

Endogenous sample selection and partial naiveté: A laboratory study*

Ignacio Esponda
(WUSTL)

Emanuel Vespa
(UC Santa Barbara)

August 26, 2015

Abstract

Accounting for sample selection is a challenge not only for empirical researchers but also for the agents populating our models. Yet most models abstract from these issues and assume that agents successfully tackle selection problems. We design an experiment where anyone who understands sample selection can easily account for it. Agents make choices under uncertainty and their choices reveal valuable information that is biased due to the presence of unobservables. We find that essentially no subject optimally accounts for endogenous selection. On the other hand, behavior is far from random but actually quite amenable to analysis: Subjects follow simple heuristics that result in a partial accounting of selection and mitigate mistakes.

*We thank Roland Bénabou, Erik Eyster, Kfir Eliaz, Guillaume Fréchet, 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. We acknowledge support from NYU's Center for Experimental Social Science and from UCSB's Center for Scientific Computing from the CNSI, MRL: NSF MRSEC (DMR-1121053) and NSF CNS-0960316. 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

ENDOGENOUS SELECTION. Accounting for sample selection is a major challenge for empirical researchers. Economic agents must also deal with selection, with the difference that they usually have more control over the process because data is endogenously generated by their own actions. Yet, a bit surprisingly, most models abstract from this difficulty and assume that agents successfully tackle selection issues.

Our main contribution is to examine how people behave in the presence of endogenous sample selection. The following examples illustrate this phenomenon.¹

1. *Bidding for procurement contracts.* Every month, a firm bids on several procurement contracts. The firm uses data on previously finished jobs to estimate its cost for a new job, but, naturally, the firm does not observe the cost of projects it was not awarded. If other firms have private information about a common value component in cost, then the average cost of *awarded* projects will be higher than the average cost of all projects. The reason is that a firm only observes costs of projects in which other firms bid above its own bid. Similarly, the more aggressively the firm bids, then the lower the average cost of projects that it is awarded.

2. *Demand estimation.* A firm wants to estimate its own-price elasticity of demand. Each period, the firm chooses a price and observes its sales. But the firm does not observe the prices of competing firms. Prices, however, are correlated, because industry costs are correlated. Thus, the firm's observed data will make demand appear less elastic than it actually is, when in fact the price increase of the firm is being mitigated by the (unobserved) price increases of other firms.

3. *Mental states and well-being.* A person is pessimistic about her life prospects, so she becomes disinterested and prefers to avoid exercising, studying, and other costly investments. As a result, she continues to obtain poor outcomes, which reinforces her pessimism. She does not realize, however, that if she were optimistic, she would feel more energetic and find it less costly to invest.

4. *Investment in risky projects.* A Hollywood studio can invest in a sequel or take a chance with a new project. The studio can easily forecast the financial return of the sequel, but assessing a new project is more involved. The standard industry

¹Example 1 is studied by Esponda (2008); Examples 2 and 4 by Esponda and Pouzo (2012, 2015), and Example 3 by Kőszegi (2010). Esponda and Pouzo (2015) show that the endogenous selection problem arises in general environments where the agent learns with a misspecified model of the world.

practice is to hire readers who, based on their experience, independently evaluate the screenplay and make a recommendation. Readers' experience is based on projects that were effectively developed, that is, they do not know what would have happened with movies that were never produced. If projects that were produced are on average better than those that were not, but readers are unaware of the selection effect, they will recommend the new project more often than optimal.

In these examples, an agent wants to learn something (a cost estimate, the elasticity of demand, her life prospects, the prospects of risky projects) in order to make decisions. People often do not know these primitives, and must learn them from experience. But data is often limited because people do not observe counterfactuals (the cost of a project that is not awarded, the sales from a price that was not chosen, the benefits of changing attitudes, the returns from a risky project that was not implemented). Moreover, observed data often comes from a selected sample due to the presence of unobservables (such as the costs, information, or choices of other agents). Finally, the agent's own decision affects the sample that is actually observed.

Economic models often disregard these issues and assume that agents choose optimal actions. The main objective of this paper is to understand how subjects make decisions in the presence of endogenous sample selection.

THE EXPERIMENT. One challenge in the experimental design is that it is difficult to distinguish a naive subject from a subject who understands selection but is unable to perfectly account for it (even professional researchers struggle here). We tackle this issue by designing a lab experiment where anyone who understands sample selection can easily account for it.

Our subjects face a toy version of the 'investment in risky projects' example. For each of 100 rounds, a subject chooses between a risky and a safe project. The project that is implemented in each round depends on the subject's recommendation and a random process (which represents the behavior of other recommenders). In the "No Selection" treatment, the random process is uninformative, and so there is no selection effect; i.e., one can correctly assess the chances that the risky project is good by simply looking at the percentage of rounds in which it was observed to be good. In the "Selection" treatment, the random process is correlated with the prospects of the risky project and, therefore, there is a selection effect; i.e., the risky project is more likely to be implemented if it is good, and so evaluating its effectiveness based

on its observed performance would lead to an upward bias in beliefs.

