Endogenous sample selection and partial naiveté in common value environments: A laboratory study*

Ignacio Esponda
(WUSTL)

Emanuel Vespa
(UC Santa Barbara)

February 20, 2015

Abstract

In environments with asymmetric information, people often need to make inferences about the information that is conveyed by others' actions. Nash equilibrium implicitly assumes that people make correct inferences. In many realistic settings, however, these inferences are complicated because primitives, such as the distribution of types, are a priori unknown, and counterfactuals, such as payoffs from alternative outcomes, are unobserved. We conduct an experiment to understand how subjects make decisions in the face of the selection problem that complicates inferences in these environments. We find that Nash is a poor predictor of behavior but that other, boundedly rational, solution concepts can rationalize the data.

*We thank Erik Eyster, Kfir Eliaz, Guillaume Fréchette, Drew Fudenberg, Philippe Jehiel, Muriel Niederle, Stefan Penczynski, Demian Pouzo, Andrew Schotter, Bernardo Silveira, Ran Spiegler, Charles Sprenger, Alistair Wilson, Leeat Yariv, and several seminar participants for helpful comments. Esponda: Olin Business School, Campus Box 1133, Washington University, 1 Brookings Drive, Saint Louis, MO 63130, iesponda@wustl.edu; Vespa: Department of Economics, University of California at Santa Barbara, 2127 North Hall University of California Santa Barbara, CA 93106, vespa@ucsb.edu.

1 Introduction

A main insight from information economics is that information is not just conveyed by prices (and only about scarcity), but also by the actions of other people. Information economics, however, maintains the old, implicit assumption of economic agents who respond optimally to their environment. This assumption seems more heroic than in traditional general equilibrium settings, such as Arrow-Debreu economies, where price is a sufficient statistic and conveys all relevant information. In environments with asymmetric information, economic agents should also worry about making inferences from the actions of other agents.

It is possible, however, that these behavioral assumptions are not as heroic as they seem because people might eventually learn to make correct inferences with enough experience. On the other hand, making inferences from others' actions is complicated by the fact that primitives (such as the distribution over types) are a priori unknown and counterfactuals (such as payoffs from alternative outcomes) are often not observed. As we explain below, agents in these environments face the task of learning to make choices based on a biased sample of previous outcomes. The challenges implied by sample selection are familiar to applied researchers, who grapple with the issue of the identification of structural parameters, such as the distribution of payoff-relevant types in the population. Yet, a bit surprisingly, most of our models neglect this issue by implicitly assuming that economic agents themselves are able to successfully tackle these challenges.

In this paper, we conduct an experiment to test the main behavioral assumptions underlying information economics and, in particular, to understand how subjects make decisions in the face of the selection problem that arises in these environments. We focus on a simple collective decision problem. Each subject in our experiment makes a decision to recommend a risky or a safe project. The risky project is best in one state of the world and the safe project is best in the other, but the subject does not know the state of the world. Two other players make recommendations *simultaneously* with the subject, and these recommendations are possibly correlated with the state of the world. These two other players are played by computers in order to guarantee that the environment faced by the subject is stationary. The project that receives a majority recommendation is then implemented. Subjects play this decision problem for 100 rounds. Feedback after each round includes past actions and payoffs but, importantly, no counterfactual information about the risky project is provided if the risky project is not implemented.

We conduct two treatments. In both treatments, subjects know neither the probability that a risky project is good nor the recommendation strategies of the computers. In the No Selection treatment, each computer recommends the safe and risky projects with equal probability, irrespective of the state of the world. In the Selection treatment, the recommendations are informative about the state of the world: If the risky project is best, then both computers recommend it; if the risky project is worst, then each computer mistakenly recommends it with probability 1/3. It is in this latter treatment where subjects face the following sample selection problem. The payoff of the risky alternative is observed when the risky alternative is implemented. Because of the informativeness of computers' recommendations, the probability that the risky alternative is best is higher when conditioning on the event that the risky alternative is implemented. Thus, subjects tend to observe the payoff of the risky alternative in situations in which the risky alternative tends to do relatively well. A subject need not realize that this is the case because she does not observe the counterfactual payoff of the risky alternative when the safe alternative is implemented. This simple setup captures essential features of environments with asymmetric information, such as auctions, health-care, antitrust policy, and elections.[1]

We focus on steady-state behavior—and we are justified in doing so because behavior converges for about 80% of the subjects. We compare the predictions of two steady-state solution concepts, Nash and naïve behavioral equilibrium; other solution concepts are discussed in Section 3.4. A Nash best response in our setting is a best response to correct beliefs about the strategies of the computers and the distribution over states of the world. In the Selection treatment, the correct best response is to choose the safe option. The reason is that a subject can only affect the project that is implemented if the computers provide different recommendations, but this event implies that the risky project must be bad (otherwise, both computers would recommend it). On the other hand, naïve behavioral equilibrium (Esponda, 2008) captures the behavior of subjects who do not take into account that the computers' recommendations convey information. Naive subjects form systematically biased beliefs by

[1]In auctions, bidders often observe the realized value of an object only if they purchase it; in health-care markets, firms do not observe how expensive it is to cover certain uninsured patients; in antitrust policy, the regulator does not observe the effects of unapproved mergers; in elections, voters do not observe the performance of the non-elected party.

focusing on the *observed* (not counterfactual) performance of the risky alternative. In the No Selection treatment, accounting for selection is irrelevant and, therefore, the Nash and behavioral predictions coincide.

The main finding is that the direction of the treatment effect is consistent with behavioral equilibrium, not with Nash equilibrium. In particular, most subjects do learn to make correct decisions in the No Selection treatment, even though they are not told the primitives of the environment. Therefore, when others' actions are independent of the state of the world, Nash equilibrium (or behavioral equilibrium) rationalizes the data. On the other hand, almost no subject learns to make correct decisions in the Selection treatment, where others' actions are correlated with the state and subjects face a biased sample. As predicted by behavioral equilibrium, subjects select the risky project more often in the Selection compared to the No Selection treatment. At the end of the experiment, we elicit subjects' beliefs and we corroborate these predictions: Reported beliefs mostly fail to account for selection and are consistent with naive (biased) beliefs, not with sophisticated beliefs.

We then examine more closely the extent to which naïve behavioral equilibrium can *quantitatively* rationalize the data. While the theory predicts behavior for the benchmark treatment (No Selection) fairly well, behavioral equilibrium tends to slightly over-predict the amount of risky behavior in the Selection treatment. The same is true regarding elicited beliefs. Thus, while subjects are clearly naïve in our setting, their behavior is consistent with a partial accounting of selection. Controlling for risk aversion, we find that subjects overestimate the benefits of playing the risky alternative, but not by the full amount predicted by the theory.

The last finding raises the puzzle of how subjects can be naive while somehow managing to partially account for selection.[2] It turns out, however, that it is rather natural for subjects to be partially naïve in our experiment. When looking at behavior in all 100 rounds, we find that subjects are more likely to change their behavior (and in the expected direction) in a given round if they were pivotal in the previous round. Thus, subjects seem to partially account for selection by placing more weight on feedback from pivotal rounds.

[2] If subjects were slightly sophisticated and placed even a small probability on the event that the computers' recommendations are informative, then one might expect them to eventually learn that this correlation exists and correctly account for it (particularly in our experiment). Thus, it would seem that, in steady-state, either subjects would be completely naïve or fully sophisticated in our setting.

Motivated by this finding, we propose a new model of partial naiveté to quantify the extent to which subjects are placing more weight on feedback from pivotal rounds. The model is a single-parameter extension of naïve behavioral equilibrium. The new parameter measures the relative weight of pivotal vs. non-pivotal observations–Nash and naïve behavioral equilibrium are special cases, where relative weights are infinite and one, respectively. We estimate the model of partial naiveté and find that the median subject places about three times more weight on pivotal vs. non-pivotal rounds. This partial overweighting, however, has a small effect on behavior, since subjects are pivotal in only about a third of the rounds. This explains why behavior is still much closer to the naïve than to the Nash prediction.

Our findings suggest that we should pay more attention to the implicit assumptions that are made regarding the inference capabilities of economics agents. Arguably, our setting is one where a subject who understands the possibility of selection should easily detect it: Every single time that a subject recommends the risky option and her decision is pivotal, the risky option turns out to do badly. Although more research is needed, the fact that almost all of our subjects fail to see this pattern suggests that people are likely to have a very hard time accounting for selection in more realistic settings. This finding is troubling to the extent that our analyses of markets with asymmetric information often rely on the assumption that people make the right inferences. The good news is that we find that behavior is not random, but actually quite amenable to analysis provided that we use other solution concepts that account for certain aspects of bounded rationality.

A large experimental literature—mostly looking at auctions and dating back to the 1980s—shows that most subjects (between 50% to 80%) fail to condition on the information content of relevant events. Kagel and Levin (2002) survey theirs and others' substantial early work and Charness and Levin (2009), Ivanov et al. (2010), and Esponda and Vespa (2014) provide more recent contributions. Because one of their main objectives is to understand whether or not subjects condition on the right event (such as having an offer accepted by a seller), it is natural for this literature to provide subjects with the primitives of the environment. In our paper, telling subjects the primitives would preclude us from our objective of understanding how subjects cope with the type of selection and identification problems that arise in these settings. While not providing the primitives changes the nature of the subject's problem, our findings are consistent with this literature. In fact, the finding that almost no subject accounts for sample selection despite extensive experience with the problem makes the standard use of Nash equilibrium even more concerning.

Our work is also influenced by the bounded rationality literature that models the failure to take into account the information content of others' actions in common value environments. The initial contributions of Kagel and Levin (1986) and Holt and Sherman (1994) in the specific context of auctions were later generalized by Eyster and Rabin's (2005) cursed equilibrium and Jehiel (2005) and Jehiel and Koeesler's (2008) analogy-based expectation equilibrium. The mistake has also been studied in the context of non-equilibrium concepts, such as level-k thinking, by Crawford and Iriberri (2007). Our subjects do not a priori know the primitives of the game nor do they observe counterfactual outcomes; thus, we focus on the notion of a naïve behavioral equilibrium, which naturally allows beliefs to depend endogenously on own actions.

Eyster and Rabin (2005) also develop a notion of partial naiveté, called partially cursed equilibrium, that includes fully cursed equilibrium and Nash equilibrium as special cases. They find that a particular value of their parameter of partial naiveté fits data from several experiments (with known primitives) better than Nash equilibrium. We also find that subjects are partially naïve, although our notion of partial naiveté is motivated by differential attention to feedback. Their model is defined for any Bayesian game, while our model of partial naiveté is intended to illustrate how it is possible to have a learning interpretation of partial naiveté in a particular context. Our model of partial naiveté could be generalized to other contexts by defining the saliency of different events and attaching weights that depend on saliency.[3]

In one aspect, our paper constitutes a return to the early days of experiments on competitive equilibrium (e.g., Smith, 1962). In these experiments, subjects traded without any information about the underlying distribution of costs or values in the population of sellers and buyers, precisely because the objective was to understand how decentralized markets could aggregate this information.[4] More generally, there are at least three reasons why it might not always be a good idea to give subjects

[3]The salient events are likely to depend on the particular context, and we view it as an interesting empirical question to understand what makes an event salient in different contexts.

[4]With the exception of the "penny jar" auctions that Bazerman and Samuelson (1983) conducted among students (although, unlike our experiment, without the chance to learn), the experimental auctions literature deviated from this premise and provided subjects with the distribution of valuations early on (e.g., Cox et al., 1982).

the primitives of the environment. First, as illustrated in our environment, the nature of the problem might change when the primitives are not known. If so, current research might be overlooking important issues (such as the selection problem). Second, knowing the primitives is neither necessary nor sufficient to justify, motivate, or give a foundation to the concept of (Bayesian) Nash equilibrium. In fact, it seems quite natural to view equilibrium as the result of a learning process (e.g., Fudenberg and Levine (1998), Dekel et al. (2004)). Finally, in many (but not all) of the realistic environments that the theory is meant to capture, people do not really know these primitives.

More recently, Fudenberg and Peysakhovich (2014) study learning in an adverse selection experiment and have a treatment in which subjects are not told the distribution over types. They find that learning models that account for recency bias provide a better fit to the data than steady-state solution concepts such as Nash, cursed or behavioral equilibrium. In particular, subjects respond more when experiencing extreme outcomes in round t-1 (and to a larger extent if the extreme outcome is bad) but much less so to experiencing extreme outcomes in earlier rounds. Their results suggest that steady-state solution concepts are not always appropriate to explain behavior and raise the broader issue of when we should apply equilibrium analysis to predict behavior.

Finally, this paper is more broadly relevant to understanding how subjects cope with selection problems that arise endogenously from their own actions. We focus on the case where the selection problem arises due to asymmetric information. Recently, Esponda and Pouzo (2014) show that the selection problem might arise in much more general environments where the agent learns with a misspecified model of the world.

We describe the experimental design in Section 2, briefly discuss the theoretical predictions in Section 3, and show the main results in Section 4. Motivated by our findings, in Section 5, we propose and estimate a new model of partial naiveté. We conclude in Section 6 and relegate the instructions of the experiment and robustness checks to the Online Appendix.

2 The experiment

Our experiment involves a single-agent problem. For presentation purposes, it is convenient to start by contrasting our experiment to a standard bandit problem.

		project A is: GOOD	project A is: BAD
Subject's choice:	A	5	1
	B	3	3

		project A is: GOOD	project A is: BAD
Majority's choice:	A	5	1
	B	x	x

Figure 1: Payoffs for a bandit problem and our experiment. The table on the left shows payoffs for a typical bandit problem, where the subject chooses between A and B without observing the state (good or bad) and with no knowledge of the probability distribution over the state. The table on the right shows payoffs for our experiment, which differs in two respects: (i) the choice is determined by a majority, and (ii) the payoff x from choosing B now varies each round from 1.25 to 4.75, although the subject observes the value of x before making a choice.

The left panel of Figure 1 depicts a typical bandit problem. In each round of a bandit problem, the decision maker must decide between investing in a project from industry A or from industry B. Project A yields 5 points if it is good and 1 point if it is bad. Project B yields a safe payoff of 3 points. The decision maker must choose A or B without observing the realized state. The probability that a project from industry A is good is the same every period, but the decision maker does not know this probability. The objectives of the theoretical and experimental bandit literatures are to characterize and test for optimal experimentation (e.g., Banks et al., 1997). Since our focus is not on experimentation, we introduce a few modifications to this standard setting.

