

Behavioral Equilibrium in Economies with Adverse Selection*

Ignacio Esponda[†]
New York University

September 12th, 2007

Abstract

I propose a new solution concept—behavioral equilibrium—to study environments with players who are naive, in the sense that they fail to account for the informational content of other players' actions. A behavioral equilibrium requires that: (i) players have no incentives to deviate given their beliefs about the consequences of deviating; (ii) these beliefs are consistent with the information obtained from the actual equilibrium play of all players; and (iii) when processing this information, naive players fail to account for the correlation between other players' actions and their own payoff uncertainty. I apply the framework to certain adverse selection settings and show that, contrary to the received literature, the adverse selection problem is exacerbated when naive players fail to account for selection. More generally, the main distinguishing feature of the framework is that, in equilibrium, beliefs about both fundamentals and strategies are jointly restricted. Consequently, whether a bias may arise or not is determined endogenously in equilibrium.

Keywords: Behavioral game theory, adverse selection, winner's curse, self-confirming equilibrium

JEL: C72, D82

*This paper is based on a chapter of my Stanford Ph.D. dissertation. I am indebted to Susan Athey, Ilya Segal, Liran Einav, and Jon Levin for their generous support and advice. I also thank Heski Bar-Isaac, B. Douglas Bernheim, Luís Cabral, Vincent Crawford, Willie Fuchs, William Greene, Yuan-Chuan Lien, Alessandro Lizzeri, Pedro Miranda, Demian Pouzo, Dan Quint, Antonio Rangel, Steve Tadelis, Adam Szeidl, three anonymous referees and several seminar participants for helpful comments and discussions. I gratefully acknowledge financial support from the Bradley Graduate and Post Graduate Fellowship.

[†]Department of Economics, NYU Stern School of Business, 44 W 4th Street, KMC 7-76, New York, NY 10012; iesponda@stern.nyu.edu, <http://www.stern.nyu.edu/~iesponda>.

A large literature presents evidence that psychological and cognitive biases affect people's behavior. This evidence raises some questions: What are the effects of these biases? And can they persist at all in economic settings where people strategically interact with each other and have the opportunity to learn from their experience?¹ In this paper, I introduce a game-theoretic equilibrium framework to study these questions in the context of a particular bias: people's failure to take into account the informational content of other people's actions. I apply the framework to a class of games that arises naturally in adverse selection settings, and then show that the adverse selection problem is exacerbated when players suffer from this bias.

To illustrate the main ideas, consider how this bias affects outcomes in a trading game in the spirit of George A. Akerlof (1970). A feature of this adverse selection setting is that lower prices select worse quality, which itself gives the buyer incentives to offer lower prices. This feature is often provided as an intuition for why markets are thinner and gains from trade are lower in the presence of information asymmetries between buyers and sellers. Given this intuition, it may be expected that the adverse selection problem will be mitigated if buyers do not realize that the seller's willingness to trade provides information about quality.

The previous literature (John H. Kagel and Dan Levin 1986; Charles A. Holt and Roger Sherman 1994; Erik Eyster and Rabin 2005) models a buyer who fails to account for selection by assuming that she incorrectly believes that the quality of traded objects is given by the unconditional expectation, rather than by the expectation conditional on the information that the seller wants to trade at the offered price. Since the unconditional expectation is higher than the conditional one when valuations of the buyer and seller are positively related, a biased buyer has a greater incentive than a non-biased buyer to choose higher prices (i.e., relative to Nash equilibrium). In a common value auction, which is an extension of the simple trading game, this overbidding phenomenon is known as the winner's curse.

Two features of the standard approach motivate the alternative approach in this paper. First, a biased buyer believes that she will, on average, obtain objects worth the unconditional expectation, while in equilibrium the quality of traded objects will be lower. But if the buyer were to face this situation repeatedly and learn from her experience, her beliefs would eventually be contradicted. A question is then raised as to whether the bias would still persist. Second, how does the buyer know what the unconditional expected quality is to begin with? In many settings, it is reasonable to instead assume that the buyer's beliefs about quality depend on her past trading experience.

The solution concept that I propose—behavioral equilibrium—makes clear that while players may adjust their behavior when experience contradicts their beliefs, they may still not realize how their learning experience would have been different had they chosen to behave differently. When the buyer receives feedback about the value of traded objects, she must then have correct beliefs about

¹See Matthew Rabin (1998) and Thomas Gilovich, Dale Griffin, and Daniel Kahneman (2002) for a review of several biases, and Drew Fudenberg (2006) for a discussion of the need of an appropriate equilibrium framework for behavioral economics.

the expected quality of objects that are traded at the price she chooses in equilibrium. Hence, the overpricing incentive described above disappears. However, a buyer who is naive, in the sense that she fails to account for selection, does not realize that higher prices would increase quality, so she has, at the margin, lower incentives to increase prices relative to a non-naive buyer. Therefore, in (a behavioral) equilibrium, the adverse selection problem is exacerbated if the buyer is naive. This example is further discussed in Section I, which also presents a dynamic learning interpretation to justify the proposed equilibrium concept.

While the above example illustrates the main intuition, the framework is applicable to a wide range of settings, and Section III discusses additional applications. Failing to account for the informational content of other peoples' actions is analogous to ignoring a potential selection problem. This problem may arise in general adverse selection settings, where the terms of the contract often *select* the type of people with whom trade will be conducted. For example, while a firm may know how different terms of trade would affect its number of customers, it may still not know whether or how different terms of trade would select different *types* of customers. More generally, this selection problem is present in *any* game that has some common value component and where players act based on privately held information.

Most of the evidence for people's failure to account for the informational content of other peoples' actions comes from experiments in auction-like environments (see Kagel and Levin 2002 for a review of the evidence), and discussions of this bias also have appeared in field settings, including the oil industry (Edward C. Capen, Robert V. Clapp, and William M. Campbell 1971), professional baseball's free agency market (James Cassing and Richard W. Douglas 1980), and corporate takeovers (Richard Roll 1986).² However, as discussed in Section IV, the consequences that the experimental literature attribute to the bias may be related to the fact that subjects are informed about the true distribution over fundamentals. Instead, I ask: given the evidence that people fail to account for adverse selection, should we expect this bias to persist in an equilibrium environment that has the realistic feature that people also have to learn the value of the object? And if so, what would its effect be?

Additional motivation is provided by the fact that the bias can be formally modeled as a failure to account for the correlation between the actions of other players and payoff-relevant uncertainty. Complexity of the environment may preclude people from understanding the relevant relationships in the data for the problem at hand.³ Certain organizational structures may also promote this bias. For example, a firm may have one division, say the research department, that produces estimates about uncertain demand conditions using past data, and a different division, say the pricing division, that keeps track of its competitors' prices. If competitors choose prices based on their own estimates of demand conditions, then these two pieces of information are likely to be

²Gary Charness and Levin's (2006) recent individual-choice experiment provides compelling evidence that people have difficulty with what they refer to as contingent reasoning, which, here, amounts to not accounting for selection.

³See Enriqueta Aragonés, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler (2005) for a justification of why people may fail to see regularities in the data.

correlated; but if these divisions do not communicate with each other, this correlation is likely to be missed. Finally, to the extent that selection problems often pose challenges for empirical researchers, it seems plausible that economic agents may not always account for selection when learning about their environment. For example, William Greene (1998) points out that credit-scoring models in the industry often fail to account that they are based on samples of individuals to whom credit has already been given.⁴

I distinguish between two types of players, each having a different and exogenously given *model of the world*: those who are not aware of the potential selection problem (naive players), and those who are aware (sophisticated players). *Within their constrained model of the world*, both types of players: (1) use available data to form beliefs about the consequences of their actions; and (2) choose actions that maximize utility subject to these beliefs.

A behavioral equilibrium is based on the idea of a self-confirming equilibrium (Eddie Dekel, Fudenberg, and David K. Levine 2004), which is a static, steady-state solution concept requiring that players have no incentives to deviate given their beliefs about the consequences of deviating, and restricting these beliefs to be consistent with the experience that results from equilibrium behavior.⁵ In contrast, Nash equilibrium is more restrictive since it requires beliefs about the consequences of deviating to *any* strategy to be correct. While sophisticated players behave as in a self-confirming equilibrium, the beliefs of naive players are restricted to being *naive-consistent*: information obtained from actual equilibrium behavior still constrains their beliefs; but, when processing this information, naive players miss the potential correlation between other players' actions and their own payoff-relevant uncertainty.⁶

I apply the new framework to a class of *monotone selection* games that satisfy two properties: (i) a *monotone selection property* (MSP), which requires “lower” actions to result in a “worse” selection of outcomes; and (ii) *complementarity between beliefs and actions*, which in turn requires beliefs about a “worse” selection of outcomes to encourage “lower” actions. These properties are present in many standard settings with adverse selection, and additional applications are discussed in Section III. Under reasonable assumptions on information feedback, I find that naive-consistent beliefs can be supported in equilibrium and that the presence of players who are not aware of the

⁴Greene (1992) shows that an additional source of bias in the industry is to ignore that different rules for selecting applicants result in applicants that may have not only different default probabilities, but also different expected expenditures. He shows that a credit card vendor that ignores the above chooses lower than optimal selection rates. More generally, in marketing, historical data are used to *learn* the relationship between product offerings and customer profitability in an attempt to avoid adverse selection problems (e.g., Yong Cao and Thomas Gruca 2005).

⁵Previous versions of self-confirming (or conjectural) equilibrium in games of complete information appear in Pierpaolo Battigalli (1987) and Fudenberg and Levine (1993a). A self-confirming equilibrium is often interpreted as the outcome of a learning process in which players revise their beliefs using observations of previous play. Fudenberg and Levine (1993b) provide explicit learning-theoretic foundations.

⁶An analogy can be drawn by thinking of a player as an empirical researcher. For example, in the context of the credit-scoring models above, a naive player ignores the fact that offering other terms would attract a different sample of individuals. A sophisticated player, on the other hand, recognizes this possibility and therefore tries to account for selection in a number of different ways (e.g., by applying the techniques suggested by Greene 1998). Finally, a player with correct beliefs (as in a Nash equilibrium) may be able to fully account for selection by, say, conducting a series of randomized trials.

