DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

NO TRADE

Juan D. Carrillo
University of Southern California and CEPR

Thomas R. Palfrey
California Institute of Technology

SOCIAL SCIENCE WORKING PAPER 1279

October 2007

No Trade *

Juan D. Carrillo
University of Southern California and CEPR

Thomas R. Palfrey
California Institute of Technology

October 2007 (Original version May 2006)

This version: October 26, 2007

Abstract

We investigate, in a simple bilateral bargaining environment, the extent to which asymmetric information can induce individuals to engage in exchange where trade is not mutually profitable. We first establish a no-trade theorem for this environment. A laboratory experiment is conducted, where trade is found to occur between 16% and 32% of the time, depending on the specific details of the environment and trading mechanism. In most cases, buyers gain from such exchange, at the expense of sellers. An equilibrium model with naïve, or "cursed," beliefs accounts for some of the behavior findings, but open questions remain.

JEL classification: O24, O26.

Keywords: bilateral bargaining, private information, no trade theorem, experimental economics.

*Part of this research was conducted while the first author was visiting Caltech. The hospitality of the institution is greatly appreciated. We also gratefully acknowledge the financial support of the National Science Foundation (SES-0450712, SES-0094800, SES-0617820), The Princeton Laboratory for Experimental Social Science, The Gordon and Betty Moore Foundation. We thank Shivani Nayyar, Uliana Popova, Stephanie Wang, Rumen Zarev and Yi Zhu for research assistance. Address for correspondence: Juan D. Carrillo, Department of Economics, University of Southern California, Los Angeles, CA 90089, <juandc@usc.edu> or Thomas R. Palfrey, Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, <trp@hss.caltech.edu>.

1 Introduction

Understanding the effects of private information on trading behavior is central to the study of markets, especially markets for risky assets. The present paper investigates whether, in the context of an extremely simple trading environment, asymmetric information can induce individuals to engage in exchange where trade is not mutually profitable. The question is interesting because, a priori, it seems unclear which trading behavior should be expected in a controlled environment. On the one hand, no-trade theories are robust and intuitive, so one might think that individuals should realize (or learn) the negative implications associated to the acceptance of an offer to trade. On the other hand, some experimental studies in other asymmetric information contexts (such as adverse selection markets and common value auctions) suggest that individuals do not fully comprehend that the actions of other players depend on their information.

To study this problem, we consider a simple two-person bargaining game with two-sided private information. One individual (the seller) is endowed with one unit of an asset. Another individual (the buyer) can acquire it if both agree on a price for the transaction. The asset is of pure common value and each individual has a private signal about this value. We assume that if agents pool their information, there is no residual uncertainty about the asset value. As a result, there cannot be trade for insurance or risk-sharing motives. Trading for pure informational reasons is not possible either, for the same reasons as in standard no-trade theorems: both agents cannot benefit from trade, so accepting the other agent's terms implies that one's end of the deal must be ex-post unfavorable. Therefore, standard theory predicts that trade should never occur in our environment.

We consider variations of the game in two dimensions: the value of the asset can be the average, the minimum or the maximum of both signals (from now on labeled "ave", "min" and "max" treatments), and the trading mechanism can be either a seller's take-it-or-leave-it price or a double auction, where trade occurs at the seller's price whenever the buyer's bid exceeds it (from now on, "price" and "auction" treatments).

Our experiment delivers two findings that are common to all treatments. First and contrary to the theoretical prediction, we observe substantial trade, with probabilities ranging from 16% to 32% depending on treatments. This amount of exchange is all the more considerable if we note that trade only occurs between 3.3% and 15.3% whenever the seller's signal exceeds the buyer's signal. Second, in almost all cases sellers suffer from adverse selection: they would earn substantial profits if the behavior of buyers did not depend on their information but, since it does, they end up making losses; given the zero sum nature of the transaction, buyers reap substantial gains from the adverse selection.

There are also interesting findings about differences in behavior across treatments. Buyers and sellers both adapt their strategy to *changes across the asset value function*. In particular, seller prices and buyer bids both increase as we move from min to ave and from ave to max. Behavior also *changes across mechanisms*. In particular, the likelihood of trade is significantly lower in the auction than in the price treatment. The difference comes mostly (although not exclusively) from buyers, who submit bids in the auction treatment which are, on average, lower than the maximum offer prices they accept in the price treatment. Interestingly, there is also little evidence of learning. This suggests the phenomena we observe are robust: sellers are consistently exploited by buyers, even though subjects gain experience in both roles and the amount of feedback is considerable.[1]

As an attempt to understand these findings, we consider an alternative model, the "cursed equilibrium" (Eyster and Rabin, 2005), where the assumption of full rationality of traders is relaxed. In particular, the model posits that traders do not correctly take into account the statistical relationship between the private information and action of their rival – in the extreme version we consider, they behave as if there is no relationship at all. In the context our model, we show that traders with such belief fallacies are vulnerable to accepting or offering unfavorable terms of trade.

In addition, we show that in a world populated by traders, all suffering this fallacy, outcomes will be systematically biased. Moreover, the severity and magnitude of the bias depends on the common value function. This is consistent with our main findings. For the price mechanism, the buyers' acceptance decision, the sellers' price function and the trade frequencies predicted by the cursed equilibrium model all match up reasonably well with the data. We conjecture the same will be true with the auction mechanism, but have been unsuccessful in solving analytically for the cursed equilibrium in that mechanism.

Related literature: Theory. Milgrom and Stokey (1982) and Tirole (1982) establish that, in equilibrium, rational individuals will not trade for purely informational reasons. More specifically, if fully rational agents have

[1]For example, in the auction treatments sellers observe the bid of the buyer they are paired with and, at the end of each round, learn the signal of that buyer. The same applies to buyers.

common prior beliefs and the existing asset allocation is Pareto optimal (say, as the result of previous trading), then new private information to some or all agents in the economy will not induce trade. The logic is simple. Traders who receive private information have the marginal valuation for their asset allocation modified. However, without insurance or transaction motives for trading, every agent realizes that a transaction beneficial for someone must necessarily be detrimental for someone else. Thus, the acceptance of the terms of a trade is evidence that the deal must be unfavorable.

This "no-trade theorem" has been extended in a number of directions. For example, Morris (1994) identifies conditions under which no-trade occurs even if individuals have heterogeneous prior beliefs. Blume et al. (2006) show that the no-trade result applies to competitive markets if and only if markets are complete. Serrano-Padial (2007) demonstrates that it holds under more general bilateral trading mechanisms.

From a theoretical viewpoint, our framework is also related to the literature on bargaining with private information. In a private value setting with two-sided private information, it has been shown that trade occurs when the seller's valuation is sufficiently lower than the buyer's valuation. This means, in particular, that full efficiency cannot be achieved and that asymmetric information prevents the realization of some profitable trades.[2] In a common value setting with sequential offers, Evans (1989) and Vincent (1989) show that one-sided private information also leads to inefficiently low trading. Instead, our experimental results imply the opposite observation: the introduction of asymmetric information leads to trade in contexts where we should observe none. Finally, our no-trade result bears some resemblance with the well-known idea that information may impede the realization of mutually beneficial agreements, as was first emphasized by Hirshleifer (1971).

Related literature: Experiments. Constant sum games with two-sided private information, of which bargaining with common values is a special case, have rarely been studied in the laboratory. A possible reason is the difficulty to find simple games that subjects can easily understand and play. Because of the signaling nature of these games, one would also want to identify games that have one or few equilibria which can be determined analytically, in order to compare the empirical behavior with the theoretical

[2] See e.g. Chatterjee and Samuelson (1983) in the context of a double auction and Myerson and Satterthwaite (1983) in a generalized bargaining game. Radner and Schotter (1989) study the Chatterjee and Samuelson model of private-values bargaining in the laboratory. Cramton et al. (1987) show that initial ownership is crucial to determine whether efficiency can be achieved.

predictions. The bargaining games we study here satisfy these criteria.

An exception is *the compromise game* (Carrillo and Palfrey, 2006), where two agents with private signals choose between two actions, "fight" and "retreat". If at least one agent fights, the agent with highest signal receives a high payoff and the other receives a low payoff. If both retreat, they each get an intermediate payoff. *The betting game* (Sonsino et al. (2001), Sovic (2004), Camerer et al. (2006)) is also related, although only the simultaneous version has been studied. In that game, agents with private information about the state of the world choose whether to bet on an asset, whose value for one agent is always the negative of the value for the other. Agents get these values if both bet and they get zero otherwise. As in our game, these studies find substantial retreating and betting, although the theoretical prediction is that it should never occur. In these two games, however, the sharing rule (intermediate payoff, payoff if betting) is exogenously fixed. Instead, we are interested precisely in how subjects set transaction prices (bids and offers) as a function of (i) their private information, (ii) their role, (iii) the value function, (iv) the trading mechanism, and (v) the timing of the game.

A number of studies have compared behavior between a game played in its extensive form and an equivalent version of the game played in (sometimes reduced) strategic form. The latter, called "the strategy method", is sometimes employed in experimental designs in order to obtain more behavioral data (Selten, 1967). The modal finding in these comparisons is that behavior is significantly different when games are played sequentially or simultaneously, but generally the differences are small.[3] Several possible explanations have been proposed for these differences, although the phenomenon remains poorly understood.[4]

Our study also relates to the winner's curse problem in common value auctions (reviewed in Kagel and Levin, 2002) and adverse selection in lemons markets (Samuelson and Bazerman, 1985).[5] Under some conditions, players do not seem to fully realize that the choices of other players in the game depend on their information, but these distortions diminish significantly with experience.

[3]See, for example, Schotter et al. (1994), Coughlan et al. (1999), Brosig et al. (2003), and Oxoby and McLeish (2004).

[4]McKelvey and Palfrey (1998) show that quantal response equilibrium behavior will generally produce some differences across "equivalent" game forms.

[5]Although related, these two environments are actually quite different. In particular, note that the first one is a simultaneous game with multi-sided private information whereas the second one is a sequential game with one-sided private information.

