Understanding Overbidding in Second Price Auctions: An Experimental Study*

David J. Cooper† Hanming Fang‡

March 3, 2006

*We are grateful to John Kagel, Dan Levin, Stephen Morris, Shyam Sunder and participants of several conferences and the IO Lunch at Yale for helpful suggestions. The experiments are made possible by financial support from Faculty Research Grant from Case Western Reserve University and Cowles Foundation for Research in Economics. We are responsible for all remaining errors.

†Department of Economics, Weatherhead School of Management, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106. Email: djc13@weatherhead.cwru.edu

‡Department of Economics and Cowles Foundation, Yale University, P.O. Box 208264, New Haven, CT 06520-8264. Email: hanming.fang@yale.edu

Abstract

This paper presents results from second price private value auction (SPA) experiments in which bidders are either given for free, or are allowed to purchase, noisy signals about their opponents' value. Even though in standard models of SPA, such information about opponents' value theoretically has no strategic use in the SPA, it provides us with a convenient instrument to change bidders' perception about the "strength" (i.e. the value) of their opponent. We argue that the empirical relationship between the incidence and magnitude of overbidding and bidders' perception of the strength of their opponent provides the key to understand whether overbidding in second price auctions are driven by "spite" motives or by the "joy of winning." The experimental data show that bidders are much more likely to overbid, though less likely to submit large overbid, when they perceive their rivals to have similar values as their own. We argue that this empirical relationship is more consistent with a modified "joy of winning" hypothesis than with the "spite" hypothesis. However, neither of the non-standard preference explanations are able to fully explain all aspects of the experimental data. We find clear evidence of learning both in avoiding costly overbidding and in subjects' choices to purchase costly information, thus lending support for the role of bounded rationality. We also find that bidder heterogeneity plays an important role in explaining their bidding behavior.

Keywords: Overbidding, Second Price Auctions, Spite, Joy of Winning, Bounded Rationality.

JEL Classification Codes: C91, C72.

1 Introduction

Second price private value auctions (SPAs) are the most easily understood auction format from a theoretical point of view.[1] In standard private value auction models of fully rational bidders with standard preferences, bidding one's own value is a weakly dominant strategy. This theoretical prediction holds irrespective of bidders' risk attitudes, the number of rival bidders, symmetry in the value distributions, and so on. In laboratory experiments, however, subjects are found to exhibit a consistent pattern of overbidding. Kagel, Levin and Harstad (1987) found that the actual bids are on average 11 percent above the dominant strategy bids. Kagel and Levin (1993) find that about 62 percent of all bids in their five-bidder SPA sessions exceed the bidder's value, while only 8 percent of all bids were below it. Both Kagel and Levin (1993) and Harstad (2000) further reports that experience has only a small effect in reducing overbidding in SPA. Another important experimental fact is that, overbidding in English auctions, a strategically equivalent mechanism for SPA in the case of private values, is known to be a short term phenomenon that subjects quickly learn not to undertake (Kagel and Levin 1993). Thus any explanation for the prevalence and persistence of overbidding in the SPA must also explain its rarity in the English auction.

Given the robustness of the findings of overbidding in SPA, it is surprising that economists have very little understanding of *why* it happens. Kagel, Levin and Harstad (1987) conjectured that bidding above one's own value in a SPA is based on the illusion that it improves the probability of winning with little cost because the winner only pays the second-highest bid.[2] Moreover, they argue that overbidding is sustainable because bidders still on average earn positive profits. Finally, overbidding is observed for experienced bidders because the negative feedback from overbidding is a rather weak mechanism in the SPA. For example, if a bidder overbids in a SPA by 10 percent above his value, there could be a high probability that he does not win at all and thus does not experience any negative feedback; even if he wins, there could still be a high probability that he obtains positive payoff from winning (in stark contrast to the first price auctions). Understanding of overbidding in the SPA is of interest in itself, but it is also of broader interest because overbidding in the SPA is a notable example of the use of dominated strategies. As such, we believe that a better understanding of overbidding in the SPA can be valuable for understanding individuals'

[1] Vickery (1961) is the first to study this auction format.

[2] Harrison (1989) used similar arguements to explain overbidding in the first price auctions.

behavior in many other games.

In a recent paper, Morgan, Steiglitz and Reis (2003) analyzed the equilibrium of standard auctions assuming that bidders care not only about her own surplus in the event that she wins the auction, but also about the surplus of her winning rival in the event that she loses the auction (the "spite" motive). They showed that, when bidder's utility function includes a spite motive component, bidders will bid more than their value in second price auctions. Andreoni, Che and Kim (2005, ACK thereafter) conducted a set of related experiments in which bidders are partitioned into groups where bidders within a group can perfectly observe each other's value.[3] They found, among other results, that overbidding is much more prevalent among "followers" – bidders whose values are known to be lower – than "leaders," thus lending support to the importance of spite motives in overbidding.

ACK, however, also found that their results from first-price auction experiments instead favor a theory of risk aversion as an explanation of the slight amount of overbidding in the FPA. The spite motive explanation thus leaves completely unexplained about when and why subjects in a short laboratory experiment will exhibit spite motives under one but not another auction format. In fact, one may argue that an *a priori* weakness of spite as an explanation for overbidding in the SPA is that it also predicts overbidding in English auctions (as Morgan, Steiglitz and Reis, 2003 showed), which as we mentioned earlier, is a rather short term phenomenon that subjects quickly learn not to undertake (Kagel and Levin 1993).

In this paper, we report results from a series of second price private-value auction experiments in which subjects either receive for free or choose to purchase *noisy* signals about their opponent's value.[4] There are several key differences between our experiments and ACK's. The first difference is that, in our experiments, bidders receive noisy but informative signals about opponent's value, whereas in ACK's experiments bidders perfectly observe rivals' values within a group but have no information about rivals outside the group. The crucial property of the noisy signals in our experiments is that, from a theoretical perspective, they are completely useless if bidders are fully rational and are motivated only by monetary payoffs because bidding their own private value remains the weakly dominant strategy regardless of their signals about opponents' value. However, such noisy

[3]The theoretical analysis in Kim and Che (2004) motivated ACK's experiments.

[4]Our experiments are motivated by the theoretical analysis in Fang and Morris (2006).

signals provide us with a convenient instrument to change bidders' perception about the "strength" (i.e. the value) of their opponent. Of course, whether or not the difference in the noisiness of the signals may be important will depend on subjects' understanding about Bayesian updating. The second difference is that, our experimental design includes treatments in which bidders have to decide whether to purchase signals about opponent's value. Because the cost of information acquisition always has a direct and obvious payoff feedback (while the cost of overbidding is typically not obvious and not direct), the endogenous information acquisition treatments allow us to more forcefully evaluate empirically the learning (and bounded rationality) hypothesis; they also allow us to evaluate the relationship between mistakes in information acquisition and overbidding, and thus the potential importance of bidder heterogeneity. The third difference is that, besides the testable hypothesis from the spite model, we also develop testable implications of a modified joy-of-winning model. Indeed our evidence suggests that, within the realm of non-standard preferences, a modified joy-of-winning seems to explain several features of the data better than a spite model.

One of the goals of this paper is to examine the empirical relationship between the incidence and magnitude of overbidding and bidders' perception of the strength of their opponent, and attempt to provide a better understanding of overbidding in second price auctions from such empirical relationships.[5] We find in our experimental data that bidders are much more likely to overbid, though less likely to submit large overbids, when they perceive their rivals to have values similar to their own. We argue that this empirical relationship is, within the framework of full rational bidders with non-standard preferences, is more consistent with a modified version of the "joy of winning" hypothesis, but inconsistent with the "spite" hypothesis. We also report direct evidence of learning and hence the importance of bounded rationality, both in avoiding overbidding through infrequent negative payoff feedbacks and the direct feedback in the costly information acquisition decisions.

Though the main purpose of our experimental design is to tease out different explanations of overbidding in the SPA, our experiments also have implications about the performance of SPAs in real world situations that have not been studied in earlier experiments. In many real-world auctions, bidders typically observe (or may have incentive to acquire) information about their opponents'

[5] Note that the information structure in ACK's setup limits the degree of variation in bidders' perceptions about their opponents' strength that can be achieved in their experiments.

values. Such information may come from gossips, espionage, or in the case of repeated auctions from observing opponents' past bids and winnings. In these situations, the assumption of the standard independent private value auction paradigm that bidder's belief about their opponent's value is independent of her own type and their opponent's true types are violated. How would such information impact bidders' bidding behavior? Is additional information a blessing or a curse for the bidders? How does such information impact the allocative efficiency and revenue? As we mentioned earlier, these questions are not interesting at all from a theoretical point of view: if bidders are fully rational and have standard preferences with no non-pecuniary concerns, then in private-value SPAs, information about opponents' value should not have impact on a bidder's bid; that is, bidding his value remains his weakly dominant strategy, just as it is in the standard SPAs when bidders do not have any information about opponents' value. In this paper, we show that such theoretical predictions do not accurately describe our experimental results. In fact, we find that bidders' bids are systematically affected by their signals about their opponents' value. Such systematic effect of opponents' value on the bids in the SPA provides an additional lens through which we can learn about the incentives for overbidding in the SPA.

The remainder of the paper is structured as follows. Section 2 describes the experimental design. Second 3 presents the theoretical predictions directly related to the experimental auction games. It includes a benchmark analysis where bidders are assumed to have standard preferences and are fully rational, as well as the derivations of testable hypotheses regarding overbidding in SPA auctions from models with "spite" and "joy-of-winning" preferences or with bidders of bounded rationality. Section 4 analyzes the experimental data. Finally Section 5 discusses our findings and concludes.

2 Experimental Design and Procedures

2.1 General Features of All Sessions

All sessions consist of 20 rounds. Subjects are anonymously and randomly matched in two-person groups for each round to play a second price auction. Subjects are not given any information about the identities of other bidders. Given that the smallest of our sessions contained 10 subjects (and sessions averaged 17 subjects), a session can be treated as a series of twenty single-shot games.

Value Distribution Prior to submitting a bid, bidders always knew their own value as well as the distribution from which values are drawn. For all sessions values were drawn from the discrete distribution shown in Table 1. All values are denominated in ECUs, which were converted to cash at a rate of 1 ECU = \$.01. In reporting results, earnings and costs are denominated in dollars rather than ECUs. Values are independent across bidders and across rounds. The distribution of values was common knowledge, although the particular values drawn were private information.

[**Table 1 about here**]

This value distribution approximates a Normal distribution with mean 5,000 ECUs and standard deviation of 2,000 ECUs.[6] We used a peaked distribution rather than a uniform distribution largely to generate more competitive auctions (e.g. auctions where the bidders' values are relatively close) without changing the range of possible values.[7]

Signal Distribution As a treatment condition, bidders are either given for free or have the choice to purchase noisy information about opponents' values. These signals are received simultaneously with values, prior to bidding. The distribution of signals is as follows: With probability K, where K is either .3 or .7 as a treatment variable, the signal that a bidder draws is exactly equal to the value of the opposing bidder. With probability $1 - K$, the signal is drawn from a uniform distribution over the other ten values. In other words, the probability of each possible incorrect signal is $(1 - K)/10$. The variable K measures the quality of the signal. The value of K is common knowledge; but the signals observed by subjects are private information.

2.2 Experimental Treatments

Our experimental design has several purposes. At the simplest level, we would like to know if bidders will respond to information about their opponent's value and whether they will be willing to pay a positive amount for such information. If bidders are rational and maximizing their expected

[6] In pilot sessions we have experimented with continuous Normal distributions, but this necessitated the inclusion of lengthy instructions on Normal distributions. We choose this discrete approximation of the Normal distribution for simplicity.

[7] The use of a peaked distribution also has important implications for comparing the efficiency of first and second price auctions, a central issue in the broader design.

monetary payoff, the signals as worthless. Bidders should therefore ignore the signals and be unwilling to pay for them. Observing otherwise extends the known anomaly of overbidding in second price auctions.

Ultimately, our goal is not to provide a laundry list of anomalous behavior but rather to understand why anomalous behavior is occurring. The two main categories of explanation that apply here are *non-standard preferences* (e.g. "spite" or "joy of winning") and *bounded rationality*. Under the rubric of bounded rationality, subjects may either be making completely random errors or making a systematic mistake. Our experimental design is intended to separate these various explanations for overbidding in second price auctions.

The experimental design has five treatments split into three categories as follow.

Control Treatment [CON]. These are basic treatments in which bidders do not observe any informative signal about opponents' values. The controls serve two purposes. Beyond serving as a baseline for comparison with the other treatments, the control treatments allow us to replicate the qualitative results of earlier experiments on overbidding in second price auctions. With this replication in hand we can be reasonably confident that our results are not driven by any peculiarities of our subject pool or other secondary features of the experimental environment.

Exogenously Provided Signals Treatments [EX3 and EX7]. In these treatments, at the same time they received their private values, subjects in each round were provided with a free signals of quality K about their opponent's values, where K is equal to .3 for the EX3 ("low quality" signals) and .7 for the EX7 ("high quality signals") treatments. The value of K was the same for all bidders and for all rounds of a session. While in neither case is the signal either completely informative or completely uninformative, the signal is substantially more informative in the high quality signal treatment. For example, suppose a bidder receives a signal of 0 about an opposing bidder. In the low quality signal treatment, the updated expected value for the opposing bidder would be \$48.41 while in the high quality signal treatment it would be \$40.87. Thus, our experimental design allows us not only to observe how bidders respond to information about their opponent's value but also how this response varies with the quality of the information.

Endogenous Signal Acquisition Treatments [END3 and END7]. In these treatments, bidders are offered an opportunity to purchase a signal about their opponents' valuations. The quality of signals offered for purchase to the subjects is respectively fixed at $K = .3$ for END3 and $K = .7$ for END7 treatments. The quality of information K is fixed and known to the subjects in each round. For each bidder in each round, a cost c, denominated in ECUs, will be drawn from a uniform $[-50, 250]$ distribution. Prior to being told their private value, each subject will be told their cost of information and asked if they wish to buy a signal. If the subject chooses to buy information, they have c deducted from their show-up fee and, prior to bidding, receive both their own value and signals on the values of the other bidder.[8] The decision to acquire information is known to all bidders before the bidding begins.