The right panel of Figure 1 depicts the main decision problem in our experiment. There are two main differences with respect to a bandit problem. First, the decision whether to invest in A or B is now a group decision. The subject and two "other players" will *simultaneously* submit a recommendation for either project A or project B, and the ultimate investment choice will be determined by the recommendation of the majority. As explained below, this is the feature that generates the selection bias that is typical (but often neglected) in these environments. The "other players" in the experiment will be played by computers that are programmed to follow specific strategies. The main reason for using computers is to make sure that the environment faced by the subject remains stationary.

The second modification is that the safe payoff from investing in project B is now

x points, where $1.25 \leq x \leq 4.75$; in particular, it is still true that project A is best in the "good" state and project B is best in the "bad" state. A new value of x is randomly drawn every round and the subject *observes* the realized value before submitting her recommendation. The reason for varying the value of x is to induce experimentation by the subject and avoid her settling for an alternative because she decides it is not worth experimenting further. The focus of this paper is to understand whether people make mistakes that are not due to lack of data but rather to the selection problem inherently present in common value problems.

A crucial feature that this experiment continues to share with a bandit problem is that, after each decision, a subject only observes whether A turned out to be good or not if the majority recommended A. In other words, subjects do not observe counterfactual payoffs.

2.1 Experimental design

In this section, we provide a summary of the instructions for each of the three parts of the experiment. Detailed instructions (with the exact wording) are provided in Appendix B.

Part I (Rounds 1-100). Summary of instructions. In each of 100 rounds:

1. You will help your company decide between investing in a new project from industry A or a new project from industry B. The chance that a project from industry A is good is fixed between 0 and 100 percent and will not change throughout the experiment.

2. Your company has programmed two computers, Computer 1 and Computer 2, to assess whether project A is good or bad. If a computer assesses project A to be good, then it recommends A; otherwise, it recommends B. The computers make two types of mistakes: recommend A when A is bad and recommend B when A is good. Computer 1 and Computer 2 make the same rates of mistakes. The chance that the computers make the first type of mistake is fixed between 0 and 100 percent and will not change throughout the experiment. The chance that the computers make the second type of mistake is fixed between 0 and 100 percent and will not change throughout the experiment.

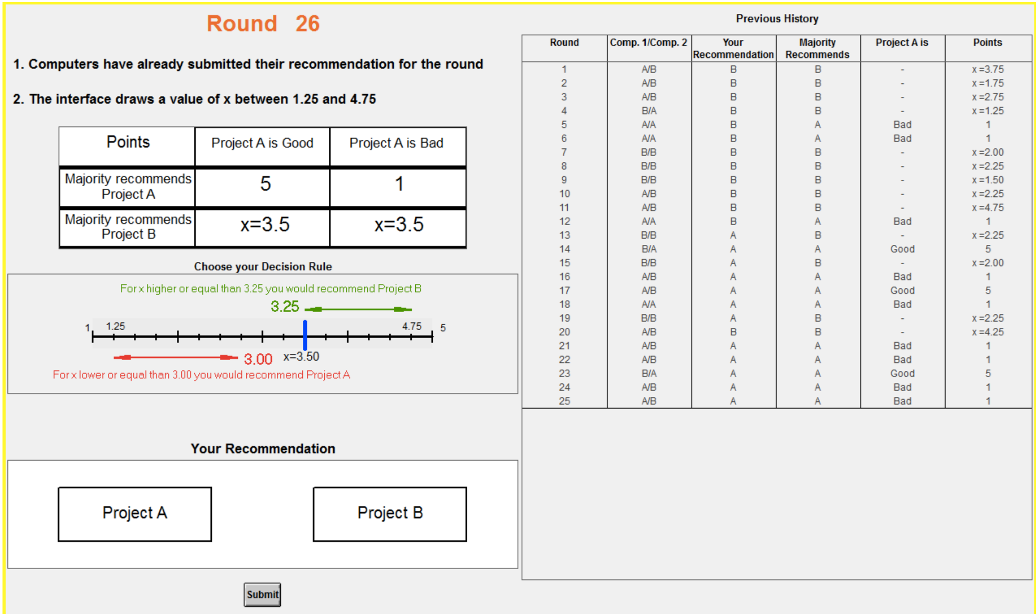

Figure 2: Screen shot for Round 26. In Rounds 1-25, the subject must submit a recommendation for a given value of x. In Rounds 26-100, the subject must first submit a threshold recommendation that indicates a choice for each value of x. She is then prompted to submit a choice for a particular value of x, as in Rounds 1-25.

3. Next, the interface draws a value of x (all values from 1.25 to 4.75, with increments in quarter points, are equally likely) that represents the payoff if the company invests in the project from industry B. You will observe the value of x but not the recommendations of the computers. You will then submit a recommendation for project A or B.

4. The payoffs for the round are given by the table in the right panel of Figure 1, where the company will invest in the project recommended by the majority.

Feedback: After each round, a subject sees the entire past history of rounds consisting of: the recommendations of the computers, her own recommendation, the recommendation of the majority, whether project A turned out to be good or not, provided it was chosen by the majority, and her payoff. Crucially, a subject does not observe whether or not A would have turned out to be good if project A is not implemented.

In the above design, we only observe a subject's decision for a particular value of x, but, ideally, we would like to know the entire strategy; i.e., a decision in each round for each possible value of x. To elicit this additional information, we introduce a novelty to our design starting in Round 26. The problem in Rounds 26-100 is exactly identical as the problem faced in the previous 25 rounds, but we now ask subjects to make one additional decision. At the beginning of the round, before the value of x is drawn, each subject must submit a threshold strategy indicating what she would recommend for each value of x. Subjects must choose a number from 1 to 5 by clicking on a slider on the screen. If they click on x^*, this means that they would recommend B for $x > x^*$ and A for $x < x^*$. After they submit their threshold strategy, the round continues as before: a value of x is drawn and they must submit a recommendation for A or B. If the recommendation submitted is not consistent with their previously selected threshold strategy, we alert them, ask them to make a consistent choice, and remind them that they can change their threshold strategy in the next round. This procedure is intended to clarify the meaning of a strategy to the subjects. We introduce the change in Round 26 to make sure that subjects are familiar with the problem before having to report a strategy. Figure 2 provides a screen shot of Round 26.[5]

Part II (Belief elicitation). After Round 100, we ask the subject to write an incentivized report for the company explaining how they reached their decision by Round 100.[6] After the report is written, we ask the subject three questions that are intended to elicit their beliefs. The subject must answer one question before moving on to read the next question. For each question, we pay $2 if the response is within 5 percentage points of the correct value.

Question 1. What is the chance that a project from industry A is good?

Question 2. What is the mistake rate of the computers when A is good? What is the mistake rate when A is bad?

[5]This design yields more (and less noisy) information in each round, compared to estimating a threshold strategy from the data (pooling data from different periods is less appealing in our setting because subjects are likely to be learning and changing their thresholds over time). Of course, without this restriction, some subjects might make a mistake and not follow threshold strategies. But this mistake is not the main focus of this paper and, more generally, implications of the strategy method have been studied elsewhere (e.g. Brandts and Charness, 2011).

[6]This part was anticipated in the instructions of Part I in order to encourage subjects to pay attention to the data. Subjects were also provided with paper and pencil in Part I to take notes about the observed data.

Question 3. What is the chance that a project from industry A is good conditional on your recommendation being pivotal?

Parts III (Risk aversion). We measure risk aversion in the following way. In the last part, the subject faces the same problem as in Rounds 1-100, but with two exceptions: there are no computers (so her decision alone determines the choice of project) and the chance that project A is good is known. The subject must make a threshold choice in each of five cases where the probability that A is good is known to be .1, .3, .5, .7, and .9.[7]

2.2 Two treatments

The primitives of the environment are given by (p, m_G, m_B), where p is the probability that project A is good, m_G is the mistake rate when A is good, and m_B is the mistake rate when A is bad. We consider two treatments. In both treatments, the probability that a project from industry A is good is $p = 1/4$, and the (unconditional) probability that a computer recommends A is $1/2$. Treatments differ by the rates of mistakes of the computers.

No Selection treatment. Each computer recommends A and B with equal probability, irrespective of whether A is good or bad, i.e., $m_G = m_B = 1/2$. The computers' recommendations in this treatment are uninformative of whether A is good or bad.

Selection treatment. Each computer correctly recommends A if A is good. Each computer mistakenly recommends A with probability $1/3$ if A is bad, i.e., $m_G = 0$, $m_B = 1/3$. The computers' recommendations in this treatment are informative.

As explained in the next section, when computers' recommendations are informative (Selection treatment) the subject must make inferences from a biased sample.

[7]At the end of the experiment, we run the experiment conducted by Holt and Laury (2002) to obtain an alternative measure of risk aversion in the population; discussed in footnote 27, the two measures are consistent with each other.

2.3 Subjects

We ran a between subjects design at NYU's Center for Experimental Social Science (CESS). We conducted three sessions per treatment (68 subjects with No Selection and 66 subjects with Selection). Part I lasted approximately 60 minutes and parts II-III lasted about 25 minutes. Average payoffs were approximately $18.

3 Theory: Nash and behavioral equilibrium

Our main objective is to explain the steady-state behavior of subjects. There are two equilibrium concepts that a priori seem particularly relevant in our setting: Nash and behavioral equilibrium. We discuss the predictions of these solution concepts for the No Selection and Selection treatments. We discuss other solution concepts at the end of this section.

3.1 Informal discussion

Table 1 shows an example of feedback from playing the first 12 rounds of the Selection treatment. What are sensible ways in which subjects could respond to this data? One reasonable approach is to figure out how many times project A has been observed to be good vs. bad. In the example, half of the observations are good and half are bad. Thus, the subject who follows this approach will estimate the probability of good to be close to 1/2 and then behave as in a decision problem where she has to choose between a risky option that delivers a payoff of 5 or 1 with equal probability and a safe option that delivers x for certain. The problem with this rule is that it does not account for the fact that the recommendations of the computers might be correlated with the state of the world. A consequence of this correlation is that the sample from which the subject is learning is biased. To see this point, note that a subject only observes whether A is good or not when a majority chooses to recommend A. But, if the computers happen to have some expertise in determining whether A is good or not (as in the Selection Treatment), then the subject will observe whether A is good or bad in those instances in which A is more likely to be good. In particular, the subject will overestimate the likelihood that A is good and choose a strategy that is more risky than optimal. We refer to this rule of behavior as a "naive rule".

There is another, "sophisticated rule", where a subject understands that the sample

Round	Comp1\Comp2	You	Majority	Project A is...	Payoff
1	A\A	B	A	Good	5.00
2	B\B	B	B	-	3.75
3	A\B	B	B	-	1.25
4	A\B	A	A	Bad	1.00
5	A\B	A	A	Bad	1.00
6	A\A	A	A	Good	5.00
7	B\B	A	B	-	3.25
8	A\A	A	A	Bad	1.00
9	A\B	A	A	Bad	1.00
10	A\A	B	A	Good	5.00
11	B\B	A	B	-	1.75
12	A\A	B	A	Good	5.00

Table 1: Example of feedback faced by a subject after 12 rounds in the Selection treatment. A naive approach is to estimate the probability of good by looking at the relative proportion of good vs. bad observed outcomes. A sophisticated approach is to look only at rounds in which a subject's decision was pivotal. In the Selection treatment, project A is always bad conditional on being pivotal.

may be biased because the recommendations of the computers might be informative about the state of the world (good or bad). A sophisticated subject can account for sample selection by learning from the subsample of rounds in which both her recommendation is pivotal and information about A is obtained (because the majority recommends A); these are rounds 4, 5, and 9 in Table 1. In the example, in all of such rounds, project A is observed to be bad. A subject following this rule will be more pessimistic about the prospects of recommending A compared to a naive subject.

Esponda and Pouzo (2012) show that the steady-state of an environment in which all players follow the sophisticated rule, where subjects account for selection by looking only at pivotal rounds, is captured by the notion of a Nash equilibrium. They also show that the steady-state of the naive rule, where subjects do not account for selection, is captured by the notion of a (naive) behavioral equilibrium (Esponda, 2008). In our setting, the strategies of the computers are fixed, so that the analog statement is that the steady-state with a sophisticated subject corresponds to playing the correct (i.e., Nash) best response and that the steady state with a naive subject corresponds to playing the (naive) behavioral best response. We now characterize the Nash and behavioral best responses in each treatment.

3.2 Steady-state behavior in No Selection treatment

In the No Selection treatment, the strategies of the computers are independent of the state of the world (good or bad). Thus, there is no selection in the data and both the naive and sophisticated rules lead to the correct belief that the probability of A being good is 1/4. Because steady-state beliefs coincide and do not depend on being pivotal, it follows that Nash and behavioral predictions are the same for this treatment.

Suppose, for example, that a subject is risk neutral. Then the steady-state belief about the expected benefit from recommending A (whether or not conditional on being pivotal) is $(1/4) \times 5 + (3/4) \times 1 = 2$. Thus, the steady-state threshold strategy is $x^* = 2$: for $x > 2$, a risk-neutral subject prefers to recommend the safe option B, and for $x < 2$ a risk-neutral subject prefers to recommend the risky option A.

In practice, it is important to account for the fact that subjects in the experiment might have different levels of risk aversion. Suppose, for concreteness, that a subject has a CRRA utility function $u_r(c) = c^{1-r}/1 - r$ with coefficient of risk aversion r, where the subject is risk neutral if $r = 0$, risk averse if $r > 0$ and risk loving if $r < 0$.[8] Then the optimal (Nash and behavioral) threshold x^* for a subject with risk aversion r is given by the solution to the following equation,

$$\frac{1}{4} \times u_r(5) + \frac{3}{4} \times u_r(1) = u_r(x^*).$$

Figure 3 plots the (Nash and behavioral) threshold $x^*_{NE}(r) = x^*_{BE}(r)$ as a function of the coefficient of relative risk aversion, r. As expected, the threshold decreases as risk aversion increases.[9]

3.3 Steady-state behavior in Selection treatment

In the Selection treatment, the strategies of the computers are correlated with the state of the world (good or bad), and Nash and behavioral best responses differ. Consider first the Nash case. Because both computers correctly recommend A if it is good, then, if a subject is pivotal, A must be bad. Thus, the best response is to

[8] For $r = 1$, we let $u(c) = \ln c$.