2 The model

The trading game can be formalized as follows. An asset is to be divided among two agents, 1 and 2. Agent 1, the seller, possesses the asset. Agent 2, the buyer, can acquire it if they mutually agree on a price. The asset has a common value to both agents, and each has a signal, denoted by s and b for the seller and buyer, respectively. The common value $v(s, b)$ is a commonly known function of the signals, s and b. There are many possible bargaining mechanisms that might apply is these environments. The simplest trade mechanism is one in which the seller sets a take-it-or-leave-it price, which is accepted or rejected by the buyer. A natural alternative, which is strategically equivalent, is a seller-price double-auction, where seller and buyer simultaneously quote price and bid, and the transaction is executed at the seller's price if and only if the bid weakly exceeds the price. In these mechanisms the total surplus $v(s, b)$ is fixed but the splitting rule (trading price) is endogenously determined. We assume s and b are private information for seller and buyer, respectively. More precisely, $s \in \mathcal{S}$ and $b \in \mathcal{B}$ with commonly known c.d.f. $F_s(s \mid b)$ and $F_b(b \mid s)$, possibly different and possibly correlated. We assume strictly positive densities $f_s(s \mid b)$ and $f_b(b \mid s)$ for all s and b. We also restrict attention to monotone value functions, i.e., $\partial v(s, b)/\partial s \geq 0$ for all b and $\partial v(s, b)/\partial b \geq 0$ for all s. Last, we assume that the utility of the seller, $u_s(x)$, and the utility of the buyer, $u_b(y)$, are strictly increasing in their own payoff x and y, that is, $u'_s(x) > 0$ and $u'_b(y) > 0$. These utility functions are not necessarily the same. Moreover, we allow for risk-averse and risk-loving utilities ($u''_s \gtrless 0$ and $u''_b \gtrless 0$).

This class of environments does *not* satisfy the conditions for no-trade described in Milgrom and Stokey (1982). In particular, the initial allocation is not Pareto optimal if the seller is more risk-averse than the buyer for all relevant levels of wealth. Therefore, we cannot apply existing no-trade theorems. Nevertheless, a no trade property can be proved, summarized in the proposition below.

Proposition. *In equilibrium, there can never be trade.*

Proof. Assume that, in case of indifference, agent 2 does not trade. Suppose there exists a price p such that for all $s \in S \subseteqq \mathcal{S}$ agent 1 offers the good at price p and for all $b \in B(S) \subseteqq \mathcal{B}$ agent 2 accepts to trade at that price. Let $\overline{s} = \max_{s \in S}$ and $\underline{b} = \min_{b \in B(S)}$. Agent 2 accepting p implies that:

$$\int_{s \in S} u_b(v(s,b)-p) dF_s(s \mid s \in S) > u_b(0) \;\; \forall\, b \in B(S) \;\; \Rightarrow \;\; u_b(v(\overline{s}, \underline{b})-p) > u_b(0)$$

Similarly, agent 1 offering p implies that:

$$u_s(p) \geqslant \int_{b \in B(S)} u_s(v(s,b)) dF_b(b \,|\, b \in B(S)) \;\; \forall\, s \in S \;\; \Rightarrow \;\; u_s(p) \geqslant u_s(v(\overline{s}, \underline{b}))$$

which contradicts the previous strict inequality. If, in case of indifference, agent 2 accepts, we might observe trade but only in trivial cases (e.g., $p = k$ and $v(s, b) = k$ for all (s, b)) or in probability-zero events (agent 1 with signal $s^* = \min_s \mathcal{S}$ sets price $p = v(s^*, b^*)$ which is accepted only by an agent 2 with signal $b^* = \max_b \mathcal{B}$). A similar proofs extends to other type of bargaining mechanisms. □

The intuition is straightforward. At the stage where individuals can trade, each agent has incomplete information, but there is no residual uncertainty about the value of the asset (formally, $v(\cdot)$ is a deterministic function of s and b). Therefore, trading for insurance or risk-sharing motives is not an option, despite the possible differences in the agents' risk-tolerance. Because of their different private information, agents will hold different beliefs about the value of the asset. This could, in principle, generate trade. However, the proposition shows that rational agents will not agree to trade *on the basis of private information alone*. The reason is the same as in standard no-trade theorems. Indeed, any deal beneficial for one player must necessarily be hurtful for the other. Both agents anticipate this simple fact and form expectations accordingly. As a result, one agent must always disapprove the terms of the trade. Because trade only occurs under mutual agreement, this is enough to break any deal.

The simplicity of the argument makes it also very robust: as long as we keep the deterministic and common value nature of the asset, extending the game in other dimensions will not change the no-trade outcome. In particular, allowing counter-offers, divisibility of the asset or more sophisticated trading mechanisms will not induce agents to trade. By contrast, it is also easy to see why the absence of residual uncertainty on the asset's value is important. Indeed, if this was not the case, incentives to trade for insurance or risk-sharing motives may be present after the revelation of information and could outweigh the adverse selection problem.[6]

[6]To grasp the intuition, imagine the limit situation where s and b provides almost no information about the value of the asset. The adverse selection effect would be minimal so if the seller were more risk-averse than the buyer, they would both gain from trading.

3 Laboratory experiments

3.1 Implementation of the game

We specialize the environment for the laboratory in the following ways. First, the private information signals, s and b, are independent draws from identical, uniform distributions. We obtain data for both the take-it-or-leave-it pricing mechanism (or "price") and the seller-price double auction (or "auction"). When there is no trade, the payoff of the buyer is 0 and the payoff of the seller is $v(s, b)$ under either mechanism. When there is trade, the payoff of the buyer is the value of the asset minus the price paid, $v(s, b) - p$, and the payoff of the seller is the price obtained, p.

Under both mechanisms the action of the buyer determines only whether there is trade, so seller behavior should be the same in equilibrium in both mechanisms. Furthermore, the buyer bidding behavior in the auction mechanism should be isomorphic to their acceptance strategy in the price mechanism. Hence the only real difference between the two mechanisms lies in the timing: sequential (price mechanism) vs. simultaneous (auction mechanism).

For the asset value function, we obtain data for three cases: average of signals ($v(s, b) = \frac{s+b}{2}$), minimum of signals ($v(s, b) = \min\{s, b\}$), and maximum of signals ($v(s, b) = \max\{s, b\}$).

3.2 Experimental design and procedures

We conduced 7 sessions with a total of 86 subjects. The subjects were registered Princeton students who were recruited by email solicitation, and all sessions were conducted at The Princeton Laboratory for Experimental Social Science. All interaction between subjects was computerized, using an extension of the open source software package, Multistage Games.[7] No subject participated in more than one session. In each session, subjects made decisions over 20 rounds. Each subject played exactly one game with one opponent in each round, with random rematching after each round.

At the beginning of each round, each subject was randomly assigned a role as either seller or buyer, and assigned a new signal, s or b. Signals were integer numbers drawn independently with replacement from a uniform distribution over $[0, 100]$. Each subject observed his own signal, but did not observe the opponent's signal. The distribution was common knowledge.

[7]Documentation and instructions for downloading the software can be found at http://multistage.ssel.caltech.edu.

The common value was computed as a deterministic function of the two signals, using either the average, minimum, or maximum. The value function was held constant within a session.[8]

In the price variant, the seller offered the asset for a price, p, which was limited to integer numbers in the range of possible values of the asset, $[0, 100]$. The buyer then decided whether to accept or reject the offer, and payoffs for that round accrued accordingly. In the auction variant, buyer and seller simultaneously quoted bid and ask prices also limited to integer numbers in the range of possible values of the good. Trade occurred at the seller's price if and only if the bid weakly exceeded the ask price. In either case, players learned at the end of each round the signal and decision of their opponent. Finally, subject computer screens included a table with the history of behavior, signals, and outcomes in previous rounds.

At the beginning of each session, instructions were read by the experimenter standing on a stage in the front of the experiment room, which fully explained the rules, information structure, and client GUI.[9] After the instructions were finished, two practice rounds were conducted, for which subjects received no payment. After the practice rounds, there was an interactive computerized comprehension quiz that all subjects had to answer correctly before proceeding to the paid rounds. The subjects then participated in 20 paid rounds, with opponents, roles (seller or buyer), and signals randomly reassigned at the beginning of each round. The trading mechanism and the common value function were held constant throughout all rounds of a session. Subjects were paid the sum of their earnings over the 20 paid rounds, in cash, in private, immediately following the session.

4 Results

4.1 Aggregate behavior and payoffs

The first cut at the data consists of comparing the prices, likelihood of trade and realized gains of buyers and sellers in the different treatments, without conditioning on the actual draws of s and b. Table 2 shows average choices for the ave, min and max value functions under the price and auction mechanisms.

---

[8]The average treatments were framed as the sum of the two signals, rather than the average, to make the instructions simpler. This only results in a rescaling of all strategies and payoffs.

[9]A sample copy of the instructions and sample subject GUI screens can be downloaded from http://www.hss.caltech.edu/~trp/notrade.

Value	ave				min				max			
Mechanism	price		auction		price		auction		price		auction	
# observations	(260)		(120)		(120)		(120)		(120)		(120)	
Average seller price	61.5	(19.7)	62.5	(21.3)	47.4	(20.4)	53.0	(20.2)	78.8	(17.7)	74.2	(17.2)
Average buyer bid	—		41.6	(14.9)	—		24.6	(20.3)	—		53.5	(22.5)
Frequency of trade (%)	31.9		26.7		24.2		16.7		30.8		21.7	
Seller gain given trade	-4.6	(14.5)	-4.2	(13.1)	-5.4	(19.7)	7.7	(12.2)	-7.8	(23.3)	-12.8	(21.9)
Seller gain	-1.5**	(8.4)	-1.1*	(7.0)	-1.3	(9.8)	1.3**	(5.6)	-2.4**	(13.3)	-2.8**	(11.3)
Gain if all traded	13.9	(22.3)	9.8	(21.6)	15.1	(25.8)	19.4	(22.3)	12.0	(25.8)	8.1	(22.6)

Table 2. Average choices and trade probabilities. Standard deviations in parenthesis. ** = significant at 5%; * = significant at 10%.

Result 1 *There is substantial trade in all treatments.*

In approximately half of our observations, the buyer's signal is below the seller's signal. These are situations with essentially no chance for trade even with completely naïve behavior, implying a natural upper bound of 50% on the amount of trade. Yet we observe trade between 16.7% to 31.9% of the time, depending on the treatment. Thus, trade occurs roughly between one-third and two-thirds of the time when $b > s$.