RELATIONSHIP TO PREVIOUS EXPERIMENTS. The specific unobservable driving sample selection in our experiment is other players' private information (represented by the random process described above). A large experimental and theoretical literature focuses on people's failure to make inferences from others' private information. Experiments find that a majority of the subjects fail to correctly make such inferences. 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. On the theory side, the initial contributions of Kagel and Levin (1986) and Holt and Sherman (1994) in an auction context were generalized by Eyster and Rabin (2005), Jehiel (2005), and Jehiel and Koessler's (2008). Esponda (2008) formalizes (the failure to account for) endogenous selection that is driven by other players' private information. These mistakes are also studied under non-equilibrium concepts (e.g., Crawford and Iriberri (2007)).

Our experiment differs from previous experiments in that subjects do not know the primitives and do not observe counterfactual outcomes. Without *either* of these features, there would be no endogenous selection problem to study. As illustrated by the introductory examples, the presence of private information is one of many possible underlying causes of endogenous selection. Thus, the problem that we study is of independent interest and has consequences for a broader range of settings.

Moreover, it is unclear how to extrapolate previous experimental findings due to the different nature of our setting. First, the right solution calls for different approaches. In previous experiments, subjects should compute an expectation conditional on some event, and this computation requires knowledge of the prior and others' strategies. Of course, if one views our experiment in a similar manner, then it seems harder because, in addition to computing a conditional expectation, subjects must also learn the prior and others' strategies. But, as we show, there is a simpler way to approach our experiment: Subjects only need to keep track of the proportion of successful projects that were implemented due to their pivotal recommendation.

Another reason why previous results are, at best, suggestive is that providing primitives might induce subjects to make mistakes. For example, in a previous paper (Esponda and Vespa, 2014), we follow the standard approach of telling subjects the chance that a project is successful. Telling a subject that this chance is, say, 75%, likely biases the subject to choose that project, even though a closer examination of

other players' strategies might reveal that, conditional on being pivotal, the chance of success is negligible. By not providing primitives, we eliminate an important mechanism underlying previous results. So it is an open question how people behave in our new environment. Of course, this comment does not detract from previous work for two reasons. First, the literature convincingly makes the important point that most people fail to compute conditional expectations in environments with known primitives. Second, there are many environments where it is natural to know the primitives. In contrast, our focus is on settings in which a priori information is not available and people need to form beliefs from endogenous data.

FINDINGS AND IMPLICATIONS. We focus on long-run, steady-state behavior for two reasons. First, we want to see if mistakes persist in the long run, after extensive experience. Second, our approach is consistent with the common focus in economics on equilibrium behavior. One important benefit of focusing on equilibrium is that many possible learning dynamics can lead to equilibrium, too many to be able to identify with just two treatments, but there are only a few reasonable candidates for (“rational” or “boundedly rational”) steady-state behavior.

The main finding is that the direction of the treatment effect is consistent with naive subjects who do not understand endogenous selection. In both treatments, subjects end up responding to the *observed* percentage of successful risky projects. In the No Selection treatment, this is an optimal response. In the Selection treatment, this is a suboptimal response that does not take into account the bias in the sample and, therefore, subjects select the risky project too often. At the end of the experiment, we elicit subjects' beliefs and corroborate these predictions: Reported beliefs mostly fail to account for selection and are consistent with naive (biased) beliefs.

We then examine the extent to which the naive theory can *quantitatively* rationalize the data. While naiveté predicts behavior in the No Selection treatment fairly well, it tends to over-predict risky behavior in the Selection treatment. Controlling for risk aversion, we find that subjects overestimate the benefits of playing the risky alternative, but not by the full amount predicted by naiveté. This finding raises the puzzle of how subjects can be so clearly naive but still manage to partially account for selection.² We discover, however, that it is rather natural for subjects to be partially

²If subjects placed even a small prior probability on the event that the random process is 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, subjects

naive in our experiment: 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 partially account for selection by placing more weight on feedback from pivotal rounds.

Motivated by this finding, we propose a new model of partial naiveté to quantify the extent to which subjects place more weight on feedback from pivotal rounds. We estimate the model and find that the median subject places about three times more weight on pivotal vs. non-pivotal rounds. This weighting, however, has a small effect on behavior, since subjects are pivotal in only a third of the rounds. This explains why behavior is still much closer to the naive than to the sophisticated prediction.

Eyster and Rabin (2005) develop a notion of partial naiveté, partially cursed equilibrium, which includes fully cursed and Nash equilibria as special cases. A particular value of their parameter of partial naiveté fits data from several experiments (with known primitives) better than Nash equilibrium. Our notion of partial naiveté is motivated by differential attention to limited feedback. While their model is defined for any Bayesian game, our model illustrates the possibility of having a learning interpretation of partial naiveté in a particular context.

There are three main implications from our results. The first is that subjects have a harder time with selection problems than documented by previous literature. The result is a bit striking, particularly because there is a very simple way to account for selection in our experiment that does not involve learning the primitives or computing difficult conditional expectations. Our finding that essentially no one understands selection can be contrasted with previous experiments in which subjects know the primitives and a non-negligible fraction—even nearing 50% in some treatments of Charness and Levin (2009) and Esponda and Vespa (2014)—become sophisticated.

Second, there are reasons not to be too pessimistic about human behavior. Although people do not understand selection, they follow certain heuristics (i.e., higher weight on pivotal observations) that help mitigate their naiveté. In addition, the experiment shows that behavior, far from being random, can be fairly accurately rationalized by sensible heuristics. The experiment also raises new questions that should ultimately help us build better models. Future work could seek to understand why people respond more to pivotal events, despite not understanding selection. More broadly, what types of events do people respond most to in general settings?

should be either completely naive or fully sophisticated in this experiment.

