Behavioral Equilibrium in Economies with Adverse Selection*

Ignacio Esponda†
New York University

February 20th, 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, Vincent Crawford, Willie Fuchs, Yuan-Chuan Lien, Alessandro Lizzeri, Dan Quint, Antonio Rangel, Steve Tadelis, Pedro Miranda, 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.

1 Introduction

A large literature presents evidence that people's behavior is affected by both psychological and cognitive biases. This evidence raises the question of what the effects of these biases are and whether they can 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 then apply the framework to a class of games that arise naturally in adverse selection settings, and 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 Akerlof (1970). A feature of this adverse selection setting is that lower prices select worse quality, which itself provides the buyer with 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 ignore that the seller's willingness to trade provides information about quality.

The previous literature (Kagel and Levin (1986), Holt and Sherman (1994), 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 an incentive to choose higher prices relative to a non-biased buyer (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 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 beliefs are contradicted by experience, 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

[1] See Rabin (1998) and Gilovich, Griffin, and Kahneman (2002) for a review of several biases, and 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 on the margin she has 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 2, where a dynamic learning interpretation is also provided to justify the proposed equilibrium concept.

While the above simple example illustrates the main intuition, the framework is applicable to a wide range of settings. 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 either ignore or not know 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 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 have appeared in field settings as well, including the oil industry (Capen, Clapp, and Campbell (1971)), professional baseball's free agency market (Cassing and Douglas (1980)), and corporate takeovers (Roll (1986)). One implication of this paper is that it is important to distinguish between the bias itself and the effect that the bias may have in different settings. While not trivial to disentangle, the previous literature does provide some support for the existence of this bias. However, as discussed further in Section 5, the results in this paper suggest that the effects attributed to the bias by the previous literature may be related to the fact that, in experiments, subjects are informed about the true distribution over fundamentals.

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 what are the relevant relationships in the data for the problem at hand.[2] 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 correlated, but if these divisions do not communicate with each other this correlation is likely to be ignored. 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

[2] See Aragones, Gilboa, Postlewaite, and Schmeidler (2005) for a justification of why people may fail to see regularities in the data.

learning about their environment.

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 (Dekel, Fudenberg, and 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.[3] 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 be *naive-consistent*: information obtained from actual equilibrium behavior still constrains their beliefs, but when processing this information naive players ignore the potential correlation between other players' actions and their own payoff-relevant uncertainty.

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 4. Under reasonable assumptions on information feedback, I find that naive-consistent beliefs can be supported in equilibrium and that the presence of players who ignore the (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 ignore 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. When players are aware of these information asymmetries, some institutions may naturally arise to mitigate this problem, initiated either by the party having more information (Spence, 1973) or the party which is less informed (Rothschild and Stiglitz, 1976). However, when players fail to account for selection, the adverse selection problem will not only be more severe but these institutions may be less likely to arise.

The most closely related work is a paper by Eyster and Rabin (2005). They provide the first

[3] Previous versions of self-confirming (or conjectural) equilibrium in games of complete information appear in 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. Explicit learning-theoretic foundations have been provided by Fudenberg and Levine (1993b).

systematic equilibrium analysis of the same bias that I study in this paper by introducing the notion of a *cursed equilibrium*. The main conceptual difference, discussed further in Section 5, 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 either naive, 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 (Stahl and Wilson (1994), Nagel (1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho and Chong (2004), Crawford and Costa-Gomes (2006)). Crawford and Iriberri (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 2, I present an example that illustrates the approach previously taken by the 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 3, 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 4, where I apply the framework to monotone selection games and characterize the set of equilibria when players are naive, sophisticated, or have correct beliefs. I conclude in Section 5 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.

2 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 $\widetilde{s}$ 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 either about 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, $ask = s$.

2.1 Nash equilibrium, the selection problem, and cursed equilibrium

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

$$\pi^{NE}(p) = \Pr(\widetilde{s} \leq p) \times \left[E\left(\widetilde{v} \mid \widetilde{s} \leq p\right) - p\right]. \tag{1}$$

Under the current assumptions, equation (1) becomes $\pi^{NE}(p) = p \times \left(x - \frac{1}{2}p\right)$ 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, $E(\widetilde{v})$. Following Eyster and Rabin (2005), call such a buyer a *cursed* buyer. The *perceived* profits of a cursed buyer are then

$$\pi^{Cursed}(p) = \Pr(\widetilde{s} \leq p) \times \left[E\left(\widetilde{v}\right) - p\right], \tag{2}$$

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 $x > 1/2$.[4]

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 $E(\widetilde{v}) \geq E\left(\widetilde{v} \mid \widetilde{s} \leq p\right)$ for all p (in general, this is true if $\widetilde{v}, \widetilde{s}$ 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.

[4] This 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.

2.2 Behavioral Equilibrium: naive and sophisticated buyers

A cursed buyer makes her pricing decision under the belief that the expected value of the object is $E(\widetilde{v})$, but if she repeatedly follows her cursed strategy, her expected valuation of objects traded in equilibrium is lower, $E\left(\widetilde{v} \mid \widetilde{s} \leq p^{Cursed}\right)$. 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 ignores 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), who knows the exact price-quality schedule in equilibrium.

A behavioral equilibrium depends on restrictions on players' beliefs that the modeler wishes to impose, which are in turn motivated by the feedback that players obtain from repeatedly playing the 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 3, 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:

$$\pi^N(p, p^*) = \Pr(\widetilde{s} \leq p) \times \left[E\left(\widetilde{v} \mid \widetilde{s} \leq p^*\right) - p\right]. \tag{3}$$

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^*.[5] 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

[5] Implicitly, the buyer does not know that $\widetilde{v} = \widetilde{s} + x$; whatever she learns about her valuation depends on the feedback she obtains. A general version of (3) is derived in Proposition 2. This equation actually holds true for $p^* > 0$. Otherwise, no objects are traded, no feedback is obtained, and therefore beliefs can be quite arbitrary.

she considers deviating (this is the failure to account for the selection problem); and (iii) are correct for $p = p^*$.

As defined in Section 3, 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.*[6]

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 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

$$\pi^S(p, p^*) = \Pr(\widetilde{s} \leq p) \left[\rho(p, p^*) - p\right],$$

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 not only knows that there might be a potential selection problem (so that $\rho(p, p^*)$ is not necessarily constant in p), but in addition also knows 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* $E\left(\widetilde{v} \mid \widetilde{s} \leq p^*\right)$. Hence, the perceived profit function of

[6] Both with a naive and 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 play their weakly dominant strategy.

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. $\rho(p, p^*) = E(\widetilde{v} \mid \widetilde{s} \leq p)$ 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 $p < p^N$ 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, p cannot be chosen in equilibrium by a sophisticated buyer. 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 $[p^N, p^{NE}]$.

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. For example, Proposition 2 can be applied to show that in the trading game the results generalize as long as $\widetilde{s}$ and $\widetilde{v}$ are affiliated, and u is nondecreasing in v and supermodular in (p, v).[7]

2.3 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.[8]

[7] In some settings, it is natural to relax the assumption that players learn the reservation values of trading partners (e.g. ask prices are revealed) with the assumption that the distribution of reservation values is known (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.