Result 2 *On average, sellers lose and buyers gain from trade.*[10]

Sellers' offer prices would, on average, earn them non-negligible profits *if* buyer acceptance decisions were uncorrelated with buyer signals. However, since buyers condition their decision on their information, and are more likely to accept when their signal is higher, sellers end up incurring net losses in 5 out of the 6 treatments. As an immediate consequence, and abstracting from endowment considerations, it would be preferable in this game to be a buyer than a seller.[11] This is especially surprising in the ave and min cases. Indeed, in those two treatments, a rational and cautious seller has at his disposal a strategy to induce a boundedly rational buyer to trade and at the same time guarantee no losses, by setting prices $p = \frac{s}{2} + 50$ and $p = s$, in the ave and min cases respectively. In fact, given the behavior of

[10] More accurately, traders lose money on average in the role of sellers, and gain money on average in the role of buyers (recall that each individual trader was a seller approximately half the time and a buyer the rest of the time).

[11] Since the seller is, by assumption, endowed with the good, his final payoff is greater than that of the buyer if there is no trade ($v(s, b)$ vs. 0). We define profit in net terms, so the comparison is between the incremental utility of buyers and sellers, which is zero in case of no trade.

buyers this would actually generate positive profits. Thus, sellers are clearly not maximizing payoffs, given the behavior of buyers. It is, of course, more difficult to evaluate whether buyers are doing as well as they can given seller behavior because the buyers are making some money.

To understand this better, we look at gains and losses as a function of the realized private information. The left column in Figure 1 displays for the price treatments, the potential –positive or negative– net gain of sellers (price minus value of the asset) as a function of the seller's signal, where each dot is one observation. It also shows whether the terms of the trade were accepted and thus the net gains realized (light circle) or not (dark triangle). The right column in Figure 1 displays the same information from the buyer's viewpoint for the auction treatments: the net loss of buyers, which is equivalent to the net gain of sellers, as a function of the buyer's signal. The treatments that are not graphically represented follow similar patterns.

[Figure 1 here]

The figures in the left column clearly illustrate the adverse selection effect. Although the expected gains would be positive if the behavior of buyers were uncorrelated with their information, they are generally negative once we condition on the buyers' actual decisions, that is, when we look only at the light circle dots. More interestingly, we notice that the biggest losses occur when the signal of sellers are high in the min treatment and low in the max treatment. Indeed, these are the cases where the dispersion of prices is highest and therefore the selective acceptance of buyers has the largest impact on payoffs.

The picture for buyers is different. Buyers rarely trade when their signal is below 50 and, if they do, they generally incur losses. Also, in the ave and max treatments under the price mechanism, buyers cannot make losses by trading whenever $p < \frac{b}{2}$ and $p < b$, respectively. In the data, 12% of the trades accepted in the ave treatment and 73% of the trades accepted in the max treatment satisfy this inequality. This explains why the gains of buyers are higher in the max than in the other two value treatments. The patterns are similar in the auction treatments as depicted in Figure 1.

Result 3 *Behavior differs across asset value treatments: asking prices and bids increase from min to ave and from ave to max.*

The seller's price and the buyer's bid all increase from min to ave and from ave to max, as expected. It is easy to see that the expected value of the asset conditional on an agent's signal also increases from min to ave and from ave to max. This suggests that players exhibit some level of rationality with respect to the asset value function. Note also that the variance is important, mainly because choices are greatly affected by signals. As a result the average differences are sizeable but not statistically significant. The average differences between bid and ask prices are of similar order across auction treatments (between 20.7 and 28.4). Also, the greatest losses for the seller are incurred in the max treatment. This is somewhat expected since it is the most difficult problem for sellers: the only way to ensure no losses is to set $p = 100$. It is also the easiest problem for buyers: by trading whenever $p < b$, buyers are sure to get positive profits.

Result 4 *Behavior differs across mechanisms: trading is consistently less frequent under the auction than under the price mechanisms.*

The difference in trade frequency between the two mechanisms is substantial. Trade under the auction mechanism is 16% less than under the price mechanism in the ave treatment; 31% less in the min treatment; and 30% less in the max treatment. In the ave and max treatments, sellers set roughly the same average prices in the two mechanisms. However, buyers are more cautious in the auction mechanism. That is, their acceptance decisions in the price mechanism are, on average, "as if" they were bidding more aggressively than we observe them doing in the auction mechanism. This behavior does not result in higher profits for buyers, possibly because sellers are equally exploited when buyers accept to trade. By contrast, sellers in the min treatment increase prices in the auction mechanism relative to the price mechanism, and buyer behavior does not change. This (barely) reverses the sign of average seller gains for the min treatment from negative in the price mechanism to positive in the auction mechanism.

4.2 Aggregate behavior and payoffs conditional on signals

4.2.1 Strategies of sellers

The picture presented so far is useful, but incomplete since it aggregates across subjects' private information. If subjects condition their decisions on their private information, then such an analysis has left out an important component of behavior.

The behavior of sellers can thus better be described as a mapping from signal to price. We can graphically display the empirical strategies of sellers and compare them across treatments. Figure 2 displays for each of the

6 treatments (ave-min-max and price-auction variants) the sellers' asking price as a function of their signal. Each dot in the graph is one observation. Figure 2 also identifies cases where the prices resulted in a trade.

[FIGURE 2 HERE]

As a natural benchmark for studying seller behavior, we use the Nash equilibrium price correspondence. That equilibrium differs across value treatments, but is the same for both mechanisms. The equilibrium price correspondences are:

$$\begin{cases} p_e(s) \in [\frac{s}{2} + 50, 100] & \text{(ave)} \\ p_e(s) \in [s, 100] & \text{(min)} \\ p_e(s) = 100 & \text{(max)} \end{cases}$$

Note that equilibrium is characterized by a range of equilibrium prices for the ave and min cases, and is a fixed constant $p_e = 100$ in the max treatment.

Result 5 *Seller pricing strategies are consistently below the Nash equilibrium.*

Seller pricing behavior coincides with Nash equilibrium play rather infrequently, particularly in view of the wide range of Nash equilibrium prices in the ave and min treatments. In those two cases, prices are in the Nash equilibrium range 26% and 57% of the time, respectively. In the max case, only 8% of the observed prices are at the Nash equilibrium ($p = 100$). All other prices are too low. Pooling across the three value treatments, sellers set prices below Nash equilibrium about 70% of the time. The lower envelope of these Nash equilibria are $p = \frac{s}{2} + 50$, $p = s$, and $p = 100$ in the ave, min, and max treatments, respectively. In equilibrium, these prices should yield no profit to the seller simply because they are too high to induce buyers to trade. At the same time, they are high enough to guarantee profits if a buyer (out of equilibrium) accepts. However, even zero or very low profits would be an improvement over the losses sellers are incurring from the lower prices they set in the experiment.

We also ran a simple OLS regression of seller price as a function of seller signal and a constant term. The results are compiled in Table 3.

Value	Mechanism	Constant	Seller signal	adjusted R^2
ave	price	42.46 (2.19)	.386 (.039)	.272
min	price	29.93 (3.24)	.356 (.057)	.241
max	price	59.61 (2.84)	.366 (.048)	.326
ave	auction	39.85 (3.06)	.433 (.051)	.376
min	auction	23.63 (2.31)	.558 (.038)	.638
max	auction	52.82 (2.53)	.419 (.044)	.432

Table 3. Seller price regression. Standard Errors in parenthesis.

Result 6 *Seller prices: (1) are increasing in their signal; (2) vary across value functions; and (3) are more responsive to signals in the auction than in the price treatment.*

Coefficients on the seller's signal are highly positive and statistically significant at the 99% level in every treatment. A one unit increase in s translates into a .35 to .56 increase in p, depending on the treatment. Second, the magnitudes of the constant term show that, as might be expected, prices are typically highest in the max treatment and lowest in the min treatment. Third, the response is greater in the auction than in the price treatments (lower intercept and higher slope for all three value treatments). This leads to prices that are, on average higher in the auction treatment, especially so for high seller signals. This effect is especially strong in the min treatment. Fourth, seller signal explains more of the variance in prices in the auction treatment than in the price treatment. That is, the prices in the auction treatment fit more tightly along the price-signal regression line.

The existence of these systematic differences suggests that sellers make pricing decisions differently under the auction and price mechanisms, especially in the min treatment. This is in spite of the fact that, theoretically, the trading mechanism should not affect their strategy. This must happen then because sellers have different expectations about how buyers are behaving in the price and auction treatments. As we will see in the next subsection, this is indeed the case.

4.2.2 Strategies of buyers

We now turn to study the behavior of buyers. Their behavior can best be described as a mapping from the pair (buyer's signal, seller's price) to a probability of accepting the terms of trade in the price treatment, and as a mapping from buyer's signal to bid in the auction treatment. We can graphically display the empirical strategies of buyers and compare them across treatments.

[Figure 3 here]

The left column in Figure 3 displays the accept/reject decision of buyers as a function of their signal and the seller's asking price, in all three price treatments. The right column displays the buyers' bid as a function of their signal in all three auction treatments. It also displays whether the bid resulted in trade or not. Finally and for comparative purposes with the price treatment, the center column displays whether trade occurred in the auction treatment at the seller's asking price given the buyer signal.

In order to understand the behavior of buyers, we ran a probit regression in the price treatments to compute the buyer's probability of accepting the trade as a function of the seller's price and the buyer's signal. For the auction treatments, we ran an OLS regression of the buyer's bid as a function of his signal. The results are compiled in Tables 4a and 4b.

Value	Mechanism	Constant	Seller price	Buyer signal	pseudo R^2
ave	price	.747 (.396)	-.058 (.0082)	.041 (.0052)	.462
min	price	-.366 (.453)	-.037 (.0094)	.023 (.0056)	.283
max	price	.673 (.692)	-.036 (.0089)	.029 (.0058)	.357

Table 4a. Dependent Variable: Buyer's acceptance decision (price mechanism).

Value	Mechanism	Constant	Buyer signal	adjusted R^2
ave	auction	23.51 (2.34)	.341 (.0392)	.386
min	auction	11.25 (3.25)	.270 (.0563)	.156
max	auction	32.09 (3.53)	.426 (.0613)	.284

Table 4b. Dependent Variable: Buyer's bid (auction mechanism).

Result 7 *Buyer decisions (acceptance or bids) are increasing in their signal.*

All six buyer signal coefficients are positive and significant. Responsiveness to buyer signals is least in the case of the min treatment, for both the price and auction mechanisms. Buyer behavior is also considerably more noisy in the min treatment (lower R^2).