(adverse) selection problem actually exacerbates this problem. This result turns out to be true with respect to both players who have correct beliefs (as in a Nash equilibrium) and players who are sophisticated (e.g., know that price and quality are positively related in a lemons market, but may not know what the exact relationship is).

The result that, in some settings, markets are thinner in the presence of naive players implies that information asymmetries may be of even greater concern for the functioning of markets than previously thought.⁷ Moreover, the standard view is that the government cannot do much to alleviate the information asymmetry underlying the adverse selection problem and, thus, must rely instead on costly regulation in insurance and credit markets that can at best achieve the second best. However, this paper suggests that the government can enhance efficiency with policies that facilitate and encourage non-biased information acquisition by less-informed parties. For example, the Federal Credit Report Act in the U.S. does not allow credit-bureau information to be provided to third parties without an individual's consent, but makes an exception in connection to 'firm offers of credit.' This exception allows credit-card companies to use individuals' credit reports to conduct large-scale randomized trials in credit card solicitations and, therefore, understand how characteristics of a customer pool change with contract offerings (Lawrence M. Ausubel 1999).⁸ In Section III, I discuss other implications in the context of particular applications.⁹

The most closely related work is a paper by Eyster and Rabin (2005), who provide the first systematic equilibrium analysis of the bias that I study in this paper by introducing the notion of a *cursed equilibrium*. The main conceptual difference, discussed further in Section IV, is that Eyster and Rabin independently place restrictions on structural and strategic beliefs. In contrast, I place restrictions directly on information feedback and on information-processing capabilities that, in turn, endogenously imply *joint* restrictions on structural and strategic beliefs. As a result, whether a bias (or incorrect model of the world) may arise is determined endogenously in my framework and depends on assumptions about the feedback that players obtain regarding equilibrium outcomes. In addition, the set of equilibria when players are naive, are sophisticated, or have correct beliefs can be unambiguously compared for a general class of settings where cursed equilibria predicts either ambiguous results or results in the opposite direction (e.g., mitigation of the adverse selection problem).

A complementary literature postulates non-equilibrium models of behavior where players follow particular decision rules characterized by a finite depth of reasoning about players' beliefs about each other (Dale O. Stahl and Paul W. Wilson 1994; Rosemarie Nagel 1995; Miguel A. Costa-

⁷When players are aware of these informational asymmetries, some institutions may naturally arise to mitigate this problem, initiated either by the party having more information (Michael Spence 1973) or the party that is less informed (Michael Rothschild and Joseph Stiglitz 1976).

⁸On the other hand, learning is more difficult in insurance industries, where firms are often not allowed to offer different set of contracts to different customers unless they can actuarially justify the difference.

⁹Since information-processing biases may arise from the choice of organizational structure or, more generally, from how data are made available to decision-makers, firms should also consider the losses from ignoring certain types of correlation when making such choices.

Gomes, Vincent P. Crawford and Bruno Broseta 2001; Colin F. Camerer, Teck-Hua Ho and Juin-Kuan Chong 2004; Crawford and Costa-Gomes 2006). Crawford and Nagore Iriberry (2005) show that behavior that arises from some of these decision rules matches the experimental evidence of overbidding in both private and common value auctions. In contrast, I focus on settings where players repeatedly face similar strategic environments and learn from this experience based on the feedback they receive.

In Section I, I present an example that illustrates the approach taken by the previous literature, the alternative approach that I propose, the result that naive players exacerbate the adverse selection problem, and a dynamic justification for the proposed steady-state solution concept. In Section II, I introduce the definition of behavioral equilibrium for general games and show that equilibrium can be characterized as a fixed point of a generalized best response correspondence. I exploit this characterization in Section III, where I apply the framework to monotone selection games and characterize the set of equilibria when players are naive, are sophisticated, or have correct beliefs. I conclude in Section IV by discussing the conceptual contribution of the new framework, implications for the experimental literature, and some extensions. All proofs are in the appendix, and additional results are provided in an online appendix available at my website.

I An Illustrative Example

Consider a trading game with one-sided asymmetric information of the sort introduced by Akerlof (1970). The seller values the object at s , while the buyer values the object at $v = s + x$, where s is the realization of a random variable S that is uniformly distributed on the interval $[0, 1]$ and $x \in (0, 1]$ is a parameter that captures gains from trade. The seller knows her valuation, but the buyer has no private information about either s or v . The buyer and seller simultaneously make offers to buy at price p and to sell at price ask , respectively. If $ask > p$, there is no trade; the seller keeps the object, and the buyer obtains her reservation utility of zero. If $ask \leq p$, the object is traded and the buyer pays p and obtains utility $u(p, v) = v - p$. I restrict attention to equilibria where the seller plays his weakly dominant strategy, $ask = s$.

A Nash equilibrium, the selection problem, and cursed equilibrium

In a Nash equilibrium, the buyer offers a price p to maximize her expected profits

$$\pi^{NE}(p) = \Pr(S \leq p) \times [E(V \mid S \leq p) - p]. \quad (1)$$

Under the current assumptions, equation (1) becomes $\pi^{NE}(p) = p \times (x - \frac{1}{2}p)$ and the optimal (i.e., Nash equilibrium) price is $p^{NE} = x$.

In this example, the buyer faces a selection problem: the price that she offers *selects* the type of objects that are traded in equilibrium. In a Nash equilibrium, the buyer accounts for selection by

conditioning her belief about the value of the object on the information that the seller is willing to trade. The literature has modeled the behavior of a buyer who fails to account for selection by assuming that such a buyer does not realize that her valuation depends on the price she offers but, rather, believes that it is given by the unconditional expectation, $E(V)$. Following Eyster and Rabin (2005), call such a buyer a *cursed* buyer. The *perceived* profits of a cursed buyer are

$$\pi^{Cursed}(p) = \Pr(S \leq p) \times [E(V) - p], \quad (2)$$

and the optimal (i.e., cursed equilibrium) price is $p^{Cursed} = \frac{1}{2}(x + \frac{1}{2})$. Hence, relative to the Nash equilibrium, a cursed buyer over-prices for $x < 1/2$ and under-prices for $x > 1/2$.¹⁰

The following intuition for the previous under-/over-pricing result is a crucial step for understanding the logic behind the main result in this paper. A cursed buyer believes her valuation to be higher than what a non-cursed buyer believes since $E(V) \geq E(V | S \leq p)$ for all p (in general, this is true if V, S are affiliated). This *level effect* increases a cursed buyer's willingness to offer higher prices in order to obtain the object. However, a cursed buyer does not realize that increasing her offer would also increase the expected quality of objects she receives. This *slope effect* provides a cursed buyer with a weaker incentive to increase her bid relative to a buyer who has correct beliefs, as in a Nash equilibrium. Depending on whether the level or slope effect dominates, a cursed buyer can either over-price or under-price relative to a buyer in a Nash equilibrium.

B Behavioral Equilibrium: naive and sophisticated buyers

A cursed buyer makes her pricing decision under the belief that the expected value of the object is $E(V)$, but if she repeatedly follows her cursed strategy, her expected valuation of objects traded in equilibrium is lower, $E(V | S \leq p^{Cursed})$. If she repeatedly obtains feedback about the value of traded objects, then it may be reasonable to expect a cursed buyer to revise her beliefs about her expected valuation and, therefore, offer a different price. In addition, a question arises as to how a cursed buyer may know the true unconditional expected value of the object to begin with.

In contrast, the solution concept proposed in this paper requires beliefs to be consistent with the feedback obtained from actual equilibrium play. By emphasizing the role of information feedback, it captures the essence of the selection problem: a buyer's pricing decision affects the average quality of objects that are traded and, therefore, the sample that she uses to form beliefs about quality. I consider two types of buyers: a *naive* buyer is not aware of the selection effect, while a *sophisticated* buyer is aware of its potential existence. However, a sophisticated buyer may still have incorrect beliefs about the quality of objects that would be traded at prices that, for example, she has not tried out in the past. This is in contrast to a buyer with correct beliefs (as in a Nash equilibrium),

¹⁰This case corresponds to Eyster and Rabin's (2005) *fully* cursed equilibrium, and in the context of the trading game, it was originally discussed by Kagel and Levin (1986) and Holt and Sherman (1994). This last paper was the first to show that both under- and over-pricing are possible.

who knows the exact price-quality schedule in equilibrium.

A behavioral equilibrium depends on feedback that players obtain from repeatedly playing their equilibrium strategies. In the context of the trading game, I assume that both naive and sophisticated players observe their own payoffs and that the auctioneer reveals the seller's ask price at the end of each encounter. According to the formal definition of equilibrium in Section II, these assumptions, in turn, imply that in equilibrium the buyer has correct beliefs both about her expected payoffs from following the equilibrium action and about the probability of trade at any possible price.

Equilibrium with a naive buyer. The following function (and its appropriate generalization) plays an important role in developing the results in this paper:

$$\pi^N(p, p^*) = \Pr(S \leq p) \times [E(V \mid S \leq p^*) - p]. \quad (3)$$

Equation (3) represents a naive buyer's equilibrium belief about her expected payoff from deviating to a price p given that in (a hypothetical) equilibrium she repeatedly chooses p^* .¹¹ Beliefs about the probability of trade at each price are correct because of the assumption that in equilibrium the distribution of ask prices is known. Beliefs about the expected value of the object: (i) are determined by the price she chooses in equilibrium, p^* (rather than just being the unconditional expected value of the object, as in a cursed equilibrium); (ii) do not depend on the price p to which she considers deviating (this is the failure to account for the selection problem); and (iii) are correct for $p = p^*$.

