have any questions please raise your hand and an experimenter will come to you.

In today's session you will have to choose between tickets. Each ticket will pay one of three possible amounts of ECU. Which of the three amounts it will pay, will be determined randomly by the computer. The computer will select each of the three amounts with a certain probability.

You will be presented with 10 sets of tickets. Each set will consist of three tickets and will be presented in the following form (only the numbers will differ):

	<i>35%</i>	<i>45%</i>	<i>20%</i>
Ticket #	32	93	65
Ticket *	69	12	50
Ticket ^	48	88	14

This is to be understood as follows: The numbers in the rows corresponding to each ticket refer to the possible amounts of ECU each ticket can pay. The percentages on top of the columns refer to the probabilities with which each ticket pays one of the three possible amounts. So in the example above:

- Ticket # pays 32 ECU with 35% probability, 93 ECU with 45% probability and 65 ECU with 20% probability
- Ticket * pays 69 ECU with 35% probability, 12 ECU with 45% probability and 50 ECU with 20% probability
- Ticket ^ pays 48 ECU with 35% probability, 88 ECU with 45% probability and 14 ECU with 20% probability

Your task will be, from each set of tickets, to choose the ticket that you prefer. For payment purposes one randomly selected set from the 10 sets of tickets will be used. Your earnings will be calculated according to the ticket that you chose from that set.

For each set of tickets you will have 90 sec to choose one of the tickets. Please do not take longer than that. Once you have confirmed an entry you cannot go back to change it.

APPENDIX D: DERIVATION OF SYMMETRIC CURSED EQUILIBRIUM FROM CH. 2

In any cursed equilibrium, no matter what a player believes about the strategies being played by others, a type $(0, H)$ never invests, and a type $(1, L)$ invests in round 1. It is easy to see that the only viable possibilities for a type $(0, L)$ are N and W , and that the only viable possibilities for a type $(1, H)$ are W and 1. Denote the probability that a type $(0, L)$ chooses N as r , and denote the probability that a type $(1, H)$ chooses 1 as q .

Two-Cost Game: Given the above probabilities, one can compute the probability that the other player chooses 1, conditional conditional on having the high common-value signal, as follows

$$pr(1|X_i = 1) = \frac{.7^2 + .3^2}{2}(1 + q) = .29(1 + q).$$

Similarly, we have

$$pr(W|X_i = 1) = .5 - .29q - .21r,$$

$$pr(N|X_i = 1) = .21(1 + r),$$

$$pr(1|X_i = 0) = .21(1 + q),$$

$$pr(W|X_i = 0) = .5 - .21q - .29r,$$

$$pr(N|X_i = 0) = .29(1 + r).$$

With a cursedness parameter χ , the objective of a type $(1, H)$ is 0.5 if she invests in round 1, and is

$$.9\chi[.29(1 + q)(.5)] + .9(1 - \chi)[(.7)(.7)\frac{1 + q}{2}(3.5) + (.3)(.3)\frac{1 + q}{2}(-6.5)] \quad (\text{A.2})$$

if she chooses W . Expression (A.2) is strictly increasing in q and decreasing in χ . For $\chi < \frac{17}{756}$, expression (A.2) is greater than 0.5 for all q , so we are at a corner solution with $q = 0$. For $\frac{517}{756} < \chi$, expression (A.2) is less than 0.5 for all q , so we are at a corner solution with $q = 1$. For $\frac{17}{756} < \chi < \frac{517}{756}$, we solve for q by setting expression (A.2) equal to 0.5, yielding $q^{2\text{cost}}$.

The objective of a type $(0, L)$ is 0 if she chooses N , and is

$$.9\chi[.21(1 + q)(-.5)] + .9(1 - \chi)[(.3)(.7)\frac{1 + q}{2}(6.5) + (.7)(.3)\frac{1 + q}{2}(-3.5)] \quad (\text{A.3})$$

if she chooses W . For all q and r , expression (A.3) is positive for $\chi < \frac{3}{4}$, and negative for $\chi > \frac{3}{4}$. Therefore, except for the knife-edge case, $\chi = \frac{3}{4}$, we are always at a corner solution. A type $(0, L)$ must choose W for $\chi < \frac{3}{4}$, and N for $\chi > \frac{3}{4}$. Combining these choices for each type yields the type-dependent strategies for the Two-Cost Game given in Table 3.