To understand the choices of buyers at a deeper level we perform the following analysis. Consider a model of buyer behavior where there exists a linear *Acceptance Threshold Function (ATF)* $\psi(b)$ such that a buyer with signal b agrees to trade if and only if the asking price is $p < \psi(b)$. For any hypothetical ATF, we use our data to construct a "misclassification score" or **MS** for that function, for each value treatment and each mechanism. This is done by adding up the number of misclassified observations (trade when $p > \psi(b)$ or no trade when $p < \psi(b)$) weighted by the magnitude of the misclassification (that is, the absolute difference between the actual price and the cutoff price such that the observation would not be misclassified) divided by the total number of observations. Table 5 reports the estimated ATF, $\widehat{\psi}(b)$, that minimizes the misclassification score. We also report the minimized **MS**. This value reflects the average amount by which observations are misclassified, with each correctly classified observation taking value 0. Last, we determine the percentage of observations that are misclassified by $\widehat{\psi}(b)$, which we call **MO**. Graphically, $\widehat{\psi}(b)$ corresponds to the best empirical dividing line between trade vs. no trade regions.

For the price treatments, this analysis involves using all the available information (buyers observe their signal and the ask price and decide whether to trade or not). In order to construct a comparable measure for the auction treatments, the only information used is whether trade occurred at the asking price or not, rather than incorporating the additional information in the buyer's bid.

Value	Treatment	$\widehat{\psi}(b)$	**MS**	**MO**
ave	price	$22.8 + 0.55\,b$	1.33	16.9%
ave	auction	$34.4 + 0.20\,b$	0.54	11.7%
min	price	$15.8 + 0.31\,b$	2.14	20.0%
min	auction	$19.2 + 0.26\,b$	1.30	16.7%
max	price	$39.4 + 0.56\,b$	2.00	18.3%
max	auction	$40.7 + 0.31\,b$	2.27	20.0%

Table 5. Linear ATF estimation results.

Result 8 $\widehat{\psi}(b)$ *is steeper in the price treatments than in the auction treatments, which results in more trade.*

For all three value treatments, the estimated classification line has a higher slope and lower constant term in the price treatment than in the auction treatment. In other words, buyers with high signals act in a more conservative way with the auction than with the price mechanism.[12] Because trade rarely occurs when buyers have low signals, this behavior tends to reduce trade. The result, combined with our previous findings about sellers' behavior, suggests that the reasons for a consistently lower likelihood of trade in the auction variant is due to a different behavior by *both* buyers and sellers. It is also worth noting that the reason for lower acceptance rates comes exclusively from a lower sensitivity of bids to own signals (coefficient on b). Finally, the linear misclassification function $\widehat{\psi}(b)$ performs quite well across all treatments and mechanisms, with a range of 80% to 88% of observations correctly classified.

4.2.3 Trading probabilities

Our next look at the data consists in describing the relation between the buyer-seller signal *combinations* and the likelihood of trade. Figure 4 plots for each treatment and each (s, b) pair whether the outcome of the game is trade (light circle) or no-trade (dark triangle).

[FIGURE 4 HERE]

Due to the deterministic and pure common value nature of the asset, the region where trade should occur consists only of the (0,1) pair. As we already know, this is not what is observed. Generally trade occurs when the seller's valuation (or signal) is sufficiently low and the buyer's valuation (or signal) is sufficiently high. The empirical likelihood of trade depending on whether the buyer's signal exceeds the seller's signal or not is reported in Table 6.

Value	Treatment	% trade given $b < s$	% trade given $b \geqslant s$
ave	price	10.6	57.6
min	price	15.3	32.8
max	price	10.9	53.6
ave	auction	3.3	50.0
min	auction	3.1	32.1
max	auction	10.2	32.8

Table 6. Likelihood of trade.

[12] This result is similar to the finding in the compromise game (Carrillo and Palfrey, 2006). In that experiment, subjects were less likely to agree to a compromise when they acted as second movers than when the game was played simultaneously, even though the Nash equilibrium prediction was identical in both variants. That paper provided an explanation based on Quantal Response Equilibrium, although they might be some others.

Result 9 *Trade rarely occurs when the seller's signal exceeds the buyer's signal. The probability of trade is increasing in the buyer's signal and decreasing in the seller's signal.*

In three out six treatments, individuals engage in trade *more than half of the time* whenever the buyer's signal exceeds the seller's signal. This is particularly striking given that the no-trade theoretical prediction does not dependent on the risk tolerance of individuals. In other words, since all that matters for our theory is that utility is increasing in the subject's monetary payoff, risk-aversion, disappointment aversion or kinks in the utility function could not account, even partially, for the observed outcomes.

We then ran a simple probit regression of the likelihood of trade as a function of the seller's and buyer's signal. The results are reported in Table 7.

Value	Mechanism	Constant	Seller signal	Buyer signal	pseudo R^2
ave	price	-1.306 (.247)	-.013* (.0036)	.029* (.0037)	.257
min	price	-1.601 (.378)	-.003 (.0047)	.019* (.0049)	.130
max	price	-1.133 (.376)	-.023* (.0057)	.033* (.0059)	.351
ave	auction	-0.051 (.450)	-.045* (.0084)	.022* (.0065)	.470
min	auction	-0.834 (.404)	-.021* (.0060)	.014* (.0052)	.211
max	auction	-1.448 (.399)	-.007 (.0052)	.018* (.0050)	.131

Table 7. Probability of trade as a function of signals; * = significant at 5%.

All slope coefficients have the expected sign, and ten out of twelve are significant at the 5% level. Trade depends more on the buyer signal than the seller signal in all the price treatments, while the reverse is true in two out of three auction treatments. However and with one exception (ave - auction), the R^2 are rather low, which suggests that a probit regression is probably not the most appropriate method for the purpose of our analysis.

To look at the relationship between buyer and seller signals more closely, we conduct a classification analysis similar to section 4.2.2. Consider a linear function $\phi(s)$ with the property that trade occurs if the pair of signals (s, b) is such that $b > \phi(s)$. As in the estimation of ATFs, for any $\phi(s)$ we empirically determine the number of misclassified observations (trade when $b < \phi(s)$ or no trade when $b > \phi(s)$) weighted by the magnitude of the misclassification (that is, the absolute difference between the actual signal of the buyer and the cutoff signal such that the observation would not be misclassified). This value divided by the total number of observations is called the misclassification score or **MS**. For each treatment, we report the estimated function, $\widehat{\phi}(s)$, that minimizes the misclassification score. We also report the percentage of misclassified observations or **MO**. Graphically, $\widehat{\phi}$ corresponds to the best dividing line between the trade and no trade regions in the (b, s) signal space. The results of the estimated functions are presented in Table 8 and included in the graphs of Figure 4.

Value (# obs.)	Treatment	$\widehat{\phi}(s)$	**MS**	**MO**
ave (260)	price	$42.3 + 0.40\,s$	3.82	23.5%
ave (120)	auction	$30.1 + 0.87\,s$	3.04	15.0%
min (120)	price	$73.7 + 0.04\,s$	5.25	31.7%
min (120)	auction	$71.1 + 0.32\,s$	3.91	20.8%
max (120)	price	$34.6 + 0.62\,s$	2.63	13.3%
max (120)	auction	$71.9 + 0.09\,s$	4.05	21.7%

Table 8. Trade vs. no-trade division.

Result 10 $\widehat{\phi}(s)$ *is an increasing function.*

The slope of the classification function is positive in all six treatments. Sellers with higher signals set higher prices for the asset, thus decreasing the likelihood of a trade. Conversely, buyers with higher signals set higher bids and are also more likely to accept a given trade. Overall, the model correctly classifies about 80% of the trade outcomes, although the slope and accuracy of classification differ substantially across both value treatments and mechanisms. The differences between the estimated functions in the price and auction treatments reinforce the argument we made previously about the impact that the trading mechanism has on the strategies selected by subjects.

4.3 Learning

A natural question to ask is whether individuals adapt their strategies over the course of a session. Clearly, the behavior is out of equilibrium, and subjects are given considerable feedback, in both roles, with 20 repetitions of the game. An adaptive player could recognize that his or her losses in the role of seller are due to the adverse selection problem and increase the price accordingly. The evolution in the response of subjects in the role of buyers is less obvious, since buyers are generally doing quite well for themselves. A simple first cut to investigate learning consists in breaking the data down into early and late plays. In each session, there were 20 rounds of play. We code the choices in the first 10 rounds as "inexperienced" and the choices in the last 10 rounds as "experienced". Table 9 presents the average choices in all six treatments broken down by experience level.

Treatment	Round	Seller price	Buyer bid	% trade	Seller gain
ave – price	inexp.	62.2 (20.61)	—	33.1	-3.8 (14.86)
	exp.	60.8 (18.71)	—	30.8	-5.4 (14.26)
ave – auction	inexp.	60.3 (21.57)	41.5 (15.11)	31.7	-4.9 (15.33)
	exp.	64.6 (20.95)	41.7 (14.85)	21.7	-3.1 (9.52)
min – price	inexp.	47.8 (20.66)	—	28.3	-7.5 (19.35)
	exp.	47.0 (20.26)	—	20.0	-2.3 (20.65)
min – auction	inexp.	55.40 (19.5)	25.8 (20.02)	15.0	5.0 (14.65)
	exp.	50.5 (20.70)	23.4 (20.62)	18.3	9.8 (9.87)
max – price	inexp.	76.1 (17.00)	—	36.7	-0.4 (21.85)
	exp.	81.5 (18.18)	—	25.0	-18.6 (22.03)
max – auction	inexp.	74.5 (16.99)	50.3 (21.87)	16.7	-24.0 (13.52)
	exp.	73.9 (17.55)	56.7 (22.77)	26.7	-5.8 (23.51)

Table 9. Average choices of sellers and buyers by level of experience.

Result 11 *There is no clear evidence of learning by either buyers or sellers.*

There is little evidence of systematic changes in the average behavior of sellers and buyers between early and late rounds. Sellers increase prices in two treatments, decrease in one and keep them roughly constant in the other three. Buyers' acceptance rate decreases in all the price treatments but their bids increase in two of the auction treatments. This is consistent with the findings of Carrillo and Palfrey (2006) in a related two-sided game of incomplete information.[13]

The absence or near absence of learning trends occurs in spite of substantial feedback after each round of play, as well as experience in both roles. For example, the buyer knows the price asked by the seller and, at the end of each round, learns the seller's signal. Therefore, in principle, subjects can partially reconstruct an average price function of sellers. The same applies when the subject is in the role of a seller, who learns the buyer's bid and signal (in the auction treatments) or the acceptance decision and signal (in the price treatments). It appears, however, that this information does not lead

[13] Little or no learning has also been emphasized in experiments on common value auctions and adverse selection (see e.g. Kagel and Levin (2002)).

to changes in individual behavior sufficiently important to produce trends at the aggregate level.