As defined in Section II, in a behavioral equilibrium, a naive buyer chooses a price p^N such that given her perceived profit function $\pi^N(\cdot, p^N)$, it is indeed optimal to choose p^N . Hence, the set of equilibria with a naive buyer is given by the set of fixed points of $H^N(p) \equiv \arg \max_{p'} \pi^N(p', p)$. A straightforward calculation yields $H^N(p) = \frac{1}{2}x + \frac{1}{4}p$, so that when the buyer is naive, there is an (essentially) unique equilibrium price $p^N = \frac{2}{3}x$ that is lower than the Nash equilibrium price $p^{NE} = x$ for any parameter value $x \in (0, 1]$. *Hence, a naive buyer offers a lower price than a buyer with correct beliefs, leading to a lower probability of trade and to lower gains from trade.*¹²

The result that the adverse selection problem is exacerbated in the presence of naive players is not a coincidence of this particular example but, rather, a more general result. Its intuition can be grasped by observing Figure 1a, which compares the perceived profit function $\pi^N(\cdot, p^*)$ of a naive buyer choosing p^* to the correct profit function in equation (1), $\pi^{NE}(\cdot)$. The restriction that

¹¹Implicitly, the buyer does not know that $V = S + x$; whatever she learns about her valuation depends on the feedback she obtains. Equation (3) holds for $p^* > 0$. Otherwise, no objects are traded, no feedback is obtained, and beliefs can be quite arbitrary.

¹²Both with a naive and a sophisticated buyer, there is a no-trade equilibrium where the buyer believes objects are worth zero and, therefore, offers a zero price; trade then takes place with zero probability, and the buyer receives no feedback, so that her beliefs are actually consistent. Some simple refinements eliminate this equilibrium when sustained by incorrect beliefs, e.g., requiring beliefs to satisfy $\pi^N(p, 0) = \lim_{p^* \downarrow 0} \pi^N(p, p^*)$ for all p . Note also that no-trade is a Nash equilibrium outcome when sellers are not restricted to playing their weakly dominant strategy.

requires a buyer to have correct beliefs about her expected payoff from playing the equilibrium price eliminates the desire to over-price (i.e., the level effect) discussed earlier. Now, due to the selection effect, only the slope effect remains: a naive player thinks that profits from deviating to a lower price are higher than they actually are, while she believes that profits from deviating to a higher price are lower than they actually are. Since only the incentives to under-price are present, naive buyers always under-price relative to buyers with correct beliefs. Figure 1a also illustrates that the price p^* in that figure cannot constitute a naive equilibrium price since a naive player choosing p^* would rather deviate to a higher price. Figure 1b depicts H^N and the (essentially) unique naive equilibrium price p^N .

Equilibrium with a sophisticated buyer. In contrast, a sophisticated buyer who chooses price p^* in (a hypothetical) equilibrium perceives her profits from deviating to p to be

$$\pi^S(p, p^*) = \Pr(S \leq p) [\rho(p, p^*) - p],$$

where $\rho(p, p^*)$ denotes her expectation about the value of objects that would be traded at price p . I assume that a sophisticated buyer knows not only that there might be a potential selection problem (so that $\rho(p, p^*)$ is not necessarily constant in p), but also that this selection problem is monotone, i.e., the quality of objects traded in equilibrium is nondecreasing in the price that she offers. Hence, $\rho(\cdot, p^*)$ is a nondecreasing function for each p^* . An implication is that when choosing p^* , a sophisticated buyer knows that the expected value of the object conditional on trading at a price higher than p^* is *at least* $E(V | S \leq p^*)$. Hence, the perceived profit function of a naive buyer who offers p^* , $\pi^N(p, p^*)$, constitutes a lower bound for the perceived profit function of a sophisticated buyer when $p \geq p^*$. Similarly, it constitutes an upper bound when $p \leq p^*$.

Furthermore, the assumption that a buyer jointly observes ask prices and realized valuations when trade occurs implies that in an equilibrium where p^* is offered, the beliefs of a sophisticated player are correct for prices lower than p^* , i.e. $\rho(p, p^*) = E(V | S \leq p)$ for $p \leq p^*$. These restrictions place tight bounds on the behavior of a sophisticated buyer:

1. *The (essentially) unique equilibrium price with a naive buyer is a lower bound for the set of equilibrium prices with a sophisticated buyer.* To see this, suppose that $p < p^N$ were an equilibrium price with a sophisticated buyer. From Figure 1b, $H^N(p) > p$, meaning that a naive buyer who offers p believes that she can do better by offering a price higher than p rather than by choosing p . Since the beliefs of a naive buyer constitute a lower bound for the beliefs of a sophisticated buyer for prices higher than p , it follows that a sophisticated buyer must also believe that she can do better by choosing a higher price. Hence, a sophisticated buyer cannot choose p in equilibrium. Note how this result uses the fact that H is monotone, a property that arises here since low prices select low quality, which in turn induce the buyer to offer low prices.

2. *The unique Nash equilibrium price constitutes an upper bound for the set of equilibrium prices*

with a sophisticated buyer. Suppose, toward a contradiction, that $p > p^{NE}$ is an equilibrium price with a sophisticated buyer. Since a sophisticated buyer must have correct beliefs in equilibrium for those prices below p , it follows that she must know that she can do better by deviating to p^{NE} . In fact, for this particular example, the set of equilibrium prices with a sophisticated buyer (excluding the no-trade equilibrium) is given by the interval $[p^N, p^{NE}]$.¹³

Therefore, the adverse selection problem is exacerbated in the presence of naive players, not only relative to a buyer with correct beliefs, as in a Nash equilibrium, but also relative to a sophisticated buyer who knows that there is a monotone nondecreasing selection effect, but who may still have incorrect beliefs about the correct price-quality schedule. In the rest of the paper, I generalize this conclusion to a wider class of games, including games where more than one player acts strategically. In particular, the previous results generalize as long as S and V are affiliated, and u is nondecreasing in v and supermodular in (p, v) .

C Dynamics leading to naive equilibrium

Naive equilibrium can be justified as the outcome of a simple dynamic learning process, hence providing further understanding and motivation for the proposed steady-state solution concept.¹⁴ Suppose that the buyer knows the probability of trade at each price (presumably because she has collected information about past trades in this market), but does not know the expected value of objects in the market. Every period, the buyer chooses a price to maximize her perceived profits, which, in the context of the example in this section, are $\pi_t^N(p) = p \times (y_t - p)$, where y_t is her expectation of the value of the object at time t . Hence, at time t she offers price $p_t^* = \frac{1}{2}y_t$. The buyer starts with a prior about the expected value of the object, $v_1 > 0$, and updates this prior as she obtains additional information about the value of the object, which occurs when she trades. No updating takes place if there is no trade, and a buyer simply offers the same price in the next period. For simplicity, the updating rule is given by $y_t = \frac{1}{t} \sum_{i=1}^t v_i$, so that a buyer's estimate of the expected value of the object is the average of all observed valuations (including her initial prior).

Proposition 1 *Under the above dynamics, price converges in probability to a naive equilibrium price.*

While the dynamics above may capture how people learn in certain situations, a closer look reveals why learning is actually naive. The buyer does not realize that by choosing different prices, she is

¹³In settings where players do not know the reservation value of trading partners (e.g., ask prices are not revealed) but still know the distribution of reservation values (e.g., demand is known), all the results continue to hold except that Nash equilibrium is now not necessarily an upper bound to the behavior of sophisticated players.

¹⁴See Fudenberg and Levine (1993b, 1998) for dynamic models with steady states that correspond to Nash equilibrium and to self-confirming equilibrium (here, equilibrium with sophisticated players). In particular, a sophisticated player will not end up with correct beliefs unless she is sufficiently patient to fully experiment.

endogenously selecting the sample from which she will learn and, therefore, pools all the information together as if arising from a common data-generating process. A sophisticated buyer, on the other hand, understands that her actions may affect the data-generating process from which realizations are drawn, and she is likely to behave differently. For example, a sophisticated player (who is patient enough) may decide to fix a price for a certain number of periods until she approximately learns the expected value of the object conditional on that price, and only then will she decide to choose a different price.

In the remainder of the paper, I abstract from the dynamics leading to equilibrium and focus instead on a steady-state definition of equilibrium.

II Definition of Behavioral Equilibrium

There is a finite set \mathcal{N} of players who simultaneously choose actions. Each player $i \in \mathcal{N}$ chooses an action from a finite, nonempty, action set \mathcal{A}_i and obtains payoff $u_i(a, v) \in \mathbb{R}$, where $a = (a_i)_{i \in \mathcal{N}} \in \mathcal{A} \equiv \times_i \mathcal{A}_i$ is an action profile, $v \in \mathcal{V}$ is the realization of a vector of random variables V that capture players' uncertainty about payoffs, and \mathcal{V} is finite. Before choosing an action, each player receives a signal realization s_i from a random variable S_i with finite support \mathcal{S}_i . The signals and the parameters of the utility function are jointly drawn according to an objective probability distribution $F \in \Delta(\mathcal{S} \times \mathcal{V})$, where $\mathcal{S} = \times_i \mathcal{S}_i$. A (pure) strategy for player i , α_i , is a function from the set of signals to the set of actions, and a strategy profile is denoted by α . Let $\Omega = \mathcal{A} \times \mathcal{V}$ denote the set of (payoff-relevant) outcomes of the stage game.¹⁵

Let \mathbb{G}_i denote the set of all joint probability distributions over (A_{-i}, V) . Each player i with signal s_i has (in equilibrium) a belief $G_{s_i} \in \mathbb{G}_i$ over the actions of other players and payoff uncertainty. The expected utility from choosing action a_i that is perceived by player i with signal s_i who has belief G_{s_i} is

$$\pi_i(a_i | G_{s_i}) \equiv Eu_i(a_i, A_{-i}, V),$$

where the expectation is taken according to the probability distribution G_{s_i} .

Let $G_{s_i, \alpha_{-i}}^0$ denote the joint probability distribution over $(\alpha_{-i}(S_{-i}), V)$ conditional on $S_i = s_i$, i.e., $G_{s_i, \alpha_{-i}}^0$ represents the true probability distribution over (A_{-i}, V) being faced by player i with signal s_i when her opponents choose strategies α_{-i} . In equilibrium, player i 's belief when her signal is s_i , G_{s_i} , need not coincide with the true distribution $G_{s_i, \alpha_{-i}}^0$.¹⁶

Feedback is formally captured through the functions $\gamma_i^{\mathcal{A}} : \Omega \rightarrow 2^{\mathcal{A}_{-i}}$ and $\gamma_i^{\mathcal{V}} : \Omega \rightarrow 2^{\mathcal{V}}$. The

¹⁵The assumptions that the game is finite and that players choose pure strategies are made for ease of exposition and can be relaxed.

