Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Regret and Feedback Information in First-Price Sealed-Bid Auctions

Richard Engelbrecht-Wiggans, Elena Katok,

MANAGEMENT SCIENCE
Vol. 54, No. 4, April 2008, pp. 808–819
ISSN 0025-1909 | EISSN 1526-5501 | 08 | 5404 | 0808

informs®

DOI 10.1287/mnsc.1070.0806

Regret and Feedback Information in First-Price Sealed-Bid Auctions

Richard Engelbrecht-Wiggans
College of Business, University of Illinois at Urbana-Champaign, Champaign, Illinois 61820,
eplus17@uiuc.edu.

Elena Katok
Smeal College of Business, Penn State University, University Park, Pennsylvania 16802,
ekatok@psu.edu.

We investigate the effect of regret-related feedback information on bidding behavior in sealed-bid first-price auctions. Two types of regret are possible in this auction format. A winner of the auction may regret paying too much relative to the second highest bid, and a loser may regret missing an opportunity to win at a favorable price. In theory, under very general conditions, being sensitive to winning and paying too much should result in lower average bids, and being sensitive to missing opportunities to win at a favorable price should result in higher bids. For example, the U.S. Government's policy of revealing losing bids may cause regret-sensitive bidders to anticipate regret and bid conservatively, decreasing the government's revenue. We test these predictions in the laboratory and find strong support for both.

Key words: auctions; competitive bidding; regret; learning; experimental economics
History: Accepted by David E. Bell, decision analysis; received June 19, 2006. This paper was with the authors 4 months for 2 revisions.

I see it all perfectly; there are two possible situations—one can either do this or that. My honest opinion and my friendly advice is this: do it or do not do it—you will regret both.
Ŝoren Kierkegaard (Danish Philosopher 1813–1855)

1. Introduction and Motivation

Bidders in first-price auction experiments overbid relative to the risk-neutral Nash equilibrium, much as if they were risk averse (see, for example, Cox et al. 1988), but there is a substantial amount of evidence that even if risk aversion is part of the explanation for overbidding, it is far from the complete explanation (see Kagel 1995 and references therein, and Isaac and James 2000). In this paper, we consider the effect of regret on bidding in first-price sealed-bid (hereafter simply "first-price") auctions.

The idea that emotions, such as disappointment (Bell 1985) or regret can play a role in decision making under uncertainty has a long history. Regret was first introduced by Savage (1951) and further explored by Luce and Raiffa (1956) and Loomes and Sugden (1982). Bell (1982) shows that incorporating regret into the utility function explains some well known behavioral anomalies, such as the Allais Paradox (Allais 1953), the coexistence of insurance and gambling, the fact that people tend to be risk averse in the domain of gains and risk loving in the domain of losses, the probabilistic insurance (Kahneman and Tversky 1979), and preference reversals (Grether and Plott 1979). Bell (1983) suggests that decision makers might be willing to pay a "risk premium" to reduce the amount of regret that they suffer; decision makers may look as if they are risk averse when in fact they are regret averse.

In addition to risk aversion, several other explanations for overbidding in first-price auctions have been advanced. Specifically, Isaac and Walker (1985) found that in auctions with four bidders, when bidders receive feedback at the end of the auction that includes all bids, the amount of overbidding relative to the risk-neutral Nash equilibrium decreases. They suggest that implicit collusion among the bidders may be responsible for the shift. Ockenfels and Selten (2005, p. 156) report a similar result in two-person auctions and propose a direction learning explanation, the impulse balance equilibrium, that they interpret "as a measure of concern for relative standing." Dufwenberg and Gneezy (2002) report a similar shift in common-value auctions that they attribute to signaling behavior. Morgan et al. (2003, p. 1) develop a theoretical model of "spiteful bidding" in which they postulate that "a bidder cares not only about her own surplus in the event she wins the auction, but also about the surplus of her rivals in the event she loses."

