

Endogenous sample selection: A laboratory study

IGNACIO ESPONDA

Department of Economics, University of California, Santa Barbara

EMANUEL VESPA

Department of Economics, University of California, Santa Barbara

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 a person who understands selection observes all the data required to account for it. Subjects make choices under uncertainty and their choices reveal valuable information that is biased due to the presence of unobservables. We find that almost no subjects optimally account 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.

KEYWORDS. Contingent thinking, learning, sample selection.

JEL CLASSIFICATION. C91, D83.

1. INTRODUCTION

ENDOGENOUS SELECTION

Accounting for sample selection is a 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 are endogenously generated by their 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

Ignacio Esponda: iesponda@ucsb.edu

Emanuel Vespa: vespa@ucsb.edu

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, Georg Weiszäcker, Alistair Wilson, Leeat Yariv, and several seminar participants and anonymous referees 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 (Grant DMR-1121053) and NSF Grant CNS-0960316.

¹Example 1 is studied by Esponda (2008), Examples 2 and 4 by Esponda and Pouzo (2016a, 2016b), and Example 3 by Kőszegi (2010). Esponda and Pouzo (2016b) show that the endogenous selection problem arises in general environments where the agent learns with a misspecified model of the world.

Copyright © 2018 The Authors. Quantitative Economics. The Econometric Society. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at <http://www.qeconomics.org>. <https://doi.org/10.3982/QE650>

1. *Bidding for procurement contracts.* Each month, a firm bids on procurement contracts. It uses data on finished projects to estimate its cost for a new job, but, naturally, it does not observe the cost of projects completed by other firms. If, at the bidding stage, other firms have private information about a common cost component, then the average cost of the projects *awarded* to the firm will be higher than the average cost of all projects. The reason is that the firm is awarded the project when other firms submitted higher bids. Similarly, the more aggressively the firm bids, then the lower is 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 either a sequel or 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 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 is 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) so as to make decisions. People often do not know these primitives, and must learn them from experience. But data are 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 come 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.

THE EXPERIMENT

So as to understand how people tackle selection, we design a lab experiment where a person who understands selection observes all the data required to account

for it.² 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 the recommendations of two computers (which represent the behavior of two other recommenders). In the “No Selection” treatment, the computers’ recommendations are uninformative, and so there is no selection effect; that is, 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 computers’ recommendations are correlated with the prospects of the risky project and, therefore, there is a selection effect; that is, 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 unobservable driving sample selection in our experiment is the private information of the computers. A large 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, due to difficulty with both Bayesian updating and pivotal thinking.³ [Kagel and Levin \(2002\)](#) survey their and others’ substantial early work, and [Charness and Levin \(2009\)](#) and [Ivanov, Levin, and Niederle \(2010\)](#) provide more recent contributions. [Esponda and Vespa \(2014\)](#) consider a setting where Bayesian updating is trivial and continue to find that a significant fraction of subjects make mistakes due to a failure of pivotal thinking. 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 \(2008\)](#). These mistakes are also studied under nonequilibrium concepts (e.g., [Crawford and Iriberry \(2007\)](#)).

Our experiment differs from previous work in that subjects do not know the primitives and, importantly, do not observe counterfactual outcomes.⁴ Without either of these modifications, there would be no endogenous selection problem to study; that is, the probability distribution over the data observed by the subject regarding the performance of the risky project would no longer depend on the subject’s decisions. To see this point, suppose that subjects do not know the primitives but observe all counterfactual outcomes. Then, asymptotically, the observed proportion of good risky projects must equal the true proportion, irrespective of the subject’s choices, and so we are back to the standard case in the literature where the subject knows the probability that the project is good.

It is unclear how to extrapolate previous experimental findings, where primitives are known, to our setting. In previous experiments, subjects should compute a conditional

²One benefit of this design is that we can distinguish a subject who does not understand selection from a subject who understands selection but is unable to perfectly account for it (even professional researchers struggle here).

³For difficulties with Bayesian updating, see [Charness and Levin \(2005\)](#) and references therein.

⁴From a theory perspective, [Esponda \(2008\)](#) formalizes (the failure to account for) endogenous sample selection that is driven by others’ private information.

expectation, and this computation requires knowledge of the primitives, which include nature's and other players' strategies. In our experiment, subjects do not even need to know these primitives to make the right decision. All they need to learn is the probability that the risky project is good conditional on being pivotal. Subjects can easily estimate this probability by keeping track of the proportion of successful projects that were implemented in the past due to their pivotal recommendation.

Another reason why previous results are, at best, suggestive is that providing primitives might actually induce 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. If this chance is say, 75%, the subject may be inclined to ignore all other information and go for the risky project, even though a deeper analysis would reveal that the chance of success conditional on being pivotal is actually very low. Thus, by not providing primitives, we eliminate an important mechanism that underlies previous results. This comment, however, 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.

While not providing primitives was common in early experimental work, it is currently underexplored.⁵ There are two reasons why it is important to examine the case where primitives are unknown and counterfactuals are unobserved. First, these assumptions match many real-world scenarios, where the agent learns from previous decisions and cannot observe the payoff from alternatives that she does not choose. Second, an important objective of experiments is to test for equilibrium behavior, but, in fact, justifications of equilibrium often do not rely on the assumption that primitives are known. 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 strategies of "nature" and other players, without the presumption that people a priori know the former.⁶

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 con-

⁵In early experiments on competitive equilibrium (e.g., Smith (1962)) and in more recent experiments on Cournot equilibrium (e.g., Huck, Normann, and Oechssler (1999) and Rassenti, Reynolds, Smith, and Szidarovszky (2000)), 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, Roberson, and Smith (1982)).

⁶See, for example, Fudenberg and Levine (1998), Dekel, Fudenberg, and Levine (2004), and Esponda (2013), who also points out that uncertainty about fundamentals and strategies are treated in the same manner by epistemic game theorists.

sistent 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. In particular, subjects are, on average, paying attention to the observed data, but not to the possibility that the data may be biased. As a result, subjects do fairly well in the No Selection treatment, but largely miss the problem in the Selection treatment.

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 overpredict risky behavior in the Selection treatment. Controlling for risk aversion, we find that subjects overestimate the benefits of choosing 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.⁷ We discover, however, that it is rather natural for subjects to be partially 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 and 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 between four and five times more weight on pivotal versus nonpivotal 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—that 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.

⁷If subjects placed even a small prior probability on the event that the random process is informative, then one might expect them eventually to learn that this correlation exists and correctly to account for it (particularly in our experiment). Thus, it would seem that, in steady state, subjects should be either completely naive or fully sophisticated in this experiment.

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 only a couple of subjects understand selection can be contrasted with previous experiments in which subjects know the primitives and a nonnegligible 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?

[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.

ROAD MAP

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 Supplementary Material, available in supplementary files on the journal website, <http://qeconomics.org/supp/650/supplement.pdf> and http://qeconomics.org/supp/650/code_and_data.zip.

2. THE EXPERIMENT

2.1 *Experimental design*

Each of our subjects participates in a single-agent decision problem. We provide a summary of the instructions in language that is familiar to economists. We include detailed instructions, with the exact wording given to subjects, in Appendix SD in the Supplementary Material.

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.

		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.