[9] For simplicity, the theory discussion assumes that both x (uniformly distributed) and the threshold can take any value in the interval $[1, 5]$. Of course, we account for the discreteness of the signal and action space when discussing the results of the experiment.

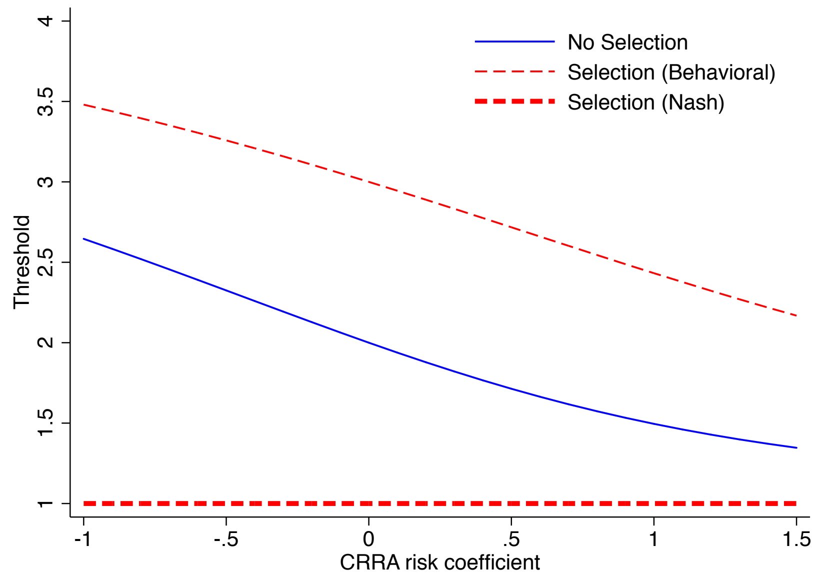

Figure 3: Theoretical prediction for Selection and No Selection treatments. For the benchmark case of the No Selection treatment, Nash and behavioral thresholds coincide. Under Selection, Nash and behavioral thresholds go in opposite direction: higher than the benchmark for behavioral and lower than the benchmark (and equal to 1) for Nash.

always recommend B, $x^*_{NE} = 1$, irrespective of the risk aversion coefficient. In terms of the sophisticated rule described above, it will be the case that *every time* that the subject is pivotal and recommends A, she will observe that A turned out to be bad. Thus, with enough experience, a sophisticated subject should stop recommending A and converge to $x^*_{NE} = 1$.[10]

Next, consider the prediction of (naive) behavioral equilibrium. The steady-state belief that A is good is given by the relative proportion of good vs. bad outcomes conditional on the event that the majority recommended A (because this is the event in which the subject observes whether A was good or bad). The probability that the majority recommends A, in turn, depends not only on the behavior of the two computers but also on the behavior of the subject, x^*. In particular, the steady-state belief of a subject who chooses strategy x^*, denoted by $z(x^*)$, is endogenous and given by the probability that A is good conditional on the event that the majority

[10]The fact that every time that the subject is pivotal she regrets recommending A makes it easiest for a subject to learn the optimal strategy. If we had chosen primitives with optimal thresholds higher than 1, then the subject would have observed both positive and negative feedback conditional on being pivotal, thus making it more difficult to determine her expected payoffs and play optimally.

recommends A (denoted as M_A),

$$\begin{aligned} z(x^*) &= \Pr(\text{good} \mid M_A;\, x^*) \\ &= \frac{\Pr(M_A|\text{good};\, x^*)p}{\Pr(M_A|\text{good};\, x^*)p + \Pr(M_A|\text{bad};\, x^*)(1-p)} \\ &= \frac{\left((1-m_G)^2 + 2m_G(1-m_G)\frac{(x^*-1)}{4}\right)p}{\left((1-m_G)^2 + 2m_G(1-m_G)\frac{(x^*-1)}{4}\right)p + \left((m_B^2 + 2m_B(1-m_B)\frac{(x^*-1)}{4}\right)(1-p)} \\ &= \frac{3}{3+x^*}, \end{aligned} \tag{1}$$

where we have used the fact that, in the Selection treatment, $m_G = 0$ and $m_B = 1/3$.[11] In particular, $z(\cdot)$ is decreasing; the intuition is that, the higher the threshold, then the more likely the subject is to vote for A, which means the more likely A is chosen when it is bad and, therefore, the lower the observed payoff from A.

Steady-state behavior (i.e., the behavioral best response) is characterized as a fixed point threshold x^* with the property that: (i) given that the subject chooses strategy x^*, then her steady-state belief is $z(x^*)$, and (ii) the strategy x^* is the optimal threshold given belief $z(x^*)$, i.e.,

$$z(x^*) \times u_r(5) + (1 - z(x^*)) \times u_r(1) = u_r(x^*). \tag{2}$$

In other words, the behavioral equilibrium threshold $x^*_{BE}(r)$ is the unique solution to equation (2).[12] For example, if the subject is risk neutral, $r = 0$, then equation (2) becomes $4/(1 + x^*/3) + 1 = x^*$ and the behavioral threshold is $x^*_{BE} = 3$. Figure 3 plots the behavioral threshold $x^*_{BE}(r)$ as a function of the coefficient of relative risk aversion, r. As expected, the threshold decreases as risk aversion increases.

To summarize, the predictions of Nash and behavioral equilibrium coincide for the No Selection Treatment. On the other hand, Nash and behavioral equilibrium differ in the predicted direction of the treatment effect: For a given level of risk aversion, the threshold increases under behavioral equilibrium and decreases under Nash equilibrium when going from the No Selection to the Selection treatment.

[11] Note that, in the No Selection treatment, $z(x^*) = 1/4$ for all x^*, as remarked earlier.

[12] The solution is unique because the LHS of equation (2) is decreasing (because $z(\cdot)$ is decreasing) and the RHS is increasing.

3.4 Other equilibrium concepts

As mentioned in the introduction, the notion of sincere bidding and its generalizations, such as fully cursed equilibrium (Eyster and Rabin (2005)) and analogy-based expectations equilibrium (Jehiel and Koessler (2008)), have usually been applied to model naiveté in common value settings. We focus on behavioral equilibrium, however, because these other solution concepts are naturally tailored to predict behavior in settings where, unlike ours, players have counterfactual information.[13] In fact, these other notions of naiveté all predict that behavior will be the same in both the Selection and No Selection treatments, since the probability that the project is good is the same in both treatments. This is a natural prediction in settings where players might know a priori the probability that project A is good or where players might observe the payoff from project A even if their own company decides not to carry out the project (say, because some other company carries it out). In our paper, in contrast, we focus on those other settings where counterfactual information is not observed, which additionally leads to a sample selection problem and endogenous beliefs for naive players.

4 Results

We organize the presentation of the results around five main findings. Motivated by the results, in Section 5 we propose and estimate a new model of partial naiveté.

Finding #1. The direction of the treatment effect is consistent with behavioral equilibrium, not with Nash equilibrium:

The first question is whether it is appropriate in our setting to use steady-state solution concepts to predict behavior, i.e,. whether or not behavior actually converges. For each round k in Part I of the experiment, we say that a subject chooses a convergent threshold if she chooses the same threshold in all remaining rounds, from k to 100. Figure 4 shows convergence rates in the population for each round in Part I, by treatment. For example, in round 30, only 18% of the subjects in the No Selection treatment and 29% in the Selection treatment choose convergent thresholds.

[13] If nature were viewed as a player, the setup of Jehiel (2005) would allow players to have erroneous perceptions about nature's strategy.

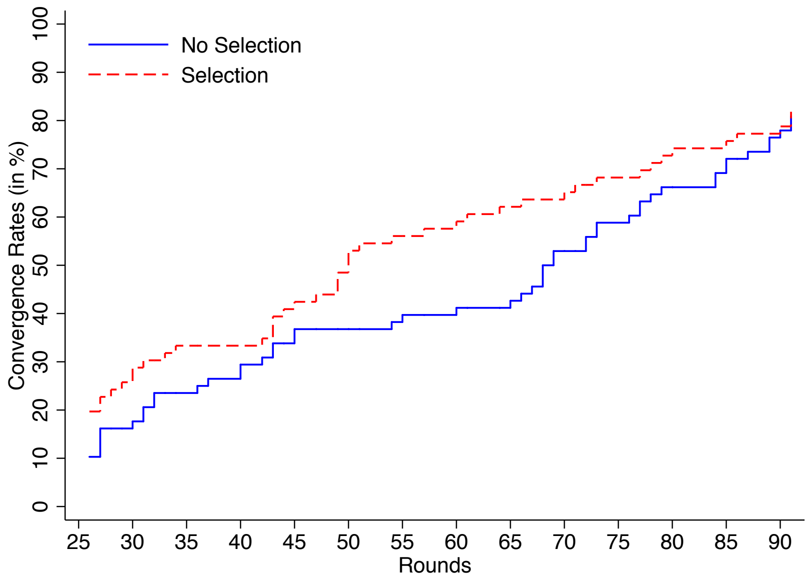

Figure 4: Convergence rates by treatment.

By round 90, however, these rates increase to 83% and 79%, respectively. Thus, we next focus on explaining behavior with steady-state solution concepts, although the figure also cautions that this is appropriate in our setting because subjects have a lot of experience (more so than in the typical experiment).

Figure 5 shows the average observed threshold choice in each round by treatment. Recall that x takes only a finite number of values, so that we can only infer that the threshold of a subject falls in an interval. For concreteness, we define the observed threshold to be the midpoint of the appropriate interval. For example, if a subject chooses A for all $\tilde{x} \leq 1.75$ and chooses B for all $\tilde{x} \geq 2$, then her preferred threshold is somewhere in the interval $[1.75, 2]$, and we code the observed threshold as $(1.75 + 2)/2 = 1.875$. The left panel of Figure 5 includes all subjects. The right panel of Figure 5 includes only subjects who choose a convergent threshold in round 91, i.e., whose behavior remains the same in the last 10 rounds (about 81% of subjects; see Figure 4). We refer to these subjects as the subjects who converge.

The patterns in the data are similar whether we look at all subjects or only those subjects who converge. Early in Round 25 (which is the first round where we observe a threshold choice), subjects have yet to receive most of their feedback and, not surprisingly, the average thresholds are similar in each treatment. As the experiment progresses and subjects observe more feedback, the average threshold in the Selec-

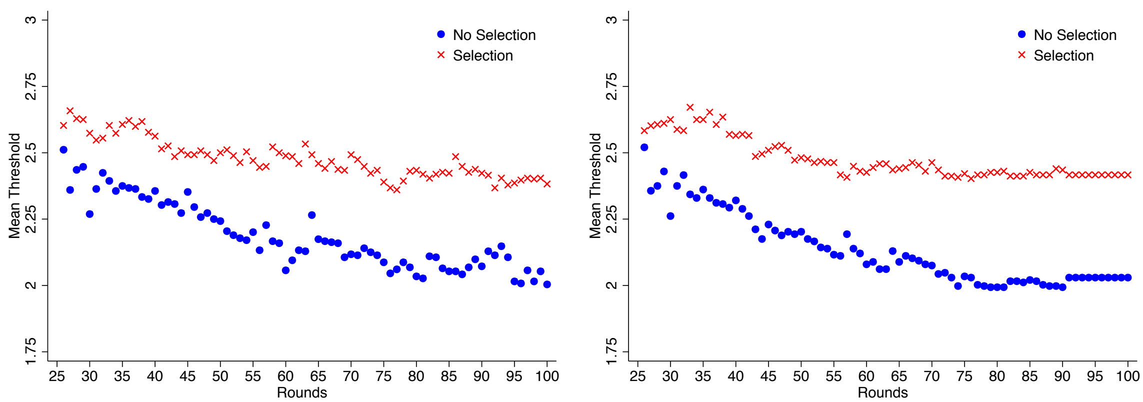

Figure 5: Mean thresholds in all rounds by treatment. The left panel shows the mean threshold for all subjects, for each round and treatment. The right panel shows the same information but only for subjects whose behavior converges in the sense that their threshold choice is constant for the last 10 rounds (approximately 80% of subjects in each treatment).

tion treatment remains above the No Selection treatment and the gap widens. Recall that in the No Selection treatment, on average, subjects will observe that A is good about 25% of the time (irrespective of their pivotality). Not surprisingly, the average threshold significantly decreases with experience in the No Selection treatment. In the Selection Treatment, in contrast, behavior depends on whether a subject is sophisticated (Nash) or not (behavioral). In the Nash case, a subject should realize that, every time she is pivotal, A is bad. Thus, the Nash threshold should converge to 1. In the behavioral case, a subject believes that the probability that project A is good is closer to 50% than to 25%, since this is what is observed in her upward-biased sample; thus, there should be a positive treatment effect. As observed in Figure 5, the direction of the treatment effect is clearly consistent with behavioral, not Nash, equilibrium.[14]

Because our objective is to explain steady-state behavior, from now on, we will focus on explaining behavior in the last rounds of the experiment, where beliefs and behavior have presumably converged and equilibrium concepts are potentially applicable. Also, it seems contradictory to use a steady-state solution concept to explain behavior that does not converge. Thus, from now on we will exclusively look at those subjects who converge, i.e., the 82% of subjects who choose the same threshold in

[14]Naiveté also explains why convergence is slower under No Selection (see Figure 4). A subject who starts with a uniform prior will take longer to converge if the observed feedback is that 25% of projects are good (No Selection) vs. about 50% (Selection).

each of the last 10 rounds, and refer to their threshold choices as their convergent thresholds. In Appendix A, we replicate the analysis with all the subjects and we find essentially the same results. For those subjects who converge, the mean convergent threshold is 2.03 under No Selection and 2.42 under Selection; the median convergent thresholds are 1.88 and 2.50, respectively. The differences in the mean (0.39) and the median (0.62) are both statistically significant at the 1% level.[15]

Finding #2. There is no shift of mass to lower thresholds under Selection compared to the No Selection treatment:

Even though average behavior is consistent with naiveté, it could still be possible that some subjects are sophisticated and choose very low thresholds in the Selection treatment. Figure 6 shows that this is not the case: the empirical distribution of convergent thresholds for the Selection treatment first-order stochastically dominates the distribution in the No Selection treatment.[16] Thus, there is essentially no evidence of sophistication in this experiment. This result can be contrasted with our finding in another paper (Esponda and Vespa, 2014), where subjects face an essentially equivalent setting but where they know both the probability over the states of the world and the strategies of the computers. In that other setting, the percentage of sophisticated subjects was between 20% and 50%, depending on the feedback.