¹⁶By assuming that players only have uncertainty about (A_{-i}, V) (that will be resolved in equilibrium), one is implicitly assuming that players know their utility function, the set \mathcal{V} of parametric uncertainty, and their set \mathcal{A}_i of feasible actions. In applications, one may wish to impose additional restrictions by restricting the set \mathbb{G}_i of feasible beliefs: for example, in Sections I and III, sophisticated players are restricted to believing that the selection effect is nondecreasing.

interpretation is that when the outcome of the stage game is $\omega \in \Omega$, player i can only observe (ex post) that the action profile played by others belongs to a subset $\gamma_i^A(\omega)$ of the actions of other players. Implicitly, players are naturally restricted to observing their own chosen actions. Similarly, player i observes that the set of payoff parameters is contained in a subset $\gamma_i^V(\omega)$ of the set of payoff parameters. To simplify, I restrict the image of these functions to be the singleton sets and the entire set (e.g., $\gamma_i^V(\omega)$ is either the payoff parameter that corresponds to ω or the entire set of payoff parameters \mathcal{V}). Hence, players observe either the realized payoff parameter when the outcome is ω , or they receive no information about the realized payoff parameter—and similarly for feedback about others' actions. In the online appendix, I generalize the setting to allow for partial information feedback.¹⁷

From feedback about actions γ_i^A and feedback about payoff parameters γ_i^V (which are treated as primitives of the game), one can infer the feedback about payoffs, γ_i^U , that player i will observe.¹⁸ Let $\gamma_i = (\gamma_i^A, \gamma_i^V, \gamma_i^U)$ summarize the feedback observed by player i .

I distinguish between players who are sophisticated when it comes to learning from feedback and players who are naive in the sense that they fail to account for the possible correlation between A_{-i} and V . Consistency requires players to have correct beliefs about the joint distribution of observable outcomes, while naive-consistency requires beliefs to be independent and to be correct only about distributions over marginal outcomes.

Definition 1 (Consistency) (i) A belief $G_{s_i} \in \mathbb{G}_i$ of player i with signal s_i is γ_i -consistent for (a_i, α_{-i}) if the **joint** probability distribution of the random variables $\gamma_i(a_i, A_{-i}, V)$ is the same whether the joint distribution over (A_{-i}, V) is given by G_{s_i} or by the true distribution $G_{s_i, \alpha_{-i}}^0$.

(ii) A belief $G_{s_i} \in \mathbb{G}_i$ of player i with signal s_i is γ_i -**naive-consistent** for (a_i, α_{-i}) if it is the product of marginal distributions over A_{-i} and V , and if the **marginal** probability distributions of $\gamma_i^A(a_i, A_{-i}, V)$, $\gamma_i^V(a_i, A_{-i}, V)$, and $\gamma_i^U(a_i, A_{-i}, V)$ are all the same whether the joint distribution over (A_{-i}, V) is given by G_{s_i} or by the true distribution $G_{s_i, \alpha_{-i}}^0$.

The consistency restriction coincides with that imposed in a self-confirming equilibrium, so that an equilibrium where all players are sophisticated coincides with a self-confirming equilibrium. The novel aspect of this paper is to allow for imperfect learning from feedback by defining the notion of naive-consistency, which requires beliefs about observable actions, payoff parameters, and payoff realizations to be correct, but does not realize that additional information could be obtained by looking at these observations jointly.

In contrast to consistency, a belief that is *naive-consistent* need not necessarily exist. While it is always possible to come up with a belief G_{s_i} such that the marginal probability distributions of γ_i^A

¹⁷While the generalization is not needed for any of the applications considered in the paper, it further clarifies what naive entails in this setting.

¹⁸Formally, $\gamma_i^U(\omega)$ is the set of payoffs $u = u_i(a_i, a'_{-i}, v')$, where $a'_{-i} \in \gamma_i^A(\omega)$, $v' \in \gamma_i^V(\omega)$, and a_i is player i 's action in state ω .

and $\gamma_i^{\mathcal{Y}}$ are correct, there may be no belief that, in addition, implies correct beliefs over $\gamma_i^{\mathcal{M}}$.¹⁹ For example, in Section I, no belief is naive-consistent for a buyer who *always* receives feedback about the value of the object, irrespective of whether or not she trades. The reason is that, by always observing the value of the object, a naive player would expect to obtain its true unconditional expectation. But with a model of the world that does not allow for correlation between the sellers' willingness to trade and the value of the object, it is not possible to rationalize the fact that, in equilibrium, she will get the conditional expectation, which is lower. Hence, whether a bias (or model of the world) arises in equilibrium is endogenous and depends on information feedback. The online appendix presents a critical sufficient condition for the existence of naive-consistent beliefs.

A behavioral equilibrium requires that players choose strategies that are optimal given their beliefs, where beliefs are required to be consistent for sophisticated players and naive-consistent for naive players.

Definition 2 (Behavioral Equilibrium) *A profile of strategies α is a behavioral equilibrium with feedback $\{(\gamma_i^{\mathcal{A}}, \gamma_i^{\mathcal{Y}})\}_{i \in \mathcal{N}}$ if for every player $i \in \mathcal{N}$ and for every $s_i \in \mathcal{S}_i$ there exists a belief $G_{s_i} \in \mathbb{G}_i$ such that:*

- (i) $\alpha_i(s_i) \in \arg \max_{a_i} \pi_i(a_i | G_{s_i})$, and
- (ii) G_{s_i} is γ_i -consistent (for sophisticated players) or γ_i -naive-consistent (for naive players) for $(\alpha_i(s_i), \alpha_{-i})$.

When considering games where either all players are naive or all players are sophisticated, I refer to a behavioral equilibrium as a *naive* or *sophisticated equilibrium*, respectively.

A feature of behavioral equilibrium that I exploit in Section III is that it can be characterized as the set of fixed points of an appropriate generalization of a best response correspondence. Let $H_i(\alpha_i, \alpha_{-i})$ denote the set of player i 's strategies that, for each s_i , maximize expected utility given a belief G_{s_i} that is γ_i -consistent (for sophisticated players) or γ_i -naive-consistent (for naive players) for $(\alpha_i(s_i), \alpha_{-i})$. The *generalized best response correspondence* is the set of fixed points of the correspondence H_i ,

$$BR_i(\alpha_{-i}) = \{\alpha_i : \alpha_i \in H_i(\alpha_i, \alpha_{-i})\},$$

and I refer to it as sophisticated-BR, naive-BR, or Nash-BR, depending on whether the player is sophisticated, is naive, or has correct beliefs. Letting $BR = \{BR_i\}_{i \in \mathcal{N}}$, a behavioral equilibrium can then be characterized as a fixed point of the BR correspondence, and standard fixed point theorems and monotone comparative statics results can be applied to characterize equilibrium.

¹⁹In contrast, if a belief about the *joint* distribution of $(\gamma_i^{\mathcal{A}}, \gamma_i^{\mathcal{Y}})$ is correct (as required by consistency), then that belief must yield the correct distribution over $\gamma_i^{\mathcal{M}}$.

III Games with Monotone Selection

I apply the equilibrium framework in Section II to a class of games with *monotone selection*, which I define below in terms of non-primitives in order to emphasize the main economic intuition behind the results. The online appendix presents conditions for the existence and uniqueness of naive-consistent beliefs. When uniqueness holds, define the naive profit function $\pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) \equiv \pi_i(a_i | G_{s_i})$, where the belief G_{s_i} of player i with signal s_i is naive-consistent for (a_i^*, α_{-i}) . To simplify, I restrict attention to cases where there exists a unique naive profit function π_i^N and where the action space is $\mathcal{A}_i \subset \mathbb{R}$ for all i .²⁰

In addition, the following four properties characterize a *game with monotone selection*: (i) players receive feedback about their own payoffs;²¹ (ii) [strict] monotone selection property (MSP): π_i^N is [increasing] nondecreasing in a_i^* , for every i, a_i, s_i, α_{-i} ; (iii) sophisticated players know that MSP is satisfied; and (iv) action-belief complementarity: π_i^N is single-crossing in (a_i, a_i^*) for every i, s_i, α_{-i} . The online appendix presents conditions on the primitives of a class of models such that the previous properties are all satisfied.

Since players receive feedback about their own payoffs, they must have, in equilibrium, correct beliefs about their expected payoff from playing their equilibrium strategies. MSP requires that higher actions lead to a selection of outcomes that are “better” for a player. When MSP holds, a naive player believes that choosing an action that is higher than her equilibrium action would result in a lower payoff than it actually would.²² In contrast, sophisticated players are aware that selection might be an issue when learning. Property (iii) is a weak restriction that allows sophisticated players substantial freedom in their beliefs about the way selection operates in the market, as long as they believe that the selection effect is monotone. While this restriction may not be appropriate for all settings, to the extent that the modeler knows players face a standard monotone adverse selection environment, it seems reasonable to assume that sophisticated players also have this information. Finally, both MSP and action-belief complementarity are standard properties in some adverse selection settings. For example, in a lemons market, the lower the price the buyer offers, the worse the quality of objects traded (this is MSP), which in turn induces the buyer to choose even lower prices. Hence, action-belief complementarity holds since the lower the price offered, the lower the price that is optimal for a naive player whose beliefs are determined by the price that she offers. As a result, H^N is monotone, hence generalizing the fact that H^N lies above the 45° line for $p < p^N$ in Figure 1b.

The following result compares generalized best responses of players who are naive, are sophisticated, or have correct beliefs. As the intuition in Section I makes clear, only properties (i) and (ii)

²⁰The online appendix shows that the results extend to a multidimensional action space as long as selection is unidimensional.

²¹Formally, $\gamma_i^{\mathcal{A}_{-i}}$ and $\gamma_i^{\mathcal{V}}$ are such that $\gamma_i^{\mathcal{U}}(\omega)$ is a singleton set for every ω, i .

²²While MSP assumes nondecreasing selection, all the results extend (but go in the reverse direction) if selection is nonincreasing.

above are used to compare naive and Nash best responses. The standard product order is used, so that a strategy α_i is higher than α'_i if $\alpha_i(s_i) \geq \alpha'_i(s_i)$ for all $s_i \in \mathcal{S}_i$.