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 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; that is, 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 to 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

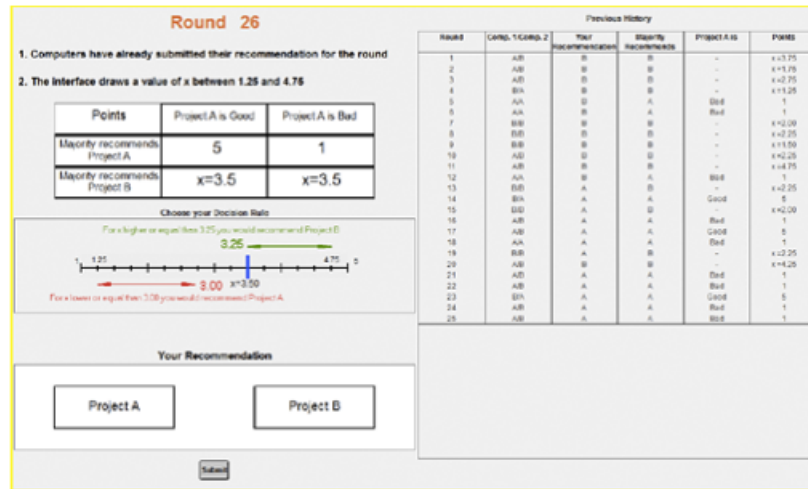


FIGURE 2. Screen shot for round 26. In rounds 1–25, the subject must submit a recommendation for a given value of x . In rounds 26–100, the subject must first submit a threshold recommendation that indicates a choice for each value of x . She is then prompted to submit a choice for a particular value of x , as in rounds 1–25. In this figure, the subject has already submitted a threshold recommendation (A if $x \leq 3$ and B if $x \geq 3.25$). A value of $x = 3.50$ is then drawn, and the subject is asked to recommend A or B. If the subject were to recommend A, she would be alerted that her choice is inconsistent with her previously submitted threshold recommendation and would be asked to submit a recommendation that is consistent (in this example, B).

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 after the subject has selected a threshold.⁸

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 100.⁹ 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?

⁸This 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)).

⁹This part was anticipated in the instructions of Part I 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 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?

Part 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 0.1, 0.3, 0.5, 0.7, and 0.9.¹⁰

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.¹¹

No Selection treatment Each computer recommends A and B with equal probability, irrespective of whether A is good or bad, that is, $m_G = m_B = 1/2$. 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, that is, $m_G = 0$, $m_B = 1/3$. The computers' recommendations in this treatment are informative.

As explained in the next section, when the 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 and III lasted about 25 minutes. Average payoffs were approximately \$18.

¹⁰At 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; as discussed in footnote 34, the two measures are consistent with each other.

¹¹An additional, atypical benefit of not providing the subjects with the primitives is that the instructions for both treatments are *exactly* the same.

TABLE 1. Example of feedback faced by a subject after 12 rounds in the Selection treatment.

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

Note: A naive approach is to estimate the probability of good by looking at the relative proportion of good versus 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.

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 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.¹²

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 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

¹²Following Esponda and Pouzo (2016b), 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 (2016) for related solution concepts.

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 1.¹³ 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 the 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 constant relative risk aversion (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 $r < 0$.¹⁴ Then the optimal (naive and sophisticated) threshold x^* for a subject with risk aversion r is given by the solution to

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

Figure 3 plots the (naive and sophisticated) threshold as a function of the coefficient of relative risk aversion, r . As expected, the threshold decreases as risk aversion increases.¹⁵

2.4.2 Steady-state behavior in the 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

¹³The importance of “pivotality” in these types of environments is highlighted by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1997). Esponda and Pouzo (2016a) show that steady-state behavior corresponds to Nash equilibrium under sophisticated learning and behavioral equilibrium under naive learning.

¹⁴For $r = 1$, we let $u(c) = \ln c$.

¹⁵For 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.

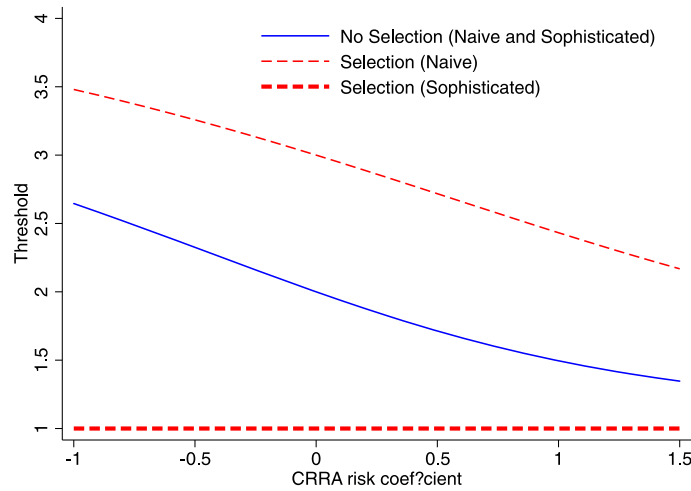


FIGURE 3. Theoretical prediction for the Selection and No Selection treatments. For the benchmark case of the No Selection treatment, naive and sophisticated thresholds coincide. Under Selection, naive and sophisticated thresholds go in opposite direction: higher than the benchmark in the naive case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

be bad. Thus, it is optimal to 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 $x_{NE}^* = 1$.¹⁶

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

$$\begin{aligned}
 z(x^*) &\equiv \Pr(\text{good} \mid M_A; x^*) \\
 &= \frac{\Pr(M_A \mid \text{good}; x^*)p}{\Pr(M_A \mid \text{good}; x^*)p + \Pr(M_A \mid \text{bad}; x^*)(1-p)} \quad (2)
 \end{aligned}$$

¹⁶In Appendix SC, we show that a sophisticated agent can actually identify the primitives (p, m_G, m_B) of the model, but that she would not be able to identify the primitives of a more general model where the votes of the computers are correlated conditional on the state of the world (in our experiment, they are conditionally independent and identically distributed (i.i.d.)). Of course, identifying these primitives is not necessary, since all that is required is that the agent can identify the probability that A is good conditional on being pivotal. This conditional probability, which is zero, can simply be identified by observing the proportion of times that A was good when the subject was pivotal and voted for A, and this is true irrespective of whether the agent thinks that votes are correlated or i.i.d. conditional on the state.

$$\begin{aligned}
&= \frac{\left((1 - m_G)^2 + 2m_G(1 - m_G)\frac{(x^* - 1)}{4} \right) p}{\left((1 - m_G)^2 + 2m_G(1 - m_G)\frac{(x^* - 1)}{4} \right) p + \left(m_B^2 + 2m_B(1 - m_B)\frac{(x^* - 1)}{4} \right) (1 - p)} \\
&= \frac{3}{3 + x^*},
\end{aligned}$$

where we have used the fact that, in the Selection treatment, $m_G = 0$ and $m_B = 1/3$. Note that the steady-state belief is above the true unconditional probability that A is good, which is $p = 1/4$; thus, the naive subject is overoptimistic about the risky project.

Equation (2) 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 is 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 $z(x^*)$, that is,

$$z(x^*) \times u_r(5) + (1 - z(x^*)) \times u_r(1) = u_r(x^*). \quad (3)$$

In other words, the naive steady-state threshold $x^*(r)$ is the unique solution to equation (3).¹⁷ For example, if the subject is risk neutral, $r = 0$, then equation (3) becomes $4/(1 + x^*/3) + 1 = x^*$ and the naive threshold is $x^*(0) = 3$. Figure 3 plots the naive threshold $x^*(r)$ as a function of the coefficient of relative risk aversion, r . As expected, the threshold decreases as risk aversion increases.

For comparison, in the No Selection treatment, the assumption that $m_B = m_G = 1/2$ implies that $z(x^*) = p$ for all x^* . This result formalizes earlier claims that, in the No Selection treatment, (i) beliefs do not depend on decisions (i.e., there is no endogenous selection problem) and (ii) the subject has a correct belief about the unconditional probability that the risky project is good.

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

¹⁷The solution is unique because the left-hand side (LHS) of equation (3) is decreasing (because $z(\cdot)$ is decreasing) and the right-hand side (RHS) is increasing.

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 (represented here by computers) 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 has 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 versus 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 the endogenous selection problem that we wish to study. If subjects knew the primitives, then the problem 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, Levin, and Niederle \(2010\)](#), and [Esponda and Vespa \(2014\)](#)). If a subject were to observe counterfactuals, then her choices would have no influence over the observed performance of the risky project. Hence, there would be no endogenous sample selection problem to study. Instead, by simply keeping track of the proportion of times that A was good, the subject would learn the true probability that project A is good. Learning this probability does not, however, imply that the subject would behave optimally. The reason is that the relevant probability is the one that is conditional on being pivotal. But whether or not the subject can carry out the pivotal calculation when the probability over the state space is known is a problem that has been studied in the previous literature.

