The Economic Journal, **118** (*October*), 1572–1595. © The Author(s). Journal compilation © Royal Economic Society 2008. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

UNDERSTANDING OVERBIDDING IN SECOND PRICE AUCTIONS: AN EXPERIMENTAL STUDY*

David J. Cooper and Hanming Fang

We present results from second price private value auction (SPA) experiments where bidders may receive noisy signals about their opponents' value. Such signals change bidders' perception about the 'strength' of their opponent, and the relationship between bidders' perception of their opponent and overbidding provides a novel channel to understand overbidding in SPA. We found that small and medium overbids are more likely to occur when bidders perceive their rivals to have similar values, supporting a modified 'joy of winning' hypothesis but large overbids are more likely to occur when bidders believe their opponents to have much higher values, consistent with the 'spite' hypothesis.

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 *et al.* (1987) found that the actual bids are on average 11% above the dominant strategy bids. Kagel and Levin (1993) found that about 62% of all bids in their five-bidder SPA sessions exceed the bidder's value, while only 8% of all bids were below it. Both Kagel and Levin (1993) and Harstad (2000) further reported that experience has only a small effect in reducing overbidding in SPA. An important related finding is that overbidding in English auctions, a mechanism that is strategically equivalent to the 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 little understanding of *why* it happens. Kagel *et al.* (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 and because the negative feedback from overbidding is a weak mechanism in the SPA. For example, if a bidder overbids in a SPA by 10% above her value, it is quite possible that she does not win at all and thus does not experience

* We are grateful to the editor, two anonymous referees, John Kagel, Dan Levin, Stephen Morris, Shyam Sunder and participants of several conferences and seminars for helpful suggestions. The experiments were made possible by financial support from Case Western Reserve University (Cooper) and the Cowles Foundation (Fang). This article is a condensed version of our Cowles Foundation Discussion Paper No. 1557 (Cooper and Fang, 2006). We are responsible for all remaining errors.

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

[2] Harrison (1989) used similar 'flat maxima' arguments to explain overbidding in the first price auctions (FPA).

any negative feedback; even if she wins, it is still possible that she obtains positive payoff from winning (in stark contrast to the FPA, first price private value auction). Understanding overbidding in the SPA is of interest in itself but is also valuable for understanding individuals' anomalous behaviour in other games because overbidding in the SPA is a notable example of the use of dominated strategies.

There are several candidate explanations for overbidding in auction experiments. The first group of explanations is based on *non-standard preferences* such as 'spite' and 'joy of winning'.[3] The second explanation is bounded rationality where systematic errors in reasoning are hypothesised to explain overbidding.[4] While there are many studies documenting overbidding in SPAs, few experimental designs are aimed at understanding *why* there is overbidding. A recent exception is Andreoni *et al.* 2007, ACK henceforth) who conducted a set of experiments where bidders are partitioned into groups such that bidders within a group can perfectly observe each other's value.[5] They found that overbidding in SPAs is much more prevalent among 'followers' – bidders whose values are known to be lower – than 'leaders', providing evidence of the role of spite in overbidding.[6]

In this article, we report results from a series of SPA experiments in which subjects either receive for free or choose to purchase noisy signals about their opponents' value. The design of these experiments draws on Fang and Morris' (2006) theoretical analysis of auctions with information about others' values. There are two differences between our experimental design and ACK's. First, bidders in our experiments receive noisy but informative signals about opponents' value, whereas ACK's bidders perfectly observe rivals' values within a group but have no information about rivals outside their own group. Our noisy signal setup provides potentially a higher degree of variation in bidders' perceptions about their opponents' strength than ACK's perfect signal setup. Second, our experimental design includes treatments in which bidders decide whether to purchase signals about opponents' value. Because the cost of information acquisition always has a direct and obvious effect on feedback about payoffs (while the cost of overbidding is typically not obvious and not direct), the endogenous information acquisition treatments allow us to evaluate the learning and bounded rationality hypothesis more forcefully; they also allow us to evaluate the potential importance of bidder heterogeneity by examining the relationship between mistakes in the domains of information acquisition and overbidding.

[3] Morgan *et al.* (2003) showed that bidders will bid more than their values in the SPA when they care not only about their own surplus in the event that they win the auction but also about the surplus of their winning rival in the event that they lose the auction (referred to as 'spite'). A weakness of spite as an explanation for overbidding in the SPA is that it also predicts overbidding in English auctions (Morgan *et al.*, 2003). Cox *et al.* (1992) attempted to explain overbidding in first price auctions with 'joy of winning'. As shown below, 'joy of winning' also implies overbidding in SPAs. Some variants of non-standard preferences that lead to overbidding in the FPAs do not explain overbidding in SPAs. Notably, Filiz-Ozbay and Ozbay (2007) and Engelbrecht-Wiggans and Katok (2007) argued that anticipated loser regret may explain overbidding with respect to risk neutral Nash equilibrium in the FPA. However, anticipated regret should have no effect on bidding in SPA (Filiz-Ozbay and Ozbay, 2007).

[4] Not all models of bounded rationality predict overbidding in SPAs. For example, Crawford and Iriberri (2007) show that level-k models cannot explain overbidding in SPAs.

[5] ACK's experiments are motivated by the theoretical analysis in Kim and Che (2004).

[6] In contrast, ACK's results from the FPA experiments favour risk aversion, instead of spite, as an explanation for the slight amount of overbidding observed there.

Even though the noisy signal about opponents' value has no strategic use in a standard SPA with fully rational bidders that care only about monetary payoffs,[7] these signals provide us with an instrument to change bidders' perception about the 'strength' (i.e., the value) of their opponents. The empirical relationship between bidders' perception of the strength of their opponents and the incidence and magnitude of overbidding provides an additional lens through which we can learn about the incentives for overbidding in the SPA.[8]

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 within the framework of fully rational bidders with non-standard preferences, the empirical relationship between small and medium overbids and bidders' perception about their opponents' strength, is more consistent with a modified version of the 'joy of winning' hypothesis than the 'spite' hypothesis but the relationship for large overbids is instead consistent with the 'spite' hypothesis. However, neither of the non-standard preference explanations fully explains all aspects of our experimental data. We find clear evidence of learning both in avoiding costly overbidding and in subjects' choices to purchase costly information, providing strong evidence of the role of bounded rationality. We also find that bidder heterogeneity plays an important role in explaining their bidding behaviour.

The remainder of the article is structured as follows. Section 1 describes the experimental design. Section 2 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 testable hypotheses regarding overbidding in the SPA from models with 'spite' and 'joy-of-winning' preferences or with bidders of bounded rationality. Section 3 analyses the experimental data. Finally Section 4 discusses our findings and concludes.