Finding #3. Reported beliefs are consistent with naive (biased) beliefs, not with sophisticated beliefs:

Recall that, after Round 100, we ask subjects to report their beliefs. While one has to be cautious when using reported beliefs to draw conclusions about behavior,

[15] To test for differences in the mean, we run a regression with the convergent threshold on the right-hand side and a dummy variable for the treatment as a control. We compute the hypothesis test using robust standard errors. To test for differences in the median we use the same dependent and control variables, but run a median quantile regression. If we use all subjects, the mean threshold in round 100 is 2.00 under No Selection and 2.38 under Selection; the median round 100 thresholds are 1.88 and 2.38. The differences are significant at the 1% level.

[16] We test for first order stochastic dominance using the test in Barrett and Donald (2003). The test consists of two steps. We first test the null hypothesis that the distribution under the Selection treatment either first order stochastically dominates or is equal to the distribution under No Selection. We cannot reject this null hypothesis, the corresponding $p-$value is 0.77. We then test the null hypothesis that the distribution under the No Selection treatment first order stochastically dominates the distribution under Selection. We reject the null in this case, with a corresponding $p-$value of 0.001.

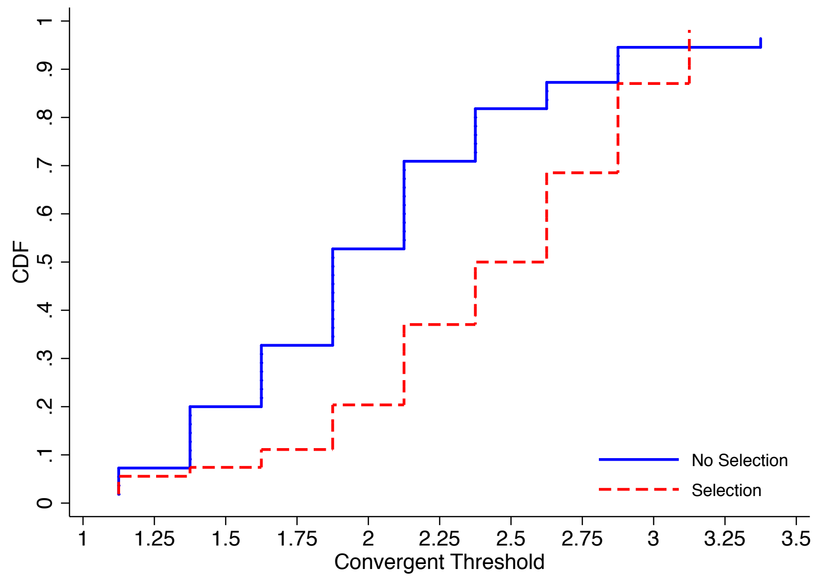

Figure 6: Distribution of convergent threshold choices, by treatment. Convergent threshold choices under Selection first order stochastically dominate choices under No Selection.

here we use the reported beliefs simply to assess what it is that subjects are paying attention to (if anything) and as a robustness check to confirm whether subjects are really being naive. Table 2 compares, for each treatment, the averages in the data and the subjects' average responses. For the averages in the data, we consider both the true, realized averages (as observed by the researchers) and the averages that would be estimated by a naive subject from the observed data. (The question on the chance A was good conditional on being pivotal was asked last but appears in the second row of the table; see Section 2 for details).

The first row in Table 2 shows the chance that A is good as observed in the data and reported by the subjects. In the No Selection treatment, the state was good 25% of the time and, of the times in which subjects got to observe whether A is good or bad (i.e., when the majority recommends A), alternative A turned out to be good 24.9% of the time (recall the true probability is 25% and that there is no selection, which explains why the true and naive estimates from the data are similar). On average, subjects report that the chance that A is good is 30.6%. For the Selection treatment, the state was good 25.6% of the time (again, the true probability is 25%). But, on average, subjects observe that, conditional on having information about A being good or bad, alternative A was good 56.1% of the time. As explained earlier,

mean values	No Selection treatment			Selection treatment		
	Data (true)	Data (naive)	Report	Data (true)	Data (naive)	Report
% Good	25.0	24.9	30.6	25.6	56.1	48.4
% Good | piv	26.1	24.9	28.0	0	56.1	44.6
% mistake | Good	49.7	49.9	43.4	0	50.1	36.1
% mistake | Bad	50.0	50.0	49.1	32.3	49.9	40.4

Table 2: Mean values of data and reported beliefs, by treatment. Reported beliefs are consistent with naive (biased) beliefs, not with sophisticated beliefs.

Legends: % Good: percentage of times that project A was good; % Good | piv: percentage of times that project A was good conditional on the subject being pivotal; % mistake | Good: percentage of times a computer mistakenly votes for B when project A is good; % mistake | Bad: percentage of times a computer mistakenly votes for A when project A is bad; Data (true): actual figure in the data; Data (naive): actual figure a naive subject would report given the data; Report: figure reported by subjects in Part II.

this higher number reflects the fact that the sample is biased because computers' strategies are correlated with the state of the world. On average, subjects report that the chance that A is good is 48.4%, which is much closer to the naive figure in the data (56.1%) than to the true figure (25.6%). In particular, it appears that subjects in both treatments are, on average, paying attention to the data, but they are doing so naively.

The second row in Table 2 shows the results when subjects are asked about the chance that A is good conditional on being pivotal. In the No Selection treatment, where the pivotal event conveys no information, the true and reported averages are similar to the unconditional case. In the Selection treatment, as explained earlier, there is not one case in which A is good when a subject is pivotal, so the realized proportion of good conditional on being pivotal is 0%. On average, subjects miss this point and report 44.6%.

Finally, the last two rows show realized rates and beliefs for the computers' mistakes. As expected, the true realized rates in the data are very close to the true rates, which are $m_G = m_B = 1/2$ under No Selection and $m_G = 0$, $m_B = 1/3$ under Selection. The naive estimates are given by the unconditional proportion of times that computers vote A, which is close to the true unconditional probability of $1/2$. In the No Selection treatment, subjects are on average correct to respond that computers' strategies are uninformative. In the Selection treatment, subjects realize that the rates of computers' mistakes is lower, but are far from realizing that computers make no mistakes when project A is good. Overall, it appears that, on average, subjects pay attention to the data, make *naive* inferences, do not realize that the computers make no mistakes when project A is good, and mostly fail to account for sample selection (though reported beliefs are slightly below naive estimates from the data).

Finding #4. Behavioral equilibrium over-predicts convergent thresholds in the Selection treatment:

While behavioral equilibrium correctly predicts the direction of the treatment effect, a more stringent test is whether it can rationalize the levels observed in the data. As discussed earlier, the average (median) convergent threshold is about 0.39 (0.62) points higher under Selection compared to the No Selection treatment, while theory predicts a difference of about 1 point (where the exact difference depends on the risk coefficient, see Figure 3). Similarly, we showed that reported beliefs are slightly lower than naive estimates from the data. There are two issues that we need to tackle, however, for a more precise comparison of the data to the theoretical predictions. First, we need to control for risk aversion. Second, the theoretical prediction is based on the assumption that beliefs have converged to steady-state beliefs, but the actual data observed by each subject is of course noisy and does not coincide with the theoretical, steady-state prediction.

To tackle the first issue, we use responses from Part III of the experiment to estimate a CRRA risk coefficient for each subject.[17] To account for the second issue, we assume that a subject's (naive) belief is determined by the observed data (i.e., the relative proportion of good vs. bad in her data) and not by the theoretical steady-state prediction. Figure 7 shows the comparison between data and theory when accounting for these two issues. On the horizontal axis, we plot, for each subject and for each round of the last 10 rounds, the threshold predicted by behavioral equilibrium (recall that this is also the Nash threshold for the No Selection treatment, but not for the Selection treatment) when using both her estimated risk coefficient and the data she observes up to that round. On the vertical axis, we plot the threshold chosen by the subject in that particular round. The size of a data point is proportional to the number of subjects characterized by that data point.

[17]We let subject i's threshold in the kth (out of five) decision in Part III be given by $x^*_{ik} = f(r_i, z_{ik}) + \varepsilon_{ik}$, where ε_{ik} is noise, z_{ik} is the known probability that A is good, r_i is the risk coefficient, and f is the optimal decision for a CRRA utility function; see Section 5.1, where we use this equation in the model of partial naiveté, for more details.

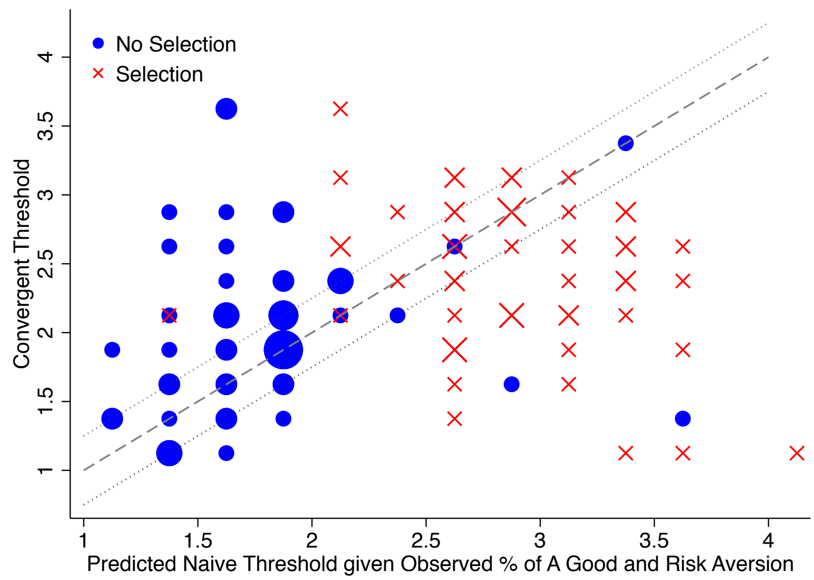

Figure 7: Observed vs. theoretical prediction of behavioral equilibrium. The figure plots the threshold predicted by behavioral equilibrium (based on the estimated risk coefficient and beliefs for each subject) against the actual choice, for rounds 91 through 100 (where the size of the plotted data point is proportional to the number of subjects). The figure confirms that behavioral equilibrium, despite being consistent with the treatment effect, over-predicts choices for the Selection treatment.

If the theoretical prediction were perfect, all data points would lie on the 45 degree dashed line in the figure. Of course, there are several reasons why the data might not perfectly fit the theory, including the fact that our estimates of risk aversion and beliefs are likely to be noisy. The main point from Figure 7, however, is that theory can rationalize a large proportion of the choices in the No Selection treatment. For example, 62% of all observations lie within 0.25 points of the 45 degree line. Instead, in the selection treatment, the predicted naive threshold is at least 0.25 points *higher* than the convergent threshold in 52% of the cases. In other words, we continue to see a systematic over-prediction of threshold choices by behavioral equilibrium in the Selection treatment.

Finding #5. Subjects are more likely to change their thresholds in a given round if they were pivotal in the previous round:

The evidence so far suggests that subjects are naive but that they partially account for the selection problem by choosing thresholds that are a bit lower than the naive behavioral equilibrium thresholds. One reasonable explanation for this behavior is that, while subjects do not know how to account for the information content of the computers' recommendations, they might be more likely to adjust their thresholds in rounds in which they are pivotal.

The top panel of Table 3 shows the results of a linear regression of an indicator variable for whether or not a subject changes her threshold in round t on two other indicator variables (and their interaction) that capture whether the subject was pivotal in the previous round (Piv_{t-1}) and whether project A was chosen by a majority in the previous period and hence she observed information about project A (Info_{t-1}). It is not surprising that observing some information (positive or negative) about project A in a previous period increases the probability that a subject will change her threshold choice; it does so by about 1.6 percentage points (from a baseline of about 4%). The key finding, however, is that the interaction effect is more than three times stronger: A subject is 5.3 percentage points more likely to change her threshold if she received information *and* was pivotal in the previous round.

While the top panel looks at the chance of adjusting the threshold, the bottom panel of Table 3 looks at the magnitude of the change. The dependent variable measures the difference between the threshold in rounds t and $t-1$ and the independent variables include indicators for whether or not a subject was pivotal and observed pos-

Dep Var: $\mathbf{1}\{T_t \neq T_{t-1}\}$	Coeff.	Std. Err.
Constant	0.040***	0.006
Piv_{t-1}	0.003	0.005
Info_{t-1}	0.016**	0.007
$\text{Piv}_{t-1} \times \text{Info}_{t-1}$	0.053***	0.013

Legends: The dependent variable and the controls are dummy variables. $\mathbf{1}\{T_t \neq T_{t-1}\}$: takes value 1 if the threshold in period t is different than the threshold in period $t-1$. Piv_{t-1}: takes value 1 if the subject was pivotal in the previous period. Info_{t-1}: takes value 1 if in the previous period the subject received feedback on whether project A was good or not.

Dep Var: $T_t - T_{t-1}$	Coeff.	Std. Err.
Constant	0.008***	0.002
$(\text{Piv and Pay 5})_{t-1}$	0.020	0.025
$(\text{Piv and Pay 1})_{t-1}$	-0.092***	0.018
$(\text{Not Piv and Pay 5})_{t-1}$	-0.002	0.005
$(\text{Not Piv and Pay 1})_{t-1}$	-0.024***	0.008

Legends: All controls are dummy variables. $(\text{Piv and Pay 5})_{t-1}$ takes value 1 if the subject was pivotal and received a payoff of 5 in the previous period. Other dummy variables are named accordingly. The excluded event is the case when the subject did not receive information in the previous period.

Table 3: Reduced Form Analysis: Reaction in Threshold to events in previous period.