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 nonstationary 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 nonstationary settings.

Mistakes of the computers In the Selection treatment, mistakes need to be asymmetric (i.e., different in the good and bad states) for the recommendations of the computers to be informative and, hence, to obtain selection effects. There are of course many choices of asymmetric mistake rates that lead to large selection effects. We choose the mistake rates that make it easiest for a subject to realize that selection is an issue as well as to be able to account for it. By choosing a zero mistake rate in the good state, it follows that *every time* a subject causes A to be implemented, she finds out that A is bad. Without this choice of mistake rates, we would be concerned about classifying as naive a subject who in fact understand selection, but, due to the noisy nature of the data, is not able to perfectly account for it.

Size of incentives The incentives to behave optimally are fairly small in our setting because subjects are pivotal with a probability of $1/3$. For example, in the Selection treatment, the naive choice gives approximately 94% of the payoff of the sophisticated choice for a risk-neutral agent. The significant treatment effects that we obtain, however, suggest that subjects are indeed responding to these small incentives. In particular, subjects do fairly well in the No Selection treatment, despite the fact that the incentives are similar in both treatments.¹⁸ Similar responses 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, but the aggregate effects of individual actions tend to have large welfare consequences.

Focus on steady-state behavior As argued in the [Introduction](#), the experimental design is intended to focus on steady-state behavior, which is typical in economics.¹⁹ We are not able to identify the exact learning dynamics with just two treatments. For example, subjects could have different incentives to experiment in the different treatments, due to different observations, and differences in behavior in initial rounds could be driven by these different incentives. Thus, we leave the important question of identifying the learning rules used by subjects for future work. The important point to keep in mind is that the steady-state predictions that we characterize and test for in this paper hold irrespective of the subjects' incentives to experiment.²⁰

3. RESULTS

We organize the presentation of the results around five main findings.

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

¹⁸For the No Selection treatment, the difference between the best and worst expected payoff is \$2.20, and the ratio of the worst divided by the best expected payoff is 80%. For the Selection treatment, the difference is \$2.67 and the ratio is 80%.

¹⁹For comparison, consider any experimental test of Nash equilibrium in a game with complete information. Subjects are initially uncertain about the strategies of other players, and researchers typically have participants face several repetitions of the game (with random matching) to provide them with experience. The main focus is often to understand whether or not beliefs and behavior eventually stabilize, and, if so, if they are consistent with a steady-state concept, such as Nash equilibrium. We follow this same approach but in a context in which subjects do not know the primitives of the environment.

²⁰This is where the assumption that x varies throughout the experiment is useful. Note that if x were constant, then the problem would be similar to a bandit problem, where the focus is instead placed on whether or not subjects experiment optimally. With variation in x , subjects will get enough information in the steady state irrespective of their initial behavior, and so the steady-state prediction will not be affected by a subjects' discount factor (hence, by her incentives to experiment).

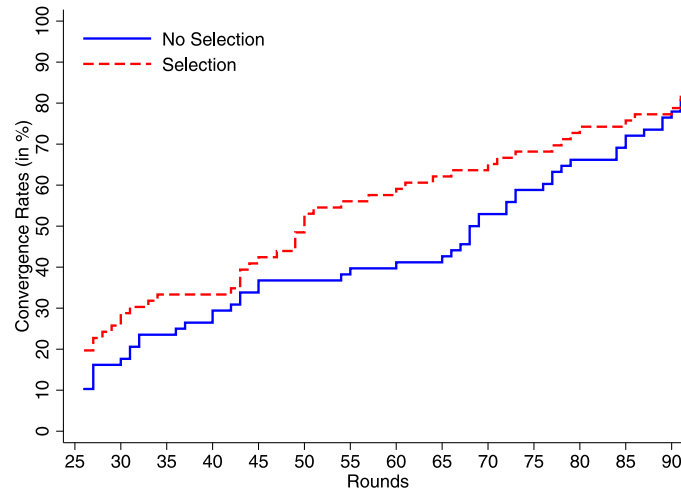


FIGURE 4. Convergence rates by treatment.

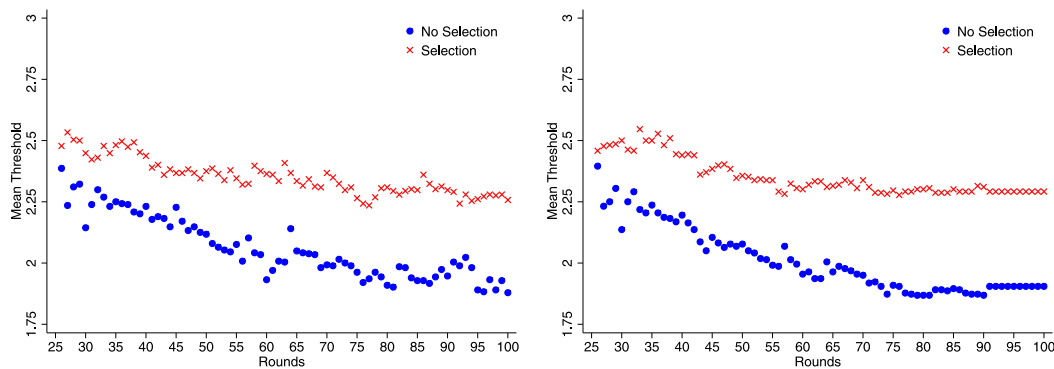


FIGURE 5. Mean thresholds in all rounds by treatment. The left panel shows the mean threshold for all subjects, for each round and treatment. The right panel shows the same information but only for subjects whose behavior converges in the sense that their threshold choice is constant for the last 10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

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 minimum number in the corresponding 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 thresh-

old is somewhere in the interval $[1.75, 2]$, and we code the observed threshold as 1.75. 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, that is, 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.

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. 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 each of the last 10 rounds, and we refer to their threshold choices as their convergent thresholds. In Appendix SA, 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 1.90 under No Selection and 2.29 under Selection; the median convergent thresholds are 1.75 and 2.38, respectively. The differences in the mean (0.39) and the median (0.63) are both statistically significant at the 1% level.²¹

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, so that there is essentially no evidence of

²¹To 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.

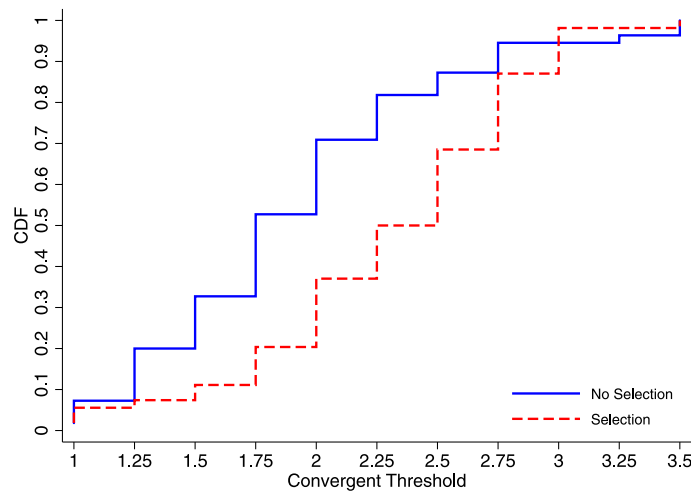


FIGURE 6. Distribution of convergent threshold choices, by treatment. Convergent threshold choices under Selection first-order stochastically dominate choices under No Selection.

sophistication in this experiment.²² Moreover, the empirical distribution of convergent thresholds for the Selection treatment first-order stochastically dominates the distribution in the No Selection treatment.²³

FINDING 3. REPORTED BELIEFS ARE CONSISTENT WITH NAIVE (BIASED) 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 beliefs are consistent with naiveté. 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

²²In Part II of the experiment, we asked subjects to provide a written report to justify their round 100 choice. In the case of the Selection treatment, only three subjects provide a correct explanation of optimal behavior. See footnote 15 in Appendix SD for further details on the reports.