[8] See Fudenberg and Levine (1998) for dynamic models with steady states that correspond to Nash equilibrium and to self-confirming equilibrium (here, equilibrium with sophisticated players).

Suppose 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 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 would therefore choose 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 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.

3 Definition of Behavioral Equilibrium

There is a finite set N of players who simultaneously choose actions. Each player $i \in N$ chooses an action from a finite, nonempty, action set A_i and obtains payoff $u_i(a, v) \in \mathbb{R}$, where $a = (a_i)_{i \in N} \in A \equiv \times_i A_i$, $v \in V$ is the realization of a random variable $\widetilde{v}$ that captures players' uncertainty about payoffs, and V is finite. Before choosing an action, each player receives a signal realization s_i from a random variable $\widetilde{s}_i$ with finite support S_i. The signals and the parameters of the utility function are jointly drawn according to an objective probability distribution $\gamma \in \Delta(S \times V)$, where $S = \times_i S_i$. A (pure) strategy for player i is a function from the set of signals to the set of actions, $\alpha_i : S_i \longrightarrow A_i$, and a strategy profile is denoted by $\alpha = \{\alpha_i\}_{i \in N}$.[9]

[9] The assumptions that the game is finite and that players choose pure strategies are made for ease of exposition and can be relaxed. Existence of equilibrium is not guaranteed under these assumptions, but I will appeal to

Players know their utility function, the set V of parametric uncertainty, and their set of feasible actions. I do not restrict beliefs about opponents' utility functions (or about their rationality). The conditional distribution over $(\widetilde{a}_{-i}, \widetilde{v})$ given a profile of opponents' strategies α_{-i} and a signal s_i is denoted by

$$p_i(s_i, \alpha_{-i})(a_{-i}, v) \equiv \sum_{s_{-i}} \frac{\gamma(s, v)}{\sum_{s'_{-i}, v'} \gamma(s_i, s'_{-i}, v')} \times 1\left\{\alpha_{-i}(s_{-i}) = a_{-i}\right\}. \tag{4}$$

Player i's conjecture when she receives signal s_i is a probability distribution $\phi_i(s_i) \in \Delta(A_{-i} \times V)$ that need not coincide with (4).[10]

The standard Nash assumption can be interpreted as requiring that in equilibrium players maximize expected utility given conjectures that are correct. However, in certain situations it may be sensible to make less restrictive assumptions on equilibrium beliefs. For example, when players form conjectures based on past experience that includes only an aggregate statistic of how other players behave, the modeler might actually want to restrict players to have correct beliefs about the distribution of this aggregate statistic, and not about each of the actions that affect the statistic but are never observed. Similarly, in games with incomplete information the realization of an uncertain variable may not be observed unless a player takes a particular action. The modeler might then want to relax the assumption that players who never take such an action have correct beliefs about the distribution of this variable.

Following the literature on self-confirming equilibrium, information feedback is formalized via partitions $\mathcal{P}_i$ for each player over the state space $\Omega = A \times V$, where $P_i(\omega)$ denotes the element of the partition that contains $\omega \in \Omega$, and $\mathcal{P} = \{\mathcal{P}_i\}$ is the collection of partitions. The interpretation is that when the (past) outcome of the game is ω, player i cannot distinguish that feedback from the outcome being any other $\omega' \in P_i(\omega)$. Letting $\overline{a}_i(\omega)$ represent the action of player i that corresponds to state ω, I assume that $\overline{a}_i$ is $\mathcal{P}_i$-measurable, i.e. $\overline{a}_i$ is constant on each element of $\mathcal{P}_i$, for every i. Hence, players at least get feedback about their own actions. Examples of information feedback partitions include: (a) fully-revealing: $P_i(\omega) = \{\omega\}$ for all ω; (b) action-revealing: $\omega' \notin P_i(\omega)$ if $\overline{a}_i(\omega') \neq \overline{a}_i(\omega)$ for some i; (c) payoff-revealing: $\omega' \notin P_i(\omega)$ if $u_i(\omega') \neq u_i(\omega)$; (d) *only* payoff-revealing: $P_i(\omega) = \{\omega' : u_i(\omega') = u_i(\omega), \overline{a}_i(\omega') = \overline{a}_i(\omega)\}$.

A behavioral equilibrium restricts conjectures to be consistent with actual equilibrium play. For sophisticated players, I adopt the definition of consistency from the literature on self-confirming equilibrium. For naive players, I propose a new definition that captures their mistake when learning from feedback.

Consistency of beliefs for sophisticated players. Consistency restricts players to have

complementarity conditions to establish existence in finite games. Allowing for mixed strategies would not guarantee existence of naive equilibrium since payoffs need not be continuous in mixed strategies if players can distinguish between the consequences that result from each of the actions in the support of the mixed strategy.

[10] A difference with the definition of self-confirming equilibrium in Dekel, Fudenberg, and Levine (2004) is that I do not necessarily restrict these conjectures to come from separate beliefs about opponents' strategies and the joint distribution of signals and payoff uncertainty.

correct beliefs about the probability of observing each of the feedback signals they obtain, here represented by elements of their partition, given the strategies that they and their opponents play. For example, if $\mathcal{P}_i$ is fully-revealing then consistency requires player i to have correct conjectures (as in a Nash equilibrium), while if $\mathcal{P}_i$ is payoff-revealing consistency requires player i to have correct beliefs about the expected payoff from following her current strategy.

Definition 1 (consistency) *The conjecture of player i when she receives signal s_i, $\phi_i(s_i)$, is $\mathcal{P}_i$-consistent for (a_i, α_{-i}) if $\phi_i(s_i)[P_i(\omega)] = p_i(s_i, \alpha_{-i})[P_i(\omega)]$ for all $\omega \in \Omega$ such that $\overline{a}_i(\omega) = a_i$.*[11]

Consistency of beliefs for naive players. The conjectures of naive players are also restricted by equilibrium feedback, but naive players fail to account for the possible correlation between the actions of other players and payoff parameters in V. Hence, naive players are modeled as drawing inferences from the observed subsamples of A_{-i} and V separately, but ignoring that a more precise picture would be learned by looking at the two samples jointly.

Formally, let

$$V_i(\omega) = \{v \in V : \exists\ a \in A \text{ s.t. } (a, v) \in P_i(\omega)\}$$

denote the set of realized payoff parameters that player i cannot rule out given that the outcome of the game is ω. Similarly, let

$$A_i(\omega) = \{a \in A : \exists\ v \in V \text{ s.t. } (a, v) \in P_i(\omega)\}$$

and

$$U_i(\omega) = \left\{u \in \mathbb{R} : u_i(\omega') = u,\ \omega' \in P_i(\omega)\right\}$$

The sets $V_i(\omega)$, $A_i(\omega)$, and $U_i(\omega)$ represent the *marginal* information feedback obtained regarding payoff parameters, actions, and payoffs, respectively. For some state ω, it is possible that $V_i(\omega)$ *reveals partial marginal information*, meaning that it is neither a singleton nor V (and similarly for $A_i(\omega)$). In that case, obtaining a belief over the probability of an element of $V_i(\omega)$ requires player i to assign a probability to such an element when all she gets feedback about is that the outcome was some element in $V_i(\omega)$. The definition of naive-consistency for this general case is given in the online appendix. The definition for the particular case where there is no partial revelation of marginal information, which is the case in all the examples in the paper, appears below.