Theorem 1 (Comparing Generalized Best Responses) *Consider a game with monotone selection where the strategies of all other players are fixed. Then:*

1. *for every naive-BR there is a Nash-BR that is higher (and if MSP is strict, then every naive-BR is lower than any Nash-BR); and*
2. *every sophisticated-BR is higher than the lowest naive-BR.*

Theorem 1 generalizes the results in Section I—it can be applied to compare equilibria in settings with only one truly strategic player. In settings with more strategic players, it is well known that a specific comparison of best responses does not necessarily extend to a comparison of equilibria, except under additional assumptions. Following the literature on modern comparative statics (e.g., Paul Milgrom and John Roberts 1990), a sufficient condition for that extension is that the game has strategic complementarities. A game with monotone selection has *Nash strategic complementarities* if $\pi_i^{NE}(a_i; s_i, \alpha_{-i}) \equiv \pi_i(a_i \mid G_{s_i, \alpha_{-i}}^0)$ is single-crossing in (a_i, α_{-i}) for all i, s_i , while it has *naive strategic complementarities* if π_i^N is single-crossing in (a_i, α_{-i}) for all i, a_i^*, s_i .²³

Theorem 2 (Comparing Equilibria) *Consider a game with monotone selection and both Nash and naive strategic complementarities.*

1. *The sets of Nash, naive, and sophisticated equilibria are each nonempty; and the sets of Nash and naive equilibria each have lowest and highest elements.*
2. *The highest naive equilibrium is lower than the highest Nash equilibrium. (If MSP is strict, then, in addition, the lowest naive equilibrium is lower than the lowest Nash equilibrium).²⁴*
3. *Every sophisticated equilibrium is (weakly) higher than the lowest naive equilibrium.*

The results in Theorem 2 can also be extended to some settings without strategic complementarities, such as symmetric games where the strategic players have no private information (see the

²³The result below differs slightly from the standard proof in that different equilibrium concepts are being compared (rather than a parameterized model under the assumption that a Nash equilibrium is played), and that the sets of best responses are not ordered by the strong set order relation. Theorem 2 also extends trivially to the case where naive, sophisticated, and “Nash” players coexist.

²⁴Naive strategic complementarities is not needed to compare naive and Nash equilibria, but without it the existence of a naive equilibrium is not guaranteed.

online appendix) and certain auctions (see below). The statements in Theorem 2 hold for these settings when restricted to the set of symmetric equilibria.

To illustrate how the previous analysis can provide additional insights in a range of familiar settings, I conclude this section by discussing additional examples of games with monotone selection.

Monopoly/monopsony.— The setting in Section I is ubiquitous in applications of adverse selection to insurance, labor, financial, credit, and used goods markets, as suggested by substituting the names ‘buyer’ and ‘seller’ with insurer and insuree, firm and worker, market maker and informed trader, venture capitalist and entrepreneur, etc. In many of these examples, a firm faces a given supply or demand, and it seems reasonable to assume that, while a firm may know the willingness-to-trade in the population, it may not know the relationship between willingness-to-trade and the type of a potential customer. One point to note is that the selection effect may also be monotone decreasing. For example, in an insurance context, the higher the price of insurance, the *lower* the ‘quality’ of the customers that the firm obtains, which, in turn, provides incentives to increase prices even further (action-belief complementarity). The result is that insurance prices are now too high relative to the Nash equilibrium price, thus exacerbating the adverse selection problem.

Oligopoly with adverse selection.— When there is competition on the less-informed side of the market, whether or not MSP holds depends on the type of customers that are attracted from competitors by improving the terms of the contract. If customers are now heterogenous in both their preferences for firms and their quality types, then MSP holds whenever these two sources of private information are independent, and the results above are applicable for symmetric firms or for asymmetric ones if there are strategic complementarities.²⁵ The adverse selection problem is still exacerbated if firms are naive, but now the terms of the contract are less favorable for consumers than in a Nash equilibrium—hence, competition is softened and firms can actually be better off if everyone is naive. In particular, firms may either implicitly coordinate or favor regulation that restricts the type of contracts that can be experimented with, such as not using certain buyer characteristics in setting insurance prices.²⁶

Symmetric first-price auctions and k th-unit auctions.— By increasing her bid, a bidder would win objects that she would have otherwise not won, and this event occurs when the highest opponent bid is between her original bid and her new, increased bid. Under standard assumptions (affiliation and nondecreasing strategies), these objects are of higher expected quality than the objects she wins at her original bid. Hence, MSP holds and, in turn, this induces bidders to choose even higher bids, so that action-belief complementarity also holds. While symmetric first-price auctions are not

²⁵Sufficient conditions on demand so that the game has strategic complementarities can be obtained following Xavier Vives (1999).

²⁶Alternatively, Amy Finkelstein and James Poterba (2006) argue that political economy considerations may explain why firms forego the use of information that could improve risk classification in markets that include U.K. annuities, U.S. long-term care insurance, and French automobile insurance.

games with strategic complementarities, it is possible to use the comparative statics results shown by Ignacio Esponda (2006) to conclude that, in a symmetric equilibrium, bidding is less aggressive when all bidders are naive.

Underbidding is also obtained in a symmetric k th-unit auction, where k identical objects are sold and the highest k bidders get an object but pay the $k + 1$ highest bid (when $k = 1$, this is the second-price auction). If everyone plays the same strategy, then the expected value of objects won by a player of type s_i is $E[u(s_i, S_{-i}) | Y_k \leq s_i]$, where Y_k is the k th-highest of the opponents' signals. Since naive bidders do not know that the expected value of the object depends on their bid, they essentially believe to be in a private-values environment, where it is a dominant strategy to bid their valuation. Hence, naive bidders bid the above conditional expectation in a symmetric equilibrium. In contrast, symmetric Nash equilibrium bidding is higher and given by $E[u(s_i, S_{-i}) | Y_k = s_i]$ (Milgrom 1981).

One implication of the underbidding result is that an auctioneer has incentives to provide public information about the ex-post value of the object in order to increase revenues (or decrease costs, in the case of procurement). For example, publishing data on the ex post production of oil and gas in U.S. offshore oil and gas lease sales (Robert H. Porter 1995) is likely to increase the revenues of the federal government.

Teamwork.— Suppose that each member of a team (say, a salesperson) decides how much (observable) effort to exert, and that if total team effort is sufficiently high, then the team is successful and shares a prize that is uncertain but for which each player has some private information (such as a percentage of a sale, where the amount of the sale depends on the unknown type of the buyer). Then, a member must take into account that: (i) her effort affects the probability of success, and (ii) since others have private information about the prize value, changing her effort increases (under standard monotonicity assumptions) the expected value of a successful outcome. A naive team member does not account for the second effect, so that in the case where efforts are strategic complements, it is possible to show that a naive team results in lower effort exertion and lower performance. This result suggests that team effort may be increased if several teams compete for the same prize since, then, teams are more likely to observe the value of the prize when another team obtains it. While competing teams may belong to other firms, there are several examples where competition takes place between teams in a same firm.²⁷

Preemption game.— Two firms simultaneously choose whether to act today or to wait and act tomorrow. A standard tradeoff exists between a first-mover advantage and the gains from coordinating to move tomorrow, when the benefits of acting are less uncertain. A firm that decides to

²⁷Julian Birkinshaw and Mats Lingblad (2005) cite examples in several industries, and Michael R. Baye, Keith J. Crocker and Jiandong Ju (1996) present a model to explain competition through divisionalization, franchising, or divestiture. The specific informational benefit that arises from my analysis complements this and other alternative explanations.

act today but evaluates whether it should wait must account for two effects: (i) the risk of being preempted, and (ii) the fact that if it waits, it will get to act tomorrow only if the other firm has not preempted it, which (assuming that the other firm has private information) is more likely to happen when the benefits from acting are lower. A naive firm does not account for the second effect and, therefore, believes that the benefit from acting is lower when it decides to wait than when it decides to act today (MSP), making it more likely to wait (action-belief complementarity). Hence, this is a game with monotone selection, and under strategic complementarities (i.e., if one firm waits, waiting becomes more attractive for the other), there is more waiting in equilibrium with naive firms. Hence, naive firms are able to coordinate better, but welfare-enhancing actions may also be delayed. For example, naivete constitutes an additional explanation to the observed delay in technology adoption (Heidrun C. Hoppe 2002).

Performance pay.— Edward P. Lazear (2000) provides evidence that performance pay increases productivity due to a well-known incentive effect and to an often-neglected selection (or sorting) effect. The latter materializes over time and results from higher turnover for less productive workers and higher productivity of new hires. A naive firm knows how its current workers respond to incentives, but it may ignore the fact that, in the long-term, performance pay further increases productivity through the selection effect. Hence, a naive firm chooses a level of performance pay that is lower than optimal. This finding suggests that part of the relatively low levels of observed performance pay are likely to be suboptimal. Indeed, the firm studied by Lazear (2000) had selected a suboptimal compensation scheme based on hourly wages before new management took over, switched to piece-rate pay, and increased profits.

IV Discussion

In this paper, I provide a framework for studying equilibrium behavior in the presence of players who fail to account for the informational content of other players' actions. I introduce the concept of a behavioral equilibrium and apply it to obtain new insights into the nature of the adverse selection problem. Contrary to what may be expected without an appropriate equilibrium framework, players who fail to account for selection actually exacerbate the adverse selection problem. In general, this result highlights the important role that specific information policies may have in mitigating adverse selection.

The distinguishing feature of the new framework is that both structural and strategic beliefs are endogenously determined in equilibrium. In contrast, the standard literature (at least since John C. Harsanyi 1967-8) makes a sharp distinction between uncertainty about fundamentals and uncertainty about the strategies of other players. While the latter is determined endogenously in equilibrium, the former is exogenous. Eyster and Rabin's (2005) cursed equilibrium can be viewed as an attempt to introduce selection bias while maintaining this standard distinction about beliefs.