Alternating One-Cost Game: Consider first the low-cost game, Game 2. We know that a type $(1, L)$ invests in round 1. For a type $(0, L)$, the probabilities that the other player chooses 1, W , and N are

$$pr(1|X_i = 0) = .42,$$

$$pr(W|X_i = 0) = .58(1 - r),$$

$$pr(N|X_i = 0) = .58r.$$

The objective of a type $(0, L)$ is 0 if she chooses N , and is

$$.9\chi[.42(-.5)] + .9(1 - \chi)[(.3)(.7)(6.5) + (.7)(.3)(-3.5)] \quad (\text{A.4})$$

if she chooses W . Therefore, except for the knife-edge case, $\chi = \frac{3}{4}$, the cursed equilibrium of Game 2 is in pure strategies. A type $(0, L)$ must choose W for $\chi < \frac{3}{4}$, and N for $\chi > \frac{3}{4}$.

Now consider the high-cost game, Game 3. We know that a type $(0, H)$ chooses N . For a type $(1, H)$, the probabilities that the other player chooses 1, W , and N are

$$\begin{aligned} pr(1|X_i = 1) &= .58q, \\ pr(W|X_i = 1) &= .58(1 - q), \\ pr(N|X_i = 1) &= .42. \end{aligned}$$

The objective of a type $(1, H)$ is 0.5 if she invests in round 1, and is

$$.9\chi[.58q(.5)] + .9(1 - \chi)[(.7)(.7)q(3.5) + (.3)(.3)q(-6.5)] \quad (\text{A.5})$$

if she chooses W . Expression (A.5) is strictly increasing in q and decreasing in χ . For $\frac{517}{756} < \chi$, expression (A.5) is less than 0.5 for all q , so we are at a corner solution with $q = 1$. For $\chi < \frac{517}{756}$, we solve for q by setting expression (A.5) equal to 0.5, yielding q^{alt} . Combining these choices for each type across Game 2 and Game 3 yields the type-dependent strategies for the Alternating One-Cost Game given in Table 3.

APPENDIX E: DERIVATION OF THE LIKELIHOOD FUNCTION FROM CH. 2

Let $j_i \in \{F, M, S\}$ denote subject i 's strategy class. Label all trials, $m = 1, \dots, M$, and let $m(1)$ and $m(2)$ be the two subjects in trial m , where $m(1)$ is the subject with the lower identification number. The likelihood function (suppressing the dependence on the realized types, realized matchings, and θ) is given by

$$\Pr(B) = \sum_{j_1, \dots, j_n} p_{j_1} \cdots p_{j_n} \Pr(B|j_1, \dots, j_n). \quad (\text{A.6})$$

Errors are independent, so behavior in one trial, conditional on the strategy classes of the subjects in that trial, is independent of behavior in any other trial. Therefore:

$$\Pr(B|j_1, \dots, j_n) = \prod_{m=1}^M \Pr(B_{m(1)}, B_{m(2)}|j_{m(1)}, j_{m(2)}). \quad (\text{A.7})$$

Now, we claim that, for subjects 1 and 2 in a particular trial, we can write

$$\Pr(B_1, B_2|j_1, j_2) = \Pr(B_1|j_1, B_2) \Pr(B_2|j_2, B_1). \quad (\text{A.8})$$

To verify this, let B_i^r be the behavior of i in round r (1 if i invests during that round, 0 otherwise). Then we have (sometimes suppressing the dependence on j_1 and j_2)

$$\begin{aligned}
\Pr(B_1, B_2 | j_1, j_2) &= \Pr(B_1^1, B_2^1) \Pr(B_1^2, B_2^2 | B_1^1, B_2^1) \Pr(B_1^3, B_2^3 | B_1^1, B_2^1, B_1^2, B_2^2) \cdot \\
&\quad \Pr(B_1^4, B_2^4 | B_1^1, B_2^1, B_1^2, B_2^2, B_1^3, B_2^3) \\
&= \Pr(B_1^1) \Pr(B_2^1) \Pr(B_1^2 | B_1^1, B_2^1) \Pr(B_2^2 | B_1^1, B_2^1) \Pr(B_1^3 | B_1^1, B_2^1, B_1^2, B_2^2) \cdot \\
&\quad \Pr(B_2^3 | B_1^1, B_2^1, B_1^2, B_2^2) \Pr(B_1^4 | B_1^1, B_2^1, B_1^2, B_2^2, B_1^3, B_2^3) \cdot \\
&\quad \Pr(B_2^4 | B_1^1, B_2^1, B_1^2, B_2^2, B_1^3, B_2^3) \\
&= \Pr(B_1 | j_1, j_2, B_2) \Pr(B_2 | j_1, j_2, B_1).
\end{aligned}$$

The behavior of one subject in a trial depends on the other subject's behavior but not on the other subject's strategy class (given the other's behavior), so the claim follows.