Let $P_i^V(s_i, \alpha_{-i})$ denote the probability, according to $p_i(s_i, \alpha_{-i})$, over the states ω such that that $V_i(\omega)$ is a singleton, i.e. the probability that i receives precise feedback about the realized payoff parameter, and similarly let $P_i^A(a_i; s_i, \alpha_{-i})$ denote the probability that $A_i(\omega)$ is a singleton, when restricted to ω such that $\overline{a}_i(\omega) = a_i$.

[11] For a probability distribution p over $A_{-i} \times V$, and for $\Omega' \subset \Omega$, let $p[\Omega'] \equiv p\{(a_{-i}, v) : \exists\ a_i \text{ s.t. } (a_i, a_{-i}, v) \in \Omega'\}$.

Definition 2 (naive-consistency) *Suppose that $\mathcal{P}_i$ is such that there is no partial revelation of marginal information. The conjecture of player i when she receives signal s_i, $\phi_i(s_i)$, is $\mathcal{P}_i$-naive-consistent for (a_i, α_{-i}) if the following conditions are satisfied:*

1. *$\phi_i(s_i) = \phi_i^V(s_i) \times \phi_i^{A_{-i}}(s_i)$, where $\phi_i^V(s_i)$ and $\phi_i^{A_{-i}}(s_i)$ are probability distributions over V and A_{-i}, respectively, that satisfy:*

   (a) for all $v \in V$,

$$\phi_i^V(s_i)(v) = \sum_{\omega : V_i(\omega) = \{v\}} \frac{p_i(s_i, \alpha_{-i})[\omega]}{P_i^V(s_i, \alpha_{-i})}, \tag{5}$$

   (b) for all $a_{-i} \in A_{-i}$,

$$\phi_i^{A_{-i}}(s_i)(a_{-i}) = \sum_{\omega : A_i(\omega) = \{(a_i, a_{-i})\}} \frac{p_i(s_i, \alpha_{-i})[\omega]}{P_i^A(a_i; s_i, \alpha_{-i})}, \tag{6}$$

2. *for all (a_{-i}, v), $\phi_i(s_i)(\mathcal{U}_i(a_i, a_{-i}, v)) = p_i(s_i, \alpha_{-i})(\mathcal{U}_i(a_i, a_{-i}, v))$, where*

$$\mathcal{U}_i(a_i, a_{-i}, v) \equiv \left\{(a'_{-i}, v') : U_i(a_i, a'_{-i}, v') = U_i(a_i, a_{-i}, v)\right\}.$$

Condition 1 requires players to believe that $\widetilde{a}_{-i}$ and $\widetilde{v}$ are independent and to form their beliefs from each of the subsamples of A and V separately. For example, the marginal probability for $v \in V$ is obtained by adding the probability that feedback $\{v\}$ is obtained, normalized by the probability that feedback about one of the elements of V is obtained. In addition, to the extent that players also receive some feedback about their own payoffs, condition 2 requires their beliefs about observed payoff-feedback to be consistent with their beliefs about the distribution of $\widetilde{a}_{-i}$ and $\widetilde{v}$ and with knowledge of their utility function. It is possible that there is no conjecture satisfying both conditions 1 and 2, so that $\mathcal{P}_i$-naive-consistent conjectures may not exist for some particular $(a_i, \alpha_{-i}, \gamma)$ (e.g. the trading game in Section 2 if the buyer always receives feedback about the value of the object). The online appendix includes a sufficient condition for existence and uniqueness of naive-consistent conjectures. When uniqueness holds, define the naive profit function

$$\pi_i^N(a_i; a_i^*, s_i, \alpha_{-i}) \equiv E_{\phi_i(s_i)} u_i(a_i, \widetilde{a}_{-i}, \widetilde{v}),$$

where $\phi_i(s_i)$ is $\mathcal{P}_i$-naive-consistent for (a_i^*, α_{-i}).

Definition and characterization of Behavioral Equilibrium. A behavioral equilibrium requires that players choose strategies that are optimal given their conjectures, where these conjectures are restricted by two sources. The first source is the consistency condition introduced above, which in turn depends on a behavioral assumption (i.e. whether the player is naive or sophisticated), and on the information feedback partition. Second, $\phi_i(s_i)$ are restricted to belong to $M_i(s_i) \subset \Delta(A_{-i} \times V)$, capturing additional restrictions on beliefs which cannot necessarily be inferred from actual equilibrium behavior. For example, I appeal to such a restriction when assuming (as in Sections 2 and 4) that sophisticated players know that selection is monotone.[12]

Consider a partition of the set of players into a set of sophisticated and a set of naive players, $N = N^S \cup N^N$, and let $M = \{M_i\}$.

Definition 3 (behavioral equilibrium) *A profile of strategies α is an $(\mathcal{P}, M, N)$ behavioral equilibrium if for every player $i \in N$ and for every $s_i \in S_i$ there exists a conjecture $\phi_i(s_i) \in M_i(s_i)$ such that*

i) $\alpha_i(s_i)$ maximizes expected utility given conjecture $\phi_i(s_i)$,

ii) $\phi_i(s_i)$ is $\mathcal{P}_i$-[naive]-consistent for $(\alpha_i(s_i), \alpha_{-i})$ if $i \in N^S$ [if $i \in N^N$].

When $\mathcal{P}$, M, and N are understood from the context, I sometimes omit explicit reference to them. In addition, 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 4 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 maximize, for all s_i, expected utility given conjectures that belong to $M_i(s_i)$ and are either $\mathcal{P}_i$-consistent or naive-consistent for $(\alpha_i(s_i), \alpha_{-i})$, depending on whether i is sophisticated or naive. The *generalized best response correspondence* is then 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, naive, or has correct conjectures. Letting $BR = \{BR_i\}_{i \in N}$, a behavioral equilibrium can then be characterized as a fixed point of the BR correspondence.

4 Games with Monotone Selection

I apply the equilibrium framework in Section 3 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. Later, I present conditions on the primitives of a class of models such that the following properties are satisfied.

[12] Implicitly, restrictions across signals are not allowed. While this assumption is easily relaxed in the general setup, it is important for the application to games with monotone selection.