In a fully cursed equilibrium, players' belief about fundamentals (such as the common valuation of an object in an auction) is exogenously fixed to be correct, while their belief about the distribution of other players' actions endogenously coincides with the equilibrium distribution. However, players do not understand the relationship between other players' private information and actions. In particular, players incorrectly believe that each type profile of the other players plays the same profile of mixed actions—which coincides with the true average distribution of actions.²⁸

There is a practical problem with this sharp distinction. The underlying assumptions about information feedback and information-processing biases that motivate restrictions on players' beliefs about the actions of other players are also likely to endogenously motivate restrictions on players' beliefs about fundamentals. By making restrictions on beliefs about fundamentals independent of restrictions on beliefs about actions, it is not clear what the underlying assumptions on feedback nor what the consequent restrictions on equilibrium beliefs are. As illustrated in Section I, a player in a cursed equilibrium may have incorrect beliefs about the expected payoff she receives *from playing her equilibrium strategy*—hence implicitly implying that players obtain no feedback about their past payoffs.²⁹

In contrast, in a behavioral equilibrium, restrictions are made directly on information feedback and on the information-processing capabilities of players, and such restrictions endogenously imply *joint* restrictions on equilibrium beliefs about both structural and strategic uncertainty. Besides obtaining results that go against the received wisdom (but that have a very clear intuition once the details are spelled out), one important implication of the proposed framework is that whether a particular bias may arise is determined endogenously in equilibrium. In addition, the framework is set up so that it should be easy to modify to study other information-processing biases.

The separation between the assumption of a bias and whether that bias may actually persist in equilibrium provides further insight into the relationship with experimental results. As argued in Section II, in the case where feedback about the value of the object is obtained irrespective of trading or not, naive-consistent beliefs do not exist and, therefore, a naive model of the world cannot persist in an equilibrium setting. There are some situations where this alternative assumption may make sense, such as a common value auction with resale, where everyone observes the resale price. An equivalent situation arises in experimental settings where subjects are told the distribution from which valuations are drawn before playing the game. The framework in this paper limits itself to predicting that a naive equilibrium will not exist in such a setting. Of course, players will still play

²⁸The same arguments that motivate the alternative approach that I follow in this paper apply to the more general frameworks of Philippe Jehiel and Frederic Koessler (2005), where players make mistakes when forecasting the type-contingent strategies of their opponents by bundling types into analogy classes, and Topi Miettinen (2007), where misperceptions of multiple correlations are allowed. In particular, the approach can be applied to these other papers to, for example, endogenize the analogy partitions that can arise in equilibrium.

²⁹See Dekel, Fudenberg, and Levine (2004) for this critique in the context of a model with no biases, and the survey by Fudenberg (2006) for a discussion of the problems with cursed equilibrium and the need for alternative equilibrium concepts for behavioral economics. Esponda (2006) provides an alternative to Bayesian Nash equilibrium and extends the present paper by allowing people to learn not only from feedback but also from introspection.

in some way, and at least two possibilities come to mind. First, players may rationalize the fact that they expect to get the unconditional expected value but always get less by thinking that they are being unlucky, or even by failing to consider payoff-feedback information altogether. In this case, players will behave as predicted by Kagel and Levin (1986), Holt and Sherman (1994), and Eyster and Rabin (2005), and it is no coincidence that the experimental evidence mostly agrees with this prediction. Second, players may eventually rationalize that the seemingly contradictory information is due to their failure to account for selection and, therefore, update their model of the world and stop being naive. A behavioral equilibrium is a steady-state concept that does not postulate how the model of the world will be revised, but rather limits itself to answering whether or not such a model of the world can persist in equilibrium. An interesting, nontrivial extension of the framework that is left for future work is to understand how players update their model of the world when it cannot rationalize what they observe.³⁰

The results in this paper also suggest that the standard intuition and experimental results in adverse selection settings are likely to be influenced by the assumption that players know a priori the true distribution over the fundamentals. This assumption may not be the most reasonable in many settings, and, in addition, it obscures the essence of the selection problem, where players' choices endogenously select the sample from which players learn about their environment. It would be interesting to extend the experimental literature by relaxing the assumption that players know the distribution over the fundamentals but must, instead, learn it.³¹

Finally, the paper provides a dynamic justification for the steady-state solution concept when only one player is engaged in learning. It would be interesting to extend these dynamics to the nontrivial case where several players are simultaneously learning in the presence of a selection problem.

Appendix

First, I present some standard terminology and results from the literature on monotone comparative statics. Then, I establish some preliminary results and use them to prove the statements in the text.

A1. Fixed points and monotone comparative statics³²

Throughout, let $\mathcal{X} \subset \mathbb{R}^K$ denote a nonempty, finite set, and let \mathcal{T} be a nonempty, partially ordered set. An element of $x \in \mathcal{X}$ is the *highest* element of \mathcal{X} if $x \geq y$ for every $y \in \mathcal{X}$; it is the

³⁰Marco Casari, John C. Ham, and Kagel's (2007) evidence linking cognitive abilities to bidders' accounting of adverse selection highlights the importance of heterogeneity for understanding people's updating of their model of the world.

³¹This is not intended as a critique of the experimental literature, which provides subjects with information about distribution of random variables in order to test for Bayesian Nash equilibrium, but rather as an assertion about the need to extend the methodology of treating both structural and strategic uncertainty as endogenous to the experimental literature, as well.

³²See, e.g., Donald Topkis (1998) and Vives (1999,2005).

lowest element if $x \leq y$ for all $y \in \mathcal{X}$. A function $g : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}$ is *single-crossing* in (x, t) if for $x > x'$ and $t > t'$, $g(x, t) \geq g(x', t')$ implies $g(x, t) \geq g(x', t)$ and $g(x, t) > g(x', t')$ implies $g(x, t) > g(x', t)$. The following results are used in the proofs.

FP1. (Alfred Tarski 1955; Milgrom and Roberts 1990) Suppose $f : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{X}$ is nondecreasing for each $t \in \mathcal{T}$. Then for each t , the set of fixed points of f is nonempty and has a lowest element $\underline{x}(t) = \inf\{x \in \mathcal{X} : f(x, t) \leq x\}$ and a highest element $\bar{x}(t) = \sup\{x \in \mathcal{X} : f(x, t) \geq x\}$. If in addition f is nondecreasing in t for all $x \in \mathcal{X}$, then $\underline{x}(\cdot)$ and $\bar{x}(\cdot)$ are nondecreasing.

MCS1. (Milgrom and Chris Shannon 1994) For each $t \in \mathcal{T}$, $h(t) \equiv \arg \max_{x \in \mathcal{X} \subset \mathbb{R}} f(x, t)$ is nonempty and has a lowest element $\underline{h}(t)$ and a highest element $\bar{h}(t)$. If f is single-crossing in (x, t) , then $\underline{h}(\cdot)$ and $\bar{h}(\cdot)$ are nondecreasing.

A2. Proof of main results

The proofs are in order of appearance in the text, except for Proposition 1, which is proved last. I start with two preliminary results. Let $\Pi_i^S(a_i^*, s_i, \alpha_{-i})$ denote the set of expected utility functions $\pi_i(\cdot | G_{s_i})$, where G_{s_i} is γ_i -consistent for (a_i^*, α_{-i}) . Define Π_i^N similarly for naive-consistent beliefs. Let BR_i^N , BR_i^S , and BR_i^{NE} denote the set of naive, sophisticated, and Nash (i.e. correct) best responses.

PR1. Fix a_i^*, s_i, α_{-i} and let $\pi_i^S \in \Pi_i^S(a_i^*, s_i, \alpha_{-i})$ and $\pi_i^N \in \Pi_i^N(a_i^*, s_i, \alpha_{-i})$. If player i receives feedback about her own payoffs, then

$$\pi_i^N(a_i^*; a_i^*, s_i, \alpha_{-i}) = \pi_i^S(a_i^*; a_i^*, s_i, \alpha_{-i}) = \pi_i^{NE}(a_i^*; s_i, \alpha_{-i}).$$

proof. Since payoffs are observed, $\gamma_i^{\mathcal{U}}(\omega)$ is a singleton for all ω . Let $G_{s_i}^N$ denote the naive-consistent belief corresponding to $\pi_i^N(\cdot; a_i^*, s_i, \alpha_{-i})$. Then:

$$\begin{aligned} \pi_i^{NE}(a_i^*; s_i, \alpha_{-i}) &= \sum_{(a_{-i}, v)} u_i(a_i^*, a_{-i}, v) \times G_{s_i, \alpha_{-i}}^0(a_{-i}, v) \\ &= \sum_{u \in \mathbb{R}} u \times G_{s_i, \alpha_{-i}}^0 \{(a_{-i}, v) : u_i(a_i^*, a_{-i}, v) = u\} \\ &= \sum_{u \in \mathbb{R}} u \times G_{s_i, \alpha_{-i}}^0 \{(a_{-i}, v) : \gamma_i^{\mathcal{U}}(a_i^*, a_{-i}, v) = \{u\}\} \\ &= \sum_{u \in \mathbb{R}} u \times G_{s_i}^N \{(a_{-i}, v) : \gamma_i^{\mathcal{U}}(a_i^*, a_{-i}, v) = \{u\}\} \\ &= \sum_{(a_{-i}, v)} u_i(a_i^*, a_{-i}, v) \times G_{s_i}^N(a_{-i}, v) \\ &= \pi_i^N(a_i^*; a_i^*, s_i, \alpha_{-i}) \end{aligned}$$

where the first, second and last equalities follow from the definitions of the corresponding expected utility functions, the third and fifth equalities follow since $\gamma_i^{\mathcal{U}}(\omega)$ is a singleton, and the fourth

equality from the definition of naive-consistency. A similar proof shows that $\pi_i^S(a_i^*; a_i^*, s_i, \alpha_{-i}) = \pi_i^{NE}(a_i^*; s_i, \alpha_{-i})$. ■

To formalize the idea that sophisticated players know MSP, extend the definition of naive-consistency by saying that G_{s_i} is γ_i -naive-consistent for $(a_i^*, G_{s_i}^S)$ if $G_{s_i, \alpha_{-i}}^0$ in the definition is replaced by $G_{s_i}^S$. Let $\Pi_i^{S-N}(a_i^*, G_{s_i}^S)$ denote the set of expected utility functions that a sophisticated player with belief $G_{s_i}^S$ thinks that a naive player may entertain, i.e. the set of $\pi_i(\cdot | G_{s_i})$, where G_{s_i} is γ_i -consistent for $(a_i^*, G_{s_i}^S)$. The equilibrium belief of a sophisticated player, $G_{s_i}^S$, must satisfy: $\Pi_i^{S-N}(\cdot, G_{s_i}^S)$ is strongly nondecreasing, i.e. for any $a_i^{*'} \leq a_i^*$, $\pi' \in \Pi_i^{S-N}(a_i^{*'}, G_{s_i}^S)$, and $\pi \in \Pi_i^{S-N}(a_i^*, G_{s_i}^S)$, it follows that $\pi'(a_i) \leq \pi(a_i)$ for all a_i .