²³We 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 0.770. 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 0.002.

TABLE 2. Mean values of data and reported beliefs, by treatment.

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	32.6	49.9	40.4

Note: Reported beliefs are consistent with naive (biased) beliefs, not with sophisticated beliefs. The label % Good denotes the percentage of times that project A was good; % Good|Piv denotes the percentage of times that project A was good conditional on the subject being pivotal; % Mistake|good denotes the percentage of times a computer mistakenly votes for B when project A is good; % Mistake|bad denotes the percentage of times a computer mistakenly votes for A when project A is bad; Data (True) denotes the actual figure in the data; Data (Naive) denotes the actual figure a naive subject would report given the data; Report denotes the figure reported by subjects in Part II.

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 that 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 the 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.²⁴

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

²⁴As shown in Appendix SC, the true primitives can be identified from the data by a sophisticated subject who believes that votes are i.i.d. conditional on the state of the world. If, for some reason, a sophisticated subject were to believe that votes are correlated conditional on the state, then the primitives are not identified. In this second case, where multiple beliefs are consistent with sophistication (all of which, importantly, make it optimal to always vote for B), the results of this section should be interpreted as a demonstration that elicited beliefs are *consistent* with naiveté, rather than as a test of naivete versus sophistication.

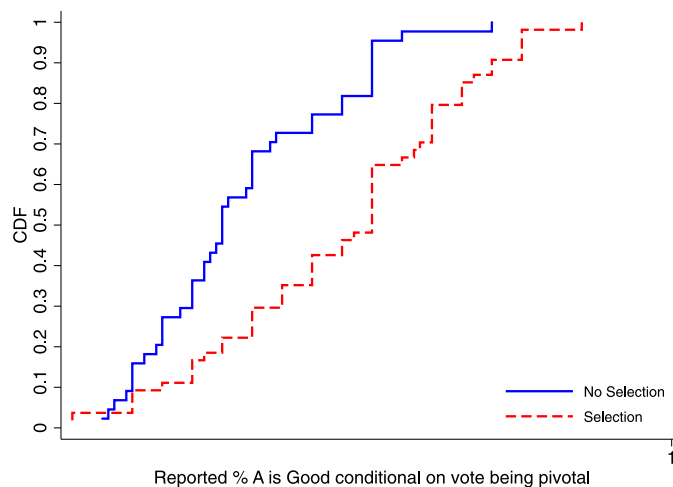


FIGURE 7. Distribution of reported beliefs on project A being good conditional on the recommendation being pivotal, by treatment. Includes only subjects with a convergent threshold.

is close to the true unconditional probability of $1/2$. In the No Selection treatment, subjects are on average correct to respond that the computers' strategies are uninformative. In the Selection treatment, subjects realize that the rates of the computers' mistakes are lower, but are far from realizing that computers make no mistakes when project A is good.

Even though the average reported belief is consistent with naiveté, it could still be possible that some subjects are sophisticated and understand that the probability that project A is good conditional on being pivotal is zero (or very low). Figure 7 shows the distributions of reported beliefs in each treatment for the probability that project A is good conditional on being pivotal. There is no mass point near zero in the Selection treatment, showing that essentially no subject realizes that the probability of the relevant event is zero.²⁵ The figure also shows that the distribution of reported beliefs for the Selection treatment first-order stochastically dominates the distribution for the No Selection treatment.²⁶ This is further evidence consistent with subjects being naive and overestimating the benefits of project A in the Selection treatment.

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).

²⁵In Appendix SA, we show that the same is true for the question about the probability that computers make a mistake when project A is good.

²⁶Following the test by Barrett and Donald (2003) (see footnote 23), we cannot reject the null hypothesis that the distribution under the the Selection treatment either first-order stochastically dominates or is equal to the distribution under No Selection (p -value of 0.731), but we reject the opposite null hypothesis (p -value of 0.001).

FINDING 4. CONVERGENT THRESHOLDS ARE 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.

For a more detailed comparison, we now contrast the observed distribution of thresholds in each treatment with the theoretical prediction. We start by describing an empirical model that we use to compute a prediction of the distribution of convergent thresholds under the assumption that subjects are naive. As discussed earlier, the predicted threshold depends on the risk coefficient, and so the distribution of naive thresholds predicted by the theory depends on the distribution of risk aversion in the population of subjects. In a first stage, we estimate the distribution of risk aversion using data from the five decisions in Part III (decision problem). In particular, we assume that subject i 's threshold choice is the optimal threshold plus some noise,

$$x_{ik}^* = u_{r_i}^{-1}(z_k \times u_{r_i}(5) + (1 - z_k) \times u_{r_i}(1)) + \varepsilon_{ik}, \quad (4)$$

where r_i is her CRRA risk coefficient, z_k is the probability that A is good, u_{r_i} is the CRRA utility function, ε_{ik} represents decision noise, and $k = 1, \dots, 5$ indexes the five decisions taken by the subject. Recall that the probability that A is good is known and given by $z_1 = 0.1$, $z_2 = 0.3$, $z_3 = 0.5$, $z_4 = 0.7$, and $z_5 = 0.9$ in each of the five decisions.²⁷ 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 decisions, and we estimate the parameters using (simulated) maximum likelihood.²⁸

For a given value of r , we can compute the naive threshold ($x^*(r)$) in the No Selection and the Selection treatments using, respectively, equations (1) and (3). We then assume that

$$x_i^* = x^*(r_i) + \varepsilon'_i,$$

where x_i^* is the convergent threshold of subject i and $\varepsilon'_i \sim N(\mu_{\varepsilon'}, \sigma_{\varepsilon'}^2)$ is an error term. In a second stage of the estimation, we use the first stage output and estimate the parameters of the distribution of ε' that maximize the likelihood that the naive threshold plus

²⁷In the data, we do not observe the exact threshold choice 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, and we account for it in the estimation. For example, if $x_{ik}^* = 2.33$, this means that a subject would choose A for all values of x lower than or equal to 2.25 and choose B for all values of x higher than or equal to 2.5.

²⁸We need to simulate the likelihood function because r_i enters nonlinearly in equation (4).

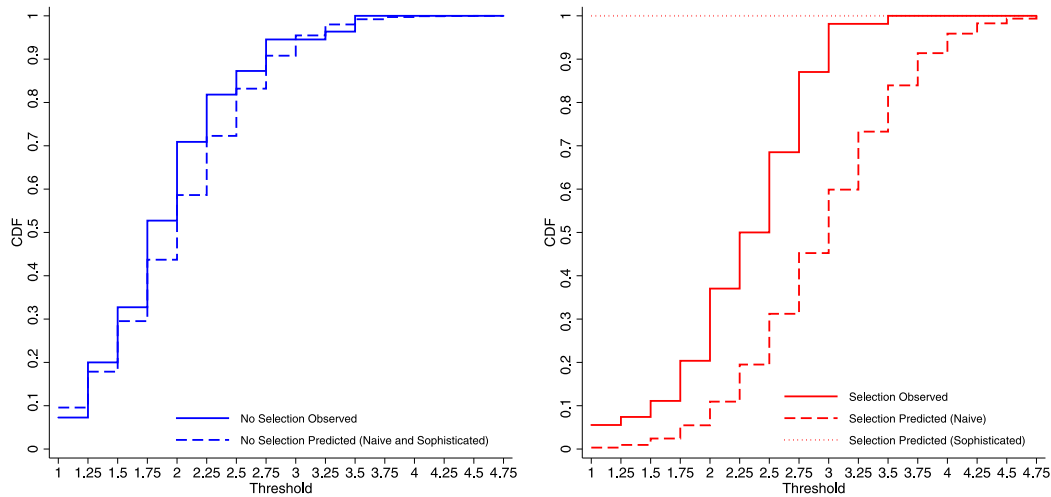


FIGURE 8. Predicted distribution of thresholds under the assumption of naive behavior versus the observed distribution of convergent thresholds, by treatment. Includes only subjects with a convergent threshold.

error is equal to the convergent threshold. Finally, we use estimates of the distributions of r and ε' to predict the distribution of convergent thresholds in each treatment.²⁹

Figure 8 depicts the observed and predicted distributions of thresholds for the No Selection (left panel) and Selection (right panel) treatments. The figure confirms that the theory provides a good fit for the No Selection treatment and that naiveté, as opposed to sophistication, correctly predicts the direction of the treatment effect. The figure shows, however, that the distribution of thresholds predicted by naiveté first-order stochastically dominates the observed distribution for the Selection treatment, thus confirming that naiveté overpredicts the treatment effect.