1. Experimental Design and Procedures

1.1. *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. Since we have a large number of subjects per session (see Table 1), a subject's data can reasonably be treated as a series of twenty single-shot games.

Value distribution. Prior to submitting a bid, a bidder privately observes her own value drawn from the discrete distribution, common for all sessions, as shown in Figure 1. All values are denominated in Experimental Currency Units (ECUs), which were

[7] Bidding one's own private value remains the weakly dominant strategy regardless of the signals about the opponent's value (see Proposition 1 in Section 2.1).

[8] This key feature of our experimental design is similar to that of Eliaz and Schotter (2006). In single-agent decision experiments, they showed that subjects are willing to pay quite a significant amount for information that in theory has no effect on their choices. They argued that this provides evidence that decision makers derive an intrinsic utility from their posterior beliefs. See also footnote 32 in Section 3.4 for more discussion about their findings.

Table 1
Summary of Treatments

Location	CON	EX3	EX7	END3	END7
Case	1 Session 16 Subjects	2 Session 34 Subjects	1 Session 20 Subjects	1 Session 18 Subjects	1 Session 24 Subjects
Yale	1 Session 20 Subjects	1 Session 24 Subjects	2 Session 30 Subjects	1 Session 12 Subjects	1 Session 10 Subjects

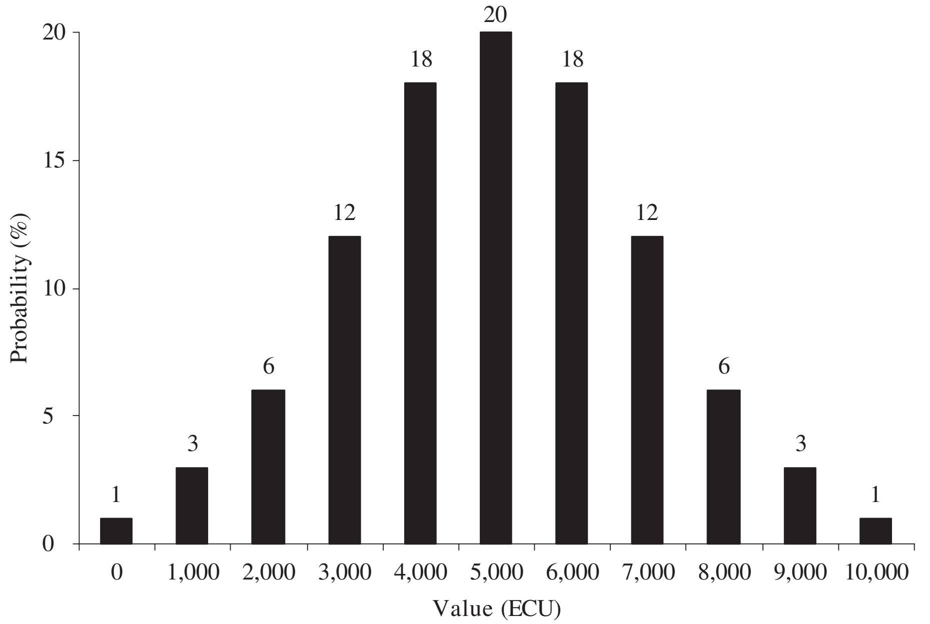

Fig. 1. *Distribution of Values*

converted to cash at a rate of 1 ECU = \$0.01.[9] The distribution of values is common knowledge; and the values are independently drawn across bidders and across rounds. This value distribution approximates a Normal distribution with mean 5,000 ECU and standard deviation of 2,000 ECU. We used a peaked distribution rather than a uniform distribution in order to generate more competitive auctions (e.g. auctions where the bidders' values are relatively close) with a sufficiently wide range of possible values.

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 0.3 or 0.7 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

9 Results related to earnings are reported in dollars rather than ECUs below.

variable K measures the quality of the signal. The value of K is common knowledge; but the signals observed by subjects are private information.

1.2. *Experimental Treatments*

Our experimental design includes 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 as a baseline for comparisons with the other treatments in our experiment, as well as the qualitative results on overbidding in SPAs in the existing literature.

Exogenously Provided Signals Treatments [EX3 and EX7]. In these treatments, at the same time they receive their private values, subjects in each round are provided with a *free* signal of quality K about their opponent's values, where K is equal to 0.3 for the EX3 ('low quality' signals) and 0.7 for the EX7 ('high quality signals') treatments. The value of K is the same for all bidders and for all rounds of a session. Since the signal is substantially more informative in the high quality signal treatment,[10] 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 $K = 0.3$ for END3 and $K = 0.7$ for END7 treatments, with K known to the subjects. *Prior* to being told her private value in each round, each subject is told her cost of information for the round, c, which is randomly drawn from a uniform $[-50, 250]$ and denominated in ECU.[11] These costs are independent across bidders and rounds. She is then asked whether she wishes to buy a signal. If she chooses to buy information, she will privately receive both her own value and a signal about the value of her opponent; her information cost c will be deducted from her show-up fee regardless of the outcome of the auction.[12] The decisions to acquire information are known to both bidders prior to the bidding.

Remark. There are two purposes for the endogenous signal acquisition treatments. First, they allow us to verify an important prediction of models of bounded rationality. If overbidding represents a mistake, rather than maximisation 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 individuals to understand the nature of the mistake since reinforcement learning is sufficient to

[10] For example, suppose a bidder receives a signal of 0 about their opponent's value. The updated expected value of their opponent is 4,841 ECU and 4,087 ECU respectively in the low and high quality signal treatments.

[11] An interesting avenue for future research is to have the bidders make a decision about information acquisition *after* they observe their own value. It is useful to investigate, both theoretically and experimentally, whether the incentives to acquire information vary with one's own value and how this can provide a further channel for understanding overbidding in the SPA.

[12] 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 SPA. 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, which we suspect would influence the results.

yield a reduction in errors. As noted by Kagel and Levin (1993), one of the reasons subjects have difficulty learning not to overbid in SPAs is that overbidding only rarely costs them money. In contrast, paying a positive price for information always costs money. If bounded rationality is a major force underlying anomalous behaviour, we expect to see subjects learning to avoid the always costly mistake of paying a positive price for information. A second purpose of the endogenous signal acquisition treatments is that they allow us to separate subjects 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 the information acquisition domain tend to make mistakes in bidding domains as well.

1.3. *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. These sessions are allocated to the five treatments as detailed in Table 1. The number of participants in each session varied between 10 to 24. Subjects were only allowed to participate in a single session.