Finally, the experiment highlights that there is much to learn from not giving primitives to subjects. While not giving primitives was common in early experimental work, it is currently under-explored.³ It need not be so, particularly since an important objective of experiments is to test equilibrium behavior. As highlighted by the learning-in-games literature, equilibrium can be viewed as the result of a learning process, and it imposes *steady-state* restrictions on what people have learned about *both* the the strategies of “nature” and other players, without the presumption that people somehow already knew one of these two objects to start with.⁴

Fudenberg and Peysakhovich (2014) highlight the importance of not giving primitives in an adverse selection experiment. They find that learning models that account for recency bias provide a better fit than steady-state solution concepts such as Nash, cursed, or behavioral equilibrium. In particular, subjects respond more to extreme outcomes in the previous round compared to much earlier rounds. Their results are an important reminder that steady-state solution concepts are not always appropriate to explain behavior.

ROADMAP. We describe the experiment and theoretical predictions in Section 2, show the results in Section 3, and propose and estimate a model of partial naiveté in Section 4. We conclude in Section 5 and relegate the instructions and robustness checks to the Online Appendix.

2 The experiment

2.1 Experimental design

Each of our subjects participate in a single-agent decision problem. 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.

³In early experiments on competitive equilibrium (e.g., Smith, 1962), subjects traded without any information about the distribution of sellers’ costs or buyers’ values, precisely because the objective was to understand how decentralized markets aggregate this information. 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).

⁴See, for example, Fudenberg and Levine (1998), Dekel et al. (2004), and Esponda (2013), who also points out that uncertainty about fundamentals and strategies are treated in the same manner

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

Figure 1: Payoffs for the experiment. The project that is implemented is determined by the choice of the subject and two other agents played by the computer. The payoff x from implementing B varies each round from 1.25 to 4.75, and the subject observes the value of x before making a choice.

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.
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 company will invest in the project recommended by the majority, and the payoffs for the round are given by the table in Figure 1.

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

by epistemic game theorists.

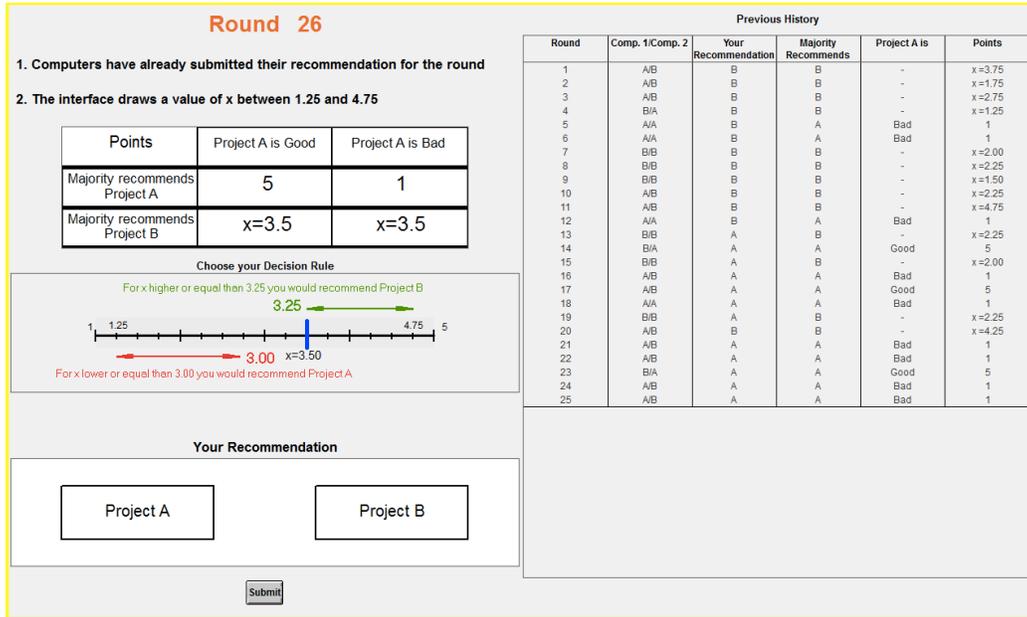


Figure 2: Screen shot for round 26. In rounds 1-25, the subject must submit a recommendation for a given value of x . In rounds 26-100, the subject must first submit a threshold recommendation that indicates a choice for each value of x . She is then prompted to submit a choice for a particular value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

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?

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

⁵This design yields more (and less noisy) information in each round, compared to estimating a threshold strategy from the data (pooling data from different rounds 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).

⁶This 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.

⁷At 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 28, the two measures are consistent with each other.

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.⁸

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.

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.

2.4 Theoretical steady-state predictions

We begin with an informal discussion of the theoretical predictions and then characterize the solutions for each treatment. Table 1 shows an example of feedback from playing the first 12 rounds of the Selection treatment. There are two natural steady-state predictions in our environment. The first prediction is that a subject will naively estimate the chance that project A is good by the proportion of times

⁸An additional, atypical benefit of not providing the subjects with the primitives is that the instructions for both treatments are *exactly* the same.