Definition 4 *In a* ***game with monotone selection****,* $A_i \subset \mathbb{R}$ *and there exists a unique naive profit function* π_i^N *for every* s_i, α_{-i}*. In addition, the following properties are satisfied:*

(i) $\mathcal{P}$ *is payoff-revealing,*

(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) kMSP: sophisticated players know that MSP is satisfied,

(iv) action-belief complementarity: π_i^N *is single-crossing in* (a_i, a_i^*) *for every* i, s_i, α_{-i}*.*

The assumptions on the action space and on the uniqueness of naive conjectures are made for simplicity and can be relaxed to some extent.[13] The assumptions driving the results in this section are (i)-(iv).

Since $\mathcal{P}$ is payoff-revealing, players are restricted to have 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.[14] kMSP is a refinement on the beliefs of sophisticated players that requires them to know that selection is monotone, and it is formally defined in the appendix. MSP and action-belief complementarity are standard properties in some adverse selection settings. For example, in a lemons market, the lower the price that is offered by the buyer, 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 45^0 line for $p < p^N$ in Figure 1b.

The following result compares generalized best responses of players who are naive, sophisticated, or have correct beliefs. As the intuition in Section 2 makes clear, only conditions (i) and (ii) 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 $s_i \in S_i$.

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.*

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

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

Theorem 1 generalizes the results in Section 2 – 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. Milgrom and 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 π_i^{NE} 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.[15]

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).*[16]

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 online appendix), and certain auctions (see Section 4.2). The statements in Theorem 2 hold for these settings, when restricted to the set of symmetric equilibria.

4.1 A class of games with monotone selection

Consider the framework in Section 3 where payoffs are $u_i^*(a_i, \theta)$ when $(a_{-i}, t) \in \Phi_i(a_i)$ and zero otherwise, where $v = (\theta, t) \in \Theta \times T$ represents payoff uncertainty, $a_i \in A_i \subset \mathbb{R}$ is player i's action, and A_i is nonempty and finite. The interpretation is that there are two possible outcomes, each of which occurs depending on players' actions and a random variable $\widetilde{t}$.[17] In addition, suppose that attention is restricted to nondecreasing strategies.

[15] The result below slightly differs 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.

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

[17] At the expense of additional notation, the setting can be slightly generalized to encompass, e.g., the k-th unit auction and the effort game discussed in Section 4.2.

Given action a_i and the profile of opponents' strategies α_{-i}, let $\varphi_i(a_i; s_i, \alpha_{-i})$ denote the probability that the non-zero outcome occurs, conditional on s_i. Suppose, for simplicity, that this probability is never zero.[18] In addition, consider the following assumptions on the economic environment and on players' beliefs.

Assumptions on fundamentals. F1. $(\widetilde{t}, \widetilde{\theta}, \widetilde{s})$ are affiliated; **F2**. u_i^* is nondecreasing in θ for all i; **F3**. u_i^* is supermodular in (a_i, θ) for all i; **F4**. for all i: $\Phi_i(a_i') \subset \Phi_i(a_i)$ whenever $a_i' \leq a_i$; i.e. the probability of the non-zero outcome is nondecreasing in a player's own action; **F5**. Φ_i is nondecreasing in the strong set order for all i.

Assumptions on beliefs. B1. $\mathcal{P}$ is *only* payoff-revealing; **B2**. every player i has correct beliefs about ρ_i for every a_i, given the equilibrium strategies of other players; **B3**. every player i believes $\widetilde{t}$ is independent of $(\widetilde{\theta}, \widetilde{a}_{-i})$; **B4.** kMSP holds.[19]

Proposition 2 *The previous environment is a game with monotone selection when F1-F5 and B1-B4 hold.*

As the proof makes clear, B1-B3 imply existence of a unique naive profit function, adding F1, F2, and F5 implies MSP, and further adding F3-F4 implies action-belief complementarity.

4.2 Additional examples

Several standard economic settings can be modeled as games with monotone selection. I discuss how the results in the paper can be applied in the context of some concrete examples.

(1) **Monopoly/monopsony.** The setting in Section 2 is ubiquitous in applications of adverse selection to insurance, labor, financial, credit, and used goods markets, as suggested by substituting the names buyer/seller with insurer/insuree, firm/worker, market maker/informed trader, venture capitalist/entrepreneur, etc. In many of these examples, a firm faces a given supply or demand, and it is reasonable to assume that while a firm may know the willingness-to-trade in the population, it may either ignore or not know the relationship between willingness to trade and the types of potential customers. Something 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,

[18] As discussed in Section 2, if the zero outcome occurs with probability one, then a player is likely to receive no feedback about $\widetilde{\theta}$, so that multiple conjectures may be entertained.

[19] B2-B4 are captured by placing restrictions on conjectures through M_i. B2 requires players, e.g., to know the demand/supply they face (i.e. willingness to pay in the population), but not necessarily the relationship between willingness to pay and the types of potential customers. In some settings (e.g. Section 2), B2 holds when $\mathcal{P}$ is action-revealing.

implying that the H^N correspondence is still monotone nondecreasing in this setting. It can be shown that the results in Section 2 are then reversed, e.g. every equilibrium with naive firms is either *above* the highest NE or is a NE itself, which confirms that the adverse selection problem is exacerbated in the presence of naive players irrespective of whether the selection effect is monotone nondecreasing or nonincreasing.

(2) **Duopoly with adverse selection.** The monopoly examples can be extended by incorporating competition in the less-informed side of the market. For concreteness, suppose two firms compete to attract workers by simultaneously offering wages w_i. Each worker has private information s about her (home) productivity and prefers to work for firm i rather than stay home whenever $s \leq g_i(w_i)$, where g is increasing. In addition, a worker prefers working for firm 1 rather than firm 2 if and only if $h(w_1, w_2) \geq t$, where h is nondecreasing in w_1 and nonincreasing in w_2 (so that higher values of t indicate a higher preference for working for firm 2). If the worker works for firm i, then firm i makes profits $u_i^*(w_i, \theta)$, where θ is the worker's productivity at work. Profits are zero when a worker is not hired. Assumptions F1-F3 and B1–B4 are natural in this context, and the corresponding Φ_i satisfies F4 and, if $\widetilde{t}$ is assumed to be independent of $(\widetilde{\theta}, \widetilde{s})$, F5 as well. If the latter independent assumption is dropped, then MSP does not necessarily hold since now an increase in wages may steal the least attractive workers from the other firm.

With the above assumptions, this is a game with monotone selection, and therefore if either firms are symmetric (see the online appendix) or if firms are asymmetric but the game has strategic complementarities,[20] then the adverse selection problem is exacerbated and both wages and the quality of hired workers are lower when firms are naive. Due to softened competition, there are examples where firms are actually better off when being naive.

(3) **Symmetric first price auctions and kth unit auctions.** Consider an auction where valuations have a common-value component. 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 the affiliation assumption, 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 games with strategic complementarities, it is possible to use the comparative statics results in Esponda (2006) to conclude that, in a symmetric equilibrium, bidding is less aggressive when all bidders are naive.