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.

Panel (a) of Table 3 shows the results of linear regressions of an indicator variable for whether or not a subject changes her threshold x_t 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,

²⁹Further details of the estimation procedure are presented Section 4, where we present a more general model of partially naive subjects, and in Appendix SB.

TABLE 3. Reduced-form analysis: reaction in threshold to events in the previous period.

(a) The dependent variable and the controls are dummy variables¹

Dep. Var.: $\mathbf{1}\{x_t \neq x_{t-1}\}$	Pooled	No Selection	Selection
Constant	0.040*** (0.006)	0.051*** (0.009)	0.029*** (0.008)
Piv_{t-1}	0.003 (0.005)	-0.007 (0.008)	0.016** (0.007)
Info_{t-1}	0.016** (0.007)	0.021* (0.012)	0.012* (0.007)
$\text{Piv}_{t-1} \times \text{Info}_{t-1}$	0.053*** (0.013)	0.058*** (0.020)	0.042** (0.013)

(b) All controls are dummy variables²

Dep. Var.: $x_t - x_{t-1}$	Pooled	No Selection	Selection
Constant	0.008*** (0.002)	0.009*** (0.003)	0.008** (0.003)
$(\text{Piv and good})_{t-1}$	0.020 (0.025)	0.020 (0.025)	- -
$(\text{Piv and bad})_{t-1}$	-0.092*** (0.018)	-0.131*** (0.029)	-0.060*** (0.023)
$(\text{Not Piv and good})_{t-1}$	-0.002 (0.005)	-0.008 (0.007)	0.001 (0.006)
$(\text{Not Piv and bad})_{t-1}$	-0.024*** (0.008)	-0.022** (0.010)	-0.032** (0.013)

Note: The asterisks *, **, and *** indicate significance at the 1, 5, and 10% levels, respectively. Standard errors are given in parentheses. 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 I (we lose one observation due to the lag). The regressions pool subjects from both treatments. Conclusions do not change if we add time dummies. ¹The term $\mathbf{1}\{x_t \neq x_{t-1}\}$ takes value 1 if the threshold in period t , x_t , is different than the threshold in period $t-1$, x_{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. ²The variable $(\text{Piv and good})_{t-1}$ 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.

and hence she observed information about project A (Info_{t-1}). The first column reports results from a regression that pools data from both treatments, while the other columns focus on each treatment separately.

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%) in the pooled data. The key finding, however, is that the interaction effect is more than three times stronger: In the pooled data, a subject is 5.3 percentage points more likely to change her threshold if she received information *and* was pivotal in the previous round. The results are similar for each of the two treatments.

While panel (a) looks at the probability of adjusting the threshold, panel (b) 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. In the pooled data, 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. The results are similar for each of the two treatments.³⁰

The findings from Table 3 confirm that subjects tend to react more to pivotal versus nonpivotal 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

In Section 2.4.2, we assumed that a naive subject forms beliefs about the probability that project A is good based on the proportion of times it was observed to be good in the past. We continue to make this assumption, but now we distinguish between periods in which a subject is pivotal or not. In particular, we assume that a subject pays more attention to whether project A was good or bad if she was pivotal than if she was not pivotal. This assumption is motivated by Finding 5 in Section 3, which suggests that subjects are more responsive to data coming from pivotal periods.

Formally, let α denote the probability that a subject recalls an observation from a period in which she was pivotal and let β denote the corresponding probability for a period in which she was not pivotal. Let y_τ denote the number of times that project A was recalled to be good in the past τ recalled periods. Note that it is indeed possible that $\tau < t$ after t periods since the subject does not necessarily recall all past data. Let $z_\tau \equiv y_\tau/\tau$ denote the proportion of times that project A was good, as recalled by the subject, and suppose that it represents the naive subject's belief about the probability that A is good. In the Appendix, we use tools from stochastic approximation to show that if the subject's

³⁰These 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).

threshold converges to x^* , then her belief z_τ converges to

$$z(x^*, \eta) = \left((1 - m_G)^2 + \eta 2m_G(1 - m_G) \frac{(x^* - 1)}{4} \right) p \\ / \left(\left((1 - m_G)^2 + \eta 2m_G(1 - m_G) \frac{(x^* - 1)}{4} \right) p \right. \\ \left. + \left(m_B^2 + \eta 2m_B(1 - m_B) \frac{(x^* - 1)}{4} \right) (1 - p) \right), \quad (5)$$

where $\eta \equiv \alpha/\beta$. The above expression, which represents the steady-state belief of the subject, is simply the probability that A is good conditional on the event that A is observed to be good and that the agent recalls it. The probability of this event, in turn, depends on the steady-state threshold choice of the subject, x^* , and the parameter η . It is important to emphasize that the naive subject is in no way required to be able to compute conditional expectations or to have an understanding of the selection problem. The subject simply follows the rules specified above, and equation (5) provides a *characterization* of the steady-state belief of a subject who follows these rules.

Assuming, once again, a CRRA utility function for convenience, the steady-state strategy x^* is the unique solution to

$$z(x^*, \eta) \times u_r(5) + (1 - z(x^*, \eta)) \times u_r(1) = u_r(x^*). \quad (6)$$

We denote the solution by $x^*(r, \eta)$.

In the [Appendix](#), we show that as time goes to infinity, the threshold converges to $x^*(r, \eta)$ provided that the subject is asymptotically myopic, meaning that there is a time after which she always chooses a threshold to maximize current expected utility. The advantage of this approach is that we do not have to make assumptions regarding how subjects behave in the early periods of the experiment, where incentives to experiment may justify deviations from myopic optimization.

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 the threshold converges to the sophisticated, optimal threshold characterized in Section 2.4. The case $\eta = 1$, which places equal weight on pivotal versus non-pivotal rounds, corresponds to what we called the naive threshold in Section 2.4 (or, equivalently, what [Esponda \(2008\)](#) calls a naive behavioral equilibrium). The parameter η captures intermediate cases where subjects are naive but account for selection by putting higher weight on feedback from pivotal rounds.

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 (5) becomes

$$z(x^*, \eta) = \frac{\left(0.25 + \eta 0.5 \frac{(x^* - 1)}{4} \right) 0.25}{\left(0.25 + \eta 0.5 \frac{(x^* - 1)}{4} \right) 0.25 + \left(0.25 + \eta 0.5 \frac{(x^* - 1)}{4} \right) 0.75} = 0.25.$$

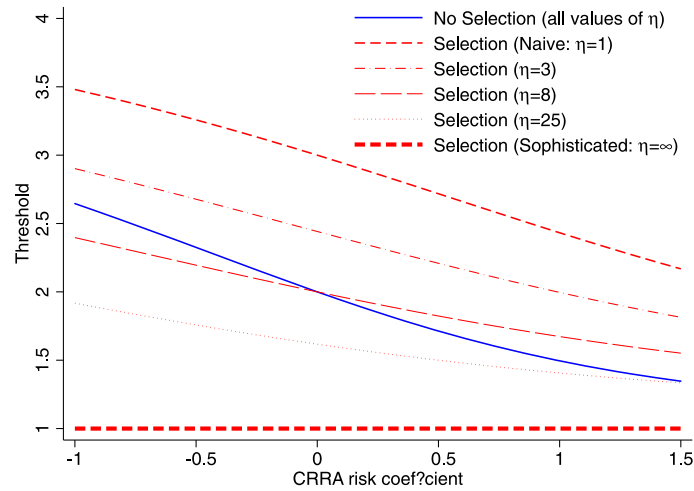


FIGURE 9. Theoretical prediction for Selection and No Selection treatments for several values of η .