All sessions were run in a computerised laboratory using the software z-Tree (Fischbacher, 2007). 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.[13] 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. 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.[14] 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.

At the end of the session, each subject was paid in cash for the earnings from a single *randomly* selected round plus their show-up fee of $12. We paid the subjects on a randomly selected round so that income effects will not be a confusion for any learning effects. Where subjects lost money for the randomly selected round, these losses were deducted from their show-up fee.[15] Finally, the average payoff was approximately $21

[13] The full text of the instructions for the EX7 treatment can be found in Appendix A of Cooper and Fang (2006).

[14] Thus, our feedback provides bidders with the information needed to determine regret as either a winner or a loser. However, recall from footnote 3 that the anticipated regret hypothesis does not lead to overbidding in SPA (Filiz-Ozbay and Ozbay, 2007).

[15] We never attempted to collect money from a subject, so losses are effectively capped at $12. Note that bidding one's value is still a dominant strategy when losses are capped at $12.

with sessions generally lasting 60–75 minutes. These payoffs were sufficient to generate a plentiful supply of subjects.

2. Theoretical Predictions

Now we derive the theoretical predictions, specific to our experimental design, 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 alternative models where bidders have either non-standard preferences (either 'spite' or 'joy of winning') or bounded rationality.

2.1. *Benchmark Theoretical Predictions: Fully Rational Bidders with Standard Preferences*

Denote bidder i's private value by v_i and her signal about opponent's value by s_i. In the benchmark model with fully rational bidders of standard preferences, the payoff for bidder i with value v_i (not counting the information acquisition cost c) when she bids b_i and her opponent bids b_j is given by:

$$U(b_i, b_j; v_i) = \begin{cases} 0 & \text{if } b_i < b_j \\ (v_i - b_j)/2 & \text{if } b_i = b_j \\ v_i - b_j & \text{if } b_i > b_j. \end{cases} \quad (1)$$

The equilibrium prediction is that a bidder's signal about her rival's values should have no effect on how much they should bid (Fang and Morris, 2006):

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

(*i*) *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.*
(*ii*) *When information acquisition is endogenous, a bidder should purchase the information only if the cost* c *is negative; and she should bid* v_i *regardless of her signal of her opponent's value.*

2.2. *Bidders with Spite Motives*

Bidders are said to be motivated by 'spite' if they care not only about their own surplus in the event of winning the auction but also about the surplus of their winning rival in the event that they lose the auction. Morgan *et al.* (2003) incorporates bidder i's spite motives into her payoff function as follows:

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

where $\alpha \in [0,1)$ represents the strength of spite with $\alpha = 0$ corresponding to standard auction model without spite motives. Morgan *et al.* (2003) derived the equilibrium of auctions in environments where bidders do not receive any noisy signals about opponents' value. For SPA, they show that the equilibrium bidding function is given by:

$$\beta(v) = v + \frac{\int_v^1 [1 - F(t)]^{(1+\alpha)/\alpha} \mathrm{d}t}{[1 - F(v)]^{(1+\alpha)/\alpha}},$$

where $F(\cdot)$ is the distribution from which bidders' values are drawn. Morgan *et al.* (2003) also showed that equilibrium level of overbidding $\beta(v) - v$ decreases with a bidder's own value. The intuition for this comparative statics result is best understood by assuming that the opponent, say bidder 2, follows the standard strategy of bidding her own value. When bidder 1 considers raising his bid marginally 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 is present, there is a third effect: by raising one's bid, one increases the price of the rival bidder in the event that bidder 2 has a higher valuation, which happens with probability $1 - F(v_1)$. The third effect, which is a marginal benefit term from overbidding, is higher the lower a bidder's own valuation. Thus, the model with spite predicts that in SPA control sessions, the overbidding should be decreasing with a bidder's own value if bidders are motivated by spite.

In an environment in which bidders also privately observe noisy signals about opponents' value (and thus have multi-dimensional private types), it is not analytically possible to derive the equilibrium of the SPA with bidders motivated by spite. However, one can 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, bids her own value v_2. When bidder 1 considers raising her bid marginally above her 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 that bidder 2 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 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 of the term $1 - F_{v_2|s_1}(v_1|s_1)$ for information accuracy $K = 0.3$ and $K = 0.7$ respectively yield the following predictions about the incentives to overbid for bidders motivated by spite in environments where bidders receive noisy signals about their opponents' values:[16]

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

SPITE HYPOTHESIS 1: *Incentives to overbid decrease in the bidders' own value v_i in all treatments.*

SPITE HYPOTHESIS 2: *Incentives to overbid tend to be lowest when bidders' own value v_i and signal s_i coincide.*

SPITE HYPOTHESIS 3: *Incentives to overbid tend to be lower when $v_i > s_i$ than when $v_i < s_i$.*[17]

It is useful to make two further remarks. First, this theory is predicated on the assumption that when subjects play in a laboratory experiment their spite is targeted towards fellow subjects rather than towards the experimenter. It is not clear whether such an assumption is valid. Second, the equilibrium in an English auction with spite-motivated bidders is identical to that of the SPA in the benchmark model where bidders do not observe noisy signals about opponents' values (see Proposition 3 of Morgan *et al.*, 2003). Thus spite cannot explain the observed difference in overbidding between SPAs and English auctions.

2.3. *Bidders with '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 hypothesis. 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* but also use heuristics in deciding how much to bid.

The implication of the simple version of the joy of winning hypothesis for 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

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

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

SIMPLE JOY-OF-WINNING HYPOTHESIS: *The amount of overbidding predicted by the simple joy-of-winning hypothesis is independent of the bidders' own valuation, their signals about opponent's valuation and the signal accuracy.*

Richer implications from the joy-of-winning hypothesis can be derived in a *modified* model where we make some additional behavioural 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 by ϵ. Her expected payoff from bidding ϵ above her value v_i is given by

[17] More specifically, for a given deviation $\Delta > 0$ of s_i from v_i, the incentives to overbid when $s_i = v_i + \Delta$ tend to be higher than when $s_i = v_i - \Delta$.

$$\int_{0}^{v_i+\epsilon} \left(v_i + t - b_j\right) \mathrm{d}G_i\left(b_j \mid v_i, s_i\right).$$

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

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

where $g_i(\cdot \mid v_i, s_i)$ is the derivative of $G_i(\cdot \mid v_i, s_i)$. This of course means that the optimal overbid is $\epsilon^* = t$. However, if we assume instead that bidders are more likely to overbid when the *marginal* benefit is higher, then we can conclude that the incentive to overbid will depend on the magnitude of $(t - \epsilon) g_i(v_i + \epsilon \mid v_i, s_i)$. In particular, it depends on $g_i(\cdot \mid v_i, s_i)$, 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 $g_i(v_i + \epsilon \mid v_i, s_i)$ 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 4,000, 5,000 or 6,000 ECUs). We restate the above discussion as two predictions based on the *modified* joy-of-winning hypothesis:

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 hypothesis suffers from the same problem as the spite hypothesis in that it cannot explain the difference in the overbidding between SPA and English auctions.