The same result can be easily obtained in a symmetric kth unit auction, where k identical objects

[20] Conditions under which a standard product-differentiated oligopoly game satisfies strategic complementarities are well known (Vives, 1999), but less is known about oligopoly games under adverse selection. In the firm-worker example, if $p(w_1, w_2)$ denotes the probability that a worker prefers to work for firm 1 rather than for firm 2 when the offered wages are (w_1, w_2), then a sufficient condition for both the Nash and naive games to satisfy strategic complementarities is that p and $1 - p$ are log-supermodular in w. (This assumption is satisfied, for example, if p is differentiable, $\frac{\partial^2 p}{\partial w_1 \partial w_2} = 0$, $\frac{\partial p}{\partial w_1} \geq 0$, and $\frac{\partial p}{\partial w_1} \leq 0$.)

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 that are won by a player of type s_i is $E\left[u(s_i, \widetilde{s}_{-i}) \mid \widetilde{y}_k \leq s_i\right]$, where $\widetilde{y}$ is the kth highest of the opponents' signals. Since naive bidders ignore 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, a symmetric equilibrium with naive bidders is given by the above conditional expectation, while the symmetric Nash equilibrium (Milgrom, 1981) is given by the more aggressive bidding strategy $\beta^{NE}(s_i) = E\left[u(s_i, \widetilde{s}_{-i}) \mid \widetilde{y}_k = s_i\right]$.

(4) **Team effort.** A member of a team must take into account that: (i) her effort affects the probability of success, and (ii) since other players choose effort based on private information about the value of a successful outcome, changing her effort also affects the expected value of a successful outcome. A naive team member ignores the second effect.

Assume that F1-F3 and B1-B4 hold, and let $c_i(a_i)$ be the cost of choosing effort a_i. If the team succeeds, each player obtains θ, while if it fails, payoffs are zero. Whether this is a game with monotone selection depends on the technology translating effort into success. Suppose that success occurs if $\min_{i \in N} x_i \geq t$, where $\widetilde{t}$ is independent of $\widetilde{\theta}$. Then F4 and F5 hold, implying that this is a game with monotone selection. Intuitively, since it is the minimum level of effort that matters for success, a player has no influence on success when (the realization of) the minimum of other players' effort level is lower than her effort, but does increase the likelihood of success when the minimum of other players' effort is higher. Hence, by increasing effort a player not only increases the probability of success, but also makes it more likely that success occurs for higher values of other players' efforts, which are indicative of higher realizations of $\widetilde{\theta}$. In addition, efforts are complements and the game has Nash strategic complementarities, implying that a team of naive players exerts less effort than in a Nash equilibrium. If in addition the game is symmetric, or if it is asymmetric but the game also has naive strategic complementarities, naive players also put lower effort compared to sophisticated ones.[21]

(5) **Preemption Game.** Two firms simultaneously choose whether to act today or to wait and act tomorrow. Firm(s) which act first are rewarded, but firms would benefit from coordinating to move tomorrow, when the benefits of acting are known. Firms get feedback about the benefits of acting only when they are the first to act. A firm which decides to act today but evaluates whether it should rather wait must account for two effects: (i) the risk of being preempted, and (ii) the fact that if it waits, it will only get to act tomorrow if the other firm has not preempted it, which (under monotone assumptions on the information structure, and given that firms have private information) is more likely to happen when the benefits from acting are lower. A naive firm ignores

[21] The naive game may not have strategic complementarities since when other players exert more effort, the marginal effect of effort on the probability of success may increase, but, on the other hand, naive players fail to realize that increasing effort would also have a positive effect on the rewards from effort.

the second effect, and therefore believes 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 standard assumptions the game also has Nash strategic complementarities. Equilibrium with naive firms involves more waiting (i.e. better coordination) than in a Nash equilibrium, and assuming either symmetry or naive strategic complementarities, the same is true with respect to sophisticated firms.[22]

5 Discussion

In this paper, I provide a framework to study 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 on 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.

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 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 ignore the relationship between other players' private information and actions. This ignorance is modeled by assuming that 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.[23],[24]

The practical problem with this sharp distinction is that 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'

[22] On one hand, if the other firm waits more often the preemption motive makes it more desirable to also wait more often. However, if a firm is naive and waits it will then believe the benefit from acting has increased (since the other firm is waiting under more favorable information), which makes acting today more attractive.

[23] More generally, Jehiel and Koessler (2005) assume players make mistakes when forecasting the type-contingent strategies of their opponents by bundling types into analogy classes. The same arguments that motivate the alternative appoach that I follow in this paper apply to their framework.

[24] Eyster and Rabin (2005) also consider intermediate cases where players "underappreciate" the selection problem: with some probability beliefs are as in a cursed equilibrium and with the remaining probability beliefs about opponents' type-contingent strategies are correct (as in a Bayesian Nash equilibrium). In contrast, in a behavioral equilibrium players are either aware or unaware of the selection problem, and it is only players who are aware (i.e. sophisticated) who might now either under or overappreciate the informational content of other players' actions.

beliefs about fundamentals. By placing restrictions on beliefs about fundamentals independently 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 2, 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.[25]

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 which 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 arise in equilibrium provides further insight into the relationship with experimental results. Consider the trading game in Section 2, where a naive equilibrium exists under the reasonable assumption that feedback about the value of the object is only obtained when the object is traded. Suppose, instead, that feedback about the value of the object is always received, irrespective of whether the object is traded or not. There are some situations where this alternative assumption may make sense, such as a common value auction with resale, where the resale price is observed by everyone. Under this alternative assumption, naive-consistent beliefs do not exist and therefore a naive model of the world cannot persist in an equilibrium setting. The reason is that knowing the correct expected quality of all objects but consistently obtaining objects of lower quality is a fact that cannot be reconciled with a naive model of the world, i.e. the seemingly contradictory feedback cannot be rationalized unless the player understands the selection problem.

An alternative assumption that leads to the same result is that the player knows a priori what the true distribution of the value of the object is. Note that this assumption is satisfied 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 in some way, and at least two possibilities come to mind. First, players can 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 ignoring to take into account payoff-feedback information altogether. In this case, players will behave as

[25] See 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) further pursues the argument in the text that no distinction should be made between structural and strategic uncertainty and presents a framework that is an alternative to Harsanyi's Bayesian Nash equilibrium and extends the present paper by allowing people to learn not only from feedback but also from introspection.

predicted by Kagel and Levin (1986), Holt and Sherman (1994), and Eyster and Rabin (2005), and it is no coincidence that the experimental evidence (reviewed by Kagel and Levin (2002)) 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 such a model of the world can persist in equilibrium or not. 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.

There are two implications for the experimental literature. First, it is likely that the intuition and experimental results developed in adverse selection settings are strongly driven 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. If naive players were to receive no information about the distribution of the value of the object, the compelling intuition provided in this paper suggests that they would actually underprice. Second, it would be interesting to extend the experimental literature by relaxing the assumption that players know the distribution over the fundamentals but must rather learn it. Of course, it is not trivial how one would make sure that players have enough opportunity to learn in an experimental setting.[26]

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.

[26] This is not so much 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 an assertion about the need to extend the methodology of treating both structural and strategic uncertainty as endogenous to the experimental literature as well.