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.

that it has been observed to be good in the past. Thus, in the example provided in Table 1, a naive subject will estimate the chance that A is good to be $1/2$ 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.⁹

The problem with this naive approach is that it does not account for the fact that the sample from which the subject makes inferences will be biased if the recommendations of the computers happen to be correlated with the state of the world. 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

⁹Following Esponda and Pouzo (2015), this form of naiveté arises from a model of misspecified learning in which subjects believe that the behavior of the computers is independent of the state of the world. This particular misspecification underlies the solution concepts of Eyster and Rabin (2005), Jehiel and Koessler (2008), and Esponda (2008). Our characterization of naive behavior follows Esponda’s (2008) behavioral equilibrium because that solution concept accounts explicitly for the lack of counterfactual information. See Kőszegi (2010) and Spiegel (2015) for related solution concepts.

is good and choose a strategy that is more risky than optimal.

The second natural steady-state prediction is that a subject is sophisticated, understands the sample may be biased, learns to account for this bias, and eventually makes optimal decisions. There are two natural ways to account for sample selection bias in our context. One way to account for the bias is for the subject to use data about the realized payoff of A only from the subsample of rounds in which her recommendation was pivotal; these are rounds 4, 5, and 9 in Table 1.¹⁰ 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. A second way to reach an optimal decision is simply to do so by trial and error. Subjects have 100 rounds to experiment with different strategy choices and settle for the one that they think maximizes their payoffs.

2.4.1 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 naive and sophisticated inferences lead to the correct belief that the probability of A being good is 1/4. Thus, the naive and sophisticated predictions coincide 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 $r < 0$.¹¹ Then the optimal (naive and sophisticated) threshold x^* for a subject with

¹⁰The importance of pivotality in these types of environments is highlighted by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1997). Esponda and Pouzo (2012) show that steady-state behavior corresponds to Nash equilibrium under sophisticated learning and behavioral equilibrium under naive learning.

¹¹For $r = 1$, we let $u(c) = \ln c$.

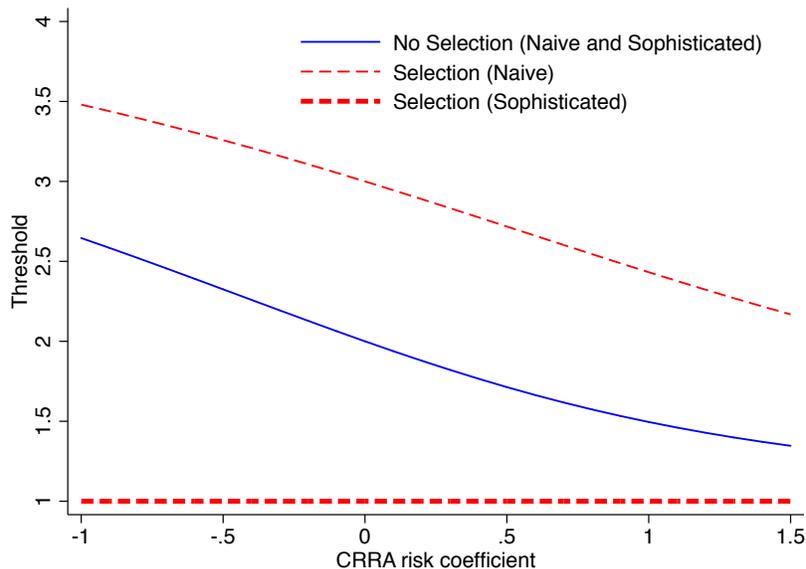


Figure 3: Theoretical prediction for Selection and No Selection treatments. For the benchmark case of the No Selection treatment, naive and sophisticated thresholds coincide. Under Selection, naive and sophisticated thresholds go in opposite direction: higher than the benchmark in the naive case and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

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 (naive and sophisticated) threshold $x_N^*(r) = x_S^*(r)$ as a function of the coefficient of relative risk aversion, r . As expected, the threshold decreases as risk aversion increases.¹²

2.4.2 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 naive and sophisticated behavior differ. Consider first the sophisticated case. Because both computers correctly recommend A if it is good, then, if a subject is pivotal, A must be bad. Thus, it is optimal to always recommend B, $x_{NE}^* = 1$, irrespective of the risk aversion coefficient. In terms of the

¹²For 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.

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 $x_{NE}^* = 1$.

Next, consider the naive steady-state prediction. The steady-state belief that A is good is given by the probability that A is observed to be good conditional on the event that the subject obtains some information about A. The latter event is equivalent to the event that the majority recommends A, which we denote by M_A in the expression below. Thus, the naive steady-state belief is

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

where we have used the fact that, in the Selection treatment, $m_G = 0$ and $m_B = 1/3$.¹³ Equation (1) makes explicit that the sample selection problem facing the subject is endogenous. The reason is that the probability that the majority recommends A 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 $z(x^*)$ is decreasing in x^* ; 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.

Because beliefs are endogenous, a naive steady-state 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.,

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

¹³Note that, in the No Selection treatment, $z(x^*) = 1/4$ for all x^* , as remarked earlier.

In other words, the naive steady-state threshold $x_N^*(r)$ is the unique solution to equation (2).¹⁴ For example, if the subject is risk neutral, $r = 0$, then equation (2) becomes $4/(1 + x^*/3) + 1 = x^*$ and the naive threshold is $x_N^* = 3$. Figure 3 plots the naive threshold $x_N^*(r)$ as a function of the coefficient of relative risk aversion, r . As expected, the threshold decreases as risk aversion increases.