2.4. *Bidders with Bounded Rationality*

Auction models in which bidders are boundedly rational have the potential to explain the difference in overbidding between SPA and English auctions if the auction formats have different effects on the *degree of perceptual biases* underlying bounded rationality. Portable models of bounded rationality have become increasing prevalent in the economics and game theory literatures over the past fifteen years, with approaches such as quantile response equilibrium (QRE) (McKelvey and Palfrey, 1995) and level-k reasoning – see Nagel (1995) and Stahl and Wilson (1995) for early examples – being particularly germane for SPAs. Both approaches have been used to explore anamolous results in other auction formats (Goeree *et al.*, 2002; Crawford and Iriberri, 2007) but neither directly predicts overbidding in SPAs.[18] Building such a model is well beyond the scope of this article. In our working paper, Cooper and Fang (2006), we sketch a simple model of bounded rationality where we hypothesise that, when a bidder contemplates her optimal bid, she *under-accounts* for the impact of an increase in her bid on her expected payoff conditional on winning while she fully accounts for the positive impact of a marginal increase of her bid on her probability of winning. *If*

[18] We are not aware of any existing papers that analyse the SPA using QRE. The QRE model predicts substantial misbidding in the SPA, as the cost of mistakes in either direction is small, but does not predict a clear direction for these mistakes (e.g. over or underbidding).

bounded rationality takes this form, it is plausible to postulate that the *key difference between SPA and English auctions* lies in the level of perceptual bias regarding the impact of overbidding on her expected payoff conditional on winning. Specifically, the format of an English auction makes it transparent to bidders that any increase in the probability of winning from overbidding will result in negative payoff conditional on winning. On the other hand, the SPA makes it less clear that bidding above one's valuation only increases the likelihood of winning while winning is not profitable. Thus, the English auction is likely to eliminate any perceptual biases with minimal experience while the SPA makes it hard to unlearn such biases.

Rather than relying on any specific model of bounded rationality, we test for the presence of bounded rationality underlying overbidding in SPAs by exploiting the simplest difference between a bounded rationality explanation and explanations based on non-standard preferences: if overbidding is driven by bounded rationality, bidders may learn to bid more accurately over time if the errors provide strong payoff feedback; but if they are driven by non-standard preferences then the overbidding will persist over time. We state this as a testable hypothesis:

Bounded Rationality Hypothesis: *If overbidding is driven by bounded rationality, bidders may learn to bid more accurately over time if the errors provide strong feedback (e.g. overbidding is costly).*

3. Experimental Results

3.1. *An Overview of Bidding Behaviour*

The bidding behaviour from our experiments is summarised in Figure 2 where we separate overbids into three categories: *low* ($0 < \text{bid} - \text{value} < \$12$), *medium* ($\$12 \leq \text{bid} - \text{value} < \$25$), and *high* ($\$25 \leq \text{bid} - \text{value}$).[19] We collectively refer to medium and high overbids as '*large*' overbids. The first cluster of bars shows data from all five treatments pooled together and the remaining clusters separate the five treatments.

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%). Many overbids cannot be characterised as small mistakes. 'Large' overbids occur for 18% of all observations and 44% of all subjects have at least one large overbid.[20]

Figure 2 also shows that the frequency and nature of overbidding differ across treatments. Overbidding is more common in the control treatment (52%) than any other treatment. Comparing EX3 and EX7 treatments, bidding one's value is more

[19] The breakpoints between these three categories are somewhat arbitrary but overbids over \$12 are sufficiently large that subjects could go bankrupt and \$25 represents the 90th percentile of overbids (rounded to the closest dollar).

[20] 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 seller revenue across all treatments is \$40.82.

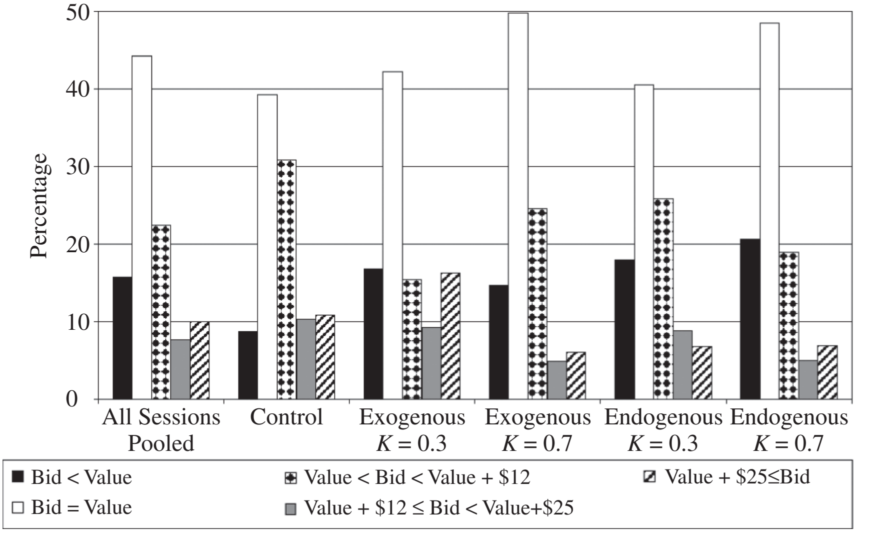

Fig. 2. *Distribution of Overbids*

common but large overbids are less frequent in EX7 treatments. This comparison is also true between END3 and END7 treatments where signal acquisition is endogenous.

Overbidding does not vanish with experience, again consistent with the existing experimental results. Compare behaviour in the first five periods with behaviour in the remaining fifteen periods.[21] Pooling across all treatments, the frequency of overbidding rises somewhat from 35% to 42%. This does not reflect a decrease in rational behaviour 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, 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 remains steady at 10% in both the first five and remaining fifteen periods.

Not only is overbidding frequent, it is costly as well. Column 1 of Table 2 lists the average bidder payoffs in the actual plays (excluding information acquisition costs/benefits if applicable), both pooled and broken down by treatments. Column 2 lists what the average payoffs would have been if all subjects had bid their values against their opponents' 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%.[22] 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 that

[21] Splitting the data unevenly into early and late periods gives a better sense of the dynamics than splitting it evenly, as most changes occurred in the early periods.