As explained earlier, there is no selection in the data, and so the belief equals the true probability that project A is good, 0.25, irrespective of the weight placed on pivotal versus nonpivotal rounds. For the Selection treatment ($p = 1/4$, $m_G = 0$, $m_B = 1/3$),

$$z(x^*, \eta) = \frac{0.25}{0.25 + \left(\frac{1}{9} + \eta \frac{4(x^* - 1)}{9 \cdot 4}\right) 0.75} \quad (7)$$

for all $x^* > 1$. As η increases, more weight is placed on pivotal rounds, where A always turns out bad, and, therefore, $z(x^*, \cdot)$ is decreasing.

Figure 9 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 is 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.³¹ 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 versus nonpivotal rounds. 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 the theoretical prediction with $\eta = 1$ is above observed values (Finding 4), Figure 9 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 use data from steady-state decisions in Part I (either No Selection or Selection treatments) and from the five decisions in Part III (decision problem).

³¹The model also allows for $\eta < 1$, which means that nonpivotal rounds receive relatively higher weight.

We estimate the model in two stages. In the first stage, we use data from Part III (decision problem) to estimate the distribution of risk coefficients. In particular, we follow the same approach that we described in the analysis of Finding 4 in the previous section. In the second stage, we use the steady-state threshold of each subject in period T of Part I of the experiment to identify the extent to which subjects are partially naive. In particular, we assume that

$$x_{iT}^* = x^*(r_i, \eta_i) + \varepsilon'_{iT}, \quad (8)$$

where x_{iT}^* is the threshold choice in period T , $x^*(r_i, \eta_i)$ is the predicted threshold for a subject with risk coefficient r_i and naiveté coefficient η_i , and ε'_{iT} is an error term. For the parameters of the distribution of risk aversion, we use the estimates from the first stage. Thus, in this second stage, we estimate the error term (identified from the No Selection treatment, since in that treatment $x^*(r_i, \cdot)$ does not depend on η) and the coefficient of naiveté (identified from the Selection treatment). We assume that η follows a normal distribution with mean μ_0 and variance σ_0^2 that is truncated to be positive, i.e., $\eta \in [0, \infty)$. We denote the corresponding mean and variance of η by μ_η and σ_η^2 , respectively.³² We also assume that the error term $\varepsilon' \sim N(\mu_{\varepsilon'}, \sigma_{\varepsilon'}^2)$, and that η and ε' are independent of each other and across subjects. We estimate the parameters using (simulated) maximum likelihood. Further details of the estimation procedure are presented in Appendix SB.

To be consistent with the steady-state model, we use data from subjects whose behavior has stabilized by round T , meaning that the threshold choice does not change after round T . For robustness, we estimate the model for several values of $T \in \{70, 75, \dots, 95, 100\}$. Note that the case $T = 100$ corresponds to using data from all subjects.

Table 4(a) presents the maximum likelihood estimates (including standard deviations and 95% confidence intervals) for $T = 90$. With these estimates, we compute the mean and median of η to be 5.45 and 4.65, respectively. Thus, the median subject puts about 4.65 times more weight on pivotal versus nonpivotal events. The result is consistent with the reduced-form results from Section 3, which suggested that subjects placed more weight on pivotal events.

Table 4(b) presents further information on the median of η .³³ Based on bootstrapping the maximum likelihood estimates, we obtain a distribution for the median of η . For $T = 90$, the 2.5th and 97.5th percentiles of the median of η are 1.69 and 10.82, respectively. This shows that the estimate of the median is concentrated around the maximum likelihood estimate (4.65) and is far from being consistent with sophisticated behavior ($\eta \approx \infty$). As explained earlier, even much higher weights on pivotal rounds are not enough to approximate sophisticated behavior, since subjects are pivotal with a relatively small probability of 1/3. Thus, the increased relative weights on pivotal events are not nearly enough to correct for mistakes.

³²Letting ϕ be the density and letting Φ be the cumulative density function (cdf) of the standard normal distribution, it follows that $\mu_\eta = \mu_0 + \frac{\phi(-\frac{\mu_0}{\sigma_0})}{1-\Phi(-\frac{\mu_0}{\sigma_0})}\sigma_0$ and $\sigma_\eta^2 = \sigma_0^2[1 + \frac{\phi(-\frac{\mu_0}{\sigma_0})}{1-\Phi(-\frac{\mu_0}{\sigma_0})} - (\frac{\phi(-\frac{\mu_0}{\sigma_0})}{1-\Phi(-\frac{\mu_0}{\sigma_0})})^2]$.

³³We focus on the median as a measure of central tendency given that the distribution of η is asymmetric. In Appendix SA, we provide statistics on the mean of η .

TABLE 4. Maximum likelihood estimation and the distribution of η .(a) Maximum likelihood estimates for $T = 90$ ¹

	Estimate	Std. Err.	95% Conf. Interval
μ_η	5.452	3.138	[1.986, 12.747]
σ_η	7.938	4.554	[2.901, 18.559]
μ_r	0.463	0.095	[0.264, 0.630]
σ_r	0.589	0.339	[0.000, 0.890]
μ_ε	0.119	0.047	[0.003, 0.193]
σ_ε	0.413	0.059	[0.322, 0.543]
$\mu_{\varepsilon'}$	0.374	0.146	[0.026, 0.595]
$\sigma_{\varepsilon'}$	0.488	0.084	[0.334, 0.649]

(b) Statistics of the median of η using the bootstrap²

Percentile	$T = 70$	$T = 75$	$T = 80$	$T = 85$	$T = 90$	$T = 95$	$T = 100$
2.5	1.62	1.68	1.58	1.54	1.69	1.95	1.94
5	1.78	1.79	1.69	1.67	1.85	2.11	2.07
25	2.49	2.40	2.17	2.16	2.70	2.99	2.85
50	3.70	3.76	2.90	3.03	4.30	4.76	4.35
75	5.84	5.21	4.38	4.59	6.05	6.40	5.68
95	9.16	7.96	6.59	6.88	9.16	9.78	8.37
97.5	10.50	9.44	7.21	7.84	10.82	11.63	9.64

Note: ¹Standard errors and the 95% confidence intervals are computed using 1000 bootstrap repetitions. ²The 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 I that were included in the estimation.

As a robustness exercise, the other columns of Table 4(b) show how the computations change depending on the choice of T . When assessing robustness, it is important to keep Figure 9 in mind, which shows that small differences in the value of η imply small differences in behavior. Thus, Table 4(b) confirms that the main conclusions are unaffected by the choice of T .

Next, we briefly comment on the other estimates of Table 4(a). The results for the risk coefficient and noise levels appear to be consistent with previous work. For example, the mean subject is risk averse with a risk coefficient of relative risk aversion of 0.463, and 95% of the population has a risk coefficient between 0.264 and 0.630, which is consistent with previous estimates (see Holt and Laury (2002), Harrison and Ruström (2008)).³⁴ 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 0.119 in the decision problem where the probabilities are known (Part III) and 0.374 where the probabilities are unknown (Part I). Naturally, the decision error is higher in the problem where the subject does not know the primitives.