Notes: (*), (**), (***) indicate significance at the 1, 5 and 10% level respectively. In both cases we report the results of fixed effects panel regressions and we cluster standard errors by subject. Both regressions include 109 subjects that converged and for each subject we use the last 74 periods of part 1 (we lose one observation due to the lag). The regressions pool subjects from both treatments. Conclusions do not change if the analysis is conducted by treatment or if we add time dummies.

itive or negative information about project A. Observing a payoff of 1 leads subjects to decrease their threshold by 0.024 points on average when they were not pivotal and by about four times this magnitude, 0.092, when they were pivotal.[18]

The findings from Table 3 confirm that subjects tend to react more to pivotal vs. non-pivotal events, which explains why their behavior can be consistent with a

[18]These numbers are small because the baseline probability of changing the threshold in a given round is small; the results are similar if we restrict the regression to rounds in which a subject changes her threshold. Also, the coefficient on being pivotal and observing a payoff of 5 is positive (as expected), but it is estimated with a higher standard error due to the fact that this coefficient is only identified from the No Selection treatment (because the event has zero probability under Selection).

partial adjustment of selection despite their inherent naiveté.

5 A Model of Partial Naiveté

Motivated by the finding that subjects are not sophisticated but seem, nevertheless, to be responding more to feedback from pivotal rounds, we now propose and estimate a model of partial naiveté.

5.1 Model

Recall from Section 3.3 that the steady-state belief in a behavioral equilibrium is given by $z(x^*)$ in equation (1), page 16, which denotes the probability that A is observed to be good conditional on having observed whether A is good or bad—which, in particular, depends on the subjects' steady-state strategy, x^*. In a behavioral equilibrium, subjects are assumed to put equal weight to an observation about A irrespective of whether or not she was pivotal. We now generalize this notion by letting $\eta \in (0, \infty)$ be a parameter that denotes the weight that a subject puts on pivotal vs. non-pivotal periods. The corresponding steady-state belief is

$$z(x^*, \eta) = \frac{\left((1-m_G)^2 + \eta 2 m_G (1-m_G) \frac{(x^*-1)}{4}\right) p}{\left((1-m_G)^2 + \eta 2 m_G (1-m_G) \frac{(x^*-1)}{4}\right) p + \left((m_B^2 + \eta 2 m_B (1-m_B) \frac{(x^*-1)}{4}\right) (1-p)}, \tag{3}$$

where η now multiplies the events that a majority recommends A and that the subject's recommendation is pivotal.

Assuming, once again, CRRA utility function for convenience, the steady-state strategy $x^*_\eta(r)$ is the unique solution to

$$z(x^*, \eta) \times u_r(5) + (1 - z(x^*, \eta)) \times u_r(1) = u_r(x^*). \tag{4}$$

Nash and behavioral equilibrium are special cases of this model. As η goes to infinity, a subject puts increasingly higher weight on pivotal periods, and it is easy to see that the threshold converges to the Nash equilibrium threshold, i.e., $\lim_{\eta\to\infty} x_\eta(r) = x_{NE}(r)$ for all risk coefficients r. And behavioral equilibrium corresponds to the case $\eta = 1$, which places equal weight on pivotal vs. non-pivotal periods, i.e., $x_1(r) = x_{BE}(r)$ for all risk coefficients r. The parameter η captures intermediate cases where subjects are naive but account for selection by putting higher weight on feedback from pivotal rounds. One way to rationalize this belief formation process is to assume that subjects pay attention to data from pivotal and non-pivotal rounds with probability α and β, respectively. Equation (3) can be interpreted as the steady-state belief where $\eta = \alpha/\beta$.[19]

We now specialize the model to the parameters chosen in each of our treatments. For the No Selection treatment ($p = 1/4, m_G = m_B = 1/2$), equation (3) becomes

$$z(x^*, \eta) = \frac{\left(.25 + \eta .5^{\frac{(x^*-1)}{4}}\right).25}{\left(.25 + \eta .5^{\frac{(x^*-1)}{4}}\right).25 + \left(.25 + \eta .5^{\frac{(x^*-1)}{4}}\right).75} = .25.$$

As explained earlier, there is no selection in the data and beliefs are always .25, irrespective of the weight placed on pivotal vs. non-pivotal rounds. For the Selection treatment ($p = 1/4, m_G = 0, m_B = 1/3$),

$$z(x^*, \eta) = \frac{.25}{.25 + \left(\frac{1}{9} + \eta \frac{4}{9}^{\frac{(x^*-1)}{4}}\right).75}.$$

As η increases, more weight is placed on pivotal rounds, where A always turns out bad, and, therefore, $z(x^*, \cdot)$ decreases.

Figure 8 plots the threshold prediction for several values of η. The prediction for the No Selection treatment is the same for all values of η and given by the solid line. The prediction for the Selection treatment is decreasing in η, with $\eta = 1$ (behavioral equilibrium) and $\eta \approx \infty$ (Nash equilibrium) representing two extreme cases in the figure.[20] The figure also illustrates that the optimal threshold is not very responsive to η; for example, a risk neutral subject, $r = 0$, would exhibit no treatment effect even if she placed $\eta = 8$ times more weight on pivotal vs. non-pivotal periods. The reason is that very high weights are needed to compensate for the fact that the probability of being pivotal is small to begin with (1/3 in this case). Because we find a positive treatment effect (Finding #1) but also that behavioral equilibrium over-predicts (Finding #4), Figure 8 already suggests that η is on average between 1

[19]Note that in this case α and β are not separately identified; only their ratio affects beliefs.

[20]The model also allows for $\eta < 1$, which means that non-pivotal rounds receive relatively higher weight.

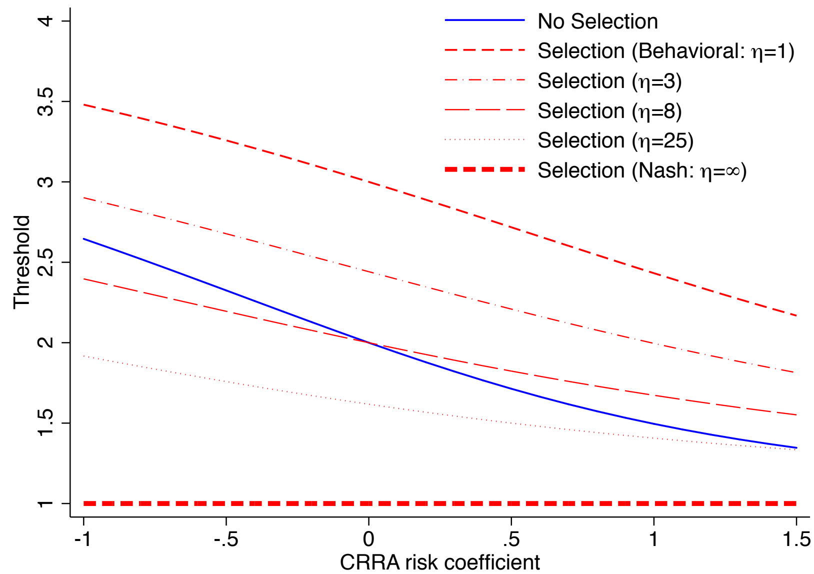

Figure 8: Theoretical prediction for Selection and No Selection treatments for several values of η.

and 8. In the next section we obtain a more precise estimate of the distribution of η in the population.

5.2 Empirical estimation and results

For each subject we have data from decisions in Part I (either No Selection and Selection treatments) and Part III (decision problem). For each case, we postulate that thresholds are chosen according to

$$x^*_{ik} = f(r_i, z_{ik}) + \varepsilon_{ik}, \tag{5}$$

where x^*_{ik} is the desired threshold choice of subject i in decision k, r_i is the CRRA risk coefficient, z_{ik} is subject i's belief that A is good, ε_{ik} is noise that affects decisions, and f is a function that maps the risk coefficient and belief to an optimal threshold choice. We assume that f is derived from a CRRA utility function, i.e., $f_{ik} = f(r_i, z_{ik})$ solves $z_{ik} \times u_{r_i}(5) + (1 - z_{ik}) \times u_{r_i}(1) = u_{r_i}(f_{ik})$, where u_{r_i} is the CRRA utility function.[21]

[21]In the data, we do not observe the exact threshold x^*_{ik} because we only observe a decision contingent on a finite number of values of x. Each value of x^*_{ik}, however, translates immediately into a choice in our environment. For example, if $x^*_{ik} = 2.33$, this means that a subject would choose A

We proceed in two stages. First, we use data from the decision problem (Part III) to estimate the distribution of risk coefficients r_i and decision noise ε_{ik}. Recall that in the decision problem, subjects face five instances of the problem from Part I, with two exceptions: (i) the outcome is determined by their own decision alone, without interference from other players/computers, and (ii) subjects are told the true probability that A is good. Each subject faces five cases, and the probability that A is good is known by the subject in each case. Thus, the value of z_{ik} is fixed for each of these five cases, $k = D_1, ..., D_5$, were D stands for decision problem. In particular, we have $z_{iD_1} = .1$, $z_{iD_2} = .3$, $z_{iD_3} = .5$, $z_{iD_4} = .7$, and $z_{iD_5} = .9$. For concreteness, we assume that the risk coefficient $r \sim N(\mu_r, \sigma_r^2)$ and the decision noise $\varepsilon \sim N(\mu_\varepsilon, \sigma_\varepsilon^2)$ are normally distributed and independent of each other and across subjects, and we estimate the parameters using (simulated) maximum likelihood.

In the second stage, we use the estimates from the previously described stage and data from Part I of the experiment to identify the extent to which subjects are partially naive. To be consistent with the steady-state model, we use data from rounds in which behavior has stabilized. In particular, we look at data from rounds T to 100, where we vary T from 70 to 100 for robustness purposes. We let $k = N_T, N_{T+1}, ..., N_{100}$ and $k = S_T, S_{T+1}, ..., S_{100}$ index data from rounds T through 100 in the No Selection and Selection treatments, respectively.

The difference with respect to the decision problem is that subjects have beliefs about z_{ik} and we do not directly observe these beliefs. We follow the model in assuming that these beliefs depend endogenously on the feedback observed by the subjects. We infer beliefs using the data actually observed by the subjects rather than the theoretical steady state. The reason is that the steady-state belief predicted by the theory is accurate provided that enough time has been spent in the steady state; according to our data, in contrast, most time has been spent outside the steady state.[22] In particular, we assume that, for observations from both treatments $k = N_t$ and S_t and all rounds $t = T, ..., 100$, a subject's belief is

$$z_{ik} = g(data_{ik}, \eta_i) + \nu_{ik},$$

for all values of x lower or equal than 2.25 and B for all values of x higher or equal than 2.5. We account for this issue when writing down the likelihood function.

[22] Alternatively, in Appendix A we estimate the model using beliefs reported in Part II and obtain similar qualitative results, although the mean estimate for partial naiveté is slightly lower. As pointed out by the literature, however, there might be issues from using reported beliefs as opposed to estimating beliefs from actions (e.g., Nyarko and Schotter, 2002).

where $g(data_{ik}, \eta_i)$ is the empirical counterpart of equation (3), which depends on the data observed by subject i up to round k and her parameter of naiveté η_i, and ν_{ik} denotes noise in the subject's estimation process in round k.[23] Intuitively, data from the No Selection treatment is used to identify the belief noise ν (since the function $g(\cdot)$ is essentially constant in η under No selection) and data from the Selection treatment is used to identify η. For concreteness, we assume that the logarithm of the naiveté coefficient $\ln \eta \sim N(\mu_{\ln \mu}, \sigma^2_{\ln \mu})$ and the belief formation noise $\nu \sim N(\mu_\nu, \sigma^2_\nu)$ are normally distributed and independent of each other and across subjects, and we estimate the parameters using (simulated) maximum likelihood.[24]

The top panel of Table 4 presents the maximum likelihood estimates (including standard deviations and 95% confidence intervals) when we use data from the last 10 rounds of part 1. With the estimates for the distribution of $\ln \eta$, we plot the resulting distribution of η (Figure 9) and obtain the mean and median of η as 6.08 and 3.14, respectively. Thus, the median subject puts 3.14 times more weight on pivotal vs. non-pivotal events.[25] The result is consistent with the reduced-form results from Section 4, which suggested that subjects were three or four times more likely to change their thresholds after facing a pivotal event.

The bottom panel of Table 4 presents further information on the median of η. Based on bootstrapping the maximum likelihood estimates we obtain a distribution for the median of η. If we use data from the last 10 rounds of part 1 ($T > 90$), the 5th and 95th percentiles of the median of η are 1.70 and 6.04, respectively. This shows that the estimate of the median is concentrated around the maximum likelihood estimate (3.14) and is far from being consistent with Nash behavior. As explained earlier, even much higher weights on pivotal rounds are not enough to approximate Nash behavior, since subjects are pivotal with a relatively small probability of 1/3. Thus, the increased relative weights on pivotal events is not nearly enough to correct for mistakes. As a robustness exercise, the other columns of the bottom panel show how the computations change depending on the data from part 1 that we include in

[23]Under this specification, z_{ik} might fall outside the $[0, 1]$ interval, in which case we set it equal to 0 or 1. This turns out not to be a serious constraint because the estimated variance of ν is fairly small.

[24]In particular, the naiveté coefficient η has a lognormal distribution with mean $E(\eta) = e^{\mu_{\ln \eta} + \frac{\sigma^2_{\ln \eta}}{2}}$ and variance $Var(\eta) = (e^{\sigma^2_{\ln \eta}} - 1)e^{2\mu_{\ln \eta} + \sigma^2_{\ln \eta}}$.

[25]Given the asymmetry in the distribution of η we focus on the median as a measure of central tendency. In the Appendix A we present more detailed information on the mean of η.