The endogenous signal acquisition treatments serve two important purposes. First, They allow us to verify an important prediction of models of bounded rationality. If overbidding represents a mistake, rather than maximization subject to non-standard preferences, subjects should learn to stop making this mistake as they gain experience *if it is costing them money.* This prediction does not require that individuals understand the nature of the mistake as reinforcement learning is sufficient to yield a reduction in errors. The critical clause here is "if it is costing them money." As noted by Kagel and Levin (1993), one of the reasons subjects have difficulty learning not to overbid is that it only rarely costs them money. By extension, it is quite difficult to observe learning in second price auctions. In contrast, paying a positive price for information always costs money. If bounded rationality is a major force underlying anomalous behavior, we ought to see subjects learning to avoid the always costly mistake of paying a positive price for information.

An additional advantage of the endogenous signal acquisition treatments is that they also allow us to separate subjects out by types. To the extent that some subjects are more rational than others, it is instructive to show that those who make wrong choices in one domain tend to make wrong choices in other domains as well. Indeed this is what we find in our empirical analysis (see Section 4.6).

[8] We are wary of using the standard Becker-DeGroot-Marschak (BDM) mechanism to elicit the willingness to pay for information because of its direct relationship to a second price auction. For the BDM technique to have any value, the instructions must carefully explain to subjects why they should bid their true value for the information. In our SPA experiments, this amounts to giving subjects detailed instructions telling them they should follow the dominant strategy. We suspect this would influence the results.

2.3 Experimental Procedures

A total of 12 experimental sessions were conducted in the Fall 2003 and Spring 2004, with subjects recruited from undergraduate students at Case Western Reserve University and Yale University using newspaper ads, posters, emails, and classroom announcements. These sessions are allocated to the five treatments as detailed in Table 2 below. The number of participants in each session varies between 10 to 24. Subjects were only allowed to participate in a single session.

[**Table 2 about here**]

For the most part, the experimental procedures were quite standard. All sessions were run in a computerized lab using the software z-Tree (Fischbacher 1999). At the beginning of each session the experimenter read the instructions aloud to the subjects, which were also displayed on the subjects' computer screens. Before beginning to play, all subjects were asked to complete a short quiz about the payoffs and the rules of the experiment.[9] All subjects were given a printed table describing the distribution of values and, where applicable, signals.

In the END3 and END7 treatments, a round began with both bidders seeing a price for information and being asked if they wished to purchase a signal. Bidders were then shown their private values and, when applicable, their signals – all other treatments began at this stage. Next, bidders simultaneously chose a bid. Negative bids were not allowed and bids were capped at 99,999 ECUs, a limit that surprisingly was reached nine times by seven different subjects. At the end of each round bidders were told whether they had the high bid for the round – we purposely did not refer to "winning" or "losing" the auction. They were also told their value and bid for the round, their opponent's bid, and their payoff for the round. When relevant, the feedback screen also reported their signal, any expenditures on information, and their payoff before and after adjusting for the cost of information. Subjects were given a record sheet to log their history of play. While keeping the record sheet was strictly voluntary, we observed that most subjects filled them out religiously.

At the end of the session, each subject was paid in cash for a single randomly selected round plus their show-up fee. We pay on a randomly selected round so that income effects will not be a confound for any learning effects. In the case that subjects lost money for the randomly selected round, these losses were deducted from their show-up fee. We never attempted to collect

[9]The full text for the instructions and quiz for the EX7 treatment is given in Appendix A.

money from a subject, so losses are effectively capped at \$12.[10] The instructions explained the bankruptcy rule to subjects multiple times, and the payoff quiz included an example with losses. Several subjects indeed left the experiment with a final payoff of \$0. Somewhat to our surprise, none of these subjects complained. The average payoff was approximately \$21 with sessions generally lasting 60 – 75 minutes. These payoffs were sufficient to generate a plentiful supply of subjects.

3 Theoretical Predictions

Now that we have described in detail our experimental design, it is useful to first derive the theoretical predictions from various models with different assumptions on bidder preferences and rationality. We first provide theoretical predictions from the benchmark model where bidders have standard preferences and are fully rational; then describe testable implications from three alternative models where bidders may have non-standard preferences (they either have spite motives or obtain additional joy from winning alone) or they may have a particular form of bounded rationality.

3.1 Benchmark Theoretical Predictions: Fully Rational Bidders with Standard Preferences

Two bidders, $i = 1, 2$, compete for an object in a second-price private value auction. Bidder $i's$ valuation for the object v_i is private and it is independently drawn from a discrete distribution with support $\{n_1, ..., n_L\}$ where $0 \leq n_1 < n_2 < \cdots < n_L$ and $\Pr(v_i = n_l) = p_l \in (0, 1)$ such that $\sum_{l=1}^{L} p_l = 1$. We consider two scenarios. In the first, bidder i is given a noisy signal s_i about her opponent's realized valuation v_j. The signal is accurate with probability $K \in (0, 1)$; and if the signal is inaccurate, it will equally likely take the other wrong values. That is,

$$\Pr(s_i = n_l | v_j = n_l) = K$$
$$\Pr(s_i = n_{l'} | v_j = n_l) = \frac{1-K}{L-1} \text{ if } l' \neq l.$$

[10] Bidding one's value is still a dominant strategy when losses are capped at \$12. Even with the losses capped, the marginal expected payoff of increasing the bid remains strictly negative at any bid greater than the value (assuming full support for the opponent's bid function).

Of course, to ensure that the signal s_i is informative about bidder $j's$ valuation, K has to satisfy $K > 1/L$. We refer to the parameter K as the precision of bidder's signal about their opponent's value.

In the second scenario, bidder i is not offered free information about her opponent's valuation v_j. Rather, they are offered to purchase a signal with precision K. The price of the signal is randomly drawn from a distribution with support $[\underline{c}, \bar{c}]$ where $\underline{c} < \bar{c}$ and $\underline{c}$ may be less than zero. The bidders are asked to make the signal purchase decision before they observe their valuations. After the bidders make the information acquisition decisions, bidders observe their private value, and the signal about their opponent's value (if they do purchase information), and they are also informed about whether their opponent has purchased information about their value. Bidders then submit bids, and the higher bidder wins the object at a price of the losing bidder. Ties are broken with a coin flip.

Suppose that bidder i receives a signal s_i about her opponent's value, her posterior about her rival's value is changed according to Bayes rule as follows: bidder i perceives that her opponent $j's$ value v_j will take on value n_l with probability

$$\Pr\left(v_j = n_l | s_i = n_m\right) = \begin{cases} \frac{Kp_l}{Kp_l + \frac{1-K}{L-1}\sum_{l' \neq l} p_{l'}} \text{ if } l = m \\ \frac{\frac{1-K}{L-1} p_m}{Kp_l + \frac{1-K}{L-1}\sum_{l' \neq l} p_{l'}} \text{ if } l \neq m. \end{cases}$$

Thus bidders' signals effectively change their perceptions about their rivals' values. However, as long as bidders' payoffs are as modelled by the standard theory, which is simply their monetary payoff:

$$U\left(b_i, b_j; v_i\right) = \begin{cases} 0 \text{ if } b_i < b_j \\ \frac{v_i - b_j}{2} \text{ if } b_i = b_j \\ v_i - b_j \text{ if } b_i > b_j, \end{cases} \tag{1}$$

their perceptions about their rivals' values should have no effect on how much they should bid.

Proposition 1 *If a bidder's payoff is given by (1), the unique equilibrium in weakly dominant strategies for the SPA is as follows:*

1. *When information is free, a bidder of type (v_i, s_i) should bid her private value v_i regardless of her signal about her opponent's value;*

2. *When information acquisition is endogenous, a bidder should purchase the information only if the cost is negative; and they should bid v_i regardless of their signal of their opponent's value.*

3.2 Spite, Joy of Winning and Bounded Rationality: Testable Implications

Before we describe our empirical results, in this section we briefly sketch the testable implications of three plausible explanations of overbidding in SPA: spite, joy of winning and bounded rationality.

3.2.1 Spite

We first consider the "spite" hypothesis recently put forward by Morgan, Steiglitz and Reis (2003) who argue that if bidders care not only about their own surplus in the event that they themselves win the auction, but also about the surplus of their winning rival in the event that the lose the auction, then bidders will overbid in the second price auctions.

We will first sketch the equilibrium deviation of the SPA auction with spite-motive à la Morgan, Steiglitz and Reis (2003) for a model in which bidders do not receive any noisy signals about opponent's value. There are two bidders $i = 1, 2$, whose values are drawn from distribution $F(\cdot)$ with PDF $f(\cdot)$ on the support $[0, 1]$. Suppose that the private valuation for bidder i and j are v_i and v_j respectively, and suppose that they bid b_i and b_j respectively. Morgan, Steiglitz and Reis' (2003) model incorporates bidder i's spite motives into her payoff function as follows:

$$W(b_i, b_j; v_i, v_j) = \begin{cases} -\alpha(v_j - b_i) \text{ if } j \text{ wins} \\ v_i - b_j \text{ if } i \text{ wins,} \end{cases} \tag{2}$$

where $\alpha \in [0, 1)$. That is, when bidder i loses the auction, she receives a negative payoff (the spite) that is proportional to her opponent's realized surplus $v_j - b_i$. Note that the standard auction model without spite motives corresponds to the case $\alpha = 0$. Now we can sketch the equilibrium for SPA with spite-motivated bidders. Suppose that bidder 2 follows a bidding strategy $\beta(\cdot)$. Bidder 1's payoff from bidding an amount b when his value is v_1 and his opponent's value is v_2 is

$$\Delta(b; v_1, v_2) = [v_1 - \beta(v_2)] \mathrm{I}_{b \geq \beta(v_2)} - \alpha(v_2 - b) \mathrm{I}_{b < \beta(v_2)}.$$

Bidder 1 takes expectation over values of v_2, assuming that bidder 2's bidder strategy $\beta(\cdot)$ is strictly increasing in v_2 :

$$\mathrm{E}_{v_2}\Delta\left(b;v_1,v_2\right)=\int\limits_0^{\beta^{-1}(b)}\left[v_1-\beta\left(v_2\right)\right]f\left(v_2\right)dv_2-\alpha\int\limits_{\beta^{-1}(b)}^{1}\left(v_2-b\right)f\left(v_2\right)dv_2.$$

Taking the first order derivative with respect to b, and imposing symmetric bidding equilibrium, we obtain

$$\left(v_1-b\right)f\left(v_1\right)\frac{1}{\beta'\left(v_1\right)}+\frac{\alpha\left(v_1-b\right)f\left(v_1\right)}{\beta'\left(v_1\right)}+\alpha\int\limits_{\beta^{-1}(b)}^{1}f\left(v_2\right)dv_2=0.$$

In a symmetric equilibrium $\beta^{-1}\left(b\right)=v_1$, thus we can rewrite the above first order condition, after some simplification, as:

$$\beta'\left(v_1\right)+\left[\frac{\left(1+\alpha\right)/\alpha f\left(v_1\right)}{F\left(v_1\right)-1}\right]\beta\left(v_1\right)=\left[\frac{\left(1+\alpha\right)/\alpha f\left(v_1\right)}{F\left(v_1\right)-1}\right]v_1$$

Thus the solution is

$$\beta\left(v\right)=v+\frac{\int_v^1\left[1-F\left(t\right)\right]^{(1+\alpha)/\alpha}dt}{\left[1-F\left(v\right)\right]^{(1+\alpha)/\alpha}}.$$

Morgan, Steiglitz and Reis (2003) showed that equilibrium level of overbidding $\beta\left(v\right)-v$ decreases with a bidder's own value (because if a bidder's value is high, it is more likely that she will win and her own bid will be less likely to decrease her rival's surplus). The intuition for this comparative statics is best revealed if we assume that the opponent, say bidder 2, is following the standard strategy of bidding her own value. When bidder 1 considers marginally raising his bid from v_1, there are three effects. First, raising one's bid leads to a marginal gain from the increase in probability of winning; second, raising one's bid also leads to a marginal cost of winning at a price in excess of one's valuation. In the absence of spite motives, these two effects exactly cancel out, and thus bidding v_1 is optimal. When spite incentives are present, there is a third effect: by raising one's bid, one increases the price of the rival bidder in the event he has a higher valuation, which happens with probability $1-F\left(v_1\right)$. Thus the third effect, which is a marginal benefit term from overbidding, is higher the lower a bidder's valuation. Thus, the spite motive model predicts that in SPA control sessions, the overbidding should be decreasing with a bidder's own value if bidders are spite-motivated.

In an environment in which bidders also privately observe noisy signals about opponents' value (and thus bidders have multi-dimensional private types), it is not analytically possible to derive the equilibrium of the SPA with spite-motivated bidders. However, it is possible to extend the above intuition to obtain some comparative statics predictions about the incentives to overbid in this environment. Suppose that the opponent, say bidder 2, is bidding her own value v_2. When bidder 1 considers marginally raising his bid above his valuation v_1, there are again three effects. The first two effects are the same as before and they again exactly cancel out each other, but the third effect – by raising one's bid, one increases the price of the rival bidder in the event he has a higher valuation – is now perceived by bidder 1 to occur with probability

$$\Pr(v_2 > v_1|s_1) = 1 - F_{v_2|s_1}(v_1|s_1).$$

In equilibrium, bidding above valuation raises the marginal cost term to just compensate for the two marginal benefit terms.