In this paper, we investigate whether the overbidding in first-price auctions could be due to regret. Engelbrecht-Wiggans (1989) looks at regret specifically in the context of auctions. He suggests that a bidder's utility depends not only on profit, but also on various forms of auction-specific regret. For example, the winner in a first-price auction typically pays more than the highest competitor's bid, thus leaving money on the table. In this case, the winner may well regret having bid too high; we will refer to this "money left on the table" regret as "winner's regret." Alternatively, winner's price may be below some losing bidder's willingness to pay. In this case, the loser has missed an opportunity to win the object at a favorable price and may regret having bid too low; we will refer to this as "loser's regret."[1] Engelbrecht-Wiggans (1989) shows that if bidders in a first-price auction weight loser's regret more heavily than winner's regret, then they should bid higher than the risk-neutral Nash equilibrium (and the converse is true as well). In other words, bidders who are more concerned with loser's regret than winner's regret may bid as if they were risk averse.

This regret-related bias has two practical implications, one related to policy and the other related to interpreting laboratory data. First, the U.S. Government has a policy of disclosing losing bids in a variety of applications, including, for example, Outer Continental Shelf mineral rights sales (Rothkopf and Park 2001). In our experiments, prices average three to five percent lower when losing bids are disclosed after the auction compared to when such information is not disclosed. Three to five percent of the total revenues from federal auctions in which losing bids are revealed would be a nontrivial amount of money. Should the Government reconsider this policy? Many other auctions do not reveal losing bids. Our work provides additional support for not revealing losing bids.

Second, many auctions do not publicize all bids after the auction, and neither does the typical laboratory setting (see Kagel 1995).[2] For example, when first-price auctions are conducted in the laboratory, the only information that is usually revealed publicly at the end of the auction is the winning price. Therefore, the winner never learns the amount of the second highest bid, he never finds out exactly how much money was left on the table, and the sensation of regret over having bid too high may not be particularly salient. However, because laboratory auctions usually do announce the winner's price, missed opportunities to win at a favorable price are quite apparent. In short, typical laboratory subjects may well be more aware of the loser's regret than the winner's regret, and this awareness might explain the observed tendency for bidders to bid above the risk-neutral Nash equilibrium.

More generally, regret can explain a variety of data from auctions, and auction-like settings, that cannot be explained by risk aversion. For example, Kagel and Levin (1993) report on sealed-bid third-price auctions and Cason (1995) reports on sealed-bid random-price auctions. In both auctions, the risk-neutral Nash equilibrium bids are above values, and risk aversion lowers bids. However, Kagel and Levin (1993) observe bids that tend to be below risk-neutral Nash equilibrium in auctions with five bidders, but above the risk-neutral Nash equilibrium in auctions with ten bidders. Cason (1995) also observes that bids tend to be above the risk-neutral Nash equilibrium. The observed behavior is inconsistent with risk aversion but it is consistent with regret, as we will discuss further in the conclusion.

Isaac and James (2000) estimate the constant relative risk-aversion utility-function parameter for individual subjects using a first-price auction and a Becker-DeGroot-Marschak (BDM) procedure (Becker et al. 1964) and find the two estimates to be inconsistent. They note that participants who act as if they were risk averse in first-price auctions act as if they were risk loving in the BDM game, and visa versa, whereas other participants behave as if they were risk neutral in both games. Isaac and James (2000) describe this result as an unsolved puzzle. In fact, the regret structure of the two games is quite different, and it turns out that the regret theory implies exactly the pattern of behavior that Isaac and James (2000) observe (see Engelbrecht-Wiggans and Katok 2007 for a formal derivation).

Finally, Ivanova-Stenzel and Salmon (2008) show that when bidders can select to enter either a first-price or an ascending-bid auction, more bidders enter ascending-bid auctions than can be explained by risk aversion. Again, these authors do not offer any specific explanation for this behavior, but they do mention regret as a general possibility. More specifically, loser's regret may explain this result; whereas both types of auctions present risk to bidders, only the first-price auction gives rise to regret, and a concern for loser's regret could shift the entry equilibrium toward one with more bidders entering the ascending auction.

[1] Other auction forms allow yet additional types of regret. For example, the third-price auctions (Kagel and Levin 1993) has a different type of winner's regret because the price may exceed the winner's value, so the winner may regret having won at an unfavorable price.

[2] However, there are some exceptions. For example, see Isaac and Walker (1985) and Ockenfels and Selten (2005).