PR2. Fix a_i^*, s_i, α_{-i} and let $\pi_i^S \in \Pi_i^S(a_i^*, s_i, \alpha_{-i})$. If player i receives feedback about her own payoffs, sophisticated players know MSP, and the game has a unique naive profit function, π_i^N , then $\pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) \leq \pi_i^S(a_i; a_i^*, s_i, \alpha_{-i})$ for $a_i \geq a_i^*$ and $\pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) \geq \pi_i^S(a_i; a_i^*, s_i, \alpha_{-i})$ for $a_i \leq a_i^*$.

proof. Let $G_{s_i}^S$ be the belief corresponding to $\pi_i^S(\cdot; a_i^*, s_i, \alpha_{-i})$. Note that $\Pi_i^N(a_i^*, G_{s_i}^S) = \Pi_i^N(a_i^*, G_{s_i, \alpha_{-i}}^0)$ since by definition naive-consistent beliefs depend only on the probability distribution over marginal feedback, and since $G_{s_i}^S$ is γ_i -consistent for (a_i^*, α_{-i}) then such distribution is the same for $G_{s_i}^S$ and $G_{s_i, \alpha_{-i}}^0$. Then $\Pi_i^{S-N}(a_i^*, G_{s_i}^S) = \{\pi_i^N(\cdot; a_i^*, s_i, \alpha_{-i})\}$ by the assumption that there is a unique naive profit function. Now let $a_i \geq a_i^*$ and consider $\bar{\pi}_i^N(\cdot; a_i, s_i, \alpha_{-i}) \in \Pi_i^{S-N}(a_i, G_{s_i}^S)$. Then

$$\begin{aligned} \pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) &\leq \bar{\pi}_i^N(a_i; a_i, s_i, \alpha_{-i}) \\ &= \pi_i^S(a_i; a_i^*, s_i, \alpha_{-i}), \end{aligned}$$

where the inequality follows since Π_i^{S-N} is strongly nondecreasing and $a_i \geq a_i^*$, and the equality follows from PR1 (i.e. $\bar{\pi}_i^N(a_i; a_i, s_i, \alpha_{-i})$ are the “correct” beliefs from playing a_i when the “true” distribution over (a_{-i}, v) is given by $G_{s_i}^S$). A similar proof establishes that $\pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) \geq \pi_i^S(a_i; a_i^*, s_i, \alpha_{-i})$ for $a_i \leq a_i^*$. ■

Proof of Theorem 1. Part 1. Let $\alpha_i^N \in BR_i^N(\alpha_{-i})$ and $\alpha_i^{NE} \in BR_i^{NE}(\alpha_{-i})$. Suppose that $\alpha_i^N(s_i) > \alpha_i^{NE}(s_i)$ for some $s_i \in \mathcal{S}_i$. Then:

$$\begin{aligned} \pi_i^{NE}(\alpha_i^N(s_i); s_i, \alpha_{-i}) &= \pi_i^N(\alpha_i^N(s_i); \alpha_i^N(s_i), s_i, \alpha_{-i}) \\ &\geq \pi_i^N(\alpha_i^{NE}(s_i); \alpha_i^N(s_i), s_i, \alpha_{-i}) \\ &\geq \pi_i^N(\alpha_i^{NE}(s_i); \alpha_i^{NE}(s_i), s_i, \alpha_{-i}) \\ &= \pi_i^{NE}(\alpha_i^{NE}(s_i); s_i, \alpha_{-i}), \end{aligned}$$

where the two equalities follow from the fact that payoffs are observed (PR1), the first inequality follows from the definition of a naive-BR, and the second inequality follows from MSP and $\alpha_i^N(s_i) \geq \alpha_i^{NE}(s_i)$.

$\alpha_i^{NE}(s_i)$. Therefore, $\alpha_i^N(s_i)$ is also a Nash-BR for i, s_i , so that $\max\{\alpha_i^N, \alpha_i^{NE}\} \in BR_i^{NE}$. Now suppose strict MSP holds: the second inequality is then strict, which contradicts the fact that $\alpha_i^{NE}(s_i)$ is a Nash-BR for i, s_i . Therefore, $\alpha_i^N \leq \alpha_i^{NE}$.

Part 2. Let

$$h_i^N(a_i^*, s_i, \alpha_{-i}) \equiv \arg \max_{a_i} \pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}),$$

so that

$$BR_i^N(\alpha_{-i}) = \{\alpha_i : \alpha_i(s_i) \in h_i^N(\alpha_i(s_i), s_i, \alpha_{-i}) \text{ for all } s_i \in \mathcal{S}_i\}.$$

Since π_i^N is single-crossing in (a_i, a_i^*) , then, by MCS1, $h_i^N(a_i^*, s_i, \alpha_{-i})$ has a lowest element $\underline{h}_i^N(a_i^*, s_i, \alpha_{-i})$ that is nondecreasing in a_i^* . Since $\underline{h}_i^N(\cdot, s_i, \alpha_{-i}) : \mathcal{A}_i \rightarrow \mathcal{A}_i$, by FP1 there is a lowest fixed point of $\underline{h}_i^N(\cdot, s_i, \alpha_{-i})$, $\underline{\alpha}_i(s_i, \alpha_{-i}) = \inf\{a_i \in \mathcal{A}_i : \underline{h}_i^N(a_i, s_i, \alpha_{-i}) \leq a_i\}$. Then $\underline{\alpha}_i$ is the lowest naive-BR.

Now fix α_i such that $\alpha_i(s_i^*) < \underline{\alpha}_i(s_i^*, \alpha_{-i})$ for some $s_i^* \in S_i$. Then $\underline{h}_i^N(\alpha_i(s_i^*), s_i^*, \alpha_{-i}) > \alpha_i(s_i^*)$, and, letting $a_i' = \underline{h}_i^N(\alpha_i(s_i^*), s_i^*, \alpha_{-i})$,

$$\begin{aligned} \pi_i^S(a_i'; \alpha_i(s_i^*), s_i^*, \alpha_{-i}) &\geq \pi_i^N(a_i'; \alpha_i(s_i^*), s_i^*, \alpha_{-i}) \\ &> \pi_i^N(\alpha_i(s_i^*); \alpha_i(s_i^*), s_i^*, \alpha_{-i}) \\ &= \pi_i^S(\alpha_i(s_i^*); \alpha_i(s_i^*), s_i^*, \alpha_{-i}), \end{aligned}$$

where the first inequality follows since sophisticated players know MSP and $a_i' > \alpha_i(s_i^*)$ (PR2), the strict inequality follows by definition of a_i' and since $\alpha_i(s_i^*)$ is not a fixed point of $\underline{h}_i^N(\cdot, s_i, \alpha_{-i})$, and the equality follows since payoffs are observed (PR1). Therefore, $\alpha_i \notin BR_i^S(\alpha_{-i})$, implying the result. ■

Proof of Theorem 2.

(a) Following the proof of part 2 of theorem 1, MCS1 and FP1 imply that there exist lowest and highest Nash and naive best responses, denoted by $\underline{BR}_i^{NE}(\alpha_{-i})$, $\overline{BR}_i^{NE}(\alpha_{-i})$ and $\underline{BR}_i^N(\alpha_{-i})$, $\overline{BR}_i^N(\alpha_{-i})$, respectively. In addition, since π_i^{NE} is single-crossing in (a_i, α_{-i}) and π_i^N is single-crossing in (a_i, α_{-i}) , each best response is nondecreasing in α_{-i} . For $n \in \{N, NE\}$, let $\underline{BR}^m(\alpha) \equiv \{\underline{BR}_i^m(\alpha_{-i})\}_{i \in N}$ be the lowest best response map. Letting $\mathcal{X} = \times_i \mathcal{A}_i^{S_i}$ denote the finite set of strategy profiles, note that $\underline{BR}^m : \mathcal{X} \rightarrow \mathcal{X}$ is nondecreasing in α , so that FP1 implies that there is a lowest fixed point of \underline{BR}^m , given by $\underline{\alpha}^m = \inf\{\alpha \in \mathcal{X} : \underline{BR}^m(\alpha) \leq \alpha\}$. For any α that is a fixed point of \underline{BR}^m , $\alpha \geq \underline{BR}^m(\alpha)$. Therefore, $\underline{\alpha}^m$ is also the lowest fixed point of \underline{BR}^m , so that $\underline{\alpha}^m$ is the lowest Nash (naive) equilibrium for $m = NE$ ($m = N$). A similar proof establishes the existence of a highest equilibrium. A sophisticated equilibrium exists since a Nash equilibrium is always a sophisticated equilibrium (this is because correct beliefs are always γ_i -consistent and because the restriction that sophisticated players know MSP is correct when MSP holds).