To explore this issue in more detail, we determine whether the behavior of buyers and sellers as a function of their own signal is different at the beginning than at the end of the experiment. Again, we divide the sample into early play (first 10 rounds) and late play (last 10 rounds). We then perform a maximum likelihood estimation in each subsample and in the full sample. For all subjects in the auction treatment and sellers in the price treatment, we run a linear regression of price (seller) or bid (buyer) on own signal and constant term, for the two experience levels separately, and compare it to the results from the pooled regression. For the case of buyers in the price treatment, we instead perform a probit estimation, and control for the seller's offer price and the buyer's signal. We then conduct a likelihood ratio test to determine whether differences in choices between early and late rounds are statistically significant. The findings are summarized in Table 10.

Player	Treatment		Likelihood estimation			χ^2-test
	value	mechanism	constrained	unconstrained	d.f.	
seller	ave	price	-1280.56	-1281.19	2	1.27
seller	min	price	-513.69	-514.45	2	1.52
seller	max	price	-488.17	-490.64	2	4.94
seller	ave	auction	-590.49	-591.09	2	1.20
seller	min	auction	-465.50	-468.93	2	6.86 *
seller	max	auction	-476.35	-476.79	2	0.88
buyer	ave	price	-86.24	-87.55	3	2.61
buyer	min	price	-44.42	-47.61	3	6.38
buyer	max	price	-43.76	-47.70	3	7.88 *
buyer	ave	auction	-543.96	-547.46	2	7.00 *
buyer	min	auction	-520.08	-520.18	2	0.19
buyer	max	auction	-520.59	-522.63	2	4.08

Table 10. Effect of experience on prices, bids and acceptance rates; * = significant at 5%.

In only 3 out of 12 treatments differences are statistically significant at the 5% level, and none are significant at the 1% level. Furthermore, in one of these three treatments (buyer's ave - auction), the change in the buyers' strategy results in lower average profits for them. Again, this reinforces the idea that subjects do not change significantly their strategy over time.

5 A behavioral theory

In this section, we consider a behavioral theory that may account for the choices of our subjects. We assume that players have an (incorrect) mutually held belief that the action of an opponent is less correlated with their information than is actually the case. This type of cognitive limitation was first discussed in Holt and Sherman (1994). Two recent theories have generalized the argument: "cursed equilibrium" (Eyster and Rabin, 2005) and "analogy based expectations" (Jehiel and Koessler, 2006).

In the extreme case, or "fully cursed", players have a mutual belief that action and information is completely uncorrelated. Applying this to our model, a fully cursed buyer in the price treatments will accept to trade if and only if the price set by the seller is less than the buyer's expected value of the asset given his own signal, $E_s[v(s,b)\,|\,b]$. Simple computations yield:

$$E_s[v(s,b)\,|\,b] = \begin{cases} 25 + b/2 & \text{(ave)} \\ b - b^2/200 & \text{(min)} \\ 50 + b^2/200 & \text{(max)} \end{cases} \tag{1}$$

The decision problem for sellers is slightly more complex. A fully cursed seller in the price treatments anticipates correctly how the buyer's probability of acceptance will depend on the offer price, p. However, the seller believes the acceptance decision is independent of b. Formally and given a price p, a cursed seller believes that his expected payoff from setting price at p, given signal s, is:

$$\Pi(p\,|\,s) = \Pr(E_s[v(s,b)\,|\,b] > p)\,p + \Pr(E_s[v(s,b)\,|\,b] < p)E_b[v(s,b)\,|\,s]$$

Denote by $p^*(s) = \arg\max_p \Pi(p\,|\,s)$, the optimal price of a fully cursed seller. After some algebra, we get:

$$p^*(s) = \begin{cases} 50 + \frac{s}{4} & \text{(ave)} \\ \frac{100}{3} + \frac{1}{3}s - \frac{1}{600}s^2 & \text{(min)} \\ \frac{1100}{18} + \frac{1}{600}s^2 + \sqrt{\frac{10000+3s^2}{81}} & \text{(max)} \end{cases}$$

Having determined the theoretical choices of cursed individuals in the price treatments, we can now compare them with the data. For the analysis of buyers, we follow the classification method employed in Table 5 of section 4.2.2. Note that the equations in (1) correspond to nonlinear theoretical ATFs for the cursed equilibrium model. We therefore consider the best quadratic (rather than linear) ATF, to make it comparable to the cursed prediction. The performance of the cursed and empirical quadratic ATFs of buyers are described in Table 11 and graphically represented in the left column of Figure 3.

Value	Strategy	$\widehat{\psi}(b)$	**M.S.**	**%M.O.**
ave	cursed	$25 + 0.5\,b$	1.33	16.9%
	empirical	$33.6 + .02\,b + .005\,b^2$	1.29	16.9%
min	cursed	$b - .005\,b^2$	2.02	20.0%
	empirical	$-14.8 + 1.86\,b - .014\,b^2$	1.60	16.7%
max	cursed	$50 + .005\,b^2$	1.88	12.5%
	empirical	$64.0 - .41\,b + .008\,b^2$	1.84	15.0%

Table 11. Classification of buyers' acceptance decision.

Result 12 *The cursed equilibrium model classifies buyer acceptance decisions as well as the best fitting quadratic ATF and better than the best linear ATF.*

Based on misclassification analysis, the dividing line for the cursed model is remarkably accurate in all price treatments. In the ave and max treatments, the cursed and empirical strategies are virtually identical in terms of the misclassification score (for the max treatment fewer observations are misclassified with the cursed function). In the min treatment, the difference in performance is more significant, but this may be due to a limited number of observations. In fact, according to the empirical strategy, the likelihood of acceptance is decreasing in the buyer's signal for all $b > 66.4$. This is the result of a few buyers with high signals of 75 and above who play the Nash equilibrium, and therefore refuse to trade even when the asking price is low (see Figure 3). When comparing with Table 5, it is also remarkable that the cursed quadratic functions perform better than the best linear fits in both the min and max treatments. Also, although the number of misclassified observations is non-negligible (up to 20%), in more than 50% of the cases, the price is within 10 units of the correctly classified value.

The cursed equilibrium strategy of sellers can also be compared to its empirical counterpart. In Table 12, we report a quadratic OLS regression of the seller's price as a function of the signal. Both the theoretical cursed function, $p^*(s)$, and the empirical quadratic estimates reported below are graphically represented in the left column of Figure 2.

Value	Mechanism	Constant	s	s^2	adjusted R^2
ave	price	41.50 (3.38)	.440 (.150)	-.001 (.001)	.270
min	price	26.33 (4.75)	.574 (.219)	-.002 (.002)	.241
max	price	66.35 (4.42)	-.001 (.192)	.004 (.002)	.342

Table 12. Seller's quadratic OLS.

Result 13 *The cursed equilibrium model implies seller pricing functions similar to what is observed in the data.*

The theoretical cursed pricing functions predict that the constant terms should be ordered max > ave > min and the linear coefficients should be ordered min > ave > max. The quadratic coefficient is predicted to be 0 in the ave treatment, small and negative in the min treatment and small and positive in the max treatment. This is the pattern we find in Table 12, with the exception that the quadratic coefficient for the min treatment is not significantly different from 0. In general, the empirical function is slightly steeper than the cursed prediction in every treatment but the overall shape is quite similar (Figure 2, left column).

Finally, we can compare the likelihood of trade and seller profits in our data with the predictions of the fully cursed model. For the ave treatments, these can easily be obtained analytically, whereas for the min and max treatments, we rely on numerical methods. The results are presented in Table 13.

	ave		min		max	
	cursed	empirical	cursed	empirical	cursed	empirical
% trade	25.0	31.9	28.9	24.2	20.7	30.8
% trade given $b < s$	0.0	10.6	0.0	15.3	0.0	10.9
% trade given $b \geqslant s$	50.0	57.6	57.7	32.8	41.3	53.6
Average profit (seller)	0	-1.5	2.41	-1.3	-1.92	-2.4

Table 13. Cursed equilibrium: % trade and seller profits in price mechanism.

Result 14 *The cursed equilibrium model implies trade frequencies similar to what is observed in the data.*

Specifically, the theoretical predictions of trade range between 20% and 30% of the time in the price mechanism, depending on the value treatment. This compares with the observed range between 24% and 32%. Furthermore, the model predicts trade only if $b > s$. In the experiment, there was almost no trade (12%) when $b < s$.

Result 15 *The cursed equilibrium model implies an ordering of seller profits (min > ave > max) that we find in the data. However, we observe seller losses in all three treatments, while the cursed model predicts losses only in the max treatment.*

In the cursed equilibrium model, expected seller profits range between 2.4 in min and -1.9 in max, whereas the corresponding numbers in our data range between -1.5 in min and -2.4 in max. The ordering is therefore correct, but the magnitudes are not. The sellers in our price mechanism lose more money on average than the expected losses in a cursed equilibrium. We have the same ranking of seller profits in the auction data as well (min > ave > max) and, in that case, seller profits in the min treatment are actually positive. Unfortunately, we have been unable to establish a theoretical solution to the cursed equilibrium model in the auction treatments.

6 Conclusion

This study addressed the question of whether asymmetric information can induce individuals to engage in exchange in environments where trade is never mutually profitable, conditions under which such trade is more or less prevalent, and the economic consequences for buyers and sellers. Despite the compelling and general logic of no-trade theories, traders trade frequently. In particular, when the buyer's signal exceeds the seller's signal, the likelihood of trade is between 32% and 58% depending on the treatment. Buyers generally outperform sellers and the difference persists even when subjects have gained experience both in the role of buyers *and* sellers. In fact, there is surprisingly little evidence of learning in all treatments of this game despite the substantial amount of feedback provided. We have also shown that a sequential mechanism (a seller's take-it-or-leave-it price) always results in more trade than a simultaneous mechanism (a seller-price double auction) even though both are strategically equivalent. Finally, we have applied the cursed equilibrium theory to our model and shown that it explains some general patterns of the data, such as the buyer's acceptance behavior and the aggregate probabilities of trade. However, it has a more difficult time accounting for the variance in the behavior of sellers and the profits of subjects in the different roles.