[22] If only observations with overbids are considered, average payoffs would have been increased 37% by bidding rationally ($7.56 vs. $10.40).

Table 2

Effects of Irrational Play on Payoffs, Subject Averages

Session	Actual Play vs. Actual Opponent (1)	Rational Play vs. Actual Opponent (2)	Actual Play vs. Rational Opponent (3)	Rational Play vs. Rational Opponent (4)
CON	8.38 (0.86)	9.80 (0.62)	10.06 (0.82)	11.25 (0.65)
EX3	6.58 (0.53)	8.45 (0.43)	8.73 (0.53)	10.40 (0.39)
EX7	10.75 (0.64)	11.63 (0.62)	10.87 (0.66)	11.78 (0.60)
END3	10.29 (0.80)	11.20 (0.71)	10.68 (0.82)	11.47 (0.73)
END7	10.27 (0.68)	11.34 (0.61)	9.35 (0.66)	10.71 (0.57)
All Sessions	9.03 (0.32)	10.32 (0.27)	9.86 (0.31)	11.08 (0.25)

Notes. These payoffs are all in dollars, and they exclude any costs/benefits from purchasing information. Standard errors are in parentheses.

reduces 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 shows that subjects' average payoffs are 24% lower than they would be if all bidders bid their values.

3.2. *How Do Signals Affect Bidding?*

As emphasised in Section 2, the 'spite' and the modified 'joy of winning' hypotheses have different predictions regarding how bidders' signals about their opponents' value affect overbidding. Here we examine how signals affect bidding in the experimental data. For this purpose 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.

Figure 3 illustrates the complex relationship between bidding behaviour and the signals received by bidders. To allow for signals both substantially larger and substantially smaller than the bidder's value, Figure 3 only includes observations with values between 4,000 and 6,000 ECUs. Regression analysis using the full data set (see Table 3) confirms the conclusions we draw from this limited dataset. We break the observations into five categories based on the difference between a bidder's signal and value. When we consider all overbids (regardless of the overbid amount) as shown in the first cluster of bars, we see that 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). However, when we distinguish overbids by their magnitudes, there seems to be stronger but different responses for low, medium and high overbids. The middle cluster of bars shows that for low and medium overbids (e.g. overbid <\$25), the relationship is strongly peaked and almost perfectly symmetric.[23] In contrast, the relationship for high overbids is U-shaped. Thus, overbidding is not a homogeneous 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.

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

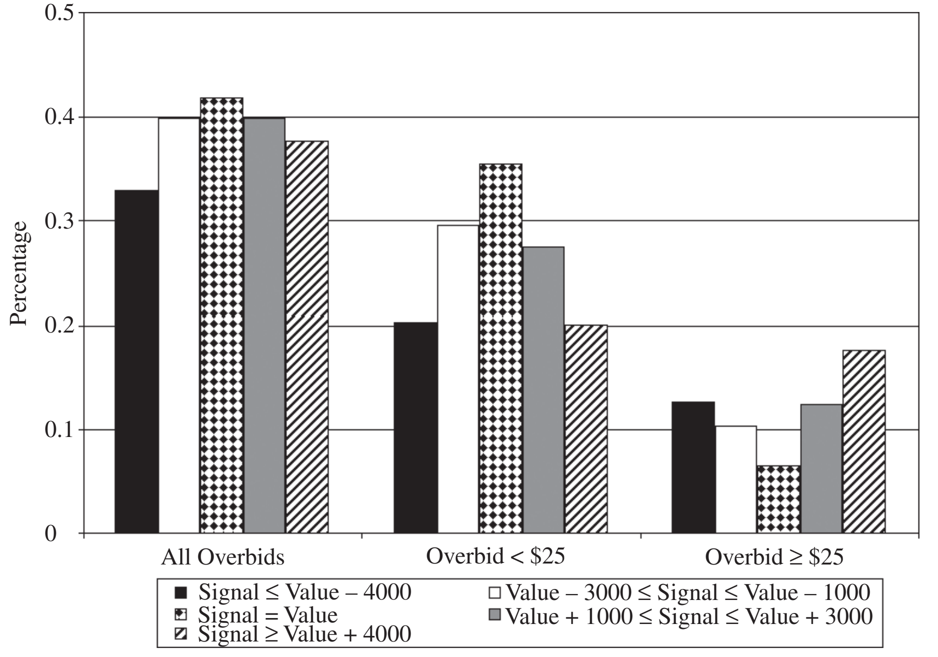

Fig. 3. *Relationship Between Overbid Incidence and Signals: Only Observations with Values Between 4,000 and 6,000 ECUs Are Included*

Table 3

The Effect of Signals on the Incidence of Overbids: Probit Regressions Using Data from CON, EX3, EX7 Treatments (144 Subjects and 2,842 Observations)

Ind. Variable	All Overbids (1)		All Overbids (2)		Overbids < \$25 (3)		Overbid ≥ \$25 (4)	
EX3	−0.318	(0.203)	−0.194	(0.206)	−0.264	(0.196)	0.117	(0.215)
EX7	−0.412**	(0.201)	−0.198	(0.204)	−0.123	(0.195)	−0.365	(0.225)
Value	0.012	(0.018)	0.017	(0.018)	0.019	(0.019)	−0.005	(0.023)
Period 6–20	0.183***	(0.057)	0.178***	(0.056)	0.159***	(0.061)	0.084	(0.071)
Yale	−0.246	(0.154)	−0.241	(0.154)	−0.027	(0.142)	−0.460***	(0.156)
Exog. × (Signal − Value) × (Signal > Value)			−0.038**	(0.019)	−0.096***	(0.021)	0.067***	(0.026)
Exog. × (Value − Signal) × (Signal < Value)			−0.052***	(0.018)	−0.055***	(0.020)	−0.013	(0.026)
Log-likelihood	−1,889.66		−1,884.58		−1,698.87		−944.37	

Notes. Standard errors in parentheses control for clustering of subjects. Constants have been suppressed. *, **, and *** denote significance at 10%, 5% and 1% respectively. Values and signals are denominated in 1000s of ECUs.

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 modified 'joy of winning' model predicts a peak-shaped function (see the discussions in Sections 2.2 and 2.3). Ignoring the highest overbids, the data is consistent with the modified 'joy of winning' model. However, the U-shaped probability function for high overbids suggests that spite is playing an important role for that class of overbids.