To summarize, the steady-state naive and sophisticated predictions coincide for the No Selection treatment. On the other hand, naive and sophisticated behavior imply different treatment effects: For a given level of risk aversion, the naive steady-state threshold increases and the sophisticated one decreases when going from the No Selection to the Selection treatment.

2.5 Discussion of experimental design

Now that we introduced the experiment and discussed the main theoretical predictions, it is easier to explain why we made certain choices in the experimental design.

Choice of environment. As illustrated by the examples in the introduction, the endogenous selection problem arises in a wide range of environments. We focus on a collective decision problem where the unobservable variable that leads to selection is the private information of other agents for three main reasons. First, as reviewed in the introduction, there is a large literature that focuses on mistakes in environments in which other players have private information. Second, our previous work (Esponda and Vespa, 2014) looked at a collective action environment but followed the more standard approach of telling the primitives to the subjects. By focusing on the same environment, we can directly contrast our results to the previous literature and understand the effect that lack of counterfactuals and primitives have on behavior. Third, to concentrate on the selection problem we wanted to make the inference problem as simple as possible. In our environment, subjects only need to learn the chance that a project is good vs. bad. In an auction environment, for example, subjects would need to learn both the value of an item and the probability of winning it.

Lack of primitives and counterfactuals. We do not provide primitives or counterfactuals to subjects because it is the lack of *both* types of information that results in

¹⁴The solution is unique because the LHS of equation (2) is decreasing (because $z(\cdot)$ is decreasing) and the RHS is increasing.

the selection problem that we wish to study. If subjects knew the primitives, then the problem exactly reduces to the problem studied in previous papers, and the source of the mistake is relatively well understood (e.g., Charness and Levin (2009), Ivanov et al. (2010), Esponda and Vespa (2014)). If subjects observed the counterfactuals, then the problem *eventually* reduces to the problem studied in previous papers, because subjects who pay attention to the data will learn the true primitives correctly (our findings already suggest that they would learn these primitives pretty well). The introduction discusses why it is interesting to depart from the standard case. Moreover, the subject would have no influence over the observed data and, hence, there would be no endogenous sample selection problem to study.

Use of computers and stationarity. The use of computers (as opposed to letting subjects interact with each other) is to make the environment stationary. This is not to downplay the importance of non-stationary environments in real life, but it seems sensible to introduce changes one at a time and to start by understanding how people respond to sample selection in stationary environments before moving on to non-stationary settings.

Mistakes of the computers. In the Selection treatment, mistakes need to be asymmetric (i.e., different in the good and bad states) in order for the recommendations of the computers to be informative and, hence, to obtain selection effects. We chose a zero mistake rate in the good state because it makes it easiest for a subject who understands the possibility of selection to account for it. Under this choice, *every time* a subject causes A to be chosen, she finds out that A is bad.

Size of incentives. The incentives to behave optimally are fairly small in our setting because subjects are pivotal with a probability of one third. The significant treatment effects that we obtain suggest that subjects are indeed responding to these small incentives. Similar response to small incentives have been found in previous work (e.g., Esponda and Vespa, 2014). Moreover, incentives are also realistically small in the type of collective action problems that our experiment represents, and the aggregate effects of individual actions tend to have large welfare consequences.

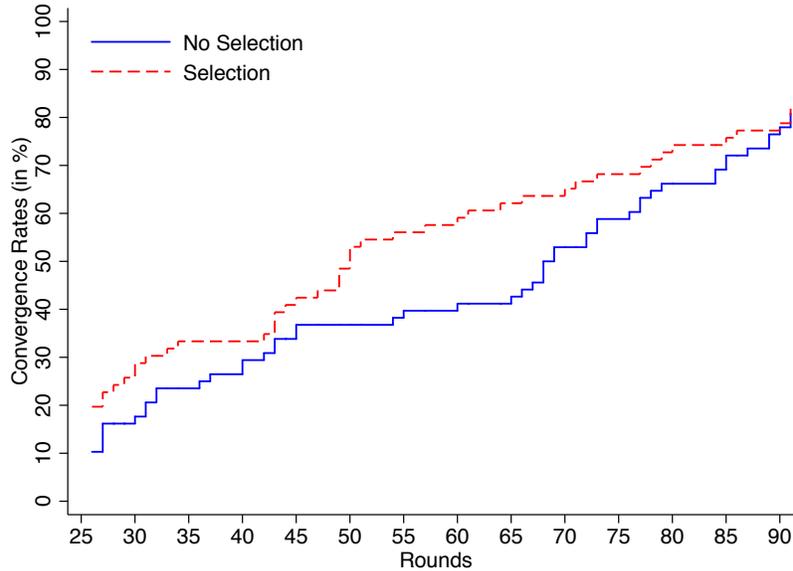


Figure 4: Convergence rates by treatment.

3 Results

We organize the presentation of the results around five main findings. Motivated by the results, in Section 4 we propose and estimate a new model of partial naivet e.

Finding #1. The direction of the treatment effect is consistent with naive, not sophisticated, steady-state behavior:

The first question is whether it is appropriate in our setting to focus on steady states, 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. By round 90, however, these rates increase to 83% and 79%, respectively. Thus, we next focus on explaining *steady-state* behavior, 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 Selection 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 or naive. A sophisticated subject realizes that, every time she is pivotal, A is bad. Thus, the sophisticated threshold converges to 1. In the naive 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 naive, not sophisticated, behavior.¹⁵

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 steady-state predictions are potentially applicable. Also, it seems contradictory to use a steady-state prediction 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

¹⁵Naiveté 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).