The effect of the trading mechanism on outcomes is particularly surprising and deserves further investigation. We have restricted our attention to two mechanisms, seller price setting and double auction, but there are many other bargaining structures that could be considered and compared. Obtaining experimental insights on how choices depend on mechanisms that are strategically equivalent could be of interest not only to improve our understanding of bilateral trading games but also to learn how to design efficient allocation mechanisms in more general economic environments such as auction and trading markets.

This approach could also be usefully applied to study bargaining between three or more parties, as in markets and auctions. It is an interesting open question whether the tendency to trade excessively is exacerbated or attenuated in environments with multiple buyers and/or multiple sellers. For example, this could provide valuable insights about the design and performance of prediction markets.

Finally, one would like to know whether alternative models could explain better the main features of the data (substantial trade, systematic advantage of buyers, importance of the order of moves, and absence of learning). Some natural candidates would be partially cursed equilibrium, quantal response equilibrium (McKelvey and Palfrey, 1995), and theories based on levels of strategic sophistication such as cognitive hierarchy (Camerer et al., 2004). Based on our earlier study of the compromise game (Carrillo and Palfrey, 2006), these theories provide only partial explanations, and there remains much to learn about behavior and outcomes in games with two-sided private information and common values.

References

[1] Brosig, J., J. Weimann, and C.L. Yang (2003) "The Hot vs. Cold Effect in a Simple Bargaining Experiment," *Experimental Economics*, 6(1), 75-90.

[2] Blume, L., Coury, T. and D. Easley (2006) "Information, Trade and Incomplete Markets," *Economic Theory* 29, 379-394.

[3] Camerer, C., T. Ho, and J. Chong (2004) "A Cognitive Hierarchy Model of Behavior in Games," *Quarterly Journal of Economics,* 119, 861-98.

[4] Camerer, C., T. Palfrey, and B. Rogers (2006) "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies," *Working Paper,* California Institute of Technology, http://www.hss.caltech.edu/~ trp/hqre051206.pdf.

[5] Carrillo, J.D. and T. Palfrey (2006) "The Compromise Game: Two-sided Adverse Selection in the Laboratory," *CEPR D.P. 6103.*

[6] Chatterjee, K. and W. Samuelson (1983) "Bargaining under Incomplete Information," *Operations Research* 31, 835-851.

[7] Coughlan, P., R. McKelvey, and T. Palfrey (1999) "An experiment on two person games with incredible threats," *Working Paper*, California Institute of Technology.

[8] Cramton, P., Gibbons, R. and P. Klemperer (1987) "Dissolving a Partnership Efficiently," *Econometrica* 55, 615-632.

[9] Evans, R. (1989) "Sequential Bargaining with Correlated Values," *Review of Economic Studies* 56(4), 499-510.

[10] Eyster, E. and M. Rabin (2005) "Cursed Equilibrium," *Econometrica*, 73: 1623-72.

[11] Hirshleifer, J. (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *The American Economic Review* 61: 561-574.

[12] Holt, C. and R. Sherman (1994) "The Loser's Curse," *American Economic Review* 84(3), 642-652.

[13] Jehiel, P. and F. Koessler (2006), "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *mimeo*, PSE, Paris.

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

[15] McKelvey, R. and T. Palfrey (1995) "Quantal Response Equilibria in Normal Form Games," *Games and Economic Behavior*, 10, 6-38.

[16] McKelvey, R., and T. Palfrey, "Quantal Response Equilibria in Extensive Form Games," *Experimental Economics*, 1 (1998): 9-41.

[17] Milgrom, P and N. Stokey (1982) "Information, trade and Common Knowledge," *Journal of Economic Theory*, 26, 11-21.

[18] Morris, S. (1994) "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62(6), 1327-1347.

[19] Myerson, R.B. and M. Satterthwaite (1983) "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* 28, 265-281.

[20] Oxoby, R. and K. McLeish (2004) "Sequential Decision and Strategy Vector Methos in Ultimatum Bargaining: Evidence on the Strength of Other Regarding Behavior," *Economics Letters*, 84, 399-405.

[21] Radner, R. and A. Schotter (1989) "The Sealed Bid mechanism: An Experimental Study," *Journal of Economic Theory*, 48, 179-220.

[22] Samuelson, W. and M. Bazerman (1985) "The Winner's Curse in Bilateral Negotiations," in *Research in Experimental Economics* (V. Smith ed.), JAI press, 105-137.

[23] Schotter, A., K. Weigelt, and C. Wilson (1994) "A Laboratory Investigation of Multiperson Rationality and Presentation Effects," *Games and Economic Behavior*, 6, 445-68.

[24] Selten, R. (1967) "Die Strategiemethode zur Erforschung des Eingeshrankt Rationalen Verhaltens im Rahmen eines Oligopolexperimentes," in *Contributions to Experimental Economics, Vol. 1* (H. Sauermann ed.), Tubingen: J.C.B. Mohr.

[25] Serrano-Padial, R. (2007) "Bilateral Trade with Pure Common Values:a No-Trade Theorem," *mimeo*, UCSD.

[26] Sonsino, D., I. Erev, and S. Gilat (2001) "On Rationality, Learning and Zero-Sum Betting – An Experimental Study of the No-Betting Conjecture," *Working Paper*, Technion.

[27] Sovic, Y. (2004) "Strength of Dominance and Depths of Reasoning," *Working Paper*, University of Oslo.

[28] Tirole, J. (1982) "On the Possibility of Speculation under Rational Expectations," *Econometrica*, 50, (5), 1163-1181.

[29] Vincent, D.R. (1989) "Bargaining with Common Values," *Journal of Economic Theory* 48, 47-62.