From (A.7) and (A.8), we have

$$\begin{aligned}
\Pr(B | j_1, \dots, j_n) &= \prod_{m=1}^M \Pr(B_{m(1)} | j_{m(1)}, B_{m(2)}) \Pr(B_{m(2)} | j_{m(2)}, B_{m(1)}) \\
&= \prod_{i=1}^n \Pr(B_i | j_i, B_{-i}).
\end{aligned} \tag{A.9}$$

Substituting (A.9) into (A.6), we have

$$\begin{aligned}
\Pr(B) &= \sum_{j_1, \dots, j_n} \left[p_{j_1} \cdots p_{j_n} \prod_{i=1}^n \Pr(B_i | j_i, B_{-i}) \right] \\
&= \sum_{j_1, \dots, j_n} \left[\prod_{i=1}^n p_{j_i} \Pr(B_i | j_i, B_{-i}) \right] \\
&= \prod_{i=1}^n \left[\sum_{j \in F, M, S} p_j \Pr(B_i | j, B_{-i}) \right] \\
&= \prod_{i=1}^n \Pr(B_i | B_{-i}),
\end{aligned}$$

which is what we wanted to show.

APPENDIX F: INSTRUCTIONS FOR THE TWO-COST TREATMENT FROM CH. 2

INSTRUCTIONS – Two Player Trials

This is an experiment on decision-making in investment markets. The National Science Foundation has provided funds for conducting this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Every participant in the experiment is guaranteed a payment of at least \$5, independent of their performance in the experiment. All monetary values in the experiment, such as investment costs, investment returns, and account balances, are written in experimental currency units (EC). Your balance of ECs at the end of the experiment will be converted to US dollars at the exchange rate of \$0.60 for each EC. Because your decisions may involve losses, we will endow you with a starting cash balance of 20 ECs. Your gains (losses) during the experiment will be added to (subtracted from) your cash balance. However, if your cash balance falls below zero, you will no longer be allowed to continue. At the end of the experiment you will receive in cash your end of experiment balance of ECs converted to US dollars, or \$5, whichever is greater.

1. In this experiment we will create a sequence of market trials. In each given market trial, the participants will act as potential investors. Each potential investor will have to decide whether, and when, s/he wishes to invest, based on the information s/he is provided (and which we will explain later).
2. In the experimental session today we will have between 20-25 market trials. Each market trial has several rounds. The initial round is round 1, the next is round 2, and so on. In each round you and the other potential investor in your market trial will have to decide (simultaneously) whether to invest in that round or not. The decision to invest is irreversible. Any potential investor who has not yet invested will be told whether the other potential investor has invested, and if so, during which round of that trial.
3. In each trial, the market in which you are a potential investor has just one more potential investor besides yourself. In a typical session we will recruit (about) 20 students. The computer will randomly match 10 pairs out of the 20 students. Each such matched pair, including the one that you are in, constitutes a separate market trial that has no relation to the other nine markets. A given market trial keeps the same matched students over the several rounds of that market trial. However, after the market trial is over, the computer randomly rematches students to form a new set of market trials. This matching procedure makes it very unlikely that you will be matched with the same student from one trial to the next.

4. The structure of information.

Information about investment cost: Each potential investor will know, before each market trial starts, his/her investment cost for that trial. There are two possible levels of investment cost: low cost, $CL=3.5$ and high cost, $CH=6.5$. Each potential investor will be assigned one of the two cost levels with equal probability ($1/2$). In other words, in your market trial, you will know your investment cost, and that the investment cost of the other potential investor is equally likely to be either 3.5 or 6.5.

Information about investment gross returns: The computer assigns a gross return to every market trial. The gross return remains the same for all rounds of the same market trial, and is completely uncorrelated with your investment cost. The computer randomly picks the gross return to be either 10 or 0, with equal probabilities. Once the gross return is picked, high or low, it is the same for both potential investors, and it remains the same for all rounds of the same trial. You will NOT observe whether the gross return for that trial is high or low. Instead, each potential investor will be given his/her own signal, which takes the value of either 0 or 1. Signals are 70% accurate, in the following sense:

If the gross return is 10, you have a 70% chance of observing signal 1 and a 30% chance of observing signal 0. If the gross return is 0, you have a 70% chance of observing signal 0 and a 30% chance of observing signal 1.

Each potential investor's signal is related to the gross return, but the computer randomizes separately for each potential investor, so the two signals can be different. For example, if the gross return is 10, there is a 49% chance that both signals are 1, there is a 42% chance that one signal is 1 and the other signal is 0, and there is a 9% chance that both signals are 0.

The signal for each potential investor is chosen at the beginning of the trial and remains the same for all rounds of that trial. Each potential investor observes his/her own signal, but not the signal of the other potential investor in that trial. Observing your signal may help you better predict the likelihood that the gross return in your market trial is high or low.