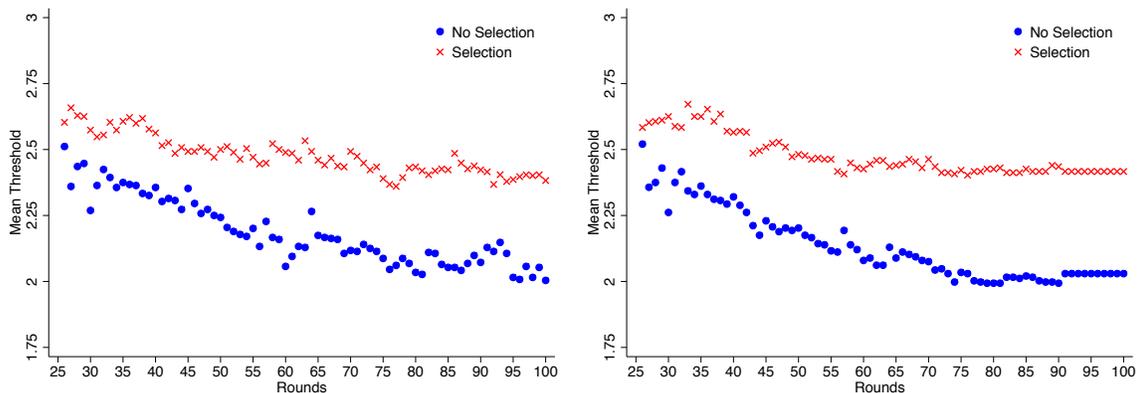


Figure 5: Mean thresholds in all rounds by treatment. The left panel shows the mean threshold for all subjects, for each round and treatment. The right panel shows the same information but only for subjects whose behavior converges in the sense that their threshold choice is constant for the last 10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

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.¹⁶

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.¹⁷ Thus, there is essentially no evidence of sophistication in this experiment.

¹⁶To 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.

¹⁷We 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 Selec-

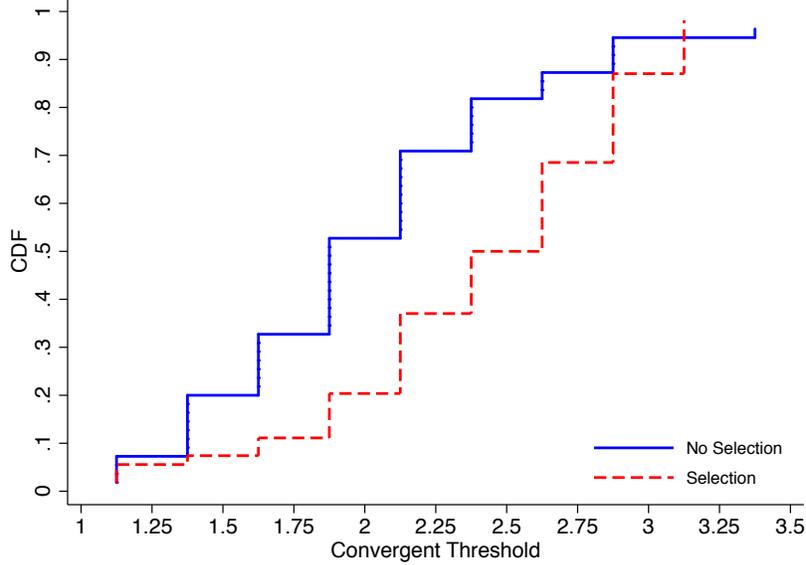


Figure 6: Distribution of convergent threshold choices, by treatment. Convergent threshold choices under Selection first order stochastically dominate choices under No Selection.

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, 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%

tion 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.

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.

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, 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 are 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. Convergent thresholds are a bit lower than predicted by the naive steady state:

While naiveté 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 the naive steady-state solution 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.¹⁸ To account for the second issue, we assume that a subject's (naive) belief is determined by the observed data (i.e., the

¹⁸We let subject i 's threshold in the k th (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 4.1, where we use this equation in the model of partial naiveté, for more details.

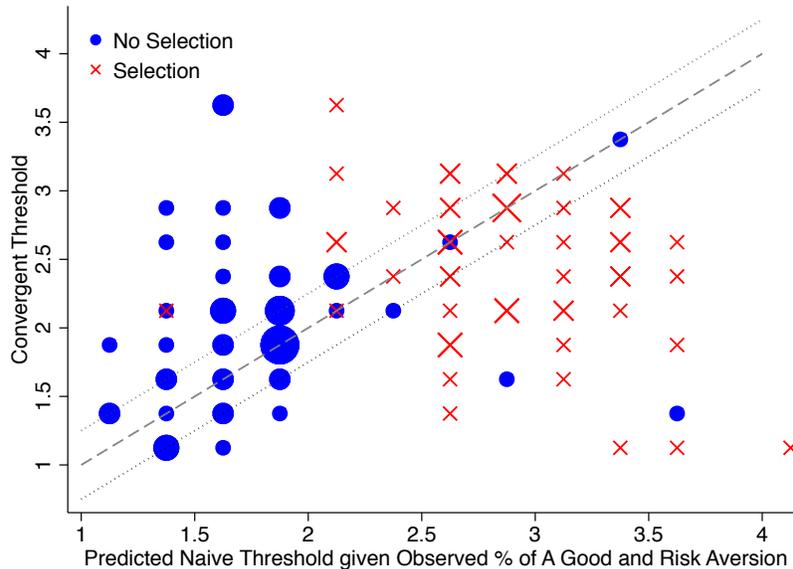


Figure 7: Observed vs. naive prediction. The figure plots the predicted naive thresholds (based on the estimated risk coefficient and beliefs for each subject) against the actual choice, for rounds 91 through 100 (where the size of the plotted data point is proportional to the number of subjects). The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