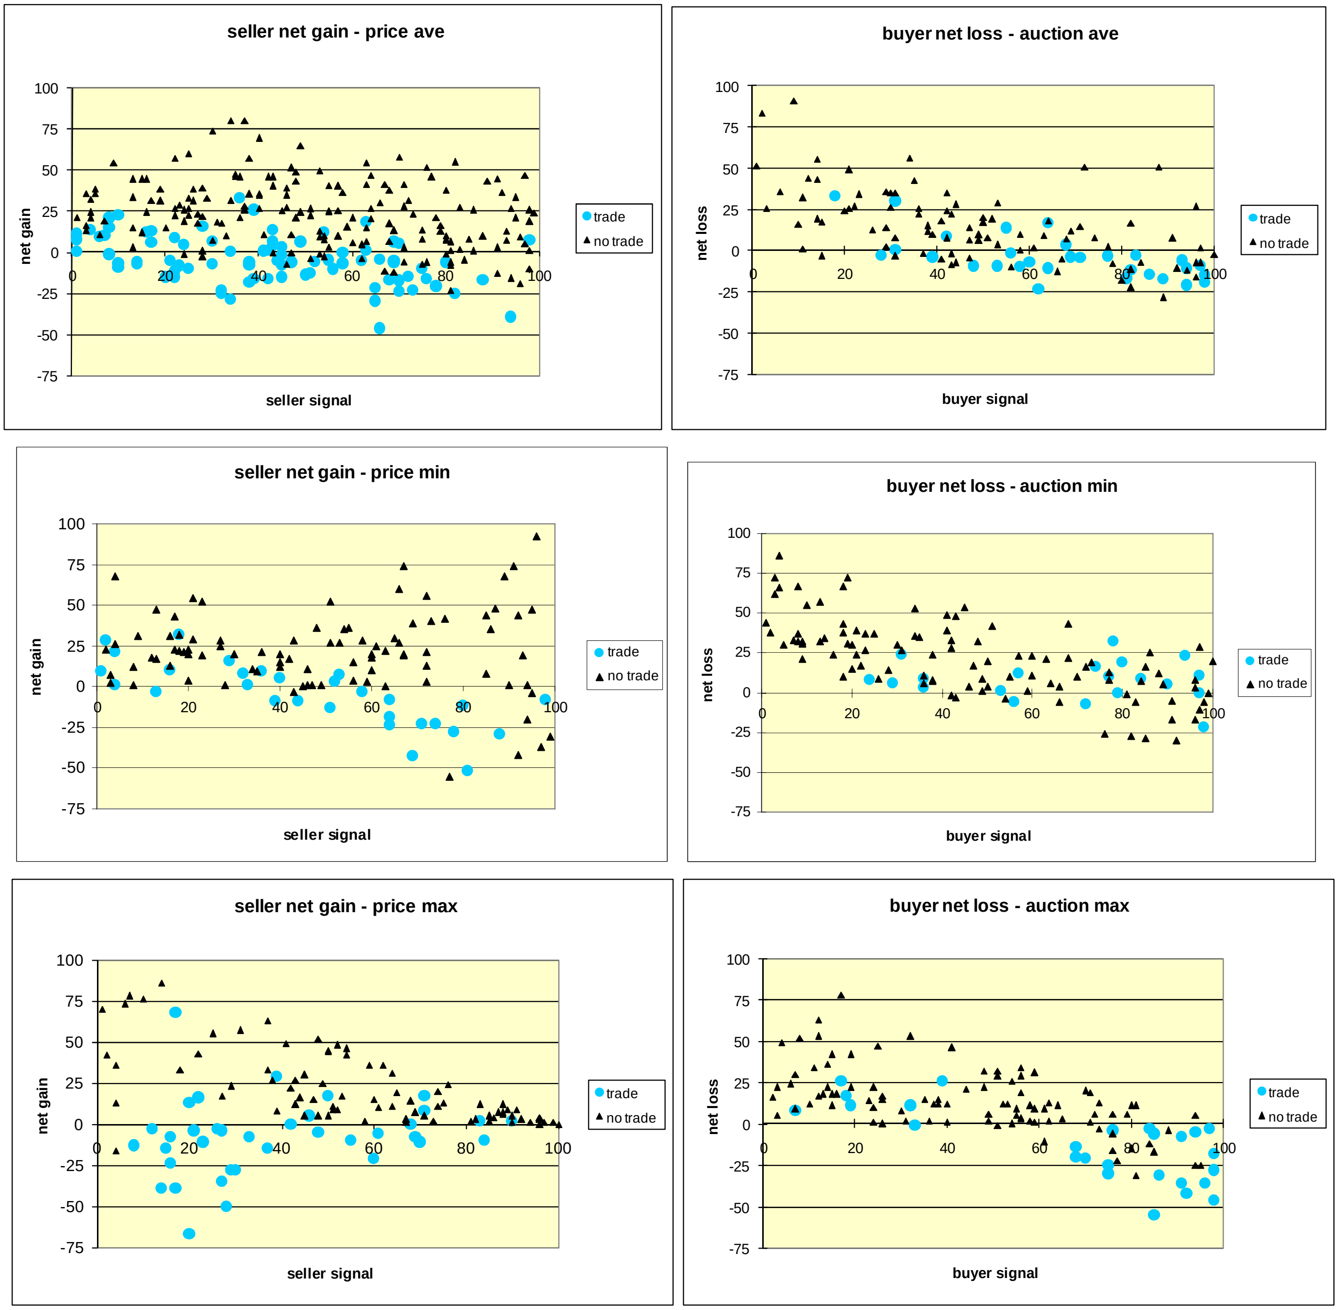

Figure 1. Sellers and buyers net gains by treatment

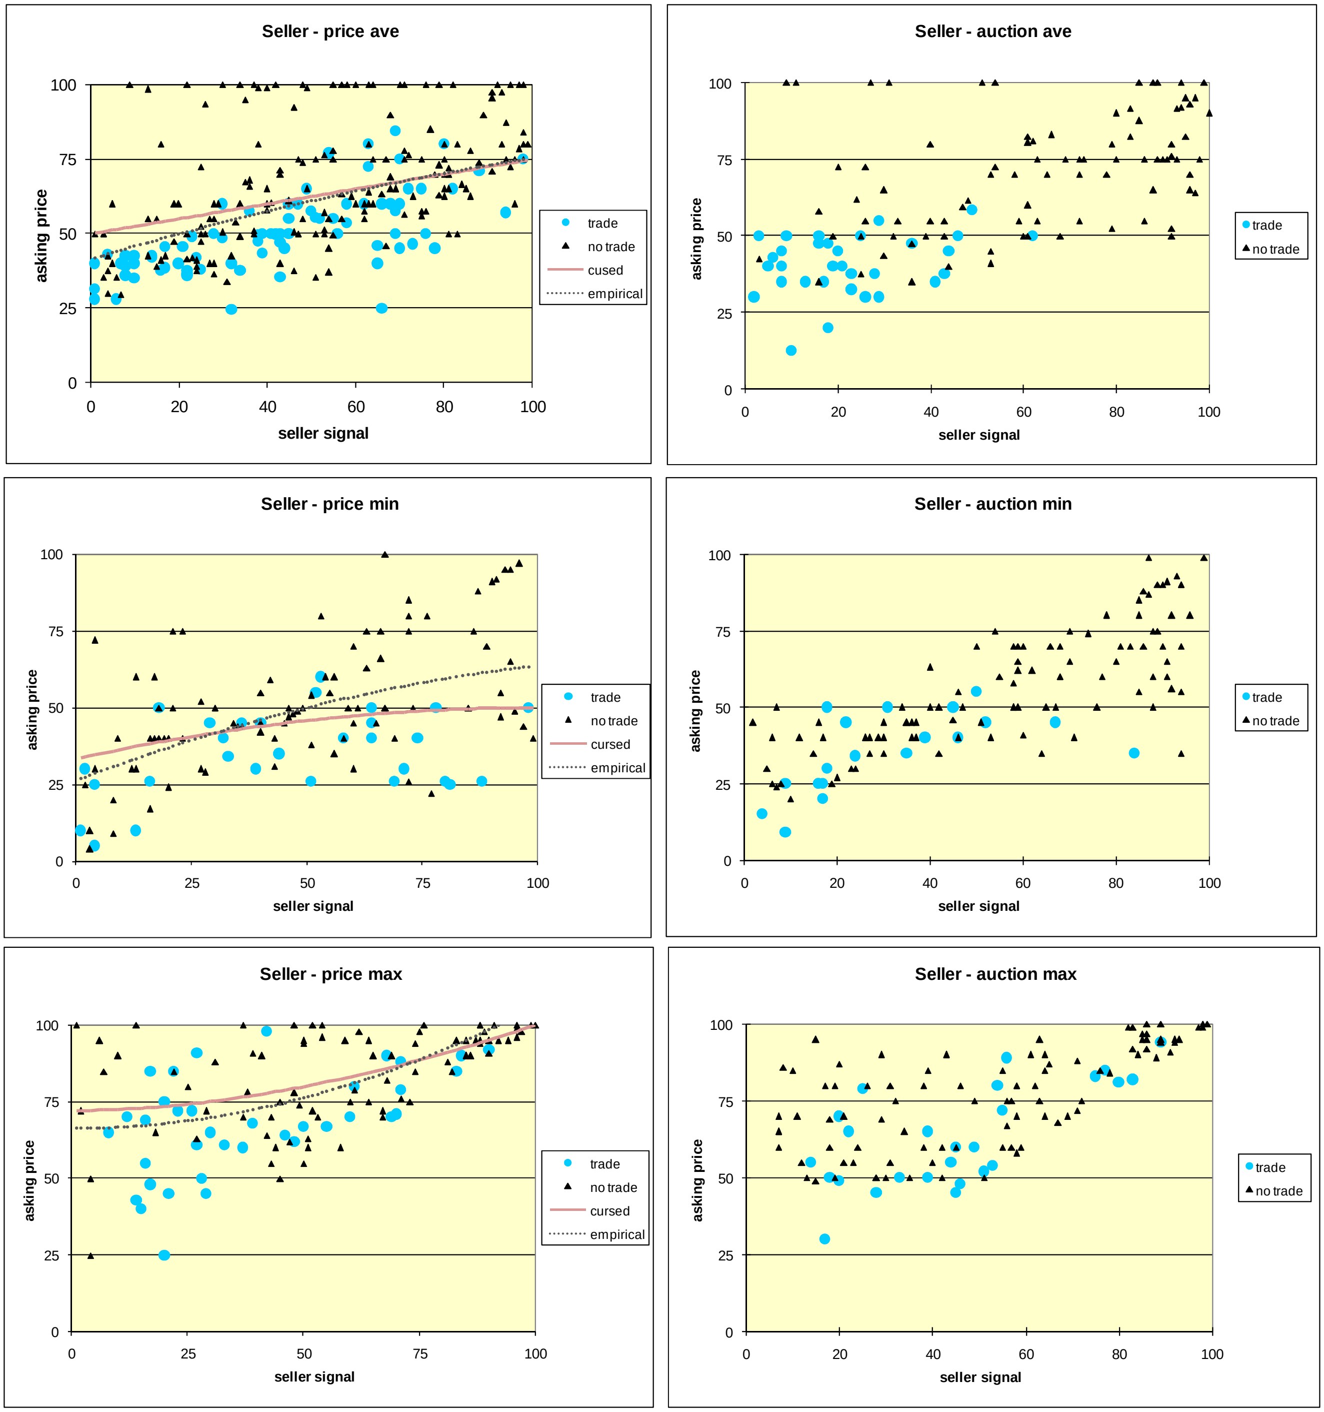

Figure 2. Seller's asking price by treatment