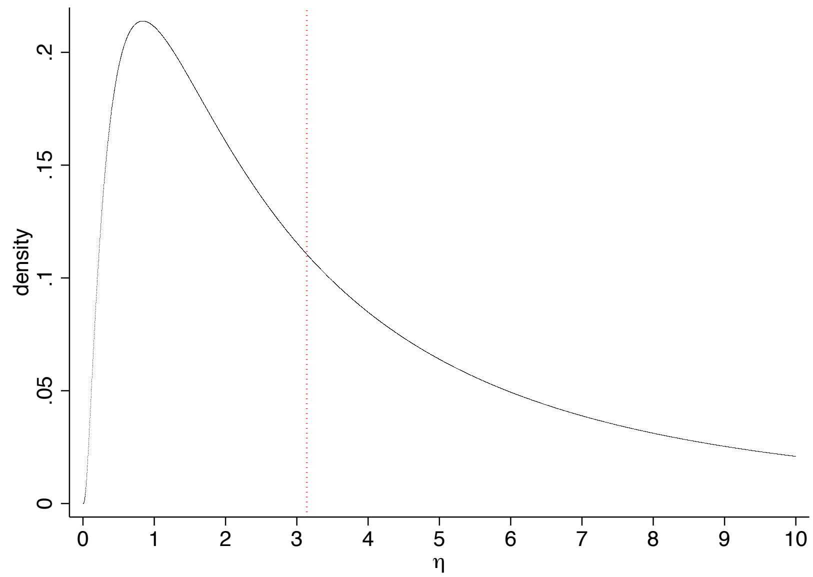

Figure 9: Estimated distribution of η. The dashed line indicates the median at 3.14.

the estimation. We confirm that the main conclusions are unaffected by the choice of T.[26]

Finally, we briefly comment on the other estimates of the top panel of Table 4. The results for the risk coefficient and noise levels appear to be reasonable and consistent with previous work. For example, the mean subject is risk averse with a risk coefficient of relative risk aversion of 0.533, and 95% of the population has a risk coefficient between 0.283 and 0.905, which is consistent with previous estimates (see Holt and Laury (2002), Harrison and Ruström (2008)).[27] The estimates also suggest that it is important to account for noise in actions to avoid biasing our results for the coefficient of partial naiveté. The mean decision noise is 0.143, which is a bit more than half the distance between two thresholds (recall that thresholds are 0.25 points apart). Noise in beliefs is estimated to be fairly small, with a mean of about 0.02 and standard deviation of 0.14.

[26] Appendix A provides detailed results on several robustness exercises.

[27] Following the maximum likelihood procedure in Harrison and Ruström (2008) we can estimate the coefficient of risk aversion using answers to the Holt-Laury choice lists that we collected at the end of the session. The coefficient equals 0.574 or 0.567 depending on whether we use only subjects who converge or all subjects. These estimates are comparable to those reported in Table 4 (which use data from Part III) and to previous estimates in the literature; for example, Harrison and Ruström (2008) report an estimate of 0.66 using data from Hey and Orme (1994).

	Estimate	Std. Err.	95% Conf. Interval
$\mu_{\ln\eta}$	1.145	3.681	[0.347, 2.337]
$\sigma_{\ln\eta}$	1.150	1.114	[0.070, 2.032]
μ_r	0.533	0.157	[0.283, 0.905]
σ_r	0.461	0.128	[0.237, 0.758]
μ_ϵ	0.143	0.071	[0.033, 0.317]
σ_ϵ	0.453	0.064	[0.325, 0.565]
μ_ν	0.024	0.042	[-0.042, 0.107]
σ_ν	0.140	0.096	[0.105, 0.227]

Maximum likelihood Estimates. Standard errors and the 95% confidence intervals are computed using 1000 bootstrap repetitions. The estimation uses data from part 1 for rounds higher than 90 ($T > 90$).

Percentile	$T = 100$	$T > 90$	$T > 80$	$T > 70$
5	1.98	1.70	0.75	0.51
25	2.53	2.35	1.73	1.78
50	3.03	3.00	2.62	2.89
75	3.67	3.69	3.41	4.58
95	4.94	6.04	5.55	21.64

Statistics of the Median of η using the Bootstrap. The bootstrap delivers 1000 estimations of the parameters of the model. For each repetition we compute the median of η and the table reports percentiles of the distribution. Each column indicates the rounds of part 1 that were included in the estimation.

Table 4: Maximum Likelihood Estimation and the Distribution of η.

6 Conclusion

The revolution in information economics that took shape in the 1970s has substantially changed our understanding of markets and strategic interactions in the presence of asymmetric information. As economists have tried to take these models to data, they have faced the challenge of identifying the primitives of their models. Most of the literature, however, implicitly assumes that subjects have no trouble choosing optimal actions given the primitives and the strategies of other players. In this paper, we conduct an experiment to understand how subjects cope with the type of endogenous selection problems that arise in many settings with asymmetric information. We find that Nash equilibrium is a poor predictor of behavior. On the other hand, the notion of behavioral equilibrium, which assumes that subjects do not account for selection at all, predicts the treatment effect fairly well.

When examining the data more closely we find, however, that subjects do tend to partially account for selection. We provide reduced form evidence that the partial accounting is likely due to the tendency to respond more to feedback from pivotal events. We then propose and estimate a new model of partial naiveté to quantify this effect and find that the median subject places about three times wore weight on pivotal vs. non-pivotal rounds. This partial accounting of selection does little to correct the mistake because of the relatively low probability of being pivotal.

While more experiments are needed to confirm behavior in these settings, our results suggest that we might want to think more seriously about the types of identification problems faced by economic agents. More generally, it would be troubling if long held beliefs about behavior in markets with asymmetric information depended crucially on unrealistic behavioral assumptions. It might then be important to revisit some of these long held beliefs under solution concepts that account for bounded rationality and data limitations.

References

Banks, Jeffrey, Mark Olson, and David Porter, "An experimental analysis of the bandit problem," *Economic Theory*, 1997, *10* (1), 55–77.

Barrett, Garry F and Stephen G Donald, "Consistent tests for stochastic dominance," *Econometrica*, 2003, *71* (1), 71–104.

Bazerman, Max H and William F Samuelson, "I won the auction but don't want the prize," *Journal of Conflict Resolution*, 1983, *27* (4), 618–634.

Brandts, Jordi and Gary Charness, "The strategy versus the direct-response method: a first survey of experimental comparisons," *Experimental Economics*, 2011, *14* (3), 375–398.

Charness, G. and D. Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–236.

Cox, James C, Bruce Roberson, and Vernon L Smith, "Theory and behavior of single object auctions," *Research in experimental economics*, 1982, *2*, 1–43.

Crawford, V.P. and N. Iriberri, "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Dekel, E., D. Fudenberg, and D.K. Levine, "Learning to play Bayesian games," *Games and Economic Behavior*, 2004, *46* (2), 282–303.

Esponda, I., "Behavioral equilibrium in economies with adverse selection," *The American Economic Review*, 2008, *98* (4), 1269–1291.

_ **and D. Pouzo**, "Learning Foundation for Equilibrium in Voting Environments with Private Information," *working paper*, 2012.

_ **and** _ , "An Equilibrium Framework for Players with Misspecified Models," *working paper*, 2014.

_ **and E. Vespa**, "Hypothetical Thinking and Information Extraction in the Laboratory," *American Economic Journal: Microeconomics*, 2014, *6* (4), 180–202.

Eyster, E. and M. Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Fudenberg, D. and D.K. Levine, *The theory of learning in games*, Vol. 2, The MIT press, 1998.

Fudenberg, Drew and Alexander Peysakhovich, "Recency, Records and Recaps: Learning and non-equilibrium behavior in a simple decision problem," *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014.

Harrison, G.W. and E.E. Ruström, 2008. "Risk Aversion in the Laboratory," in *Risk Aversion in Experiments (Research in Experimental Economics, Volume 12)*, ed. by J.C. Cox, G.W. Harrison. Emerald Group Publishing Limited.

Hey, John D and Chris Orme, "Investigating generalizations of expected utility theory using experimental data," *Econometrica: Journal of the Econometric Society*, 1994, pp. 1291–1326.

Holt, Charles A and Roger Sherman, "The loser's curse," *The American Economic Review*, 1994, pp. 642–652.

_ **and Susan K Laury**, "Risk aversion and incentive effects," *American economic review*, 2002, *92* (5), 1644–1655.

Ivanov, A., D. Levin, and M. Niederle, "Can relaxation of beliefs rationalize the winner's curse?: an experimental study," *Econometrica*, 2010, *78* (4), 1435–1452.

Jehiel, P., "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

_ **and F. Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Kagel, J.H. and D. Levin, "The winner's curse and public information in common value auctions," *The American Economic Review*, 1986, pp. 894–920.

_ **and** _ , *Common value auctions and the winner's curse*, Princeton Univ Pr, 2002.

Nyarko, Yaw and Andrew Schotter, "An experimental study of belief learning using elicited beliefs," *Econometrica*, 2002, *70* (3), 971–1005.

Smith, Vernon L, "An experimental study of competitive market behavior," *The Journal of Political Economy*, 1962, pp. 111–137.

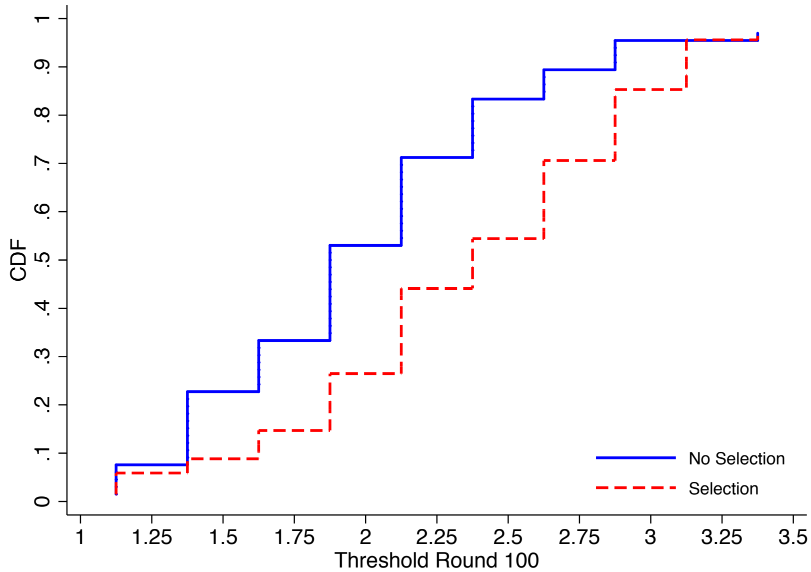

Figure 10: Distribution of round 100 thresholds, by treatment using All Subjects

Online Appendix A

In this appendix we report several robustness exercises to the findings presented in the main text.

Figure 10 updates Figure 6 by including all subjects. Since not all subjects have a convergent threshold we focus on their Round 100 choices. There are no qualitative differences and thresholds in the Selection Treatment still first order stochastically dominate thresholds in the No Selection treatment.

Table 5 expands the information in Table 2. It reports standard errors and the figures when all subjects are included in the analysis. When we include all subjects there are only minor quantitative deviations and the findings reported in the paper still hold.

Figure 11 is the counterpart to Figure 7, when we include all subject in the analysis and focus on their Round 100 choices. We observe the same pattern that we reported in the text for subjects who converge. In the No Selection treatment 55% of the subjects have round 100 thresholds that are within a 0.25 points band of the predicted naive threshold. In the Selection treatment for 54% of the subjects the predicted naive threshold is at least 0.25 points higher than their round 100 choice.

The difference between Figure 12 and Figure 7 resides in how beliefs are computed. For Figure 7 we compute beliefs using the proportion of times that subjects observed project A to be good. Figure 12 uses the beliefs that subjects reported in Part II of the experiment. Although less salient, the reported pattern remains. In the No Selection treatment 46% of subjects have thresholds that are within a 0.25 points

	No Selection treatment						Selection treatment					
	Subjects who Converged			All Subjects			Subjects who Converged			All Subjects		
	Data (true)	Data (naive)	Report	Data (true)	Data (naive)	Report	Data (true)	Data (naive)	Report	Data (true)	Data (naive)	Report
% Good	25.0	24.9	30.6	24.9	24.5	29.3	25.6	56.1	48.4	25.5	56.2	47.8
	(4.3)	(6.6)	(15.3)	(4.3)	(6.5)	(14.8)	(4.6)	(7.5)	(18.1)	(4.7)	(7.9)	(18.4)
% Good | piv	26.1	24.9	28.0	25.6	24.5	27.3	0	56.1	44.6	0	56.2	43.9
	(6.1)	(6.6)	(15.6)	(5.9)	(6.5)	(15.5)	(0)	(7.5)	(20.8)	(0)	(7.9)	(21.1)
% mistake | Good	49.7	49.9	43.4	50.4	50.2	43.9	0	50.1	36.1	0	50.5	35.4
	(7.5)	(3.1)	(16.4)	(7.3)	(3.2)	(17.3)	(0)	(4.3)	(18.7)	(0)	(4.3)	(18.4)
% mistake | Bad	50.0	50.0	49.1	49.9	49.7	49.1	32.3	49.9	40.4	32.3	49.5	39.3
	(4.4)	(3.1)	(17.1)	(4.1)	(3.1)	(18.9)	(3.9)	(4.3)	(17.9)	(3.8)	(4.3)	(17.8)

Table 5: Mean values in the data and reported beliefs, by treatment for subjects who converge and all subjects.

Legends: % Good: percentage of times that project A was good; % Good | piv: percentage of times that project A was good conditional on the subject being pivotal; % mistake | Good: percentage of times a computer mistakenly votes for B when project A is good; % mistake | Bad: percentage of times a computer mistakenly votes for A when project A is bad; Data (true): actual figure in the data; Data (naive): actual figure a naive subject would report given the data; Report: figure reported by subjects in Part II. Standard errors between parentheses.

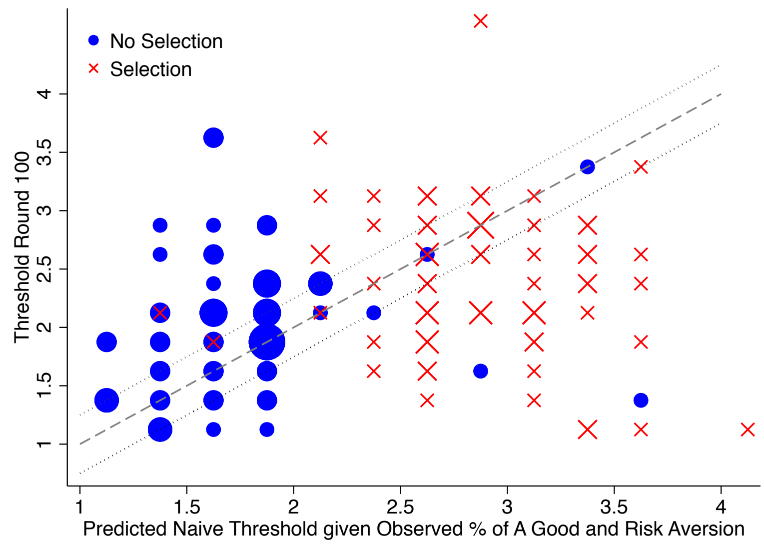

Figure 11: Observed vs. theoretical prediction of behavioral equilibrium. (Using All Subjects) The figure plots the threshold predicted by behavioral equilibrium (based on the estimated risk coefficient and beliefs for each subject) against the actual choice, for round 100 (where the size of the plotted data point is proportional to the number of subjects).

band of the predicted naive threshold (25% above the band and 29% below). In the selection treatment for 41% of the subjects the predicted naive threshold is at least 0.25 points higher than the round 100 choice (33% inside the band, 26% above).