Thus the incentives to overbid in our treatments in which bidders receive noisy signals is proportional to $1-F_{v_2|s_1}(v_1|s_1)$, which is bidder 1's belief that bidder 2's value is above v_1 given v_1 and s_1. This perceived probability can be calculated from the Bayes' rule, and not surprisingly, it depends on bidder 1's own valuation v_1, his signal about opponent's value s_1, and the signal accuracy K. Numerical simulations for the term $1 - F_{v_2|s_1}(v_1|s_1)$ for information accuracy $K = .3$ and $K = .7$ respectively yield the following key predictions of the overbidding incentives for spite-motivated bidders in environments where bidders receive noisy signals about opponents' valuation:[11]

Spite Hypothesis 1: *Overbidding incentives decrease in bidders' own valuation v_i in all treatments.*

Spite Hypothesis 2:

(a) *Overbidding incentives is lowest when bidders' own value v_i and signal s_i coincide;*

(b) *Overbidding incentives are lower when $v_i > s_i$ than when $v_i < s_i$;*

(c) *Overbidding incentives increase in s_i when a bidder's signal s_i is higher than her own value v_i.*

Two remarks about the spite theory of overbidding in SPA are worth making. First, this theory predicates that when subjects play in a lab experiment their spite is targeted toward fellow subjects

[11] Details about the calculations are available from the authors upon request.

rather than toward the experimenter. Whether or not this assumption is valid is not clear. Second, the equilibrium in an English auction with two spite-motivated bidders is identical to that of the second price auction, in the benchmark model where bidders do not observe noisy signals about opponent's values (see Proposition 3 of Morgan et. al. 2004). Thus spite motive can not explain the observed difference in overbidding between SPA and English auctions.

3.2.2 Joy of Winning

An alternative hypothesis is that bidders overbid because they derive positive utility from winning, *over and beyond* any monetary payoffs, which we will call the "Joy of Winning" hypothesis. We will distinguish between two versions of the joy of winning theory. In the *simple* version, we assume that other than the additional positive utility from winning, the bidders are able to figure out the equilibrium bidding strategy with full rationality; in the *modified* version, we assume that the bidders not only care about winning *per se*, they also use heuristics in deciding how much to bid.

The implication of the simple version of the joy of winning theory on overbidding is easy to establish. Suppose a bidder's valuation of an object is v_i, then she receives a utility of $v_i + t_i$ from winning the object, and 0 otherwise, where $t_i > 0$ denotes the additional joy from winning the object. Let $G_i(b_j|v_i, s_i)$ be bidder $i's$ belief about her opponent's bid given her own type (v_i, s_i). Then bidder $i's$ problem is

$$\max_{\{b_i\}} \int_0^{b_i} (v_i + t_i - b_j)\, dG_i(b_j|v_i, s_i).$$

The optimal bid $b_i^* = v_i + t_i$. Thus in equilibrium, a simple joy-of-winning theory predicts that bidders overbid by the amount of their joy t_i. That is,

Simple Joy-of-Winning Hypothesis 0: *The amount of overbidding in the simple joy-of-winning theory is independent of the bidders' own valuation, their signals about opponent's valuation and the signal accuracy.*

Richer implications from the joy-of-winning theory can be derived in a *modified* model where we make some additional behavioral assumptions about overbidding incentives. For simplicity, suppose that $t_i = t$ for all i. Again let $G_i(b_j|v_i, s_i)$ be bidder $i's$ belief about her opponent's bid given her own type (v_i, s_i). Consider bidder i who is contemplating overbidding his valuation v_i by ϵ. Her expected payoff from bidding ϵ above her value v_i is given by

$$\int_0^{v_i+\epsilon} (v_i + t - b_j)\, dG_i(b_j|v_i, s_i).$$

The marginal benefit from overbidding is thus (taking derivative with respect to ϵ) :

$$(t - \epsilon)\, g_i(v_i + \epsilon|v_i, s_i)$$

where $g_i(\cdot|v_i, s_i)$ is the derivative of $G_i(\cdot|v_i, s_i)$. This of course means that the optimal overbidding is $\epsilon^* = t$. However, if we assume instead that bidders are more likely to overbid when the marginal benefit is higher, than we can conclude that the incentives to overbid will depend on the magnitude of $(t - \epsilon)\, g_i(v_i + \epsilon|v_i, s_i)$. In particular, it depends on $g_i(\cdot|v_i, s_i)$, which measures bidder $i's$ belief about opponent $j's$ bid. Just as we did heuristically for the spite motive model earlier, if bidder i imagines that the other bidders are bidding their values, then $g_i(v_i + \epsilon|v_i, s_i)$ is higher when s_i is close to v_i; and when the base probabilities of v_i are higher (i.e. when v_i is 4000, 5000 or 6000 ECUs). We restate the above discussion as two hypotheses for the *modified* joy-of-winning theory:

Modified Joy-of-Winning Hypothesis 1: *Overbidding is more likely when a bidder's signal s_i is close to her own value v_i;*

Modified Joy-of-Winning Hypothesis 2: *Overbidding is more likely for values with higher base probabilities.*

Finally, it is useful to point out that the joy-of-winning story suffers from the same problem as the spite story in that it can not explain the difference in the overbidding between SPA and English auctions.

3.2.3 Bounded Rationality

To start with, we should note that overbidding is clearly *not* a purely random bidding error, as overbidding occurs systematically more frequently than underbidding. However, what is the right model for bounded rationality in auctions is a challenging question that is beyond the scope of this paper. Here we sketch a simple model of how bounded rationality may impact overbidding.

As before, let $G_i(b_j|v_i, s_i)$ be bidder $i's$ belief about opponent $j's$ bid distribution where (v_i, s_i) is bidder $i's$ type. Bidder $i's$ expected payoff from bidding b_i is given by

$$\int_0^{b_i} (v_i - b_j)\, dG_i(b_j|v_i, s_i)\,. \tag{3}$$

A perfectly rational bidder will of course choose b_i to maximize the expected payoff above, which will yield a first order condition

$$(v_i - b_i)\, g_i(b_i|v_i, s_i) = 0, \tag{4}$$

implying an optimal bid of $b_i^* = v_i$, as predicted by standard theory.

The bounded rationality story we propose hypothesizes that bidders put a higher weight on the impact of their bid on the probability of winning, and a lower weight on its impact on the expected payoff conditional on winning. To be more precise, let us rewrite the expected payoff (3) as

$$G_i(b_i|v_i, s_i)\, \mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right].$$

The form of bounded rationality that we model here assume that, while bidders can fully appreciate the positive impact of a marginal increase of b_i on her probability of winning, she *under-appreciates* the impact of an increase in b_i on her expected payoff conditional on winning $\mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right]$. More specifically, we assume that she only appreciate $1-\omega \le 1$ of the impact of b_i on her expected payoff conditional on winning $\mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right]$. Under this assumption her first order condition is

$$g_i(b_i|v_i, s_i)\, \mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right] + (1-\omega)\, G_i(b_i|v_i, s_i)\, \frac{\partial \mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right]}{\partial b_i} = 0,$$

where the first term captures the marginal effect of bidding on the probability of winning, and the second term captures the under-weighted marginal effect of bidding on the expected payoff conditional on winning. After some manipulations, the above first order condition is simplified to

$$(1-\omega)(v_i - b_i)\, g_i(b_i|v_i, s_i) + \omega g_i(b_i|v_i, s_i)\, \mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i\right] = 0. \tag{5}$$

Note that when $\omega = 0$, i.e., when bidders are fully rational, the first order condition is identical to (4). The above first order condition immediately implies overbidding, i.e. $b_i^* > v_i$, since the second term is always positive. This can be seen directly by eliminating $g_i(b_i|v_i, s_i)$ from both terms and rewriting:

$$b_i^* - v_i = \frac{\omega}{1-\omega} \mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i^*\right]. \tag{6}$$

It can be easily shown that the amount of overbid $b_i^* - v_i$ increases in ω, the perception bias.

It is plausible to postulate that the *key difference between SPA and English auctions* lie in the level of perception bias ω. specifically, the format of English auction makes it salient to the bidders when they bid above their valuation that any increase in the probability of winning the objects will result in negative payoff conditional on winning, that is, in English auctions the perception bias ω is likely to be zero, at least after a few rounds of bidding. On the other hand, SPA never makes it clear (and learning may be slow) that bidding above one's valuation only increases winning when winning is not profitable.

In terms of testable implications, the key difference between a bounded rationality explanation and the explanations based on non-standard preferences is as follows. If overbidding is driven by bounded rationality, bidders may learn to bid more accurately overtime if the errors provide strong feedbacks; but if they are driven by non-standard preferences then the overbidding will persist overtime. We state this as a testable hypothesis:

Bounded Rationality Hypothesis 1: *If overbidding is driven by bounded rationality, bidders may learn to bid more accurately overtime if the errors provide strong feedbacks.*

It would also have been useful to derive how overbidding under this particular form of bounded rationality relates to the signals. However, even though the right side of formula (6) can be written as

$$\frac{\omega}{1-\omega}\mathrm{E}\left[(v_i - b_j)\,|b_j \le b_i^*\right] = \frac{\omega}{1-\omega}\frac{\int_0^{b_i^*}(v_i - b_j)\,dG_i\left(b_j|v_i, s_i\right)}{G_i\left(b_i^*|v_i, s_i\right)},$$

where indeed bidder i's belief about her opponent's bids conditional on v_i and s_i, namely $G_i\left(\cdot|v_i, s_i\right)$, appears, unfortunately formula (6) has b_i^* on both sides, and as a result it does not provide a clear-cut prediction about the relationship between the amount of overbid $v_i - b_i^*$ and (v_i, s_i).

4 Experimental Results

4.1 An Overview of Bidding Behavior

Figure 1 summarizes bidding behavior from our experiments. The first cluster of bars shows data from all five treatments pooled together and the remaining clusters break out the five treatments. Overbids are separated into three categories: low ($0 <$ bid – value $< \$12$), medium ($\$12 \le$

bid – value < \$25), and high (\$25 ≤ bid – value) overbids. The breakpoints between these three categories are somewhat arbitrary. Overbids over \$12 are sufficiently large that subjects could go bankrupt and \$25 represents the 90th percentile of overbids (rounded to the closest dollar). We collectively refer to medium and high overbids as "large" overbids.

[**Figure 1 about here**]

Consistent with previous experimental findings from Kagel and Levin (1993), there is frequent overbidding in all five treatments. Pooling across all treatments, 40% of all observations are overbids (compared with 64% in Kagel and Levin, 1993) and 76% of the subjects overbid at least once. Overbids are more than twice as common as underbids (16%), making it unlikely that overbidding can be explained as purely random noise. Many overbids cannot be characterized as small mistakes. Large overbids (e.g. overbids > \$12) occur for 18% of all observations and 44% of all subjects have at least one large overbid.[12]

Overbidding does not vanish with experience, again consistent with the existing experimental results. Compare behavior in the first five periods with behavior in the remaining fifteen periods.[13] Pooling across all treatments, the frequency of overbidding rises somewhat from 35% to 42%. This does not reflect a decrease in rational behavior as the proportion of observations for which the bid and value are equal also rises, from 37% to 47%. Instead, there is a dramatic decrease in the proportion of underbids (e.g. bid < value), which decrease from 28% of the observations to 11%. This suggests that underbids are largely being driven by mistakes. The growth in overbids takes place primarily for the low overbids which grow from 19% to 24% of all observation. In contrast, the proportion of high overbids (e.g. overbid ≥ \$25) remains steady at 10% in both the first five and remaining fifteen periods.

A quick visual inspection of Figure 1 shows that the frequency and nature of overbidding differs across treatments. Overbidding is more common in the control treatment (52%) than any other treatment. Comparing the two treatments with $K = .3$ with the corresponding treatments with

[12] Our focus is on overbidding, but some other statistics are worth noting. Ignoring observations in which the two bidders have equal values, 87% of all auctions are efficient (e.g. the bidder with the higher value wins). The average revenue across all treatments is \$40.82.

[13] Splitting the data unevenly into early and late periods gives a better sense of the dynamics than splitting it evenly, as most changes are confined to the early periods.

$K = .7$, bidding one's value is more common with better information and large overbids (e.g. bid − value ≥ \$12) are less frequent. This is true both with exogenously provided signals and endogenous signal acquisition. Beyond Figure 1, behavior also differed between the two locations. Overbids were more likely among subjects from Case (45% of all observations) than Yale (34% of all observations). This difference was even more noticeable if we focus on large overbids which are more than twice as likely for Case subjects (23% of all observations) than Yale subjects (11% of all subjects).

Not only is overbidding frequent, it is costly as well. Table 3 shows that irrational bidding, including both over- and under- bids, causes substantial payoff losses for bidders. The first column lists the average bidder payoffs, both pooled and broken down by treatments.[14] The second column lists what the average payoffs would have been if all subjects had bid rationally (e.g. bid = value) against their opponent's actual bids. The difference between column 1 and 2 measures how much bidders could have benefited by unilaterally changing to rational bidding. Pooling across all treatments, bidding rationally would have increased subjects' average payoffs by 15%. If only overbids are considered, average payoffs would have been increased 37% by bidding rationally (\$7.56 versus \$10.40). Column 3 shows what the average payoff would have been if subjects, using their actual bids, had faced opponents who bid rationally. Irrational bidding, particularly overbidding, generates a negative externality for other subjects reducing average payoffs by 10%. Column 4 reports what the average payoffs would have been if all subjects had bid rationally and faced others who bid rationally. Comparing columns 1 and 4, subjects' average payoffs are 24% lower than they would be if all bidders behaved rationally.

[**Table 3 about here**]

4.2 How Do Signals Affect Bidding?

We now turn to one of the central questions raised by our experimental design: how do signals affect bidding? To answer this question we focus on data from the treatments with exogenously provided signals, EX3 and EX7, as the effect of signals on bidding is confounded with the decision to purchase information in the endogenous signal acquisition treatments.

[14] Throughout this table, costs of acquiring information are not included in the payoff calculation.

Figure 2 illustrates the complex relationship between bidding behavior and the signals received by bidders. The data for this figure is drawn from observations in EX3 and EX7 where \$40 $\leq$ value $\leq$ \$60. By focusing on these intermediate values we allow for signals both substantially larger and substantially smaller than the bidder's value. The conclusions we draw from this limited dataset extend to the full dataset as will be demonstrated in the regression analysis presented below. Observations are broken into five categories based on the difference between signal and value. The first cluster of bars, on the left of the figure, shows the proportion of overbids as a function of the difference between signal and value. The probability of overbidding is a weakly-peaked function, with overbids most likely when the auction is perceived to be competitive (e.g. signal and value are relatively close).