³⁴Following 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 have converged by $T = 90$ or all subjects. These estimates are comparable to those reported in Table 4 (which

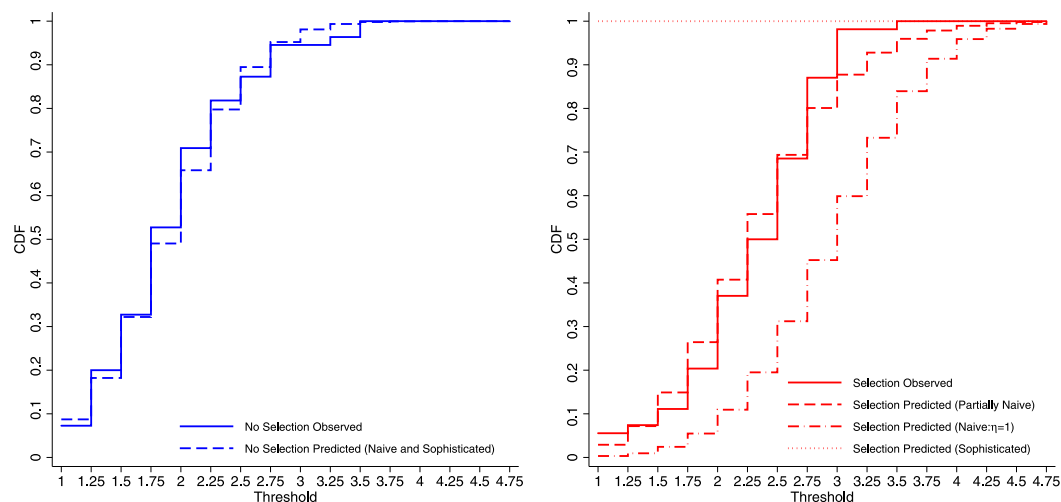


FIGURE 10. Goodness of fit using subjects who converged starting at $T = 90$. In the Selection treatment, the distribution “Selection Predicted (Naive: $\eta = 1$)” corresponds to the prediction under the assumption of full naiveté previously reported in Figure 8.

Finally, in Figure 10, we report goodness of fit for the estimated model, both for the No Selection (left panel) and Selection (right panel) treatments. For the No Selection treatment, the fit was already good under the naive model, where the parameter value is exogenously fixed at $\eta = 1$, and continues to be good now that a distribution of η is being estimated.³⁵ For the Selection treatment, however, the model where the distribution of η is estimated provides a much better fit than the model where we fix $\eta = 1$ (and which leads to overprediction of threshold choices, as discussed in Section 3). Thus, the model of partial naiveté does a good job of rationalizing decisions in both treatments simultaneously.³⁶

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 a subject who understands sample selection has all the available data necessary to ac-

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).

³⁵This is not surprising since the theoretical prediction is that η does not affect decisions in the No Selection treatment.

³⁶Appendix SA shows that the results are robust to several different assumptions, including using a log-normal distribution for η rather than a truncated normal, adding mass points to the distribution of η , and assuming that the belief of a subject is given not by the model's steady-state prediction, but rather by the observed data. We also conduct an out-of-sample prediction exercise that estimates the model by excluding one experimental session and then uses the parameter estimates to predict the results for the excluded session.

count 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 almost no subjects optimally account for endogenous selection. On the other hand, behavior is far from random, but actually quite amenable to analysis. Subjects respond to the observed data and partially account for selection by placing between four and five times more weight on pivotal versus nonpivotal 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.

APPENDIX: CONVERGENCE TO THE STEADY STATE

In the text, we characterized steady-state behavior and beliefs for the model with partial naiveté (which includes full naiveté and perfect sophistication as special cases). In this Appendix, we use tools from the theory of stochastic approximation to show that if a subject eventually chooses a threshold to maximize her perceived current-period payoff, then her beliefs and actions converge almost surely to those characterized in the text.

Recall that y_τ is the number of times in the past that a subject recalls project A being good and that $z_\tau \equiv y_\tau/\tau$, where τ is the number of times that the subject recalls project A being either good or bad. By simple algebra,

$$z_{\tau+1} = z_\tau + \frac{1}{\tau+1}(\xi_{\tau+1} - z_\tau), \quad (9)$$

where $\xi_{\tau+1} = 1$ if the subject registers project A to be good and $\xi_{\tau+1} = 0$ if the subject registers project A to be bad. Without loss of generality, we only keep track of periods in which a subject pays attention to the data she observes, so that, in period t , $\tau \leq t$. In other words, if the subject does not register project A to be good or bad—either because she is not paying attention or because project B is implemented in that period—then we do not advance time from τ to $\tau + 1$.³⁷

Let $p : [1, 5] \rightarrow [0, 1]$ be the function that maps a threshold choice $x_{\tau+1} \in [1, 5]$ into the probability that project A is registered to be good, $\xi_{\tau+1} = 1$, where $\tau + 1$ is a period in which project A is implemented and the subject registers information about project A. In our context,

$$\begin{aligned} p(x_{\tau+1}) &= \frac{\beta \times 0.25}{\beta \times 0.25 + \left(\frac{1}{9}\beta + \frac{4}{9}\frac{(x_\tau - 1)}{4}\alpha\right) \times 0.75} \\ &= \frac{0.25}{0.25 + \left(\frac{1}{9} + \frac{4}{9}\frac{(x_{\tau+1} - 1)}{4}\eta\right) \times 0.75}, \end{aligned} \quad (10)$$

³⁷Note that τ goes to infinity as t goes to infinity because there is a strictly positive probability that project A is implemented in any period and, therefore, that the subject receives information about project A.

where $\eta \equiv \alpha/\beta$, and $\alpha \in (0, 1]$ and $\beta \in (0, 1]$ are the probabilities of paying attention to an observation that project A is good or bad conditional on being pivotal and not pivotal, respectively. Note that $x \rightarrow p(x)$ is decreasing.

We assume that $\{z_\tau\}_\tau$ is the naive subject's sequence of beliefs and that there is a period after which the subject always chooses her threshold to maximize current expected payoff, as perceived given her current belief.³⁸ Thus, for all sufficiently large τ , the threshold $x_{\tau+1}$ solves

$$z_\tau \times u(5) + (1 - z_\tau) \times u(1) = u(x_{\tau+1}),$$

where u is the utility function, assumed to be increasing, continuous, and bounded. Therefore, the threshold given belief z_τ is given by

$$x_{\tau+1} = x^*(z_\tau) \equiv u^{-1}(z_\tau \times u(5) + (1 - z_\tau) \times u(1)),$$

where u^{-1} is the inverse of the utility function. Note that $z \rightarrow x^*(z)$ is increasing.

Next, define the function $q : [0, 1] \rightarrow [0, 1]$ by letting

$$q(z) = p(x^*(z))$$

for all $z \in [0, 1]$. The function q specifies the probability that project A is registered to be good, given that some information about A is registered (i.e., $\xi_{\tau+1} = 1$), if the subject's belief is $z_\tau = z$ and if she chooses a threshold to maximize current expected payoff. Note that $z \rightarrow q(z)$ is decreasing. We can use this function to rewrite expression (9) as

$$z_{\tau+1} = z_\tau + \frac{1}{\tau+1}(q(z_\tau) - z_\tau) + \frac{1}{\tau+1}M_{\tau+1}, \quad (11)$$

where $M_{\tau+1} \equiv \xi_{\tau+1} - q(z_\tau)$ is a martingale difference sequence. The above equation can be thought of as a noisy discretization for the ordinary differential equation (ODE)

$$\dot{z}(\tau) = q(z(\tau)) - z(\tau). \quad (12)$$

The ODE in (12) has a unique steady state that solves $q(z^*) = z^*$ and, moreover, $\dot{z} > 0$ for all $z < z^*$ and $\dot{z} < 0$ for all $z > z^*$.³⁹ Thus, the trajectories of z_τ converge to z^* for any initial condition $z_0 \in (0, 1)$. A standard result from stochastic approximation (see, e.g., Borkar (2008)) says that as τ goes to infinity, the trajectories of $\{z_\tau\}$ in (11) are almost surely given by the trajectories of the ODE in (12). Thus, the sequence of beliefs $\{z_\tau\}_\tau$ converges to z^* with probability 1. Moreover, by continuity of $x^*(\cdot)$, it follows that the sequence of thresholds $\{x_\tau\}$ converges to $x^*(z^*)$ with probability 1. Finally, note that the steady-state belief and thresholds z^* and $x^*(z^*)$ are those that solve equations (6) and (7) in the text.