(b) By part 1 of Theorem 1, $\overline{BR}^N(\alpha) \leq \overline{BR}^{NE}(\alpha)$ for all $\alpha \in \mathcal{X}$. Let $\mathcal{T} = \{0, 1\}$ and define $f : \mathcal{X} \times \mathcal{T} \rightarrow \mathcal{X}$ such that $f(\cdot, 0) = \overline{BR}^N(\alpha)$ and $f(\cdot, 1) = \overline{BR}^{NE}(\alpha)$. Then f is nondecreasing in $t \in \mathcal{T}$, and from FP1, the highest fixed point of \overline{BR}^N (i.e. the highest naive equilibrium) is (weakly) lower than the highest fixed point of \overline{BR}^{NE} (i.e. the highest Nash equilibrium). With the additional assumption that MSP is strict, part 1 of Theorem 1 implies that $\overline{BR}^N(\alpha) \leq \underline{BR}^{NE}(\alpha)$ for all $\alpha \in \mathcal{X}$, so that a similar application of FP1 yields that the lowest fixed point of \underline{BR}^{NE} (i.e. the lowest Nash equilibrium) is (weakly) higher than the lowest fixed point of \overline{BR}^N , which is itself (weakly) higher than the lowest naive equilibrium.³³

(c) Let α^S be a sophisticated equilibrium, i.e. $\alpha^S \in BR^S(\alpha^S)$. I show that there exists a naive equilibrium α^N such that $\alpha^N \leq \alpha^S$. Consider $\mathcal{X}' = \{\alpha \in \mathcal{X} : \alpha \leq \alpha^S\}$. For any $\alpha \in \mathcal{X}'$, $\underline{BR}^N(\alpha) \leq \underline{BR}^N(\alpha^S) \leq \alpha^S$, where the first inequality follows from \underline{BR}^N being nondecreasing and the second inequality follows from the ordering of best responses established in part 2 of Theorem 1 (i.e. $\underline{BR}^N(\alpha) \leq \alpha'$ for any $\alpha' \in BR^S(\alpha)$). Hence, $\underline{BR}^N(\mathcal{X}') \subset \mathcal{X}'$ and it follows from FP1 that there exists a naive equilibrium $\alpha^N \in \mathcal{X}'$.

Proof of Proposition 1. First note that given that $v_t \leq 1$ almost surely for any t and the fact that y_{t+k} can be written as $\frac{t}{t+k}y_t + \frac{k}{t+k} \left(\frac{1}{k} \sum_{i=1}^k v_{t+i} \right)$ it can be shown that

$$|y_{t+k+j} - y_{t+k}| \leq \frac{2j}{t+k+j} \quad (4)$$

almost surely for any t , k and j . Now take t and k such that both go to infinity but t goes faster than k . Notice that $\zeta_{t+i} \equiv v_{t+i} - (x + \frac{1}{4}y_{t+i-1})$ is a bounded martingale difference, as $E_{t+i-1}[\zeta_{t+i}] = 0$. Hence by invoking Azuma-Hoeffding inequality, standard probability theory arguments and some algebra it follows that $k^{-1} \sum_{i=1}^k \zeta_{t+i}$ converges to zero almost surely. Thus this fact and equation (4) yield that $\lim_{k,t \rightarrow \infty} \left| \frac{1}{k} \sum_{i=1}^k v_{t+i} - (x + \frac{1}{4}y_t) \right| = 0$ with probability one. Finally, by noting that $y_{t+k} = y_t + \frac{k}{t+k} \left(\frac{1}{k} \sum_{i=1}^k v_{t+i} - y_t \right)$, using the two aforementioned results over convergent subsequences (which we still denote as y_t) it follows that $|y_t - \frac{4}{3}x| \rightarrow 0$ and thus the desired result follows by invoking the fact that a sequence converges in probability if and only if for every sub-sequence there exists a further sub-sub-sequence that converges almost surely. ■

References

- [1] Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84: 488-500.

³³When MSP is strict, the proof of part (b) can be established along the lines of the proof of part (c) and does not require that the naive game have strategic complementarities.

- [2] Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler. 2005. "Fact-Free Learning." *American Economic Review*, 95(5): 1355-1368.
- [3] Athey, Susan. 1998. "Comparative Statics Under Uncertainty: Single-Crossing Properties and Log-Supermodularity." MIT working paper 96-22.
- [4] Athey, Susan. 2001. "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information." *Econometrica*, 69: 861-890.
- [5] Athey, Susan. 2002. "Monotone Comparative Statics Under Uncertainty." *Quarterly Journal of Economics*, 117: 187-223.
- [6] Ausubel, Lawrence M. 1999. "Adverse Selection in the Credit Card Market." Working Paper, Department of Economics, University of Maryland.
- [7] Battigalli, Pierpaolo. 1987. "Comportamento Razionale Ed Equilibrio Nei Giochi e Nelle Situazioni Sociali." unpublished undergraduate dissertation, Bocconi University, Milano.
- [8] Baye, Michael R., Keith J. Crocker, and Jiandong Ju. 1996. "Divisionalization, Franchising, and Divestiture Incentives in Oligopoly." *American Economic Review*, 86(1): 223-236.
- [9] Birkinshaw, Julian and Mats Lingblad. 2005. "Intrafirm Competition and Charter Evolution in the Multibusiness Firm." *Organization Science*, 16 (6): 674-686.
- [10] Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics*, 119: 861-898.
- [11] Cao, Yong and Thomas Gruca. 2005. "Reducing Adverse Selection Through Customer Relationship Management." *Journal of Marketing*, 69: 219-229.
- [12] Capen, Edward C., Robert V. Clapp, and William M. Campbell. 1971. "Competitive Bidding in High Risk Situations." *Journal of Petroleum Technology*, 23: 641-53.
- [13] Casari, Marco, John C. Ham, and John H. Kagel. 2007. "Selection Bias, Demographic Effects and Ability Effects in Common Value Auction Experiments." forthcoming, *American Economic Review*.
- [14] Cassing, James and Richard W. Douglas. 1980. "Implications of the Auction Mechanism in Baseball's Free Agent Draft." *Southern Economic Journal*, 47: 110-21.
- [15] Charness, Gary and Dan Levin. 2006. "The Origin of the Winner's Curse: A Laboratory Study." working paper.
- [16] Costa-Gomes, Miguel A., Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica*, 69: 1193-1235.

- [17] Crawford, Vincent P. and Miguel A. Costa-Gomes. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review*, 96: 1737-1768.
- [18] Crawford, Vincent P. and Nagore Iriberri. 2005. "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" working paper.
- [19] Dekel, Eddie, Drew Fudenberg, and David K. Levine. 2004. "Learning to Play Bayesian Games." *Games and Economic Behavior*, 46: 282-303.
- [20] Esponda, Ignacio. 2006. "Information Feedback in First Price Auctions." working paper.
- [21] Esponda, Ignacio. 2006. "Robust Equilibrium Analysis in Games with Uncertainty." working paper.
- [22] Eyster, Erik and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica*, 73: 1623-1672.
- [23] Finkelstein, Amy and James Poterba. 2006. "Testing for Adverse Selection with 'Unused Observables'." NBER Working Paper 12112.
- [24] Fudenberg, Drew. 2006. "Advancing Beyond Advances in Behavioral Economics." *Journal of Economic Literature*, 44: 694-711.
- [25] Fudenberg, Drew and David K. Levine. 1993a. "Self-confirming Equilibrium." *Econometrica*, 61: 523-545.
- [26] Fudenberg, Drew and David K. Levine. 1993b. "Steady-State Learning and Nash Equilibrium." *Econometrica*, 61: 547-573.
- [27] Fudenberg, Drew and David K. Levine. 1998. *The Theory of Learning in Games*. The MIT Press, 1998.
- [28] Gilovich, Thomas, Dale Griffin, and Daniel Kahneman, ed. 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press.
- [29] Greene, William. 1992. "A Statistical Model for Credit Scoring." Working paper No. EC-92-29, Department of Economics, Stern School of Business.
- [30] Greene, William. 1998. "Sample Selection in Credit-Scoring Models." *Japan and the World Economy*, 10: 299-316.
- [31] Harsanyi, John C. 1967-8. "Games with Incomplete Information Played by 'Bayesian' players." Parts I, II, and III, *Management Science*, 14: 159-182, 320-334, and 486-502.

- [32] Holt, Charles A. and Roger Sherman. 1994. "The Loser's Curse." *American Economic Review*, 84: 642-652.
- [33] Hoppe, Heidrun C. 2002. "The Timing of New Technology Adoption: Theoretical Models and Empirical Evidence." *The Manchester School*, Vol. 70 No. 1 January.
- [34] Jehiel, Philippe and Frederic Koessler. 2005. "Revisiting Games of Incomplete Information with Analogy-Based Expectations Equilibrium," forthcoming *Games and Economic Behavior*.
- [35] Kagel, John H. and Dan Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review*, 76: 894-920.
- [36] Kagel, J. H. and Dan Levin (2002): *Common Value Auctions and the Winner's Curse*, Princeton, NJ: Princeton University Press.
- [37] Lazear, Edward P. 2000. "Performance Pay and Productivity." *American Economic Review*. 90(5): 1346-1361.
- [38] Miettinen, Topi. 2007. "Learning Foundation for the Cursed Equilibrium." Max Planck Institute of Economics Papers on Strategic Interaction.
- [39] Milgrom, Paul. 1981. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica*, 49: 921-943.
- [40] Milgrom, Paul and John Roberts. 1990. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica*, 58: 1255-77.
- [41] Milgrom, Paul and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica*, 62: 157-80.
- [42] Milgrom, Paul and Robert Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 50: 1089-1122.
- [43] Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85: 1313-1326.
- [44] Porter, Robert H. 1995. "The Role of Information in US Offshore Oil and Gas Lease Auctions." *Econometrica*, 63: pp. 1-27.
- [45] Rabin, Matthew. 1998 "Psychology and Economics." *Journal of Economic Literature*, 36: 11-46.
- [46] Roll, Richard. 1986. "The Hubris Hypothesis of Corporate Takeovers." *Journal of Business*, 59: 197-216.

- [47] Rothschild, Michael and Joseph Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, 90: 629-649.
- [48] Samuelson, William F. and Max H. Bazerman. 1985. "The Winner's Curse in Bilateral Negotiations," in *Research in Experimental Economics*, vol. 3, Vernon L. Smith, ed., Greenwich, CT: JAI Press.
- [49] Spence, Michael. 1973. "Job Market Signaling." *Quarterly Journal of Economics*. 87: 355-374.
- [50] Stahl, Dale O. and Paul W. Wilson. 1995. "On Players' Model of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior*, 10: 218-254.
- [51] Tarski, Alfred. 1955. "A Lattice-theoretical Fixpoint Theorem and its Applications." *Pacific Journal of Mathematics*, 5: 285-308.
- [52] Topkis, Donald M. 1998. *Supermodularity and Complementarity*. Princeton University Press.
- [53] Vives, Xavier. 1999. *Oligopoly Pricing. Old Ideas and New Tools*. Cambridge: MIT Press.
- [54] Vives, Xavier. 2005. "Complementarities and Games: New Developments." *Journal of Economic Literature* 43, 437-479.




Figure 1a




Figure 1b