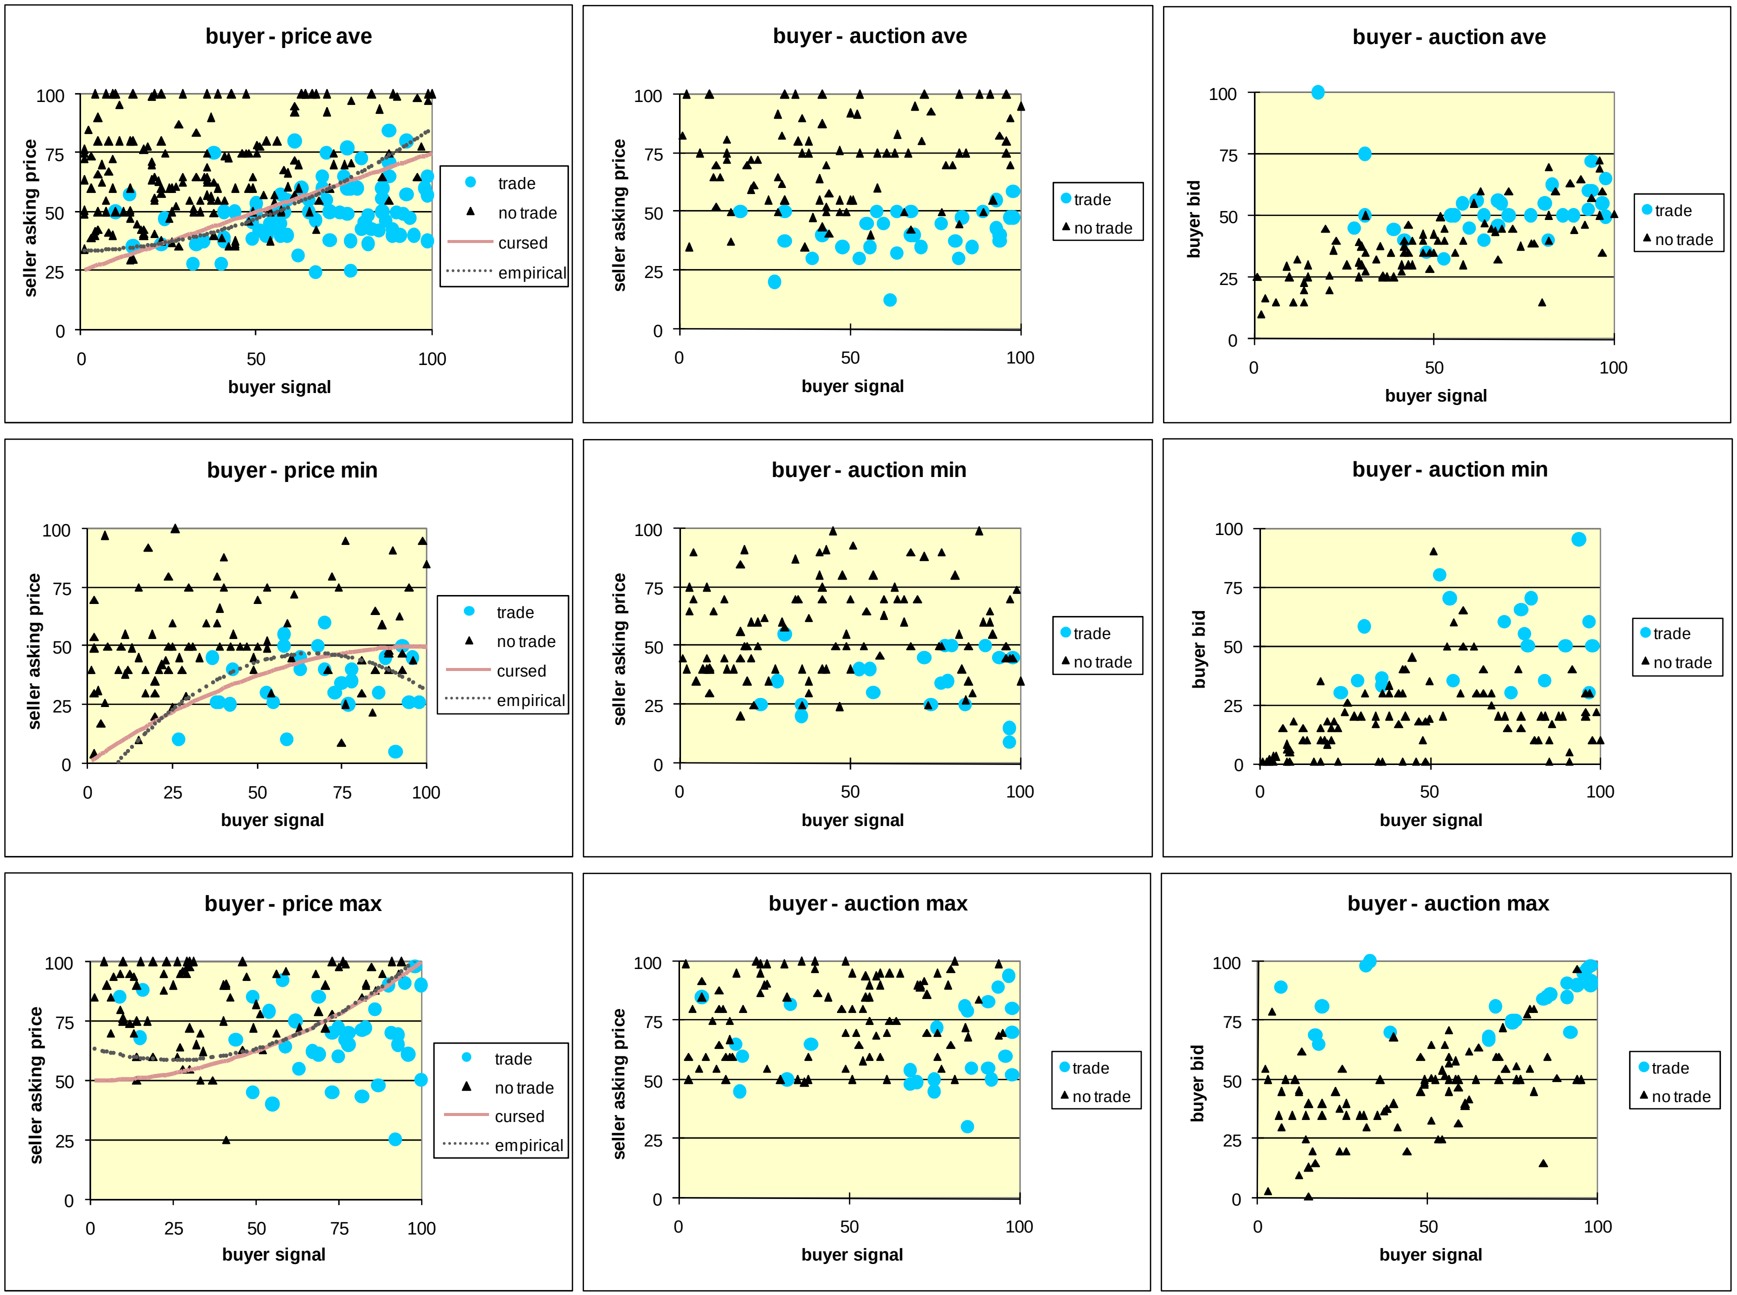

Figure 3. Buyer’s acceptance or bid by treatment

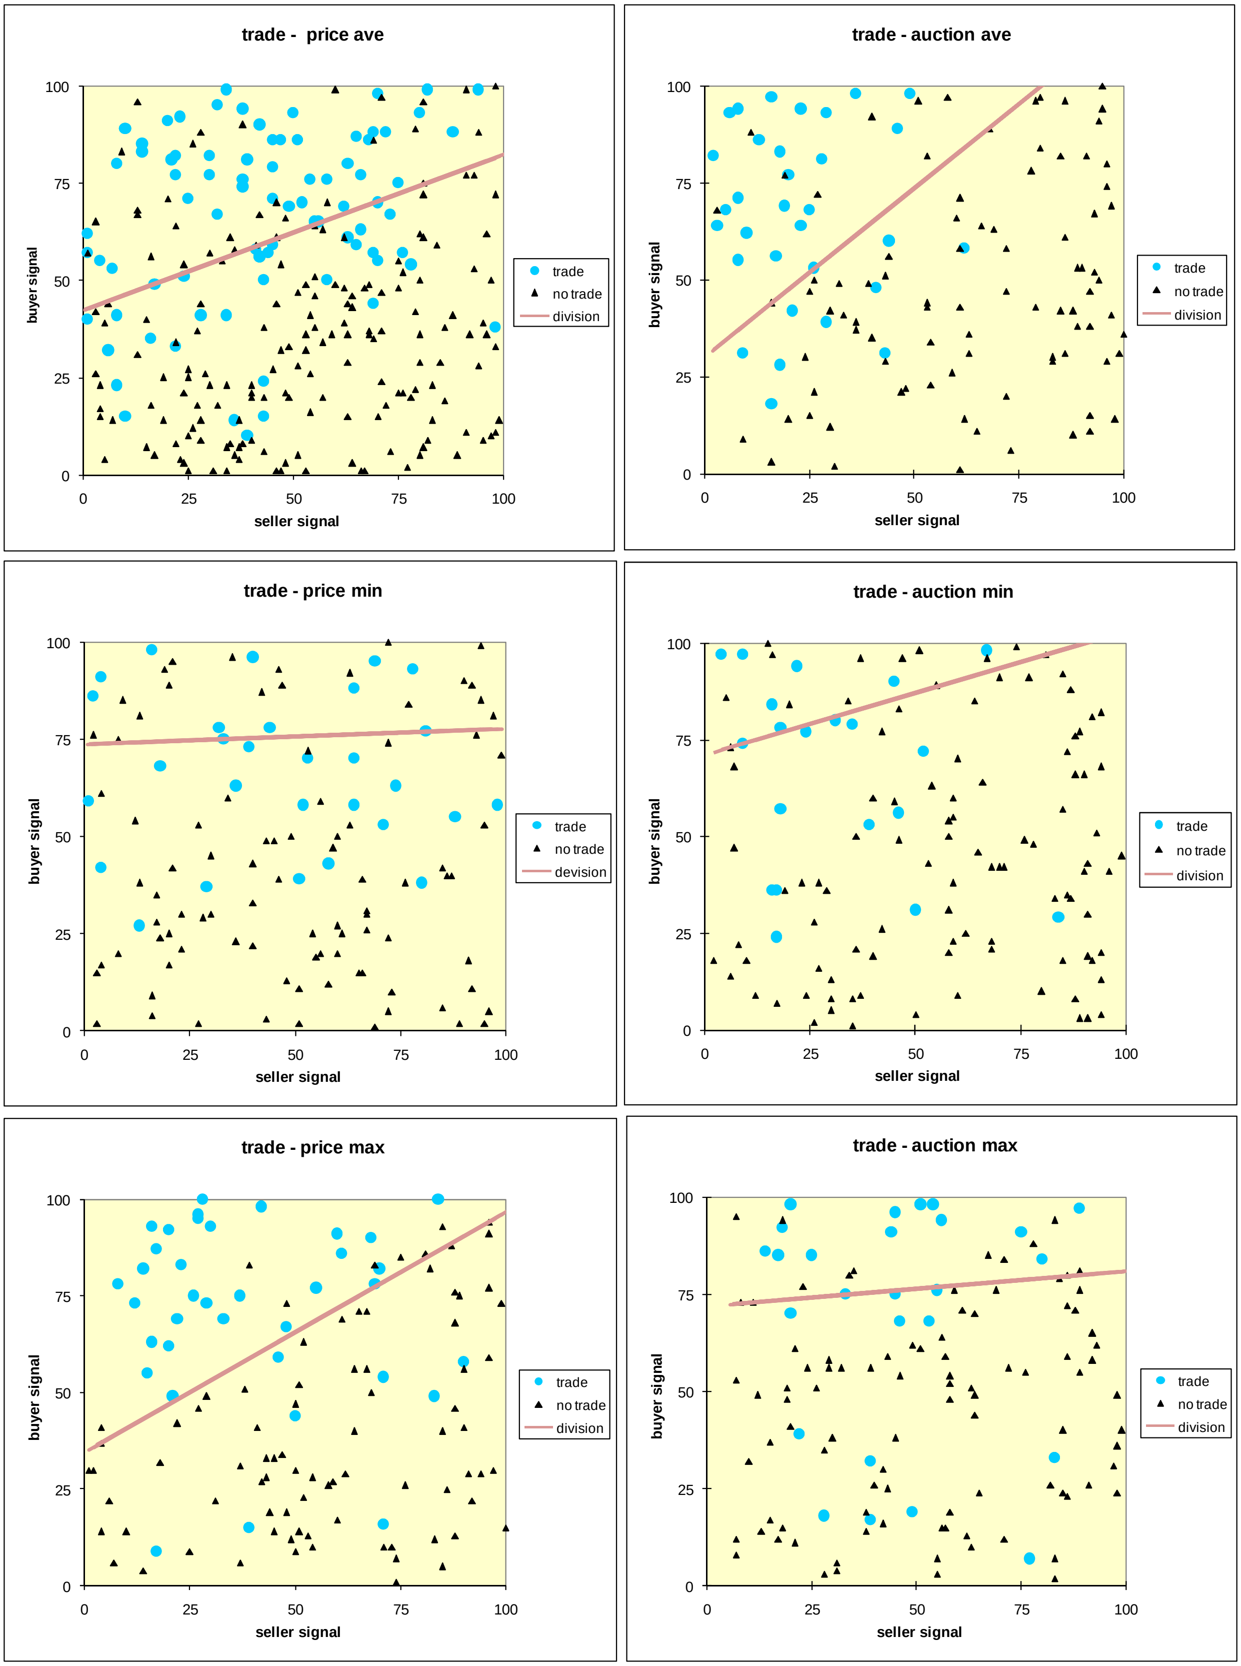

Figure 4. Likelihood of trade as a function of signals