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 predicted naive threshold (recall that this is also the predicted sophisticated 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.

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 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, the naive prediction systematically over-predicts threshold choices 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 threshold. One reasonable explanation for this (admittedly, unexpected) 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 probability 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 positive or negative information about project A. Observing that A was bad 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.¹⁹

¹⁹These 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 that A was good 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).

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
Piv _{t-1} × 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
(Piv and Good) _{t-1}	0.020	0.025
(Piv and Bad) _{t-1}	-0.092***	0.018
(Not Piv and Good) _{t-1}	-0.002	0.005
(Not Piv and Bad) _{t-1}	-0.024***	0.008

Legends: All controls are dummy variables. (Piv and Good)_{t-1} takes value 1 if the subject was pivotal, the company invested in A, and it turned out to be Good. Other dummy variables are named accordingly. The excluded event is the case when the subject did not receive information in the previous period because the company invested in B.

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 rounds 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.

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 partial adjustment of selection despite their inherent naiveté.

4 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.

4.1 Model

Recall from Section 2.4.2 that the steady-state belief for a naive subject is given by $z(x^*)$ in equation (1), page 14, 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^* . The implicit assumption underlying this definition of naive beliefs is that a subject puts 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 rounds. The corresponding steady-state belief is

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

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

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

The naive and sophisticated predictions discussed in Section 2.4 are special cases of this model. As η goes to infinity, a subject puts increasingly higher weight on pivotal rounds, and it is easy to see that the threshold converges to the sophisticated, optimal

threshold, i.e., $\lim_{\eta \rightarrow \infty} x_\eta(r) = x_S(r)$ for all risk coefficients r . And the case $\eta = 1$, which places equal weight on pivotal vs. non-pivotal rounds, corresponds to what we called the naive threshold (or, equivalently, what Esponda (2008) calls a naive behavioral equilibrium), i.e., $x_1(r) = x_N(r)$ 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 $\eta = \alpha/\beta$.²⁰

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

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

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 ($p = 1/4, m_G = 0, m_B = 1/3$),

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

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 $\eta = 1$ (naive behavior) and $\eta \approx \infty$ (sophisticated behavior) representing two extreme cases in the figure.²¹ 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 rounds. The reason is that very high weights are needed to compensate for the fact that the probability

²⁰Note that in this case α and β are not separately identified; only their ratio affects beliefs.

²¹The model also allows for $\eta < 1$, which means that non-pivotal rounds receive relatively higher weight.

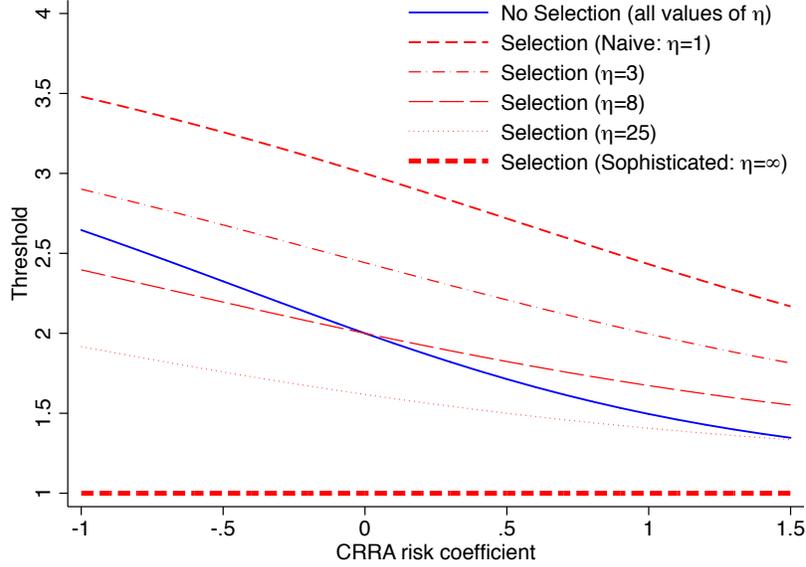


Figure 8: Theoretical prediction for Selection and No Selection treatments for several values of η .

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 the theoretical prediction with $\eta = 1$ is above observed values (Finding #4), Figure 8 already suggests that the average η in the population is between 1 and 8. In the next section we obtain a more precise estimate of the distribution of η .

4.2 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

$$x_{ik}^* = f(r_i, z_{ik}) + \varepsilon_{ik}, \quad (5)$$

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.²²

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, $k = D_1, \dots, D_5$, where 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 $k = N_T, N_{T+1}, \dots, N_{100}$ and $k = S_T, S_{T+1}, \dots, S_{100}$ 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.²³ In particular, we assume that, for observations from both treatments $k = N_t$

²²In 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 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.