³⁸This is a form of asymptotic myopia (Fudenberg and Kreps (1993)). Alternatively, we could assume that the threshold choice of the agent converges, which is what we see in the data for most of the subjects, and that the convergent threshold maximizes current expected payoff given beliefs.

³⁹This is because $q(\cdot)$ is decreasing, is continuous, and satisfies $q(0) > 0$ and $q(1) < 1$.

The result that the threshold converges relies on the continuity of $x^*(\cdot)$, which in turn relies on the assumption that the agent can choose from a continuum of thresholds, $x \in [1, 5]$. In the experiment, the subject was restricted to choose from a finite set of thresholds. In that case, the above result that the belief converges with probability 1 continues to be true, but now it is not necessarily the case that the threshold converges to a unique threshold. For those cases where z^* is such that the agent is indifferent between two thresholds, then the threshold choice will converge to a probability distribution over two contiguous thresholds (such as, for example, 2.5 with probability 1/2 and 2.75 with probability 1/2). We account for this issue in the estimation, as described in Appendix SB.⁴⁰

REFERENCES

- Austen-Smith, D. and J. S. Banks (1996), “Information aggregation, rationality, and the Condorcet jury theorem.” *American Political Science Review*, 90, 34–45. [193]
- Barrett, G. F. and S. G. Donald (2003), “Consistent tests for stochastic dominance.” *Econometrica*, 71 (1), 71–104. [200, 202]
- Bazerman, M. H. and W. F. Samuelson (1983), “I won the auction but don’t want the prize.” *Journal of Conflict Resolution*, 27 (4), 618–634. [186]
- Borkar, V. S. (2008), *Stochastic Approximation*. Cambridge University Press. [213]
- Brandts, J. and G. Charness (2011), “The strategy versus the direct-response method: A first survey of experimental comparisons.” *Experimental Economics*, 14 (3), 375–398. [190]
- Charness, G. and D. Levin (2005), “When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect.” *American Economic Review*, 95 (4), 1300–1309. [185]
- Charness, G. and D. Levin (2009), “The origin of the winner’s curse: A laboratory study.” *American Economic Journal: Microeconomics*, 1 (1), 207–236. [185, 188, 196]
- Cox, J. C., B. Roberson, and V. L. Smith (1982), “Theory and behavior of single object auctions.” *Research in Experimental Economics*, 2, 1–43. [186]
- Crawford, V. P. and N. Iriberry (2007), “Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions?” *Econometrica*, 75 (6), 1721–1770. [185]
- Dekel, E., D. Fudenberg, and D. K. Levine (2004), “Learning to play Bayesian games.” *Games and Economic Behavior*, 46 (2), 282–303. [186]

⁴⁰In the data, a few subjects show convergence to contiguous thresholds. For example, out of 134 subjects, 109 of them converge to a single threshold and 10 of them (5 in each treatment) converge to contiguous thresholds at $T = 90$. Note that the case $T = 100$ in the text includes all subjects, whether or not they converge, and the findings continue to be the same.

- Esponda, I. (2008), “Behavioral equilibrium in economies with adverse selection.” *American Economic Review*, 98 (4), 1269–1291. [183, 185, 192, 207]
- Esponda, I. (2013), “Rationalizable conjectural equilibrium: A framework for robust predictions.” *Theoretical Economics*, 8 (2), 467–501. [186]
- Esponda, I. and D. Pouzo (2016a), “Conditional retrospective voting in large elections.” *American Economic Journal: Microeconomics*, 9 (2), 54–75. [183, 193]
- Esponda, I. and D. Pouzo (2016b), “Berk–Nash equilibrium: A framework for modeling agents with misspecified models.” *Econometrica*, 84 (3), 1093–1130. [183, 192]
- Esponda, I. and E. Vespa (2014), “Hypothetical thinking and information extraction in the laboratory.” *American Economic Journal: Microeconomics*, 6 (4), 180–202. [185, 186, 188, 196, 197]
- Eyster, E. and M. Rabin (2005), “Cursed equilibrium.” *Econometrica*, 73 (5), 1623–1672. [185, 187, 192]
- Feddersen, T. and W. Pesendorfer (1997), “Voting behavior and information aggregation in elections with private information.” *Econometrica*, 65, 1029–1058. [193]
- Fudenberg, D. and D. M. Kreps (1993), “Learning mixed equilibria.” *Games and Economic Behavior*, 5 (3), 320–367. [213]
- Fudenberg, D. and D. K. Levine (1998), *The Theory of Learning in Games*, Vol. 2. The MIT press. [186]
- Fudenberg, D. and A. Peysakhovich (2014), “Recency, records and recaps: Learning and non-equilibrium behavior in a simple decision problem.” In *Proceedings of the 15th ACM Conference on Economics and Computation*. [188]
- Harrison, G. W. and E. E. Ruström (2008), “Risk aversion in the laboratory.” In *Risk Aversion in Experiments* (J. C. Cox and G. W. Harrison, eds.), Research in Experimental Economics, Vol. 12, 41–196, Emerald Group Publishing Limited. [210, 211]
- Hey, J. D. and C. Orme (1994), “Investigating generalizations of expected utility theory using experimental data.” *Econometrica*, 62, 1291–1326. [211]
- Holt, C. A. and S. K. Laury (2002), “Risk aversion and incentive effects.” *American Economic Review*, 92 (5), 1644–1655. [191, 210]
- Holt, C. A. and R. Sherman (1994), “The loser’s curse.” *American Economic Review*, 84, 642–652. [185]
- Huck, S., H.-T. Normann, and J. Oechssler (1999), “Learning in Cournot oligopoly—An experiment.” *The Economic Journal*, 109 (454), 80–95. [186]
- Ivanov, A., D. Levin, and M. Niederle (2010), “Can relaxation of beliefs rationalize the winner’s curse?: An experimental study.” *Econometrica*, 78 (4), 1435–1452. [185, 196]
- Jehiel, P. (2005), “Analogy-based expectation equilibrium.” *Journal of Economic Theory*, 123 (2), 81–104. [185]

Jehiel, P. and F. Koessler (2008), “Revisiting games of incomplete information with analogy-based expectations.” *Games and Economic Behavior*, 62 (2), 533–557. [185, 192]

Kagel, J. H. and D. Levin (1986), “The winner’s curse and public information in common value auctions.” *American Economic Review*, 76, 894–920. [185]

Kagel, J. H. and D. Levin (2002), *Common Value Auctions and the Winner’s Curse*. Princeton Univ Pr. [185]

Kőszegi, B. (2010), “Utility from anticipation and personal equilibrium.” *Economic Theory*, 44 (3), 415–444. [183, 192]

Rassenti, S., S. S. Reynolds, V. L. Smith, and F. Szidarovszky (2000), “Adaptation and convergence of behavior in repeated experimental Cournot games.” *Journal of Economic Behavior & Organization*, 41 (2), 117–146. [186]

Smith, V. L. (1962), “An experimental study of competitive market behavior.” *Journal of Political Economy*, 70, 111–137. [186]

Spiegler, R. (2016), “Bayesian networks and boundedly rational expectations.” *Quarterly Journal of Economics*, 131, 1243–1290. [192]

Co-editor Petra E. Todd handled this manuscript.

Manuscript received 8 December, 2015; final version accepted 23 January, 2017; available online 6 February, 2017.