References

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

[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] Battigalli, Pierpaolo. 1987. "Comportamento Razionale Ed Equilibrio Nei Giochi e Nelle Situazioni Sociali." unpublished undergraduate dissertation, Bocconi University, Milano.

[7] Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *Quarterly Journal of Economics,* 119: 861-898.

[8] Capen, Edward C., Robert V. Clapp, and William M. Campbell. 1971. "Competitive Bidding in High Risk Situations." *Journal of Petroleum Technology,* 23: 641-53.

[9] Cassing, James and Richard W. Douglas. 1980. "Implications of the Auction Mechanism in Baseball's Free Agent Draft." *Southern Economic Journal,* 47: 110-21.

[10] 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.

[11] 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.

[12] 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.

[13] Dekel, Eddie, Drew Fudenberg, and David K. Levine. 2004. "Learning to Play Bayesian Games." *Games and Economic Behavior,* 46: 282-303.

[14] Esponda, Ignacio. 2006. "Information Feedback in First Price Auctions." working paper.

[15] Esponda, Ignacio. 2006. "Robust Equilibrium Analysis in Games with Uncertainty." working paper.

[16] Eyster, Erik and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica,* 73: 1623-1672.

[17] Fudenberg, Drew. 2006. "Advancing Beyond Advances in Behavioral Economics." *Journal of Economic Literature,* 44: 694-711.

[18] Fudenberg, Drew and David K. Levine. 1993a. "Self-confirming Equilibrium." *Econometrica,* 61: 523-545.

[19] Fudenberg, Drew and David K. Levine. 1993b. "Steady-State Learning and Nash Equilibrium." *Econometrica,* 61: 547-573.

[20] Fudenberg, Drew and David K. Levine. 1999. *The Theory of Learning in Games.* The MIT Press, 1999.

[21] Gilovich, Thomas, Dale Griffin, and Daniel Kahneman, ed. 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment.* New York: Cambridge University Press.

[22] 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.

[23] Holt, Charles A. and Roger Sherman. 1994. "The Loser's Curse." *American Economic Review,* 84: 642-652.

[24] Jehiel, Philippe and Frederic Koessler. 2005. "Revisiting Games of Incomplete Information with Analogy-Based Expectations Equilibrium," working paper.

[25] Kagel, John H. and Dan Levin. 1986. "The Winner's Curse and Public Information in Common Value Auctions." *American Economic Review,* 76: 894-920.

[26] Kagel, J. H. and D. Levin (2002): *Common Value Auctions and the Winner's Curse*, Princeton, NJ: Princeton University Press.

[27] Kagel, John H. and Alvin. E. Roth, ed. 1995. *Handbook of Experimental Economics.* Princeton, NJ: Princeton University Press.

[28] Milgrom, Paul. 1981. "Rational Expectations, Information Acquisition, and Competitive Bidding." *Econometrica,* 49: 921-943.

[29] Milgrom, Paul and John Roberts. 1990. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica,* 58: 1255-77.

[30] Milgrom, Paul and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica,* 62: 157-80.

[31] Milgrom, Paul and Robert Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica,* 50: 1089-1122.

[32] Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *American Economic Review,* 85: 1313-1326.

[33] Rabin, Matthew. 1998 "Psychology and Economics." *Journal of Economic Literature,* 36: 11-46.

[34] Roll, Richard. 1986. "The Hubris Hypothesis of Corporate Takeovers." *Journal of Business,* 59: 197-216.

[35] 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.

[36] 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., Greenwhich, CT: JAI Press.

[37] Spence, Michael. 1973. "Job Market Signaling." *Quarterly Journal of Economics.* 87: 355-374.

[38] Stahl, Dale and Paul Wilson. 1995. "On Players' Model of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior,* 10: 218-254.

[39] Tarski, Alfred. 1955. "A Lattice-theoretical Fixpoint Theorem and its Applications." *Pacific Journal of Mathematics,* 5: 285-308.

[40] Topkis, Donald. 1978. "Minimizing a Submodular Function on a Lattice." *Operations Research,* 26: 305-321.

[41] Topkis, Donald. 1998. *Supermodularity and Complementarity.* Princeton University Press.

[42] Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance,* 16: 8-27.

[43] Vives, Xavier. 1999. *Oligopoly Pricing. Old Ideas and New Tools.* Cambridge: MIT Press.

[44] Vives, Xavier. 2005. "Complementarities and Games: New Developments." *Journal of Economic Literature* 43, 437-479.

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.

A. Fixed points and monotone comparative statics[27]

Throughout, let $X \subset \mathbb{R}^K$ denote a nonempty, finite set, and let T be a nonempty, partially ordered set. An element of $x \in X$ is the *highest* element of X if $x \geq y$ for every $y \in X$; it is the *lowest* element if $x \leq y$ for all $y \in X$. A correspondence $\phi : T \longrightarrow X$ is *increasing in the strong set order* if when $t > t'$, then for each $x \in \phi(t)$ and $y \in \phi(t')$, $\sup(x, y) \in \phi(t)$ and $\inf(x, y) \in \phi(t')$. A function $f : X \longrightarrow \mathbb{R}$ is *supermodular* if for all $x, y \in X$, $f(\inf(x, y)) + f(\sup(x, y)) \geq f(x) + f(y)$. A function $g : X \times T \longrightarrow \mathbb{R}$ has *increasing differences* in its arguments (x, t) if $g(x, t) - g(x, t')$ is nondecreasing in x for all $t \geq t'$. A function $g : X \times T \longrightarrow \mathbb{R}$ is the *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 $g(x, t) > g(x', t)$. A function with increasing differences in (x, t) is also single-crossing in (x, t), but the reverse need not hold. The following results are used in the proofs.

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

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

MCS2. (Athey, 1998) Let $u : A \times S \longrightarrow \mathbb{R}$ be a function where $A \subset R^K$ and $S \subset R^K$ is the support of a vector of affiliated random variables $\widetilde{s}$. Let $\Phi : X \longrightarrow S$ be a correspondence that is nondecreasing in the strong set order. Define $U(a, x) \equiv E\left[u(a, \widetilde{s}) \mid \widetilde{s} \in \Phi(x)\right]$. If $u(a, \cdot)$ is nondecreasing in s, then $U(a, \cdot)$ is nondecreasing in x. If u is supermodular in (a, s), then U has increasing differences in (a, x).

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 (profit) functions $E_{\phi_i(s_i)} u_i(\cdot, \widetilde{a}_{-i}, \widetilde{v})$ given a conjecture $\phi_i(s_i)$ that is $\mathcal{P}_i$-consistent for (a_i^*, α_{-i}) and that belongs to

[27] See, e.g., Topkis (1998) and Vives (1999,2005).