[**Figure 2 About Here**]

Underlying the modest response of the frequency of overbidding to the signals are strong but differing responses by low, medium, and high overbids. The remaining two clusters of bars in Figure 2 compare pooled low and medium overbids (e.g. overbid < 25) with high overbids (e.g. overbid $\geq$ \$25). For the lower two categories of overbids, the probability function is strongly peaked and almost perfectly symmetric.[15] In contrast, the probability function for high overbids is U-shaped. The right arm of the U, representing cases where the signal is much greater than the bidder's value, appears higher than the left arm. Overbidding is not a heterogenous phenomenon. Overbids, particularly low overbids, are most frequent when the auction is perceived to be competitive, but the largest overbids tend to occur when the auction is perceived to be non-competitive, especially when the bidder seems to have little chance of winning.

These differing responses suggest differing motives underlying overbids. Ignoring bounded rationality for the time being, a spite-based model predicts that the probability function of overbidding should be an increasing function of the difference between the signal and value while a joy-of-winning model predicts a peak-shaped function (see the discussions in Section 3.2.1 and 3.2.2). Ignoring the highest overbids, the data is consistent with the joy-of-winning model. However, the asymmetry of the probability function for high overbids suggests that spite is playing an important role for this class of overbids.

[15] If the low and medium overbids are considered separately,the probability function is peaked in both cases (but more so for the low overbids).

Finding 1: (Incidence of Overbidding) *Subjects are more likely to overbid, but overbid to a lesser extent, in seemingly competitive auctions.*

4.3 Statistical Analysis of the Impact of Signals on Overbidding

Thus far we have only used a subset of the data to support the preceding finding and our analysis has been limited to a strictly visual examination of the data. To put ourselves on firmer ground we now turn to formal statistical analysis. Table 4 reports the results of Probit regressions studying the effect of signals on overbidding.

[**Table 4 About Here**]

In using Probits, we focus on the probability of overbidding (or of particular types of overbids) rather than on trying to explain the magnitude of overbids. This necessarily involves discarding a great deal of information from the dataset. However, any statistical model that attempts to treat overbids as a continuous variable is going to be fraught with difficulties because of the large spike at an overbid of zero. Our use of Probits also makes it simple to consider different types of overbids separately.

For the regressions shown in Table 4, the dataset consists of all bids from the CON, EX3, and E7 treatments. Note that an observation is a bid, not an auction, so there is one observation for each subject in each period. For Models 1 and 2, the dependent variable is a dummy for an overbid (e.g. bid $>$ value). In Model 3 the dependent variable is a dummy for low and medium overbids (e.g. overbid $<$ 25) and in Model 4 the dependent variable is a dummy for high overbids (e.g. overbid $\geq$ \$25). The base for all regressions is the control session. As independent variables, all of the regressions include dummies for the EX3 and EX7 treatments, the bidder's value (denominated in 1000s of ECUs), a dummy for periods 6 – 20, and a dummy for whether the session took place at Yale. Models 2 – 4 also include the absolute value of the difference between the bidder's signal about his opponent's value and his value. This is interacted with a dummy for the treatments with exogenously provided signals as the variable is not defined in the control sessions. The difference between the signal and value is then interacted with a dummy for the difference being positive and a dummy for the difference being negative, creating two independent variables. Separating positive and negative differences between a bidder's signal and value allows for asymmetric responses to information.

Model 1 allows us to identify basic treatment effects. It partly confirms our observation that the control session yields more overbidding as both the EX3 and EX7 parameters are negative. Only the EX7 parameter is statistically significant in Model 1, but if we replace the EX3 and EX7 dummies with a single dummy for the treatments with exogenous signal provision, the resulting parameter is negative and statistically significant at the 5% level. Unlike our impression from Figure 1, the difference between EX3 and EX7 is negative but not statistically significant.[16] Contrary to the prediction of a spite model, the parameter estimate for "Value" is positive, albeit not significantly so. The significant positive estimate for "Periods 6 – 20" should not be taken as evidence that subjects are not learning as it largely reflects the sharply decreased likelihood of underbidding with experience rather than a move away from bidding one's value.

Model 2 explores the effect of signals of the likelihood of overbidding. The parameter estimates for "Exogenous*(Signal – Value)*(Signal > Value)" and "Exogenous*(Signal – Value)*(Signal < Value)" are both negative and statistically significant at least at the 5% level. The difference between these two parameters is not statistically significant. The same general pattern is observed in Model 3, with a dummy for low and medium overbids as the dependent variable, but the measured effect is stronger. Removing the interaction with the sign of the difference, the marginal effect of difference between the signal and the value is 46% more negative for low and medium overbids than for all overbids combined.[17] Once again the difference between positive and negative differences is not statistically significant. The results of Model 4, with a dummy for high overbids as the dependent variable, are quite different from those for Models 2 and 3. The parameter estimate for "Exogenous*(Signal – Value)*(Signal > Value)" is positive and statistically significant while the estimate for "Exogenous*(Signal – Value)*(Signal < Value)" is actually slightly negative (although nowhere close to statistical significance). The difference between these two parameter estimates is statistically significant.[18]

[16] If we consider high overbids (e.g. overbid $\geq$ \$25) rather than all overbids, the difference between EX3 and EX7 is significant at the 1% level.

[17] Comparing Models 2 and 3 directly, the marginal effect is 125% larger for positive differences and virtually the same for negative differences.

[18] As secondary results, the parameter estimate for periods 6 - 20 is no longer statistically significant and high bids are significantly less likely in the Yale population.

For the most part the regressions contained in Models 2 – 4, which use all of the data rather than only a subset with intermediate values, confirm our impressions from Figure 2. As a function of the difference between a bidder's signal and value, the probability of overbidding is peaked, consistent with a "joy of winning" model. This pattern is more extreme if attention is restricted to low and medium overbids, but breaks down for high bids. The one surprise from the regressions is that our observation that the probability of high overbids is a lop-sided U-shaped is not quite correct. The left arm of the U is actually non-existent while the right arm is quite robust. This is the one case in which the data is consistent with a model of spite as high overbids are most likely when they are likely to harm the bidder's opponent without affecting his own payoff.

The regression analysis shown in Table 4 can be extended to answer additional questions about the data. Consider interacting the two variables measuring the difference between signals and values with a dummy for the EX7 treatment. We have rerun Models 2 – 4 with these two interaction terms added to the models. None of the interaction terms are statistically significant even at the 10% level. It does not appear that bidder's responses to information depends on the quality of this information. This should not come as a great surprise given the general difficulty experimental subjects have with Bayesian updating.

The results are more interesting if we interact the two variables measuring the difference between signals and values with the dummy for periods 6 – 20. Once again we rerun Models 2 – 4 with these two interaction terms added. In Models 2 and 3 the new variables have little impact, failing in all cases to achieve statistical significance at any standard level. However, in Model 4 the interaction term with "(Signal – Value)*(Signal > Value)" is *negative* and significant at the 5% level. This negative effect offsets 60% of the positive marginal effect for "(Signal – Value)*(Signal > Value)." With experience, the probability of subjects choosing a high overbid no longer depends strongly on the signal being greater than the value (although the effect is still statistically significant at the 10% level). This instability should be counted as a mark against models of spite as an explanation of overbidding.

4.4 Are Subjects Learning to Not Overbid?

Understanding how subjects' behavior changes with experience can be critical for separating explanations of overbidding that rely on non-standard preferences from those based on bounded rationality. If overbidding is largely a mistake, subjects should learn to stop making this mistake *to the extent it is costly*. As established in Table 4, overbidding increases with experience. However, in most cases subjects face no cost from overbidding. Only in 7% of all observations (and only 18% of all observations with overbids) do subjects overbid and lose money. In this subsection we explore whether *costly* overbidding causes subjects to learn not to overbid.

[**Figure 3 About Here**]

The data shown in Figure 3 indicates that subjects do learn from costly overbids. This data is drawn from all five treatments. The left pair of bars is based on data from all observation in which the bidder overbid in the previous period. In other words, the lagged bid was greater than the lagged value for these observations. The data is split into observations where this lagged overbid did not cause a loss and observations where, in the previous period, the bidder overbid, won the auction, and lost money. Only the latter case provides the correct experience for subjects to learn to not overbid. The graph reports the proportion of overbids in the current period for each of these cases. Not surprisingly, subjects who overbid in period $t-1$ also tend to overbid in period t. However, they are 13% less likely to do so if overbidding led to a loss. Breaking the data down by treatment, a subject who overbid in period $t-1$ is less likely to overbid in period t in all five treatments conditional on losing money in period $t-1$. This suggests that this relationship is not likely to be a coincidence. The right pair of bars reports data from all observations where the bidder had a high overbid (e.g. bid − value $\geq$ \$25) in the previous period. These bars show the proportion of high overbids in the current period, once again split by whether the bidder lost money in the previous period. High overbids are 12% less likely to be repeated if the bidder lost money. This pattern is again present across all five treatments.

[**Table 5 About Here**]

The regressions shown in Table 5 put our observations from Figure 3 on a firmer statistical basis. For both regressions the data set is all observations from all treatments except for observations from period 1. These are discarded to allow the use of lagged variables. For Model 1 the dependent variable is a dummy for whether an overbid occurred in the current period and for Model 2 the dependent variable is a dummy for whether a high overbid was observed. The control treatment serves as the base. As independent variables, both regressions include dummies for the other four treatments (EX3, EX7, END3, and END7), the value (denominated in 1000s of ECUs), a dummy for periods $6 - 20$, and a dummy for the location. A number of lagged dependent variables are included in both regressions. The critical variable in Model 1 is a dummy for whether the bidder overbid in the previous round and lost money. Similarly, the variable of interest in Model 2 is a dummy for whether the bidder submitted a high overbid in the previous period and lost money. If subjects are learning from negative experience to avoid a mistake, the estimates for these critical parameters should be negative. The other lagged dependent variables in Models 1 and 2 play an important role as well. Suppose we rerun Model 1 with no lagged dependent variables other than the dummy for having overbid and lost money in the preceding round. The resulting parameter estimate for this dummy is positive and statistically significant at the 1% level. This does not indicate that subjects are somehow learning to overbid more following a negative experience, but instead reflects the strong individual effects in the data. Essentially, we are regressing on the fixed effects. Including a dummy for the lagged overbid takes care of this problem. The parameter of interest now measures the effect of losing money *subject to having overbid*. Dummies for lagged large overbids (bid − value $\geq$ \$12) and lagged high overbids (bid − value $\geq$ \$25) allow for the possibility that the magnitude of the lagged overbid could drive a negative estimate for "Lagged Overbid and Lose Money." Losing money is more likely as the overbid increases. If there is regression to the mean in overbids and no variables controlling for the magnitude of the lagged overbid are present in the regression, a negative estimate for "Lagged [High] Overbid and Lose Money" may result even if no learning is taking place.[19] The dummy for "Lagged Overbid and Win" allows for the possibility that winning, rather than winning and losing money, drives a negative estimate for "Lagged Overbid and Lose Money." If subjects have a taste for winning, satiation could lead to a negative estimate.

The results for Models 1 and 2 are consistent with our observations from Figure 3 as the estimate for "Lagged [High] Overbid and Lose Money" is negative and statistically significant in both cases. When given the correct experience, subjects are less likely to either overbid or choose a high overbid.

[19] We have considered a variety of alternative specifications to control for the magnitude of the lagged overbid, including continuous function rather than the step function used here. The main results are robust to these alternative.

This indicates that any explanation of overbidding must include a bounded rationality component.

Finding 2: (Learning to not Overbid) *The evidence is consistent with subjects learning from costly overbidding to avoid mistakes. The apparent stability of overbidding is due to a paucity of opportunities to learn the costs of overbidding rather than a failure to learn from relevant experience.*

4.5 The Demand for Signals

We begin our discussion of the sessions with endogenous signal acquisition by examining when signals are acquired. Recall that the standard theoretical prediction about the demand for signals is very simple: a bidder should acquire a signal only if it has a negative price (see Proposition 1). Moreover, the theory predicts that bidders should pay no attention to the signals even if one is acquired as a result of a negative price. Given our earlier observations that subjects respond to their signals in the treatments with exogenously provided signals, our realistic expectation is that at least some subjects will pay for signals. Our goal is to determine whether there is a systematic pattern to when signals are acquired. We are particularly interested in whether signal acquisition fades with experience as this would be clear evidence of subjects learning to avoid a mistake.

[**Figure 4 About Here**]

Figures 4 graphs the demand curve for signals. Purchases of information are quite common. The signal is purchased for 35% of all observations, including 22% of observations where the cost of information is strictly positive. Most subjects (72%) purchase information at a positive cost at least once. Ignoring the fact that the information subjects purchase is intrinsically useless for their monetary payoffs, subjects are fairly rational in their purchase decisions. The likelihood of purchasing information decreases monotonically in the cost and subjects are more likely to purchase high quality signals at a positive cost (25%) than low quality signals (18%). Subjects at Case were slightly more likely to purchase signals at positive price than those at Yale (23% vs. 18%). Consistent with learning, the likelihood of purchasing information at a positive price decreases with experience, dropping from 28% in the first five periods to 19% for the remaining fifteen periods.

There is a clear link between overbidding and purchasing signals. There are sixteen subjects (out of 64 in the endogenous signal acquisition sessions) who never overbid. These subjects only purchase signals for 7% of the observations with a positive cost, compared with 26% for subjects who overbid at least once.

[**Table 6 About Here**]

The preceding conclusions are made more formally by the Probit regression shown in Table 6. The data set for these regressions is all observations from sessions with endogenous signal acquisition (END3 and END7). The dependent variable is a dummy for whether a signal was purchased. As independent variable, both regressions include a dummy for observations with a negative cost, the cost of a signal interacted with a dummy for observations with a (weakly) positive cost, a dummy for periods $6 - 20$, a dummy for the location, and a dummy for the quality of the signal. Model 2 also includes a dummy that equals 1 if the subject never overbid in any of the twenty periods.