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

The basic pattern of the incidence of overbidding summarised in Finding 1 above is confirmed by Probit regressions reported in Table 3,[24] where we use all bid observations from the CON, EX3, and EX7 treatments. The dependent variable in models 1 and 2 is a dummy for an overbid (e.g. bid > value); in model 3, it is a dummy for low and medium overbids (e.g. overbid < \$25); and in model 4 it is a dummy for high overbids (e.g. overbid $\geq$ \$25). The base for all regressions is the control session.

Model 1 estimates the basic treatment effects. It partially 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; however, 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.[25] 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 sharp decrease in the probability of underbidding with experience rather than a move away from bidding one's value.

Models 2–4 include two additional independent variables generated from the interaction of bidders' absolute signal/value difference with a dummy for whether the difference is positive or negative.[26] In both models 2 and 3 (which differ in the dependent variables, see above), the parameter estimates for 'Exog. $\times$ (Signal – Value) $\times$ (Signal $>$ Value)' and 'Exog. $\times$ (Signal $-$ Value) $\times$ (Signal $<$ Value)' are both negative and statistically significant at least at the 5% level,[27] though the measured effect is stronger in model 3 where the dependent variable is low or medium overbids, especially for positive signal/value differences.[28] 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 'Exog. $\times$ (Signal – Value) $\times$ (Signal $>$ Value)' is positive and statistically significant

[24] 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 certainly discards a great deal of information from the dataset. However, any statistical model that attempts to treat overbids as a continuous variable will 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.

[25] Unlike our impression from Figure 2, the difference between EX3 and EX7 is negative but not statistically significant. 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.

[26] They are further interacted with 'Exog.' because the signal/value difference is not defined in CON treatments. Separating positive and negative signal/value differences allows for asymmetric responses to information.

[27] The difference between these two parameters is not statistically significant in either model 1 or model 2.

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

while the estimate for 'Exog. × (Signal – Value) × (Signal < Value)' is actually slightly negative (although not statistical significant). Thus, the left arm of the U-shape we saw for the third cluster of bars in Figure 3 is actually non-existent while the right arm is robust.

To summarise, Table 3 shows that as a function of the difference between a bidder's signal and value, the probability of 'any overbids' is peaked, consistent with a 'joy of winning' model; in particular, this pattern is more extreme if attention is restricted to low and medium overbids but breaks down for high bids. The finding in model 4 that the probability of high overbids increases with the positive signal/value difference is the one case in which the data are consistent with a model of spite.

3.3. *Are Subjects Learning Not to Overbid?*

Understanding how subjects' behaviour 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*. However, in most cases subjects experience no cost from overbidding: in our data 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 4 plots the probability of overbidding in period t conditional on having overbid in period $t-1$, using data from all five treatments. The difference between the left and right pair of bars in Figure 4 is the size of the overbid in the previous period. The data is split into observations where this lagged overbid did not cause a loss, and observations where a lagged overbid led to winning the auction at a monetary loss. 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, overbidding is 13% less likely if the lagged overbid led to a monetary loss; and high overbids are 12% less likely if the lagged high overbid led to a monetary loss. The qualitative pattern also holds in all five treatments when we break down the data by treatment. This suggests that this relationship is not likely to be a coincidence.

Table 4 presents regression results with the dependent variable being a dummy for whether an overbid occurred in the current period (model 1) or a dummy for whether a high overbid was observed in the current period (model 2). For both regressions the data set includes all observations from all treatments except for observations from period 1[29] and the control treatment serves as the base. Besides the standard independent variables, we include a number of lagged dependent variables in both regressions. The critical variable of interest in model 1 (respectively, model 2) is a dummy for whether the bidder overbid (respectively, submitted a high overbid) in the previous round *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. For

[29] These are discarded to allow the use of lagged variables.

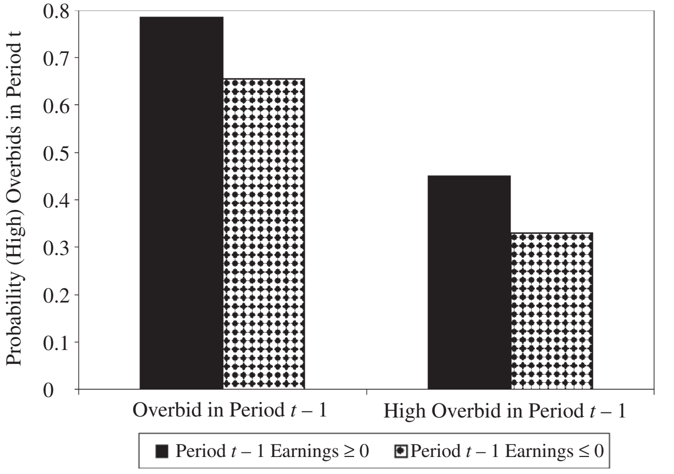

Fig. 4. *Learning Not to Overbid*

Table 4

Learning: Probit Regressions Using Data from All Treatments (208 Subjects and 3,914 Observations)

Ind. Variable	All Overbids (1)		Overbid ≥ 25 (2)	
EX3	−0.186	(0.128)	0.228	(0.164)
EX7	−0.229*	(0.126)	−0.138	(0.171)
END3	−0.193	(0.146)	−0.190	(0.186)
END7	−0.380***	(0.143)	−0.184	(0.192)
Value	0.019	(0.019)	−0.030	(0.020)
Periods 6–20	−0.042	(0.041)	−0.116*	(0.059)
Yale	−0.207**	(0.081)	−0.272***	(0.101)
Lagged Overbid	1.685***	(0.103)	0.336***	(0.092)
Lagged Overbid ≥ \$12	0.190	(0.117)	0.524***	(0.114)
Lagged Overbid ≥ \$25	−0.218*	(0.116)	0.666***	(0.191)
Lagged [High] Overbid and Win	−0.034	(0.090)	−0.058	(0.207)
Lagged [High] Overbid and Lose Money	−0.389***	(0.090)	−0.351***	(0.129)
Log Likelihood	−1914.76		−1074.19	

Notes. See Table 3.

example, 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. Including a dummy for whether the lagged bid was an overbid takes care of this problem.[30] The parameter of interest now measures the effect of losing money *conditional on having overbid*. Dummies for lagged large overbids (bid – value ≥ \$12) and

[30] Without including the dummy for 'Lagged Overbid,' estimating a parameter for 'Lagged Overbid and Lose Money' is akin to regressing on the fixed effects.

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.[31] The dummy for 'Lagged [High] 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 in Table 4 are consistent with our observations from Figure 4 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. 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.*

3.4. *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 bidders acquire signals. We are particularly interested in whether signal acquisition fades with experience as this would be clear evidence of subjects learning to avoid a mistake.