Information about other investors in your market: You will NOT be told the signal of the other potential investor in your market trial. However, you will be informed about whether the other potential investor has invested, and if so, during which round. If this information reveals something about his/her signal, it could improve your decision about if and when to invest.

You are not allowed to reveal or discuss your information with other students or look at another student's screen (this will be strictly monitored and violators will be removed from the experiment).

5. **The structure of the game.**

The computer randomly matches you with another potential investor to form a market trial. Once you are assigned to a market trial, you privately observe your cost and your signal, which remain constant for that market trial. The other potential investor observes his/her cost and signal. In round 1, you are asked to decide if you wish to invest. If you do not invest in round 1, you are informed about whether the other potential investor invested in round 1, and you are asked if you wish to invest in round 2. If you have not invested by round 2, we move to round 3, and so on. Once you have decided to invest, there are no more decisions to make in that market trial. That is, an investment decision in a given trial is irreversible. You cannot disinvest or invest a second time. After two consecutive rounds in which no one in your trial invests, that trial is over.

In order to make good decisions, you must understand how your gains and losses are determined. This will be carefully explained below.

Once a market trial is over, the whole process starts again. The computer matches you to another potential investor to form a new market trial, you will be assigned an investment cost and a signal, etc.

Your screen will inform you of the trial number, and the round number within the trial.

How your gains (discounted net returns) or losses are determined.

If you invest, your gains from that trial are the discounted difference between the gross return and your investment cost. Let us illustrate what this means by using a simple example. Suppose that in the current market trial your investment cost is 3.5. If you decide to invest in round 1, then your gains are: 6.5 if the gross return is 10 ($10 - 3.5 = 6.5$) or -3.5 , a loss of 3.5, if gross return is 0 ($0 - 3.5 = -3.5$). Note that gains or losses in round 1 are not discounted; they are just the difference between the market gross return and your investment cost. For each round that you wait, your gains or losses are discounted by a factor of 0.9, as shown in the following table.

Discounted Net Returns when Cost is 3.5		
Round that you Invest	If return is 10 (high)	If return is 0 (low)
1	6.5	-3.5
2	5.85	-3.15
3	5.26	-2.84

There are several important things to note here:

(i) If for whatever reason you have decided not to invest at all in a particular market trial, you will earn zero for that market trial.

(ii) You will not be told the actual gross return during a market trial. After each trial is over, the gross return is revealed and you will learn your discounted net gains or losses, which will be added to, or subtracted from, your cash balances.

(iii) It is up to you to decide if and when to invest. Clearly, your investment cost and your signal can affect your decision. Observing the activity of the other potential investor in your trial might indirectly yield useful information about your gross return, because his/her behavior might tell you something about his/her signal.

6. Information on the computer screen. Throughout the experimental session, the computer screen will show your ID number and current cash balances, in the upper left corner. The upper left corner of the screen will also remind you of the number of potential investors in each trial (2), the discount factor (0.9), and the “accuracy parameter” of your signal (70%).

At the beginning of each round of each market trial, you will see the number of the market trial, your cost of investment (either 3.5 or 6.5), and your signal (0 or 1). This information stays the same during the trial. In the middle of the screen, you will see the current round number. At the bottom of the screen, you will see a “history” of investment in previous rounds of that trial. (If it shows all zeros, no one has invested; if it shows a 1 under some round, the other potential investor invested during that round.)

You will have 25 seconds to think about whether to invest in that round. At that time, boxes marked “YES” and “NO” will appear, and you should mark a box to indicate whether you want to invest or not. Please make your choice within 5 seconds.

At the end of the market trial, you will see a screen that tells you the market trial number, your investment cost, your signal, the actual gross return, and your net discounted gains or losses from that trial. You will also see your current cash

balance and your personal statistics from your previous trials. (If you are listed as investing in round -1, this means that you never invested during that trial.)

7. We will start the session with two practice “dry runs” that do not count towards your earnings, at which point we will stop and answer additional questions. At the end of the experiment, while we are calculating your earnings, we ask that you answer the short questionnaire on your computer.

8. Are there any questions?

APPENDIX G: SCREEN PRINTOUT FROM ALTERNATING ONE-COST TREATMENT FROM CH. 2

ID 1
Accuracy parameter 70%
Discount factor 0.90
Investors in Market trial 2
YOUR CASH BALANCES 10.00

Remaining time (sec) 4

MARKET TRIAL 3

Your cost 6.5
Your signal 1

Current round 2
Do you wish to invest in this round?

☐ Yes
☐ No

OK

Round: 1
Number of investors: 1

Market trial	Your Cost	Your Signal	Gross return	Round you invested in	Your Profit
-1	6.5	1	10	1	3.50
0	3.5	0	0	1	-3.50
1	6.5	1	0	1	-6.50
2	3.5	0	0	1	-3.50