The results from Model 1 provide mixed support for our preceding observations. The parameter estimate for the cost of information (subject to the cost being negative) is negative and significant at the 1% level, indicating that the demand curve for signals is downward sloping. The coefficient for periods $6 - 20$ is also negative and is significant at the 5% level. With experience subjects are significantly less likely to pay for information. This gives further credence to the idea that any explanation of overbidding must include a component of bounded rationality. Although the estimate for END7 is positive, it surprisingly fails to achieve statistical significance. It appears that any response by subjects to the quality of information is weak at best. The weak location effect we observed in the descriptive statistics also turns out (non-surprisingly) not to be significant. In Model 2, the dummy for subjects who never overbid is negative and significant at the 1% level. Overbidding and purchasing signals are closely connected phenomena, suggesting that these "mistakes" share a common cause. It is worth noting that the dummy for END7 becomes significant at the 5% level in Model 2.

Finding 3: (Purchasing Costly Information) *Subjects decisions to purchase costly information are consistent with rational choice, but, critically, subjects learn with experience to not purchase costly information.*

4.6 The Connection Between Signal Purchase and Overbidding

[**Figure 5 About Here**]

As suggested above, there is a strong link between purchasing information and overbidding. This relationship is illustrated by Figure 5. The bidding data shown in this figure is drawn from the sessions with endogenously acquired signals (END3 and END7) and only includes observation with a positive cost for information. As in Figure 1, bids have been broken down into five categories: underbids, bids equal to the value, and low, medium, and high overbids. The first cluster of bars is drawn from observations where a signal is purchased. The second cluster shows observations where a signal was not purchased. The final cluster shows observations from the eighteen subjects who never purchase information at a positive cost. Given that all of these subjects have numerous opportunities to purchase information at a positive price (all eighteen have at least twelve observations with a positive cost), they can be classified as strongly following the theoretical prediction of no costly signal purchases.[20]

Subjects who pay a positive cost for information are far more likely to overbid (and underbid as well) than those who do not. Subject to overbidding, subjects who purchase information at a positive cost are more likely to have large overbids. The relationship between purchasing information and overbidding becomes especially clear when those subjects who never purchase information at a positive cost are considered. These subjects bid their value for 72% of the observations while large overbids (e.g. bid - value $\geq$ \$12) are chosen for only 3% of all observations. This is as close to the theoretical prediction as we could ever hope to see. The subject population appears to be heterogeneous, consisting of types whose behavior, both in purchasing information and in bidding, is consistent with the theoretical predictions and types who violate the theoretical predictions across the board.

[**Table 7 About Here**]

[20]The maximum over these eighteen subjects of their respective minimum positive costs for information was 54 ECUs (with an average of 15 ECUs), so it is difficult to argue that they never had the opportunity to purchase information at a reasonable price. We have run probit regressions testing whether the number of observations with a positive price or the minimum positive price has predictive value for whether a subject ever purchases costly information. While both parameter estimates have the correct sign, neither even approaches statistical significance. As such it must be considered more than a coincedence that these subjects never paid a positive cost for information.

The Probit regressions reported in Table 7 explore how purchasing information affects subjects' likelihood of overbidding. For all of these regressions the dependent variable is a dummy for overbidding. As independent variables, all of the regression include a dummy for high quality information (END7), the value (denominated in 1000s of ECUs), a dummy for periods $6-20$, and a location dummy.

The data set for Model 1 is all observations from END3 and END7 with a positive cost for information. This regression includes a dummy for purchasing information as an independent variable. The parameter estimate associated with this variable is positive and statistically significant at the 1% level – purchasing information is strongly associated with overbidding. The problem with interpreting this result is causality – do informed bidders behave less rationally because they are informed or because they are the same types who think the (worthless) information has value.[21]

The inclusion of negative costs for information in the experimental design helps us to sort out this question. Model 2 includes only observations for which the cost of a signal was negative. Since information is purchased in 99% of these observations, our sample can be treated as if subjects are exogenously informed. The key variable in Model 2 is a dummy is for whether a subject ever purchased information with a positive cost, either before or after the current observation. If overbidding and purchasing costly information are driven by a common type, this dummy should be highly correlated with this type. Since subjects are (essentially) exogenously informed in the restricted sample for Model 2, "Ever Purchase at Cost ≥ 0" is not correlated with subjects' current information. Thus, its coefficient measures the influence of type separate from any direct effect of being informed. The parameter estimate for "Ever Purchase at Cost ≥ 0" is negative and significant at the 5% level. The marginal effect of this variable is almost identical to the marginal effect for "Signal Purchased" in Model 1 (21% versus 17%). This result suggests that most of the effect of becoming informed in Model 1 is driven by a subject's type rather than being informed *per se*.[22]

[21] As an alternative method of answering this question, we have run probit regressions which instrument for buying information. Specifically, we use a dummy for negative costs and the cost of information interacted with a dummy for positive costs as instruments. This takes advantage of the exogeneity of information costs. The resulting parameter estimate has just about the same marginal effect as shown in Model 1 but no longer achieves statistical significance.

[22] We have rerun Models 1 and 2 using high overbids (e.g. overbid $\geq$ \$25) as the dependent variable. The conclusions are virtually the same. The estimates for "Signal Purchased" in Model 1 and "Ever Purchase at Cost ≥ 0" are positive and significant at least at the 5% level, and the marginal effects for these two variables are almost

Model 3 looks at the question of whether giving information to subjects has any effect independent of subjects' types. We restrict the sample to those subjects whose actions indicate they believe the information is worthless, subjects who never paid a positive cost for information. This yields only a small dataset (18 subjects). As such, any results reported in Model 3 should be considered as suggestive at best. The reader particularly should be aware that the correction for clustering yields biased estimates of the standard errors when the number of clusters is small as is the case here (see Wooldridge, 2003, for example). The central variable in the regression is a dummy for whether the cost of information is negative. Given that information is almost always purchased in this case, this is equivalent to estimating the effect of exogenously provided information. The coefficient for this variable is positive but only weakly significant, suggesting that giving subjects information leads to more overbidding independent of their type.[23]

Finding 4: (Heterogeneity) *There is a strong relationship between signal purchase and overbidding. This relationship appears to be based on a common type rather than a causal relationship where becoming informed leads to overbidding.*

5 Discussion and Conclusions

This paper reports results from a series of second price auction experiments, where bidders are presented with exogenous signals about opponents' value, or with opportunities to purchase signals about opponents' value. Such signals are theoretically useless for the bidders if they are only concerned about their monetary payoffs, as assumed by standard auction models; but it provides a convenient way to change the bidders' perceptions about the value of their rivals. We examine how subjects' incidence and magnitude of overbidding varies with their perceptions about how their own value compares with that of their opponent, and use the empirical findings to shed new light to the question of why bidders overbid in second price auctions.

Our central goal in designing these experiments was to separate out various explanations for overbidding in second price auctions. Ex ante, the scale was tilted in favor of explanations that

identical (8% in both cases).

[23] If Model 3 is redone with high overbids as the dependent variable, the parameter estimate for "Cost < 0" becomes tiny and statistically insignficant. This reinforces our impression that any impact of information on overbidding, independent of type, is weak.

involve bounded rationality. Otherwise the differences between sealed bid second price auctions and English auctions are quite troublesome. Indeed, our experiments provide clear evidence in support of bounded rationality, as we find clear evidence of learning both in avoiding costly overbidding and in subjects' choices to purchase costly information. As to the nature of this bounded rationality, it is unlikely that a single cause for overbidding can be identified. Random errors do not appear to be the dominant explanation – the well-behaved demand for costly information argues strongly otherwise.

We also find that non-standard preferences may be partly responsible for the overbidding. Our experimental design provides us with the opportunity to see how the incidence and magnitude of overbidding reacts to bidders' perceptions about how their own value compares with that of their opponent. We find that bidders are more likely to overbid, though they are less likely to submit large overbids (e.g. overbid $\geq$ \$25), when they perceive that their own values are relatively close to that of their opponents. This is inconsistent to the "spite" hypothesis of overbidding, but lends support to a modified "joy of winning" hypothesis (see the hypothesis listed in Section 3.2).

Finally, we find that bidder heterogeneity is playing an important role in our data. There is a group of subjects in our dataset whose behavior is almost completely in line with the standard theoretical predictions. These subject do not purchase costly information and rarely overbid. Other subjects get everything wrong, both purchasing costly information and overbidding. These results are important for predicting the external validity of the experimental results. In the laboratory, overbidding can only be extinguished through learning, but in the field selection can play an equally important role. Given that there exist subjects who bid according to the theory, forces of selection may quickly drive out those subjects who are prone to overbidding.

References

[1] Andreoni, James, Yeon-Koo Che and Jinwoo Kim (2005). "Asymmetric Information about Rivals' Types in Standard Auctions: An Experiment." mimeo, University of Wisconsin-Madison.

[2] Fang, Hanming and Stephen Morris (2006). "Multidimensional Private Value Auctions." *Journal of Economic Theory*, Vol. 126, No. 1, 1-30.

[3] Fischbacher, Urs (1999). "z-Tree: Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual." Working Paper No. 21, Institute for Empirical Research in Economics, University of Zurich. harstad : Harstad, Ronald (2000). "Dominant Strategy Adoption and Bidders' Experience with Pricing Rules," Experimental Economics, Vol. 3, 261-280.

[4] Harrison, Glenn W. (1989). "Theory and Misbehavior of First Price Auctions." *Amerian Economic Review*, Vol. 79, No. 5, 749-762.

[5] Kagel, John, Dan Levin and Ronald Harstad (1987). "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica*, Vol. 55, No. 6, 1275-1304.

[6] Kagel, John and Dan Levin (1993). "Independent Private Value Auctions: Bidder Behavior in First-, Second- and Third Price Auctions with Varying Number of Bidders." *Economic Journal*, Vol. 103, 868-879.

[7] Kim, Jinwoo and Yeon-Koo Che (2004). "Asymmetric Information about Rivals' Types in Standard Auctions." *Games and Economic Behavior*, 46, 383-397.

[8] Morgan, John, Ken Steiglitz and George Reis (2003). "The Spite Motive and Equilibrium Behavior in Auctions." *Contributions to Economic Analysis & Policy*, Vol. 2, No. 1, Article 5.

[9] Vickery, William (1961). "Counterspeculation, Auctions and Competitive Sealed Tenders." *Journal of Finance*, Vol. 16: 8-37.

[10] Wooldridge, Jeffery (2003). "Cluster-Sample Methods in Applied Econometrics." *American Economic Review Papers and Proceedings*, Vol. 93: 133-138.

Table 1
Distribution of Values

Value	Probability (%)
0 ECUs	1 %
1,000 ECUs	3 %
2,000 ECUs	6 %
3,000 ECUs	12 %
4,000 ECUs	18 %
5,000 ECUs	20 %
6,000 ECUs	18 %
7,000 ECUs	12 %
8,000 ECUs	6 %
9,000 ECUs	3 %
10,000 ECUs	1 %

Table 2
Summary of Treatments

Location	Treatment				
	Control	Exogenous K = .3	Exogenous K = .7	Endogenous K = .3	Endogenous K = .7
Case	1 Session 16 Subjects	1 Session 34 Subjects	2 Sessions 20 Subjects	1 Session 18 Subjects	1 Session 24 Subjects
Yale	1 Session 20 Subjects	2 Sessions 24 Subjects	1 Session 30 Subjects	1 Session 12 Subjects	1 Session 10 Subjects

Table 3

Effect of Irrational Play on Payoffs, Subject Averages
(Standard Errors in Parentheses)

Session Type	Actual Play vs. Actual Opponent *	Rational Play vs. Actual Opponent	Actual Play vs. Rational Opponent*	Rational Play vs. Rational Opponent
Control Sessions	$8.38 (0.86)	$9.80 (0.62)	$10.06 (0.82)	$11.25 (0.65)
Exogenous K = .3	$6.58 (0.53)	$8.45 (0.43)	$8.73 (0.53)	$10.40 (0.39)
Exogenous K = .7	$10.75 (0.64)	$11.63 (0.62)	$10.87 (0.66)	$11.78 (0.60)
Endogenous K = .3	$10.29 (0.80)	$11.20 (0.71)	$10.68 (0.82)	$11.47 (0.73)
Endogenous K = .7	$10.27 (0.68)	$11.34 (0.61)	$9.35 (0.66)	$10.71 (0.57)
All Sessions Pooled	$9.03 (0.32)	$10.32 (0.27)	$9.86 (0.31)	$11.08 (0.25)

* These payoffs exclude any costs/benefits from purchasing information.

Table 4
Response to Signals: Probit Regressions on Data from CON, EX3, and EX7
Standard Errors Controlled for Clustering, 144 Subjects, 2842 Observations

	(1)	(2)	(3)	(4)
Dependent Variable	All Overbids	All Overbids	Overbids < \$25	Overbids ≥ \$25
EX3 (Exogenous × K = .3)	-.318 (.203)	-.194 (.206)	-.264 (.196)	.117 (.215)
EX7 (Exogenous × K = .7)	-.412** (.201)	-.298 (.204)	-.123 (.195)	-.365 (.225)
Value	.012 (.018)	.017 (.018)	.019 (.019)	-.005 (.023)
Periods 6 – 20	.183*** (.057)	.178*** (.056)	.159*** (.061)	.084 (.071)
Yale	-.246 (.154)	-.241 (.154)	-.027 (.142)	-.460*** (.156)
Exogenous * (Signal – Value)*(Signal > Value)		-.038** (.019)	-.096*** (.021)	.067*** (.026)
Exogenous * (Value - Signal)*(Signal < Value)		-.052*** (.018)	-.055*** (.020)	-.013 (.026)
Log-Likelihood	-1889.66	-1884.58	-1698.87	-944.37

Notes: Constants have been suppressed. *, **, *** denote significance at 10%, 5% and 1% respectively. Values and signals are denominated in 1000s of ECUs.