²³Alternatively, 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

and S_t and all rounds $t = T, \dots, 100$, a subject's belief is

$$z_{ik} = g(\text{data}_{ik}, \eta_i) + \nu_{ik},$$

where $g(\text{data}_{ik}, \eta_i)$ 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 .²⁴ 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 $\ln \eta \sim N(\mu_{\ln \mu}, \sigma_{\ln \mu}^2)$ and the belief formation noise $\nu \sim N(\mu_\nu, \sigma_\nu^2)$ are normally distributed and independent of each other and across subjects, and we estimate the parameters using (simulated) maximum likelihood.²⁵

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.²⁶ The result is consistent with the reduced-form results from Section 3, 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 sophisticated behavior. As explained earlier, even much higher weights on pivotal rounds are not enough to approximate

to estimating beliefs from actions (e.g., Nyarko and Schotter, 2002).

²⁴Under 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.

²⁵In particular, the naiveté coefficient η has a lognormal distribution with mean $E(\eta) = e^{\mu_{\ln \eta} + \frac{\sigma_{\ln \eta}^2}{2}}$ and variance $Var(\eta) = (e^{\sigma_{\ln \eta}^2} - 1)e^{2\mu_{\ln \eta} + \sigma_{\ln \eta}^2}$.

²⁶Given 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 .

	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 .

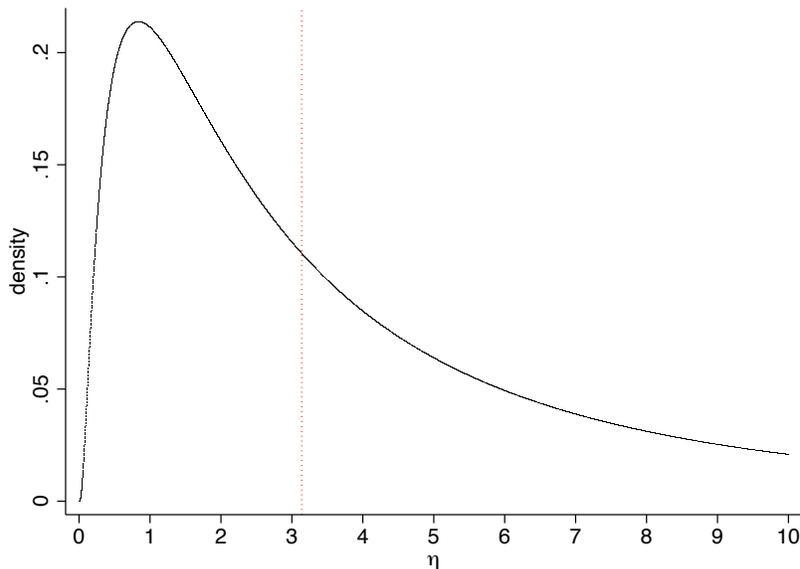


Figure 9: Estimated distribution of η . The dashed line indicates the median at 3.14.

sophisticated 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 the estimation. We confirm that the main conclusions are unaffected by the choice of T .²⁷

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)).²⁸ The estimates also suggest that it is important to account for noise in actions to avoid biasing our results for the

²⁷Appendix A provides detailed results on several robustness exercises.

²⁸Following 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).

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.

5 Conclusion

Accounting for selection is a challenge not only for empirical researchers but also for economics agents in a wide range of important environments. Yet most models assume that agents successfully tackle selection problems. We design an experiment where anyone who understands sample selection can easily account for it. The design incorporates assumptions, such as the provision of no information about primitives and counterfactual outcomes, that are nonstandard in the literature but are crucial to study endogenous selection. We find that essentially no subject optimally accounts for endogenous selection. On the other hand, behavior is far from random but actually quite amenable to analysis. Subjects partially account for selection by placing three times more weight on pivotal vs. non-pivotal observations, thus mitigating losses from their (suboptimal) risky behavior. While more experiments are needed to confirm behavior in these types of settings, our results suggest that we might want to think more seriously about the types of identification problems faced by economic agents.

References

- Austen-Smith, D. and J.S. Banks**, “Information aggregation, rationality, and the Condorcet jury theorem,” *American Political Science Review*, 1996, pp. 34–45.
- 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. Iriberry**, “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 –**, “Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models,” *working paper*, 2015.
- **and E. Vespa**, “Hypothetical Thinking and Information Extraction in the Laboratory,” *American Economic Journal: Microeconomics*, 2014, 6 (4), 180–202.
- Esponda, Ignacio**, “Rationalizable conjectural equilibrium: A framework for robust predictions,” *Theoretical Economics*, 2013, 8 (2), 467–501.
- Eyster, E. and M. Rabin**, “Cursed equilibrium,” *Econometrica*, 2005, 73 (5), 1623–1672.
- Feddersen, T. and W. Pesendorfer**, “Voting behavior and information aggregation in elections with private information,” *Econometrica*, 1997, pp. 1029–1058.
- 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.
- Kőszegi, Botond**, “Utility from anticipation and personal equilibrium,” *Economic Theory*, 2010, *44* (3), 415–444.
- 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.
- Spiegler, R.**, “Bayesian Networks and Boundedly Rational Expectations,” *working paper*, 2015.