$M_i(s_i)$. Define Π_i^N similarly for naive-consistent conjectures, and let $\pi_i^{NE}(\cdot; s_i, \alpha_{-i})$ denote the profit function for correct conjectures $p_i(s_i, \alpha_{-i})$. 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 $\mathcal{P}_i$ is payoff-revealing, 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 $\mathcal{P}_i$ is payoff-revealing, $U_i(\omega)$ is a singleton for all ω. Let $\phi^N(s_i)$ denote the naive-consistent conjecture corresponding to $\pi_i^N(\cdot; a_i^*, s_i, \alpha_{-i})$. Then:

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

where the first, second and last equalities follow from definitions, the third and fifth equalities follow since $U_i(\omega)$ is a singleton, and the fourth equality from condition 2 in the definition of naive-consistency. A similar proof shows that $\pi_i^S(a_i^*; a_i^*, s_i, \alpha_{-i}) = \pi_i^{NE}(a_i^*; s_i, \alpha_{-i})$ (since condition 2 in the definition of naive-consistency is implicitly required by the definition of consistency). ■

Abusing notation, let $\Pi_i^N(a_i^*, p)$ be the union of $\Pi_i^N(a_i^*, s_i, \alpha_{-i})$ over (s_i, α_{-i}) that induce $p \in \Delta(A_{-i} \times V)$. Condition kMSP is formally defined as follows: for all $i \in N^S$ and $s_i \in S_i$, $M_i(s_i)$ is the set of all $\phi \in \Delta(A_{-i} \times V)$ such that $\Pi_i^N(\cdot, \phi)$ is *strongly nondecreasing*, i.e. for any $a_i^{*\prime} \leq a_i^*$, $\pi' \in \Pi_i^N(a_i^{*\prime}, \phi)$, and $\pi \in \Pi_i^N(a_i^*, \phi)$, 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 $\mathcal{P}_i$ is payoff-revealing, kMSP holds, 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 $\phi_{a_i^*}^S$ be the conjecture corresponding to $\pi_i^S(\cdot; a_i^*, s_i, \alpha_{-i})$ and let $p_i \in \Delta(A_{-i} \times V)$ be induced by (s_i, α_{-i}). Note that $\Pi_i^N(a_i^*, \phi_{a_i^*}^S) = \Pi_i^N(a_i^*, p_i)$ since by definition naive-consistent beliefs depend only on the probability distribution over marginal feedback, and since $\phi_{a_i^*}^S$ is $\mathcal{P}_i$-consistent for (a_i^*, α_{-i}) then such distribution is the same for $\phi_{a_i^*}^S$ and p_i. Then $\Pi_i^N(a_i^*, \phi_{a_i^*}^S) = \left\{\pi_i^N(\cdot; a_i^*, s_i, \alpha_{-i})\right\}$

by the assumption that there is a unique naive profit function. Now let $a_i \geq a_i^*$ and consider $\overline{\pi}_i^N(\cdot; a_i, s_i, \alpha_{-i}) \in \Pi_i^N(a_i, \phi_{a_i^*}^S)$. Then

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

where the inequality follows from kMSP and $a_i \geq a_i^*$, and the equality follows from PR1 (i.e. $\overline{\pi}_i^N(a_i; a_i, s_i, \alpha_{-i})$ are the "correct" beliefs from playing a_i when the "true" distribution over (a_{-i}, v) is given by $\phi_{a_i^*}^S$). 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 $s_i \in S_i$. 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 $\mathcal{P}_i$ is payoff-revealing (PR1), the first inequality follows from the definition of a naive-BR, and the second inequality follows from MSP and $\alpha_i^N(s_i) \geq \alpha_i^{NE}(s_i)$. 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

$$BR_i^N(\alpha_{-i}) = \left\{\alpha_i : \ \alpha_i(s_i) \in h_i^N(\alpha_i(s_i), s_i, \alpha_{-i}) \text{ for all } s_i \in S_i\right\}.$$

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 $\underline{h}_i^N(\cdot, s_i, \alpha_{-i}) : A_i \longrightarrow A_i$, by FP1 there is a lowest fixed point of $\underline{h}_i^N(\cdot, s_i, \alpha_{-i})$, $\underline{\alpha}_i(s_i, \alpha_{-i}) = \inf\left\{a_i \in A_i : \underline{h}_i^N(a_i, s_i, \alpha_{-i}) \leq a_i\right\}$. 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 from kMSP 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 $\mathcal{P}_i$ is payoff-revealing (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 $X = \times_i A_i^{S_i}$ denote the finite set of strategy profiles, note that $\underline{BR}^m : X \longrightarrow X$ is nondecreasing in α, so that FP1 implies that there is a lowest fixed point of $\underline{BR}^m$, given by $\underline{\alpha}^m = \inf\{\alpha \in X : \underline{BR}^m(\alpha) \leq \alpha\}$. For any α that is a fixed point of BR^m, $\alpha \geq \underline{BR}^m(\alpha)$. Therefore, $\underline{\alpha}^m$ is also the lowest fixed point of 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 $\mathcal{P}_i$-consistent and because correct beliefs belong to M_i^S – i.e. the set of beliefs that satisfy kMSP – given that it is actually true that MSP holds).

(b) By part 1 of Theorem 1, $\overline{BR}^N(\alpha) \leq \overline{BR}^{NE}(\alpha)$ for all $\alpha \in X$. Let $T = \{0, 1\}$ and define $f : X \times T \longrightarrow X$ such that $f(\cdot, 0) = \overline{BR}^N(\alpha)$ and $f(\cdot, 1) = \overline{BR}^{NE}(\alpha)$. Then f is nondecreasing in $t \in T$, 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 $\alpha \in X$, 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.[28]

(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 $X' = \{\alpha \in X : \alpha \leq \alpha^S\}$. For any $\alpha \in X'$, $\underline{BR}^N(\alpha) \leq \underline{BR}^N(\alpha^S) \leq \alpha^S$, 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

[28] When 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.

1 (i.e. $\underline{BR}^N(\alpha) \le \alpha'$ for any $\alpha' \in BR^S(\alpha)$). Hence, $\underline{BR}^N(X') \subset X'$ and it follows from FP1 that there exists a naive equilibrium $\alpha^N \in X'$.

Proof of Proposition 2. By B2, beliefs about $\varphi_i(a_i; s_i, \alpha_{-i})$ are correct for any a_i, so consider beliefs about $Eu_i^*(a_i, \widetilde{\theta})$ when a_i^* is played. Since *only* payoffs are revealed by $\mathcal{P}_i$ (B1), it follows that players observe the exact realization of $\widetilde{\theta}$ if and only if $(\alpha_{-i}(s_{-i}), t) \in \Phi_i(a_i^*)$, and get no feedback otherwise. In addition, player i believes that this conditional expectation does not depend on their action a_i^* since: i) she believes $\widetilde{t}$ is independent of $\widetilde{\theta}$ (B3), and ii) she is naive, so she ignores that opponents' actions might be correlated with $\widetilde{\theta}$. Therefore, naive-consistency requires that beliefs about $Eu_i^*(a_i, \widetilde{\theta})$ be given by the conditional expectation

$$\nu(a_i; a_i^*, s_i) \equiv E\left(u_i^*(a_i, \widetilde{\theta}) \mid (\alpha_{-i}(\widetilde{s}_{-i}), \widetilde{t}) \in \Phi_i(a_i^*),\ \widetilde{s} = s_i\right)$$

establishing uniqueness of a naive profit function $\pi_i^N = \varphi_i(a_i; s_i, \alpha_{-i}) \times \nu(a_i; a_i^*, s_i)$.