Table 5
Learning: Probit Regression on All Data
Standard Errors Controlled for Clustering
208 Subjects, 3914 Observations

	(1)	(2)
Dependent Variable	All Overbids	Overbids ≥ \$25
EX3 (Exogenous × K = .3)	-.186 (.128)	.228 (.164)
EX7 (Exogenous × K = .7)	-.229* (.126)	-.138 (.171)
END3 (Endogenous × K = .3)	-.193 (.146)	-.190 (.186)
END7 (Endogenous × K = .7)	-.380*** (.143)	-.184 (.192)
Value	.019 (.019)	-.030 (.020)
Periods 6 – 20	-.042 (.041)	-.116* (.059)
Yale	-.207** (.081)	-.272*** (.101)
Lagged Overbid	1.685*** (.103)	.336*** (.092)
Lagged Overbid ≥ \$12	.190 (.117)	.524*** (.114)
Lagged Overbid ≥ \$25	-.218* (.116)	.666*** (.191)
Lagged[High] Overbid and Win	-.034 (.090)	-.058 (.207)
Lagged [High] Overbid and Lose Money	-.389*** (.097)	-.351*** (.129)
Log Likelihood	-1914.76	1074.19

Notes: Constants have been suppressed. *, **, *** denote significance at 10%, 5% and 1% respectively. Values and signals are denominated in 1000s of ECUs.

Table 6

Demand for Information: Probit Regressions on Data from END3 and END7
Standard Errors Controlled for Clustering
64 subjects, 1280 observations

	(1)	(2)
Cost < 0	2.328*** (.263)	2.629*** (.302)
(Cost ≥ 0) * Cost	-6.159*** (1.108)	-6.777*** (1.098)
Periods 6 – 20	-.259** (.131)	-.324** (.141)
Yale	-.166 (.278)	.029 (.288)
END7 (Endogenous × K = .7)	.293 (.234)	.582** (.261)
Never Overbid		-1.186*** (.307)
Log-likelihood	-514.14	-472.03

Notes: Constants have been suppressed. *, **, *** denote significance at 10%, 5% and 1% respectively. Costs are denominated in 1000s of ECUs.

Table 7
Information Purchase and Overbidding: Probit Regressions on Data from END3 and END7
Standard Errors Controlled for Clustering

	(1)	(2)	(3)
Data Set	Cost ≥ 0 64 subjects, 1063 obs.	Cost < 0 63 subjects, 217 obs.	Never Purchase 18 subjects, 360 obs.
END7 (Endogenous × K = .7)	-.392* (.223)	-.556** (.275)	-.655 (.502)
Value	.033 (.027)	-.060 (.043)	.072** (.036)
Periods 6 – 20	.267** (.119)	-.014 (.248)	-.063 (.203)
Yale	-.557** (.229)	-.755*** (.279)	-.496 (.448)
Signal Purchased	.460*** (.154)		
Ever Purchase at Cost ≥ 0		.593** (.279)	
Cost < 0			.465* (.246)
Log-likelihood	-645.48	-130.60	-155.98

Notes: Constants have been suppressed. *, **, *** denote significance at 10%, 5% and 1% respectively. Costs are denominated in 1000s of ECUs.