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 increasing their monetary payoffs, subjects are otherwise fairly rational in their purchase decisions. Figure 5 graphs the demand curve for signals. The likelihood of purchasing information decreases monotonically in the cost. Subjects are more likely to purchase high quality signals at a positive cost (25%) than low quality signals (18%). The likelihood of purchasing information at a positive price decreases with experience,

[31] 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 alternatives.

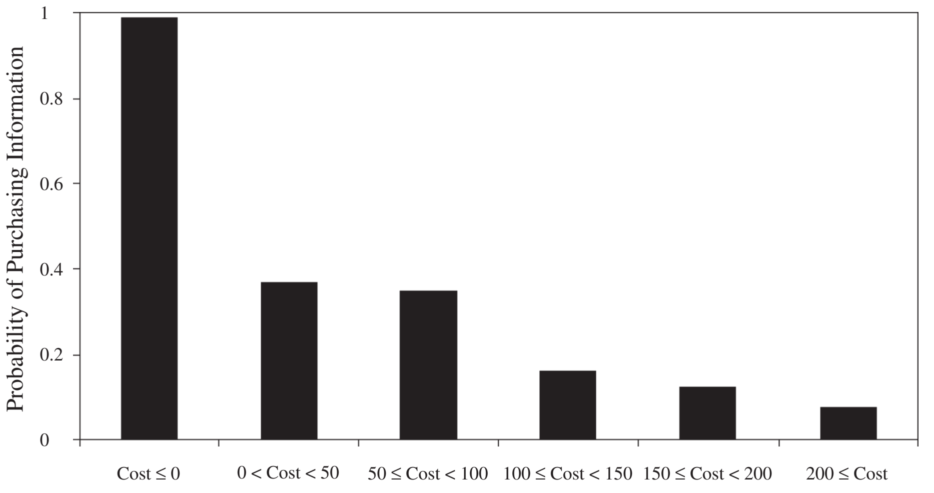

Fig. 5. *Demand for Signals*

dropping from 28% in the first five periods to 19% for the remaining fifteen periods. This is consistent with learning.

There seems to be a clear link between overbidding and signal purchases. 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 5 formalises the preceding conclusions using Probit regressions. These regressions use all observations from END3 and END7 sessions. The dependent variable is a dummy for whether a signal was purchased. As independent variables, 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 20 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)

Table 5

Demand for Information: Probit Regressions Using Data from END3 and END7 Treatments (64 Subjects and 1,280 Observations)

Ind. Variable	(1)		(2)	
Cost < 0	2.328***	(0.263)	2.628***	(0.302)
(Cost $\geq$ 0) $\times$ Cost	−6.159***	(1.108)	−6.777***	(1.098)
Period 6–20	−0.259**	(0.131)	−0.324**	(0.141)
Yale	−0.166	(0.278)	0.029	(0.288)
END7	0.293	(0.234)	0.582**	(0.261)
Never Overbid			−1.186***	(0.307)
Log Likelihood	−514.14		−472.03	

Notes. See Table 3.

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.[32] 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. 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 from experience to not purchase costly information.*

3.5. *The Connection Between Signal Purchase and Overbidding*

As suggested above, there is a strong link between purchasing information and overbidding. This relationship is illustrated by Figure 6. 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 2, 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 and 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 18 have at least 12 observations with a positive cost), they can be classified as strongly following the theoretical prediction of no costly signal purchases.[33]

Figure 6 shows that subjects who pay a positive cost for information are far more likely to overbid (and underbid as well) than those who do not. Conditional on 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 are chosen for only 3% of all observations. The subject population appears to be heterogeneous, consisting of types whose behaviour (both in signal

[32] Our finding that subjects do learn over time not to purchase costly information is at odds with Eliaz and Schotter's (2006) finding that decision makers have intrinsic preferences over beliefs. In their experimental design, however, subjects are not presented with opportunities to learn. It would be interesting to see how subjects' information purchase decisions change over time in their experimental setting.

[33] The maximum over these eighteen subjects of their respective minimum positive costs for information was 54 ECU (with an average of 15 ECU), 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 coincidence that these subjects never paid a positive cost for information.

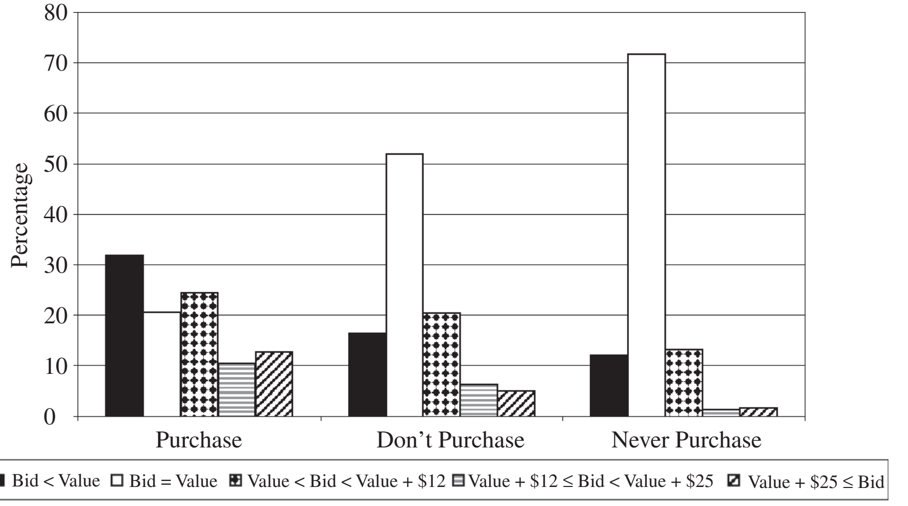

Fig. 6. *Signal Purchase at a Positive Cost and Bidding Behaviour*

purchasing and in bidding) are consistent with the theoretical predictions and types who violate the theoretical predictions across the board.

Table 6 reports Probit regression results regarding how information purchase affects subjects' likelihood of overbidding. All three regression models in Table 6 has the same dependent variable, which is a dummy for overbidding but they differ in the estimation sample, as described by the column head. The data set for model 1 includes all observations from END3 and END7 with a positive cost for information; and this regression includes a dummy 'Signal Purchased' as an explanatory variable. The parameter estimate associated with this variable is positive and statistically significant at the 1% level, suggesting that purchasing information at positive cost is strongly *associated* with overbidding.

Model 2 helps us to answer whether the positive association found between 'Signal Purchased' and overbidding in model 1 reflects a common unobserved type that induces both information purchasing at positive cost and subsequent overbidding, or reflect a causality – informed bidders behave less rationally because they purchased signals.[34] In model 2, we only use 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 the dummy variable 'Ever Purchased at Cost $\geq$ 0', 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 $\geq$ 0' is not correlated with subjects' current information. Thus, its coefficient

[34] 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.