Since $(\widetilde{\theta}, \widetilde{s}, \widetilde{t})$ are affiliated (F1), u_i^* is nondecreasing in θ (F2), and Φ_i is nondecreasing in the strong set order (F5), it follows from (MCS2) and from the assumption that α is nondecreasing that $E\left(u_i^*(a_i, \widetilde{\theta}) \mid (\alpha_{-i}(\widetilde{s}_{-i}), \widetilde{t}) \in \Phi_i(a_i^*),\ \widetilde{s} = s_i\right)$ is nondecreasing in a_i^*, so that MSP holds.

Since $(\widetilde{\theta}, \widetilde{s}, \widetilde{t})$ are affiliated (F1), u_i^* is supermodular in (a_i, θ) (F3), and Φ_i is nondecreasing in the strong set order (F5), it follows from (MCS2) and from the assumption that α is nondecreasing that $E\left(u_i^*(a_i, \widetilde{\theta}) \mid (\alpha_{-i}(\widetilde{s}_{-i}), \widetilde{t}) \in \Phi_i(a_i^*),\ \widetilde{s} = s_i\right)$ has increasing differences in (a_i, a_i^*). Since $\varphi_i(a_i; s_i, \alpha_{-i})$ is nonnegative and nondecreasing in a_i (F4) and since ν is nondecreasing in a_i^* (see above), it then follows that π_i^N has increasing differences in (a_i, a_i^*), implying that it is single-crossing in (a_i, a_i^*).

Together with kMSP (B4), the properties in definition 4 are then established. ■

Proof of Proposition 1. Let t denote a time period where trade takes place, so that the realized value of the object at time t is a draw from a random variable $\widetilde{h}_{\widetilde{y}_t}$ that is uniformly distributed in the interval $[x, x + \frac{1}{2}\widetilde{y}_t]$ (since trade occurs when $\widetilde{s}_t \le \widetilde{p}_t^* = \frac{1}{2}\widetilde{y}_t$). Then $\widetilde{y}_t$ can be written as

$$\widetilde{y}_t = \frac{t-1}{t}\widetilde{y}_{t-1} + \frac{1}{t}\widetilde{h}_{\widetilde{y}_{t-1}},$$

and the proof proceeds by showing that $\widetilde{y}_t$ converges in probability to $\overline{y}$, where $\overline{y}$ is a solution to $\overline{y} = E\widetilde{h}_{\overline{y}} = x + \frac{1}{4}\overline{y}$, i.e. $\overline{y} = \frac{4}{3}x$. Hence, $\widetilde{p}_t$ converges to $\overline{p} = \frac{1}{2}(\frac{4}{3}x) = \frac{2}{3}x$, which is the naive equilibrium price.

The proof of convergence of $\widetilde{y}_t$ to $\overline{y}$ is as follows (alternatively, a law of large numbers for the sum of asymptotically independent random variables can be applied). Write

$$\widetilde{y}_{T+k} = \frac{T}{T+k}\widetilde{y}_T + \frac{1}{T+k}\sum_{j=1}^{k}\widetilde{h}_{\widetilde{y}_{T+k-j}} \tag{7}$$

and pick T, $k(T)$, and $\varepsilon > 0$ such that $\lim_{T\to\infty} k(T) = \infty$ and T is large enough relative to $k(T)$ so that $\Pr(|\widetilde{y}_T - \widetilde{y}_{T+k-j}| > \varepsilon) = 0$ for $j = 1, ..., k-1$. Hence, (7) implies that $\Pr\left(\left|\widetilde{y}_T - \frac{1}{k}\sum_{j=1}^{k} \widetilde{h}_{\widetilde{y}_T}\right| > \varepsilon\right) = 0$. Now take $T \to \infty$, so that $k \to \infty$ and then by the law of large numbers $\frac{1}{k}\sum_{j=1}^{k} \widetilde{h}_{\widetilde{y}_T} \xrightarrow{p} E\widetilde{h}_{\widetilde{y}_T} = x + \frac{1}{4}\widetilde{y}_T$. Hence, $\lim_{T\to\infty} \Pr(|\widetilde{y}_T - (x + \frac{1}{4}\widetilde{y}_T)| > \varepsilon) = 0$, so that $\lim_{T\to\infty} \Pr(|\widetilde{y}_T - \frac{4}{3}x| > \frac{4}{3}\varepsilon) = 0$. ■

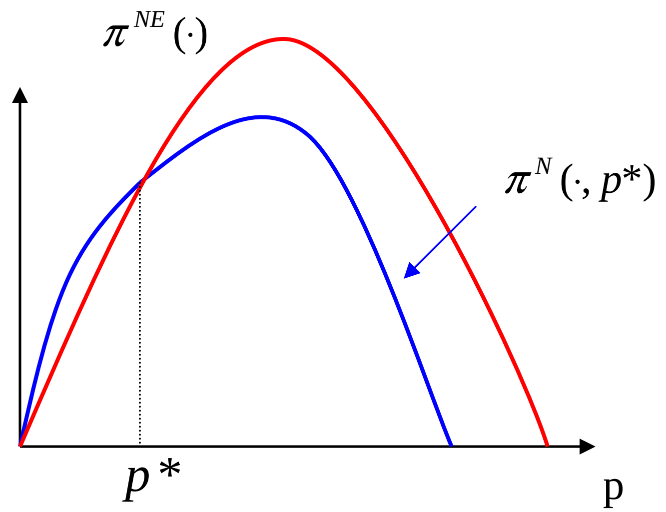

Figure 1a

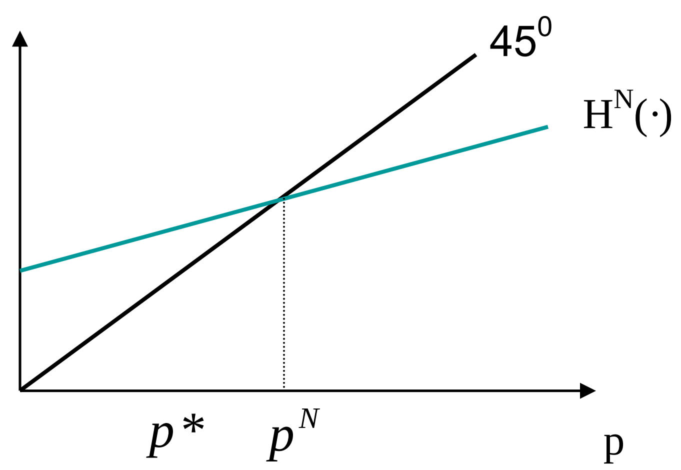

Figure 1b