Figure 1

Distribution of Overbids

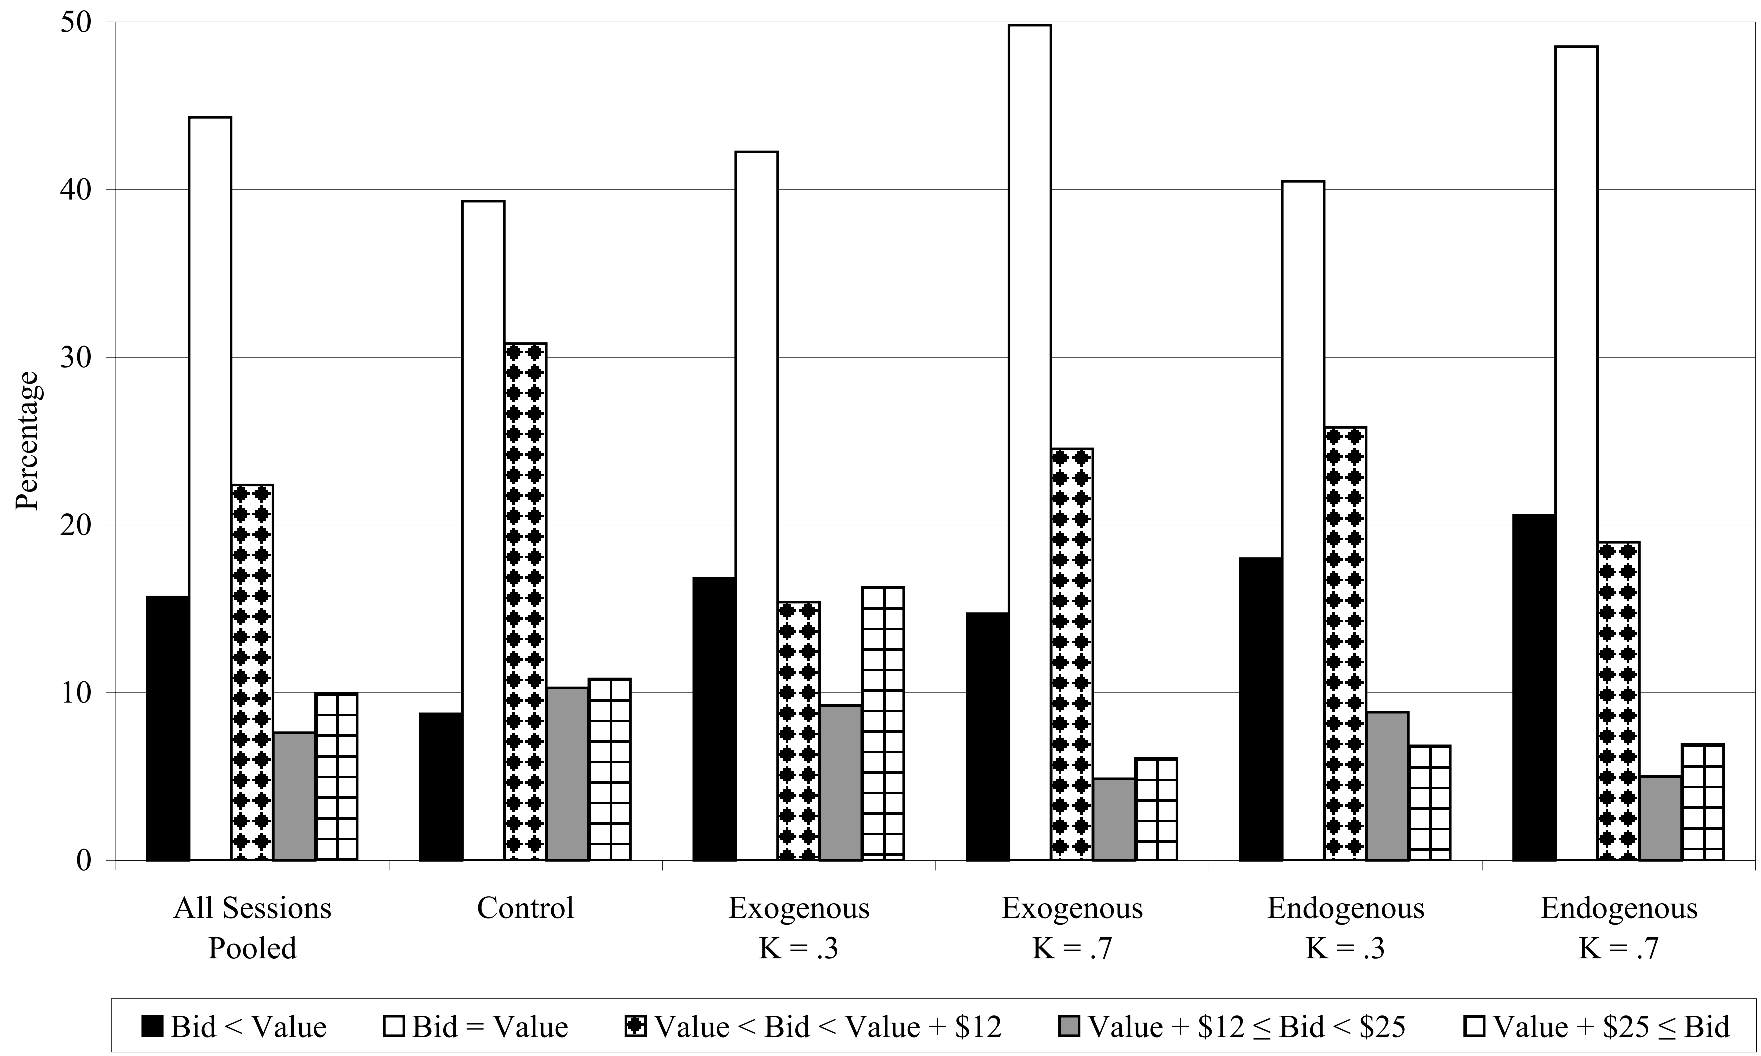

Figure 2
Relationship of Overbids and Signals

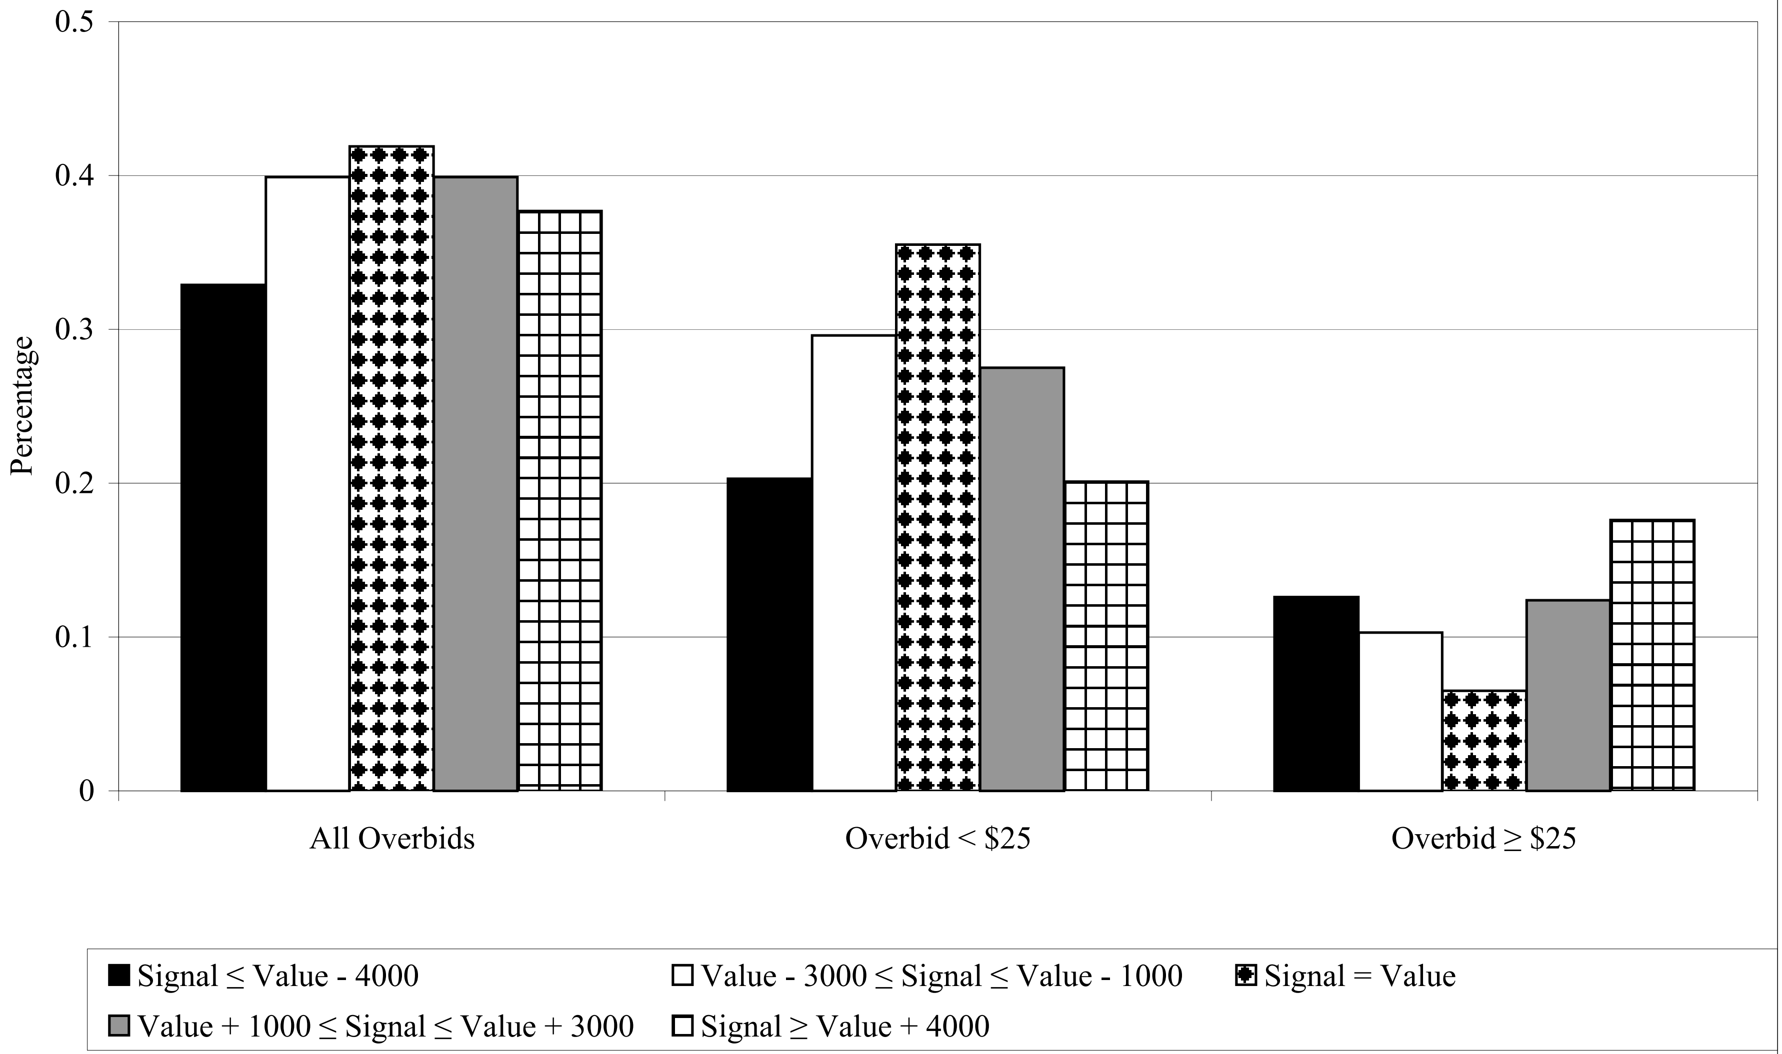

Figure 3
Learning to Not Overbid

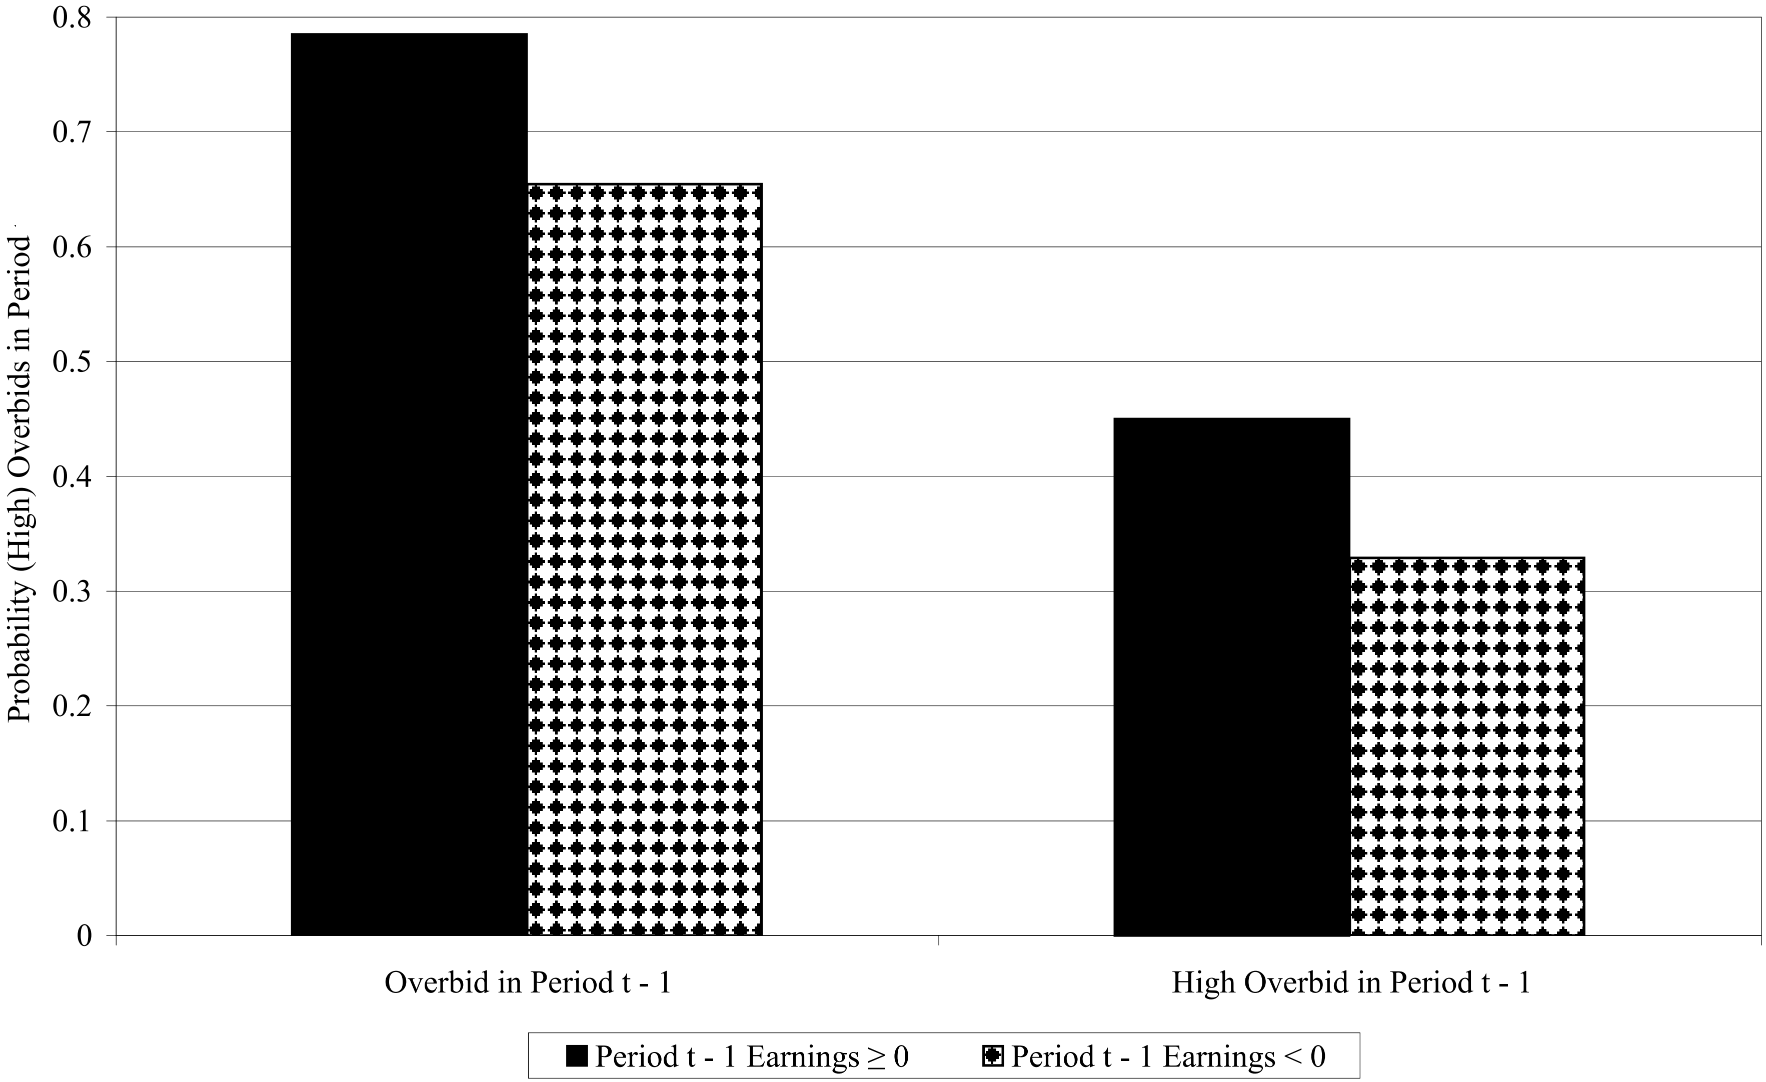

Figure 4
Demand for Signals

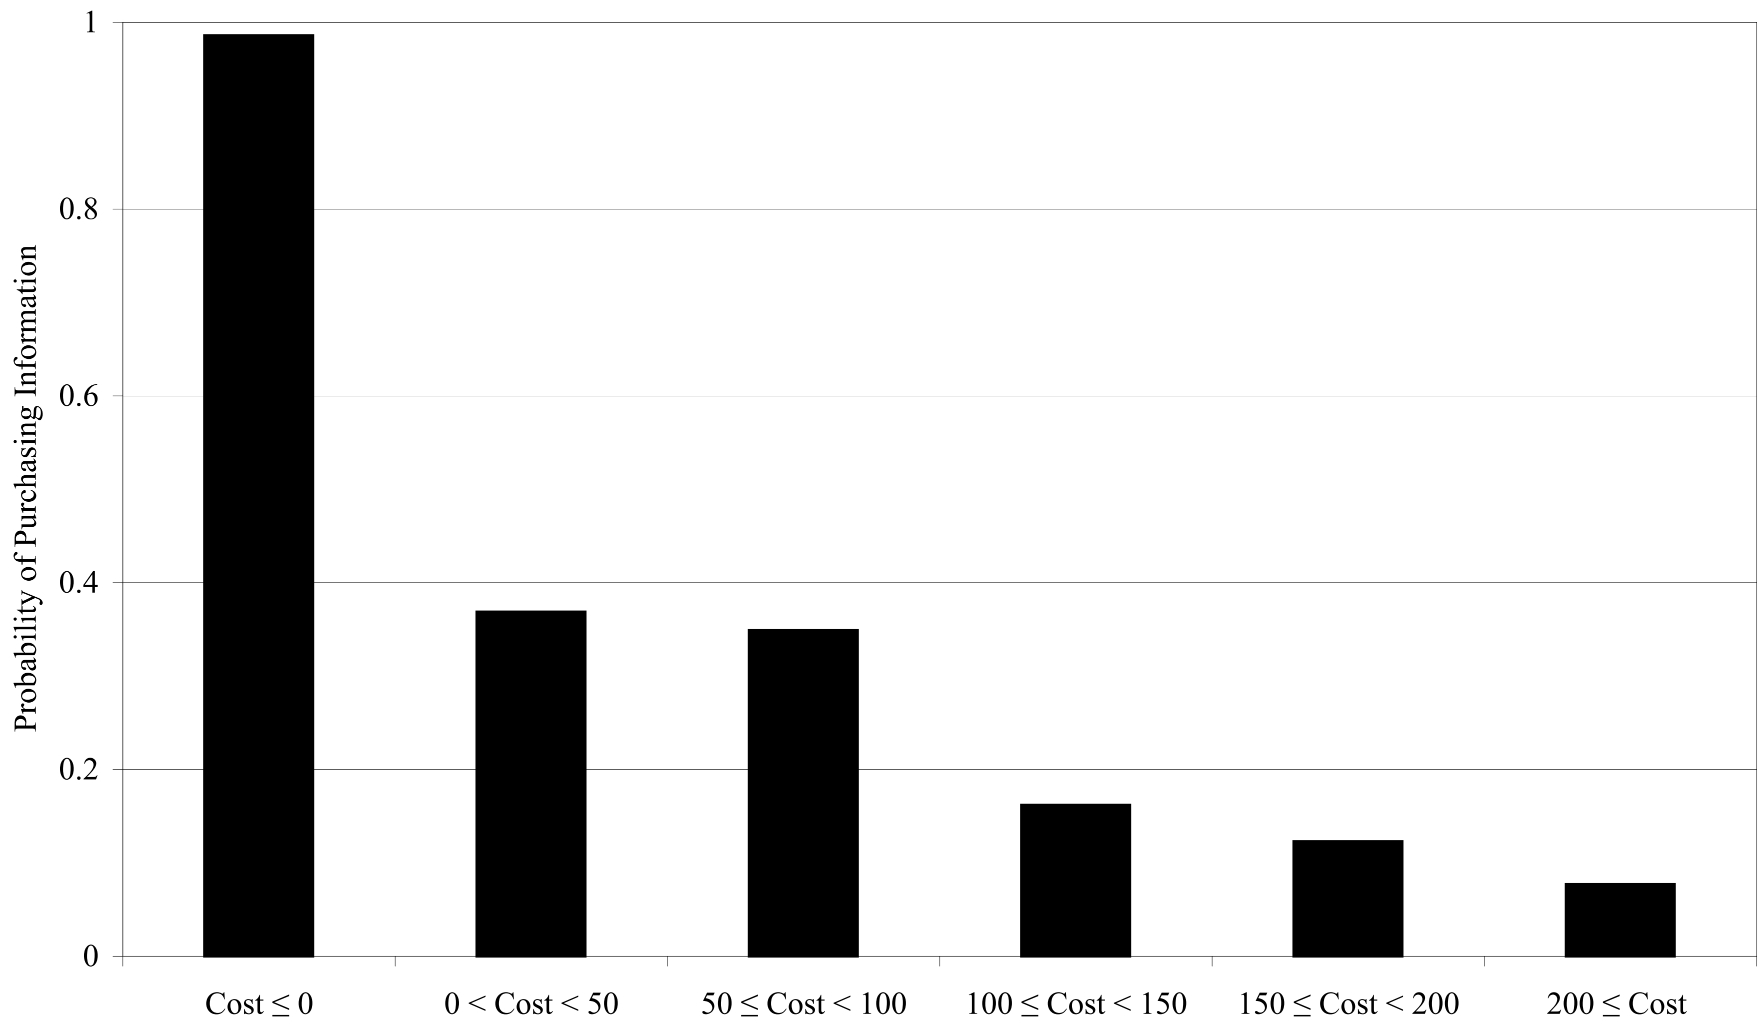

Figure 5
Signal Purchase at a Positive Cost and Bidding Behavior

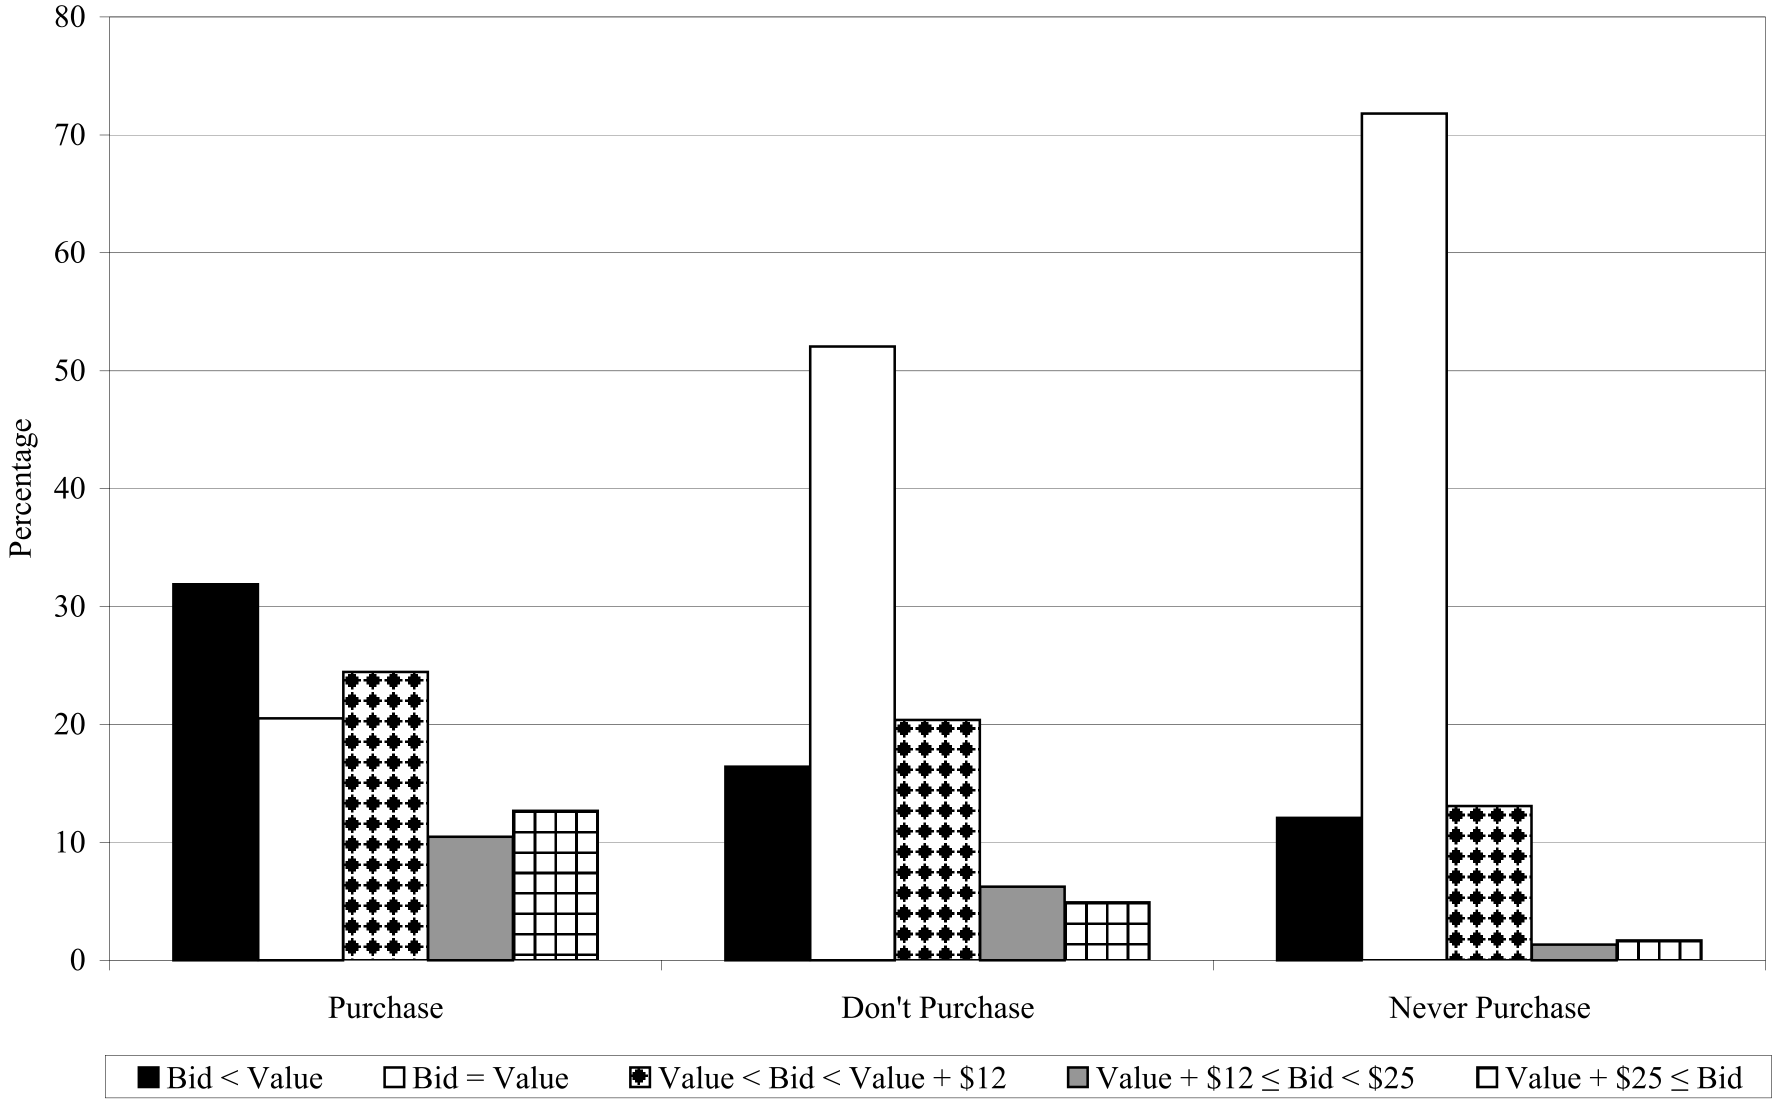

A Appendix: Instructions and Quiz for EX7 Treatment

Introduction

Today you are participating in a decision making experiment. These instructions describe a game that you will play 20 times. To make money in this experiment you must follow the instructions closely. Your payoffs in this experiment will depend on the choices made by you and the other players you are matched with. You will be given $12.00 for coming on time. This $12.00 and any money that you earn during the experiment will be paid to you, in cash, at the end of the experiment. It is possible to lose money in this experiment. Any losses will be deducted from your $12.00. Your total payoff will always be non-negative. You should feel free to make as much money as possible. Money has been provided for this experiment by Cowles Foundation for Research in Economics at Yale University (Faculty Research Fund at Case Western Reserve University).

If you have any questions while these instructions are being read, please raise your hand and we will attempt to answer your questions. *Please do not talk with the other subjects, even to ask questions about the instructions. If we hear you talking at any point in the experiment other than to talk with me or one of my assistants, you will be removed from the room and will not receive any payment. You will be barred from participating in any future economics experiments at Yale University (Case Western Reserve University).*

Matching

For each round of play you will be randomly and anonymously paired with another player. You will not know whom you are matched with nor will any other bidder know whom you are matched with while you are playing the games. Further, no bidder will know whom he or she was matched with after the experiment is finished. To repeat, you are not being matched against the same individual in each round. The matching is randomly redone at the beginning of each round.

Experimental Currency

All experimental payoffs are denominated in experimental currency units (ECUs). Your ECU earnings will be converted to dollars at the end of the experiment at a conversion rate of 100 ECUs equal one dollar. For example, suppose you have earned 1525 ECUs in the experiment. This would be divided by one hundred to give you a monetary payoff of $15.25. With your show-up fee of 12 dollars, this would give you total earnings of $27.25 for the session.

Auction Rules: Part One

In each round of the experiment, you will have the opportunity to bid on a single unit of a fictitious commodity. This fictitious commodity will have some "value" to you - you can think of this value as being the amount of money (in ECUs) that we will pay you for the item if you obtain it in the auction. Before any bidding takes place in any round, you will be told the value of the fictitious commodity to you. The other bidder you are matched with will know his or her own value for the fictitious commodity. You will also have some information about the value of the fictitious commodity for the bidder you are matched with. The bidder you are matched with will also have information about your value. How the values for you and the other bidder will be generated is described in detail below. How the information you have about others' values (and they have about your value) is generated is also described in detail below.

You and the bidder matched with you will submit bids for the fictitious commodity. (The computer screen will show an abbreviated summary of the instructions, your value, and your information about the values of others. You will then be prompted to submit a bid.) Like the values, these bids will be denominated in ECUs. Your bid must be in whole numbers (no fractions or decimals will be allowed). Negative bids will not be allowed. You will not know the bid of the bidder you are matched with at the time you submit a bid, nor will the bidder matched with you know your bid when they are choosing a bid.

Auction Rules: Part Two

The high bidder in an auction will obtain the unit of the fictitious commodity. The high bidder pays the second highest bid for the fictitious commodity. The high bidder's payoff (in ECUs) for the auction is then the difference between their value for the fictitious commodity and the second highest bid for the fictitious commodity:

$$\text{Payoff of High Bidder} = (\text{Value of Commodity}) - (\text{Second Highest Bid}).$$

The low bidder does not obtain a unit of the fictitious commodity and has a payoff of zero.

If there is a tie for the high bid, one of the high bidders is randomly selected to obtain the unit of the fictitious commodity. In this special case, the second highest bid is equal to the highest bid. The individual who is not picked as the winner will receive a payoff of zero in that round.

For example, you have a value of 5,000 ECUs. Suppose the other bidder matched with you for the round bids 4,000 ECUs. If you bid 4,500 ECUs, you obtain the unit of the fictitious commodity at a price of 4,000 ECUs (the second highest bid) and earn a payoff of 1,000 ECUs (5,000 - 4,000 = 1,000). If you bid 3,500 ECUs, you do not obtain the unit of the fictitious commodity and earn a payoff of 0 ECUs.

All of the rounds of the experiment are independent of each other. Your bids and payoffs in one round have no impact on your payoffs in any other round.

How Are the Values Generated?

Your value and the value of the bidder you are matched with is generated in each round from a random distribution. The probability of each possible value is given by the table below. You have been given a printed copy of this table as well.

[Table 1 Here]

The values that you receive in any two rounds are "independent." This means that knowing your value for any one round gives you no additional information about your values for other rounds. Likewise, the values of the bidders matched with you are independent. Knowing the value of the bidder matched with you in one round tells you nothing about the values of the bidders matched with you in other rounds.

How are the Signals Generated?