Table 6 shows that the reported findings are unaffected if we extend the reduced form analysis to include all subjects.

Table 7 provides details on the Maximum Likelihood Estimates for the model of partial naiveté. Table 7a reproduces the estimates of the coefficients related to risk and noise in actions for subjects who converged (presented in the top panel of Table 4), and provides the estimates when we include all subjects in the estimation. The table shows that there are only small differences in the estimates. For example, the estimate of the mean of risk aversion being slightly higher for all subjects (0.653) than for subjects who converge (0.533), but both estimates are close to each other and to previous estimates in the literature.

Table 7b presents the estimates of the coefficients related to the distribution of η and noise in beliefs (ν). The table shows how the estimates change depending on the sample of part 1 rounds included in the estimation and on whether we use only subjects who converge or all subjects. There are no large differences in the coefficients in either direction. We do note that the confidence intervals are generally larger when we include all subjects than when we focus only on subjects who converge. To capture the effect of the estimates on η, Table 7c presents information on the related distributions of the median of η. In all cases the 50th percentile of the median

Dep Var: $\mathbf{1}\{T_t \neq T_{t-1}\}$	Coeff.	Std. Err.
Constant	0.083***	0.007
Piv_{t-1}	0.005	0.008
Info_{t-1}	0.022***	0.007
$\text{Piv}_{t-1} \times \text{Info}_{t-1}$	0.063***	0.013

Legends: The dependent variable and the controls are dummy variables. $\mathbf{1}\{T_t \neq T_{t-1}\}$: takes value 1 if the threshold in period t is different than the threshold in period $t-1$. Piv_{t-1}: takes value 1 if the subject was pivotal in the previous period. Info_{t-1}: takes value 1 if in the previous period the subject received feedback on whether project A was good or not.

Dep Var: $T_t - T_{t-1}$	Coeff.	Std. Err.
Constant	0.016***	0.004
$(\text{Piv and Pay 5})_{t-1}$	-0.020	0.035
$(\text{Piv and Pay 1})_{t-1}$	-0.158***	0.031
$(\text{Not Piv and Pay 5})_{t-1}$	-0.001	0.006
$(\text{Not Piv and Pay 1})_{t-1}$	-0.025**	0.011

Legends: All controls are dummy variables. $(\text{Piv and Pay 5})_{t-1}$ takes value 1 if the subject was pivotal and received a payoff of 5 in the previous period. Other dummy variables are named accordingly. The excluded event is the case when the subject did not receive information in the previous period.

Table 6: Reduced Form Analysis: Reaction in Threshold to events in previous period. (All Subjects)

Notes: (*), (**), (***) indicate significance at the 1, 5 and 10% level respectively. In both cases we report the results of fixed effects panel regressions and we cluster standard errors by subject. Both regressions include all 134 subjects and for each subject we use the last 74 periods of part 1 (we lose one observation due to the lag). The regressions pool subjects from both treatments. Conclusions do not change if the analysis is conducted by treatment or if we add time dummies.

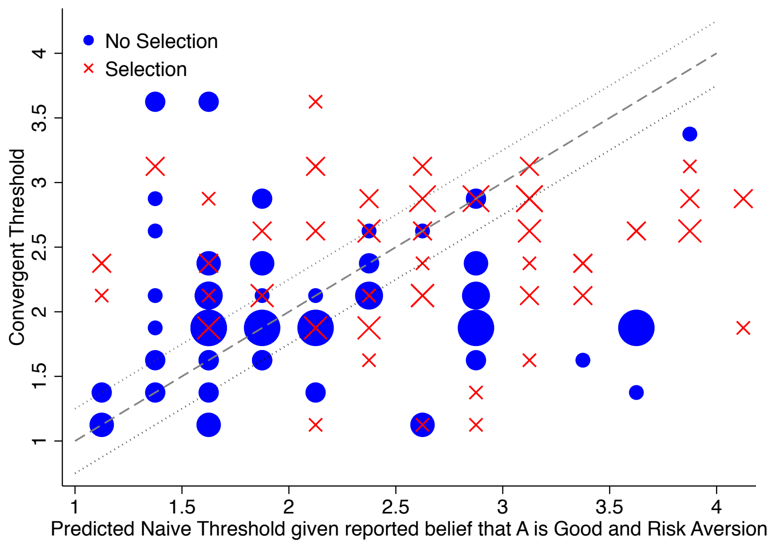

Figure 12: Observed vs. theoretical prediction of behavioral equilibrium. (Subjects who Converged) The figure plots the threshold predicted by behavioral equilibrium (based on the estimated risk coefficient and beliefs for each subject) against the actual choice, for rounds 91 through 100 (where the size of the plotted data point is proportional to the number of subjects). Beliefs are computed using reports in Part II of the experiment.

distribution is between 2.6 and 4.2, consistent with our main finding reported in the text.

In the text we focus on the median as a measure of central tendency for η. Table 8 allows to compare the reports on the distribution of the median of η to the distribution of the mean. Given the shape of the distribution (see 9), the mean is expected to be higher than the median and this is verified in 8 for all cases. The difference may be large as the mean can be severely affected by outliers.

As explained in Section 5, in the model that we estimate we compute beliefs (z) as a function of observed data, the coefficient of naiveté (η) and noise in beliefs (ν) (i.e. $z = g(data, \eta) + \nu$). Alternatively, we could use Maximum Likelihood to find the estimates of η that make the beliefs under which behavior in round 100 is optimal closest to the beliefs reported in Part II of the experiment. In particular, the estimation is comparable to the $T = 100$ estimates of Table 7b. For subjects who converged, Table 9 reproduces the estimates for round $T = 100$ of the model of partial naiveté and the estimates using reported beliefs.[28] The estimates of the median (mean) of η are 3.10 and 2.05 (4.49 and 2.61) using the model of partial naiveté and reported beliefs respectively. In other words, using reported beliefs leads to lower estimates of the mean and median of η.

[28]There are only small quantitative changes when we estimate the model using reported beliefs using all subjects.

	Subjects who converged			All Subjects		
	Estimate	Std. Err.	95% Conf. Interval	Estimate	Std. Err.	95% Conf. Interval
μ_r	0.533	0.157	[0.283, 0.905]	0.653	0.171	[0.382, 1.042]
σ_r	0.461	0.128	[0.237, 0.758]	0.519	0.131	[0.278, 0.774]
μ_ϵ	0.143	0.071	[0.033, 0.317]	0.186	0.077	[0.061, 0.365]
σ_ϵ	0.453	0.064	[0.325, 0.565]	0.475	0.050	[0.370, 0.566]

(a) Maximum likelihood Estimates of coefficients related to risk (r) and noise in actions (ϵ). Standard errors and the 95% confidence intervals are computed using 1000 bootstrap repetitions.

	Subjects who converged				All Subjects			
	$T = 100$	$T > 90$	$T > 80$	$T > 70$	$T = 100$	$T > 90$	$T > 80$	$T > 70$
$\mu_{\ln \eta}$	1.134	1.145	1.222	1.094	1.668	1.144	1.310	1.221
	[0.602, 1.740]	[0.347, 2.337]	[-2.347, 3. 323]	[-2.162, 5.402]	[0.684, 1.806]	[0.040, 8.783]	[-2.836, 6.047]	[-0.940, 7.738]
$\sigma_{\ln \eta}$	0.857	1.150	1.113	1.430	1.274	0.977	1.066	0.968
	[0.099, 1.385]	[0.070, 2.032]	[0.086, 9.834]	[0.103, 10.67]	[0.382, 1.559]	[0.038, 6.118]	[0.099, 10.86]	[0.061, 13.53]
μ_ν	0.035	0.024	0.033	0.028	0.115	0.024	0.049	0.043
	[-0.018, 0.093]	[0.283, 0.905]	[-0.111, 0.164]	[-0.121, 0.252]	[-0.013, 0.092]	[-0.051, 0.193]	[-0.103, 0.220]	[-0.101, 0.234]
σ_ν	0.079	0.140	0.141	0.148	0.156	0.154	0.139	0.069
	[0.000, 0.136]	[0.237, 0.758]	[0.095, 0.273]	[0.094, 0.373]	[0.000, 0.145]	[0.114, 0.255]	[0.120, 0.275]	[0.108, 0.387]

(b) Maximum Likelihood Estimates of coefficients related to the weight on pivotal rounds (η) and noise in beliefs (ν). 95% confidence intervals computed with bootstrapping between brackets. Columns in each case indicate the rounds of part 1 that were included in the estimation.

Percentile	Subjects who converged				All Subjects			
	$T = 100$	$T > 90$	$T > 80$	$T > 70$	$T = 100$	$T > 90$	$T > 80$	$T > 70$
5	1.98	1.70	0.75	0.51	2.19	1.77	0.76	0.84
25	2.53	2.35	1.73	1.78	2.77	3.09	2.61	2.90
50	3.03	3.00	2.62	2.89	3.32	4.07	3.42	4.17
75	3.67	3.69	3.41	4.58	4.07	5.70	4.71	6.47
95	4.94	6.04	5.55	21.64	5.47	66.65	42.45	104.68

(c) Statistics of the Median of η using the Bootstrap. The bootstrap delivers 1000 estimations of the parameters of the model. For each repetition we compute the median of η and the table reports percentiles of the distribution.

Table 7: Maximum Likelihood Estimates: Subjects who Converged and All Subjects.

	Median η				Mean η			
Percentile	$T = 100$	$T > 90$	$T > 80$	$T > 70$	$T = 100$	$T > 90$	$T > 80$	$T > 70$
5	1.98	1.70	0.75	0.51	2.40	2.19	1.35	1.28
25	2.53	2.35	1.73	1.78	3.59	3.03	2.68	3.08
50	3.03	3.00	2.62	2.89	4.58	4.87	4.17	5.23
75	3.67	3.69	3.41	4.58	5.82	6.27	6.03	8.66
95	4.94	6.04	5.55	21.64	8.73	11.54	20.35	180.6

Table 8: Mean and Median of η using the Bootstrap. The bootstrap delivers 1000 estimations of the parameters of the model. For each repetition we compute the mean and median of $\eta$and the table reports percentiles of the distribution of the mean and median respectively. Each column indicates the rounds of part 1 that were included in the estimation.

	Partial Naiveté Model ($T = 100$)			Using Reported Beliefs		
	Estimate	Std. Err.	95% Conf. Interval	Estimate	Std. Err.	95% Conf. Interval
$\mu_{\ln \eta}$	1.134	0.291	[0.602, 1.740]	0.720	0.170	[0.383, 1.056]
$\sigma_{\ln \eta}$	0.857	0.333	[0.099, 1.385]	0.689	0.282	[0.001, 1.168]
μ_ν	0.035	0.028	[-0.018, 0.093]	0.049	0.015	[0.016, 0.074]
σ_ν	0.079	0.038	[0.000, 0.136]	0.139	0.017	[0.096, 0.159]

Table 9: Partial Naiveté Model v. Reported Beliefs: Maximum Likelihood Estimates of coefficients related to the weight on pivotal rounds (η) and noise in beliefs (ν). 95% confidence intervals and standard errors computed with bootstrapping between brackets.

Online Appendix B: Instructions[29]

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation with cash, privately at the end of the session. What you earn depends partly on your decisions and partly on chance.

The entire session will take place through computer terminals and there will be no interaction with participants seated at other terminals. Please do not communicate with other participants during the session. Please turn off cell phones now.

Please remember that the experiment will last 90 minutes. Remember also that you will be compensated for being in the lab for the next 90 minutes, and that the better your performance, then the higher the amount of money you are likely to earn.

We will start with a brief instruction period. Please pay attention. When I finish reading the instructions, you will be asked questions regarding these instructions. If you have any questions, please wait until we finish reading the instructions.

Instructions: Part 1

You work for a company and your job is to help decide which investment projects to undertake. Projects come from two industries: industry A and industry B.

In every round, the company has the option of investing in a project from industry A or in a project from industry B. The company will invest in one of these two projects, but not in both of them. For simplicity, we will often refer to these projects as project A and project B, respectively. You will help the company decide between investing in project A or B.

The projects from industry A can be either good or bad. The company does not know the chance that a project from industry A is good.

Investing in a good project A results in a payoff of 5 points and investing in a bad project A results in a payoff of 1 point. The payoff from investing in project B is higher than 1 point but lower than 5 points. Therefore, if project A is good, it is best to invest in project A; and, if project A is bad, it is best to invest in project B.

You can perfectly assess the performance of a project from industry B. However, you do not know whether a particular project from industry A is good or bad.

The company has programmed two computers, Computer 1 and Computer 2, to evaluate whether the project from industry A is good or bad. Each computer performs its own evaluation of the project from industry A and submits a recommendation. If a computer assesses project A to be good, then it recommends project A. If a computer assesses project A to be bad, then it recommends project B.

[29]Part 1 in the instructions corresponds to what we call Part I in the text; Part 2 to Part II; and Parts 3 and 4 to Part III.

The software of each computer is in beta mode, so it is possible that the computers make mistakes in their recommendations. Each computer can potentially make two types of mistakes when assessing whether project A is good or bad: it can mistakenly recommend project A when project A is bad, and it can mistakenly recommend project B when project A is good. The company does not know the rates of either type of mistake. However, it is known that Computer 1 and Computer 2 have the same rate for each of these two types of mistakes, although of course they might make different recommendations.

The company wants you to submit a recommendation for project A or a recommendation for project B. You will make this recommendation without knowledge of the recommendation of the computers. Together with the recommendations of the two computers, the company will then have received a total of 3 recommendations. The company will invest in project A if a majority of the recommendations is for project A (that is, 2 or 3 of the 3 recommendations are for project A). And the company will invest in project B if a majority of the recommendations is for project B (that is, 2 or 3 of the 3 recommendations are for project B). In other words, the company will follow the recommendation of the majority.