Table 6

Information Purchase and Overbidding: Probit Regressions of Using Data from END3 and END7 Treatments

	Data Sets					
	Cost $\geq$ 0 64 subjects, 1,063 obs.		Cost $<$ 0 63 subjects, 217 obs.		Never Purchased 18 subjects, 360 obs.	
Ind. Variable	(1)		(2)		(3)	
END7	-0.392^{*}	(0.223)	-0.556^{**}	(0.275)	-0.655	(0.502)
Value	0.033	(0.027)	-0.060	(0.043)	0.072^{**}	(0.036)
Period 6–20	0.267^{**}	(0.119)	-0.014	(0.248)	-0.063	(0.203)
Yale	-0.557^{**}	(0.229)	-0.755^{***}	(0.279)	-0.496	(0.448)
Signal Purchased	0.460^{***}	(0.154)				
Ever Purchased at Cost $\geq$ 0			0.593^{**}	(0.279)		
Cost $<$ 0					0.465^{*}	(0.246)
Log Likelihood	-645.48		-130.60		-155.98	

Notes. See Table 3.

measures the influence of type separate from any direct effect of being informed. The parameter estimate for 'Ever Purchase at Cost $\geq$ 0' is positive 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 'Signal Purchased' in model 1 is driven by a subject's type rather than being informed *per se*.[35]

In model 3, we address the question of whether giving information to subjects has any effect independent of subjects' types. For this purchase we restrict the sample to those subjects whose actions indicate they believe the information is worthless, i.e., subjects who never paid a positive cost for information. This yields only a small data set (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 may weakly lead to more overbidding independent of their type.[36]

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.*

[35] 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 $\geq$ 0' are positive and significant at least at the 5% level; the marginal effects for these two variables are almost identical (8% in both cases).

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

4. Discussion and Conclusions

This article 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 they provide a convenient way of changing 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 favour of explanations that 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 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 hypotheses listed in Section 2).

Finally, we find that bidder heterogeneity is playing an important role in our data. There is a group of subjects in our data set whose behaviour is almost completely in line with the standard theoretical predictions. These subjects 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 subjects exist who bid according to the theory, forces of selection may quickly drive out those subjects who are prone to overbidding.

Florida State University
Duke University, SHUFE and CEMA

Submitted: 8 March 2006
Accepted: 28 June 2007

References

Andreoni, J., Che, Y.-K. and Kim, J. (2007). 'Asymmetric information about rivals' types in standard auctions: an experiment', *Games and Economic Behavior*, vol. 59(2), pp. 240–59.

Cooper, D.J. and Fang, H. (2006). 'Understanding overbidding in second price auctions: an experimental study', Cowles Foundation Discussion Paper 1557, Yale University.

Cox, J.C., Smith, V.L. and Walker, J.M. (1992). 'Theory and misbehavior of first-price auctions: comment', *American Economic Review*, vol. 82 (5) (December), pp. 1392–412.

Crawford, V.P. and Iriberri, N. (2007). 'Level-*k* auctions: can a non-equilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?', *Econometrica*, vol. 75 (6) (November), pp. 1721–70.

Eliaz, K. and Schotter, A. (2006). 'Paying for confidence: an experimental study of preference over beliefs', mimeo, New York University.

Engelbrecht-Wiggans, R. and Katok, E. (2007). 'Regret in auctions: theory and evidence', *Journal of Economic Theory*, vol. 33, pp. 81–101.

Fang, H. and Morris, S. (2006). 'Multidimensional private value auctions', *Journal of Economic Theory*, vol. 126 (1) (January), pp. 1–30.

Filiz-Ozbay, E. and Ozbay, E.Y. (2007). 'Auctions with anticipated regret: theory and experiment', *American Economic Review*, vol. 97(4), pp. 1407–18.

Fischbacher, U. (2007). 'z-Tree: Zurich toolbox for readymade economic experiments', *Experimental Economics*, vol. 10(2), pp. 171–8.

Goeree, J.K., Holt, C.A. and Palfrey, T.R. (2002). 'Quantal response equilibrium and overbidding in private-value auctions', *Journal of Economic Theory*, vol. 104 (1) (May), pp. 247–72.

Harrison, G.W. (1989). 'Theory and misbehavior of first price auctions', *American Economic Review*, vol. 79 (5) (October), pp. 749–62.

Harstad, R. (2000). 'Dominant strategy adoption and bidders' experience with pricing rules', *Experimental Economics*, vol. 3 (3) (December), pp. 261–80.

Kagel, J.H. (1995). 'Auctions: a survey of experimental research', in (J.H. Kagel and A.E. Roth, eds.) *The Handbook of Experimental Economics*, pp. 501–85, Princeton, NJ: Princeton University Press.

Kagel, J.H., Harstad, R. and Levin, D. (1987). 'Information impact and allocation rules in auctions with affiliated private values: a laboratory study', *Econometrica*, vol. 55 (6) (November), pp. 1275–304.

Kagel, J.H. and Levin, D. (1993). 'Independent private value auctions: bidder behavior in first-, second- and third price auctions with varying number of bidders', ECONOMIC JOURNAL, vol. 103 (419) (July), pp. 868–79.

Kim, J. and Che, Y.-K. (2004). 'Asymmetric information about rivals' types in standard auctions', *Games and Economic Behavior*, vol. 46 (2) (February), pp. 383–97.

McKelvey, R.M. and Palfrey, T.R. (1995). 'Quantal response equilibria in normal form games', *Games and Economic Behavior*, vol. 10 (1) (July), pp. 6–38.

Morgan, J., Steiglitz, K. and Reis, G. (2003). 'The spite motive and equilibrium behavior in auctions', *Contributions to Economic Analysis & Policy*, vol. 2 (1), Article 5.

Nagel, R. (1995). 'Unraveling in guessing games: an experimental study', *American Economic Review*, vol. 85 (5), (December), pp. 1313–26.

Stahl, D. and Wilson, P. (1995). 'On players' models of other players: theory and experimental evidence', *Games and Economic Behavior*, vol. 10 (1), (July), pp. 218–54.

Vickery, W. (1961). 'Counterspeculation, auctions and competitive sealed tenders', *Journal of Finance*, vol. 16 (1) (March), pp. 8–37.

Wooldridge, J. (2003). 'Cluster-sample methods in applied econometrics', *American Economic Review Papers and Proceedings*, vol. 93 (2) (May), pp. 133–8.