In each round you will receive a signal that gives you information about the value of *the other bidder you are matched with*. The bidder you are matched with will also receive a signal about your value.

Your signal is correct 70% of the time. In other words, there is a 70% chance that your signal is exactly equal to the realized value of your opponent. If you signal is incorrect, then each of the other values (different from the realized value of your opponent) is equally likely to be drawn. In other words, there is a 3% chance of drawing each of the incorrect values.

Your opponent also receives a signal about your realized value. His or her signal about your value is generated in an analogous manner.

The signals that you receive in any two rounds about the values of other bidders are independent. This means that knowing your signal for any one round gives you no additional information about what signals will be in other rounds. Likewise, the signals of other bidders about your value are independent across rounds.

Feedback

At the end of each round, the computer will give you a summary of the outcomes for the round. In particular, it will tell you your value for the round, your bid, the bid of the other bidder you are matched with and your payoff for the round. You have a record sheet that you can use to record this information.

Payoff in Dollars

At the end of the experiment, the computer will select one of the rounds at random (all 20 rounds are equally likely to be selected). Your ECU earnings for this round will be converted to dollars. The conversion rate of ECU's to the dollar is 100 to 1. You will then be paid your converted earnings plus your 12 dollar show-up fee. It is possible to lose money in a round of this experiment. If you have losses for the selected round, these losses will be deducted from your 12 dollar show-up fee. In the extremely unlikely event that you lose more than 12 dollars in the selected round, you will receive a monetary payoff of $0.

You will be paid in cash at the conclusion of the experiment. You will be paid privately, and no other bidder will be told what you earned for the experiment.

Summary

The experiment will consist of 20 rounds. In each round you will be randomly matched with another bidder. In each round you and the bidder you are matched with will be given values for a unit of a fictitious commodity. You will also be given a signal about the other bidder's value for this unit. Bidders will then bid on the unit. The high bidder obtains the unit at a price equal to the second highest bid. The high bidder's payoff is equal to the difference between his value and the second highest bid. All other bidders get a payoff of zero. At the end of the experiment, one of the 20 rounds will be randomly selected. You will be paid your earnings for this round (converted from ECUs to dollars) plus your $12 show-up fee. You will be paid in cash at the end of the experiment. All information about your choices and payoffs in this experiment will be kept strictly confidential.

Please do not talk with the other subjects at any point during the experiment, even to ask questions about the instructions. If we hear you talking at any point in the experiment other than to talk with me or one of my assistants, you will be removed from the room and will not receive any payment. You will be barred from participating in any future economics experiments at Yale University (Case Western Reserve University).

If you have any questions about any part of the instructions, this would be a good time to raise your hand. We want everyone to understand the instructions before we begin the experiment.

Quiz on Instructions: Part 1

Before we begin the experiment, we would like to confirm that everyone understands the instructions. Please complete the following questions. If you have any questions, please raise your hand. The computer will prompt you if you make an incorrect answer.

The first part of this quiz asks you questions about the rules of the experiment. Assume that you have drawn a value of 8,000 ECUs in round 1.

1. If you submit a bid of 5000 ECUs, and the other bidder you are matched with in round 1 has placed a bid of 9,000 ECUs, what will your payoff be in ECUs in round 1?

2. If you submit a bid of 6,500 ECUs, and the other bidder you are matched with in round 1 has placed a bid of 6,000 ECUs, what will your payoff be in ECUs in round 1?

3. If you submit a bid of 10,000 ECUs, and other bidder you are matched with in round 1 has placed a bid of 9,000 ECUs, what will your payoff be in ECUs in round 1?}

Quiz on Instructions: Part 2

Reminder: You have a table giving you the probability of drawing each value. The probability of receiving a correct signal about your opponent's value is **70%**.The probability of receiving each incorrect possible signal is **3%**.

1. What is the probability in percentages that you draw a value of 4,000 ECUs?

2. What is the probability in percentages that you draw a value of 7,000 ECUs?

3. If you have a value of 2,000 ECUs, what is the probability in percentages that the bidder you are matched with has a value of 2,000 ECUs?

4. If you have a value of 8,000 ECUs in the current round, what is the probability in percentages that you draw a value of 8,000 ECUs in the next round?

5. If the bidder you are matched with has a value of 3,000 ECUs, what is the probability in percentages that you receive a signal of 3,000 ECUs?

6. If the bidder you are matched with has a value of 7,000 ECUs, what is the probability in percentages that you receive a signal of 3,000 ECUs?

Quiz on Instructions: Part 3

1. Does the bidder you are matched with know your value for sure? Yes ____ No ____
2. Does the bidder you are matched with know your signal about his or her value for sure? Yes ____ No ____
3. You will be matched with the same bidder in all 20 rounds. True _____ False _____
4. We will pay you your converted payoffs from all 20 rounds. True _____ False _____
5. After the experiment has finished, we will not tell any other bidder about the choices you have made. True _____ False _____
6. We will never let any other bidder know what payoff you have received for participation in this experiment. True _____ False _____

We will now begin the experiment. If you have any questions at any point in the experiment, please raise you hand and we will answer your question as fully as possible. Please feel free to make as much money as you can!!!