Your payoffs in the round are given by the following table:

	Project A is good	Project A is bad
Majority recommends Project A	5 points	1 point
Majority recommends Project B	x points	x points

In other words, if a majority recommends project A and project A turns out to be good, then your payoff for the round is 5 points. If a majority recommends project A and project A turns out to be bad, then your payoff for the round is 1 point. Finally, if a majority recommends project B, then your payoff for the round is x points. Here, x represents the payoff of investing in project B and x is equally likely to take values from 1.25 to 4.75, with increments in quarter points, that is 1.25, 1.50, 1.75, 2, and so on, all the way to 4.75. In each round, you will know the value of x before making your decision. The computers, on the other hand, have no information about industry B, and, therefore, their recommendations will not depend on the value of x.

This decision problem will be repeated for a total of 100 rounds. In each round, you will have to decide between a new project from industry A and a new project from industry B. The chance that a project from industry A is good is fixed between 0 and 100 percent. The chance that the computers make the first type of mistake (recommend A when A is bad) is fixed between 0 and 100 percent, and the chance that the computers make the second type of mistake (recommend B when A is good) is fixed between 0 and 100 percent.

You do not know the chance that a project from industry A is good, the chance that the computers make the first type of mistake, and the chance that the computers make the second type of mistake. But these chances are fixed and will not change throughout the experiment. In particular, notice that the problem that you face is the same in every round. In every round, the interface will display on the screen information from past rounds, including the round number, the recommendation of the computers, your recommendation, the majority recommendation (which is the project in which the company invests), whether the project turned out to be good or bad in those cases where the company invested in a project from industry A, and the payoff for that round.

Your total payment for Part 1 will be as follows. We will randomly select 12 out of 100 rounds (with each round having equal chance of being chosen) and we will pay you the total number of points that you made in these 12 rounds. We will then convert points into dollars at the rate of $1 dollar for every 3 points.

Here is a brief reminder of what happens in each of the 100 rounds:

1. You will help decide between a new project from industry A and a new project from industry B. The chance that a project from industry A is good is fixed between 0 and 100 percent and will not change throughout the experiment.

2. Each of two computers evaluates whether the project from industry A is good or bad and submits a recommendation for project A or project B. Computer 1 and Computer 2 make the same rates of mistakes. The chance that the computers make the first type of mistake is fixed between 0 and 100 percent and will not change throughout the experiment. The chance that the computers make the second type of mistake is fixed between 0 and 100 percent and will not change throughout the experiment.

3. The interface draws a value of x (all values from 1.25 to 4.75, with increments in quarter points, are equally likely) that represents the payoff if the company invests in the project from industry B. You, but not the computers, will observe the value of x. You will then submit a recommendation for project A or B.

4. The payoffs for the round are given by the following table:

	Project A is good	Project A is bad
Majority recommends Project A	5 points	1 point
Majority recommends Project B	x points	x points

Now, you will be asked to answer some questions about these instructions. If you have any questions, please raise your hand.

Questions on the instructions [Subjects read and answer these questions on the screen. After subjects submit answers to each set of questions they are provided with feedback. For each question they are informed if their answer was correct or not and the interface highlights the correct answer. If all answers are correct, subjects start with part 1 of the experiment. If any answer is incorrect, subjects face the set of questions that corresponds to the incorrect answer again until they answer correctly.]

Set of questions #1:[30]

1. The chance that project from industry A is good:
   (a) is the same in every round.
   (b) can be different in each round.
2. Each computer can potentially make two types of mistakes when assessing whether project A is good or bad: it can mistakenly recommend project A when project A is bad, and it can mistakenly recommend project B when project A is good. Select all true statements:
   (a) Computer 1 and Computer 2 have the same rates of mistakes.
   (b) Computer 1 and Computer 2 can end up making different recommendations.
   (c) The rates of mistakes of the computers are the same in every round.
3. The recommendations of the computers:
   (a) can not depend on the value of x.
   (b) can depend on the value of x.

Set of questions #2:[31]

1. Suppose that Computer 1 recommends Project A and Computer 2 recommends Project A. Suppose that you recommend Project B. What is the recommendation of the majority?

[30] The correct answers is the first alternative for the first and third questions in this set, and all alternatives should be selected for the second question. Out of 134 subjects 48, 56 and 21 answer correctly on their first, second and third try, respectively. Fewer subjects needed more tries: 3 subjects needed 5 tries, 4 subjects needed 5, one subject needed 6 and one subject needed 8. The findings reported in the paper do not change if we exclude subjects who needed several tries to answer this set of questions correctly.

[31] Out of 134 subjects, 129 answer correctly on the first try. Four subjects need 2 tries and 1 subject needed 3.

(a) Project A.

(b) Project B.

2. Suppose that Computer 1 recommends Project B and Computer 2 recommends Project A. Suppose that you recommend Project B. What is the recommendation of the majority?

    (a) Project A.

    (b) Project B.

Set of questions #3:[32]

1. Suppose that you are in a round in which the majority recommends Project A and Project A is Good. Your payoff in that round is:

    (a) 5 points.

    (b) 1 point.

    (c) x points.

2. Suppose that you are in a round in which the majority recommends Project A and Project A is Bad. Your payoff in that round is:

    (a) 5 points.

    (b) 1 point.

    (c) x points.

Instructions: Part 1, rounds 26-100

The company introduces these additional instructions with the purpose of helping you make better decisions.

The problem in rounds 26 through 100 is exactly the same as before, with the only difference that now the company decides to introduce an additional task. This task should help you make better choices.

At the beginning of each round, the company will ask you to submit a decision rule. A decision rule indicates which option you would recommend for each possible value of x from 1.25 to 4.75.

You will submit your decision rule in each round as follows. On your screen, you will see a slider with the numbers from 1 through 5, with increments in quarter points. You will then click on the slider. Your choice is interpreted in the following way:

- *For any value of x lower than the value on which you clicked, you would recommend project A.*

[32] Out of 134 subjects, 132 answered correctly on the first try. The rest needed 2 tries.

- *For any value of x higher than the value on which you clicked, you would recommend project B.*

Once you submit your decision rule, the round continues exactly in the same way as you played rounds 1 through 25. In particular, the interface will draw a value of x and you will observe the value of x. You will then be asked to recommend project A or B. If your recommendation goes against the decision rule that you submitted in the round, you will be alerted and you will be asked to make a recommendation that is consistent with your decision rule. If you would like to change your decision rule, you will be able to do so in the following round. At the beginning of each round, the interface will display the decision rule you selected in the previous round. You are free to change the decision rule as desired.

Please click on Continue to begin playing round 26, but please wait for my instruction before submitting your choice. Please go ahead and click on the slider to practice. Once you click, you will see a clear indication of the values of x for which you would recommend project A or B. You can adjust your decision rule by clicking on a new value on the slider. Please go ahead and try adjusting your choice by clicking on the slider. Once you click the submit bottom, your decision will be final.

Company report[33]

When you finish playing 100 rounds, the company will ask you to report how you came to your decisions. Your report will explain the following:

1. A description of the aspects of the data that informed your decision.
2. A quantitative assessment of those aspects of the data that informed your decision.
3. Using your responses to (1), (2), and any additional argument you find relevant, you will be asked to justify your decision rule in round 100.

In order to make better decisions and earn more money in the following rounds, it is important that you pay attention to the data displayed on your screen. For this reason, we have provided you with scrap paper in case you want to take notes about your data. Paying attention to the data on your screen will also help you write a better report. The data on your screen will not be available to you when you write the report, you will only have your notes. All of you will have the chance to make an additional $8 from the report, and one of you will have the chance to make up to an additional $20.

Please go ahead now and make your decision for round 26 and then continue to play all rounds until round 100. If you have any questions, please raise your hand.

[33] We mention that we will ask for a report later to highlight incentives for paying attention to the data that subjects gather as part 1 of the experiment evolves.

Instructions: Part 2—(10 minutes)[34] [Subjects are asked to provide a report on paper. After providing the report they answer questions on beliefs on the screen]

Please take the next few minutes to complete the report on the back of this page.

At the end of the experiment, we will randomly select one of you to play the role of the CEO of the company. For the rest of you, we will randomly select 1 out of every 4 reports. This means that, if you are not selected as the CEO, your report has a 25% chance of being selected.

The CEO will then be given these selected reports and he/she will have to pick what he/she considers to be the best report. The experimenter will publicly announce the best report among the reports that were selected. The person who wrote this best report will get an additional payment of $12 and the person selected as the CEO will get an additional payment of $5.

Please enter Lab # ________________________

Please answer the following questions in the space provided.

Question 1: Your decision rule in round 100:

Please draw a circle around the values of x for which you recommend project A:
1.25 1.5 1.75 2 2.25 2.5 2.75 3 3.25 3.5 3.75 4 4.25 4.5 4.75
Please draw a circle around the values of x for which you recommend project B:
1.25 1.5 1.75 2 2.25 2.5 2.75 3 3.25 3.5 3.75 4 4.25 4.5 4.75
(please check that every value of x is circled exactly once; that is, for each value of x you must indicate a recommendation for either A or B, but not both)

Question 2: Please explain the logic behind your decision rule in round 100 by answering the following questions.

(a) Describe which aspects of the data informed your decision and provide a quantitative assessment of what you believe to be the relevant aspects of the data.

[34] The report had two main objectives. First, as mentioned earlier, we use it as a way to incentivize subjects to pay attention to the data in part 1. Second, it serves as a device to check for possible misunderstandings. For example, it is possible that some subject would follow the logic required for a optimal behavior in the Selection treatment, but that it was not properly reflected in their choices. Given that the answers subjects provide are free form, they are challenging to standardize and we do not base any findings on these data. However, some points are worth mentioning. First, 50 out of 66 subjects who participate in the No Selection treatment explicitly mention that they experienced project A not to be good a majority of the rounds as the basis for their choices. In the case of the Selection treatment, only three subjects provide a correct explanation of optimal behavior. Eight subjects provide a correct explanation, but end up selecting thresholds higher than 1. This indicates that at most 16% (11/68) of subjects are making a decision guided by optimal behavior. The vast majority (50 subjects) provide an explanation that either explicitly mentions a computed probability of A being good as the basis of their choice, or either implicitly assumes the probability to be 50% and make an argument for the choice based on risk aversion. (10 of these 50 subjects acknowledge pivotality, but their choices are not based on optimal behavior.) Finally, there are 7 subjects whose report does not fall in either of these classifications.

(b) Use your answer to part (a) and any other relevant arguments to provide a justification for your choice of decision rule in round 100.

Questions on Beliefs [Subjects read and answer the following questions on the screen]

Question 1: What is the chance that a project form industry A is good? Enter a number between 0 and 100. You will receive $2 if your answer is within 5 points of the correct percentage.

Question 2: Computers' Mistakes.

1. Suppose a Project from Industry A is Good. What is the chance that a computer will mistakenly recommend Project B? Enter a number between 0 and 100. You will receive $2 if your answer is within 5 points of the correct percentage.

2. Suppose a Project from Industry A is Bad. What is the chance that a computer will mistakenly recommend Project A? Enter a number between 0 and 100. You will receive $2 if your answer is within 5 points of the correct percentage.

Question 3: Suppose that one computer recommends project A and the other recommends project B. Suppose that you recommend project A, which implies that the company will then invest in Project A. What is the chance that this investment in project A turns out to be good? Enter a number between 0 and 100. You will receive $2 if your answer is within 5 points of the correct percentage.

Instructions: Part 3 (10 minutes)

In part 3, you will participate in an environment that is very similar to the 100 rounds in part 1, but with the following differences:

1. You will face 5 different cases. The chance that a project from industry A is good will now be 10%, 30%, 50%, 70%, or 90%, depending on the case. You will now know which case you are facing, so you will know the chance that project A is good. For each case, you will have to choose a decision rule by clicking on the slider below the case.

2. The second difference is that there are no computers submitting recommendations. Therefore, your recommendation is the only one that counts.

How do you make your choice for each case? In the same screen, you will see all 5 cases at once. You will then have to select a decision rule for each of the 5 cases. The interpretation is the same as in part 1:

For any value of x lower than the value on which you clicked, you recommend project A for that case.

For any value of x higher than the value on which you clicked, you recommend project B for that case.

Once you make your choices for each of the 5 cases, you can submit your choices by clicking on the “Submit” button. You can change your choices as many times as you want before clicking the “Submit” button.

How your payment in Part 3 is determined: Your payment will be determined by selecting one of the 5 cases with equal probability. We will then select a value of x (with any number between 1.25 and 4.75 with equal probability) and use your decision rule in the selected case to determine your vote. We will then pay you according to the following table, which is similar to the payoff table in part 1 of the experiment, except that now there are no computers submitting recommendations and your recommendation alone determines the project on which the company invests.

	Project A is good	Project A is bad
You recommend Project A	5 points	1 point
You recommend Project B	x points	x points

You will now be asked some brief questions about these instructions. Please raise your hand if you have any questions.

[Subjects read and answer the following questions on the screen. The questions are repeated until all are answered correctly.]

Question 1: In part 3 you have to submit 5 decision rules. True or False?

Question 2: In Part 3 your recommendation and the recommendations of two computers will determine your payoffs. True or False?

Question 3: The chance that a project from industry A is good will be the same in all cases of Part 3. True or False?

Instructions: Part 4 (final part of the experiment, 5 minutes)

This final part of the experiment is NOT related to any of the previous parts. In these instructions, we will explain how to answer part 4.

To illustrate, consider a choice between Column A and Column B in the table below.

Column A You will receive			Column B You will receive
\$2 with chances 50/100; \$1.6 with chances 50/100			\$3.85 with chances 50/100; \$0.1 with chances 50/100

If you choose Column A, then your payoff will be \$2 with chances 50/100 and \$1.6 with chances 50/100. If you choose Column B, then your payoff will be \$3.85 with chances 50/100 and \$0.1 with chances 50/100.

In Part 4, you will actually observe a table with 10 rows. Each row of the table contains a choice between Column A and Column B. The full table will contain 10 rows. For each row, you have to choose whether you prefer the lottery described in Column A or the lottery described in Column B. If you prefer the lottery in Column A, you must click on the center column next to Column A; if you prefer the lottery in Column B, you must click on the center column next to Column B. Only one option can be selected for each row, and you can change your selection as many times as you would like before clicking the "Submit" button.

How your payment in Part 4 is determined: Your payment will be determined by selecting one of the 10 rows in the table with equal probability. We will pay you according to the choice you made for that row (either Column A or Column B).

Please raise your hand if you have any questions.