The goal of this paper is to test directly the Engelbrecht-Wiggans (1989) concept of regret in first-price auctions. We specifically want to focus on the effect of regret in isolation of other factors. In particular, we want to control for interpersonal factors such as collusion, signaling, and spite.[3] We also recognize that having several bidders simultaneously trying to discover how to bid makes it all that much more difficult for any one bidder to figure out how to bid, and we are interested in the effect of regret on bidding rather than in subject's ability to discover an equilibrium bidding strategy. Therefore, each human subject bids against several computerized opponents in our experiments. Note that this at least partially controls for the possible explanations presented earlier for bids above the risk-neutral equilibrium; human subjects cannot collude with the computerized rivals, and human subjects may well be less concerned with how they do compared to computerized opponents than compared to human opponents.

Having subjects bid against computerized opponents is the appropriate way to investigate "best reply" behavior motivated by regret independently of other effects that might emerge from strategic interactions. In the next section, we present a general argument (from Engelbrecht-Wiggans and Katok 2007) that regret moves best replies in the same direction as Engelbrecht-Wiggans (1989) previously argued for equilibrium bids. We also derive the best reply for the specific setting used in our experiments. This theory explicitly allows different types of regret to have different weights in a decision maker's utility function. Our experimental design (§3) manipulates those weights by varying the feedback information. It turns out that, in our setting, the regret theory organizes the data well—it predicts two different shifts, and we observe both in the data (§4). In §5, we offer discussion and summary of how this work fits into the overall literature on bidding behavior in first-price auctions.

2. Theoretical Predictions

Consider a setting in which one human bidder competes against $N-1$ computerized opponents in a first-price auction without reserve. The human bidder is risk neutral, but factors in addition to profit affect the utility that the bidder derives from the outcome of the auction. Specifically, let v denote the bidder's value, and let b denote his bid. On winning, the bidder realizes a profit (or loss) of $v-b$. Additionally, the bidder may suffer from one of two possible types of regret. The first type occurs when the bidder wins. Typically, the winner pays strictly more for the object than the next highest bid; the winner leaves an amount $b-z$ "on the table," where z denotes the highest bid made by the competitors. The winner may regret doing so, and we refer to this as "winner's regret." The second type of regret occurs when $b \le z \le v$; the bidder loses, but has a value above the price paid by the winning competitor. In this case, the loser has missed an opportunity to win at a favorable price and may regret doing so; we refer to this as "loser's regret."

To illustrate the effect of regret, consider a simple example. Specifically, imagine that the bidder is risk neutral with regard to profits, that regret enters additively into the bidder's utility function, and that the effect of regret on the bidder's utility is proportional to the amount of regret suffered. Thus, the bidder's utility suffers by an amount $\alpha(b-z)$ (where $\alpha \ge 0$) due to winner's regret when the bidder wins. If the bidder loses with a bid $b \le z \le v$, then the bidder's utility suffers by an amount $\beta(v-z)$ (where $\beta \ge 0$) due to loser's regret. Therefore, the ex ante expected utility of a bidder who has a value v and who bids b (where $b < v$) is

$$\Pi(b; v) \equiv (v-b)F(b) - \int_{z:\, z\le b} \alpha(b-z)\, dF(z) - \int_{z:\, b\le z\le v} \beta(v-z)\, dF(z),$$

where z denotes the highest of the $N-1$ computerized bidders' bids and F denotes the cumulative distribution function of z.

Now consider what happens to such a bidder in our experimental setting. The $N-1$ computerized bidders' values are independent draws from a uniform distribution (and are independent of the human bidder's value). Each of the computerized bidders uses the multiplicative bidding strategy $b(v) = Av$. Therefore, the computerized bids are independent and uniform on $[0, A]$. In this specific setting, we have

Proposition 1. *Imagine N-1 opponents who have i.i.d. Uniform$(0,1)$ values and bid a constant fraction A of their values. A regret-sensitive bidder's expected utility-maximizing bid in the face of such competition is*

$$b^*(v) = \min\left\{A, \frac{(1+\beta)v(N-1)}{(1+\beta)(N-1)+(1+\alpha)}\right\} = \min\left\{A, \frac{v(N-1)}{(N-1)+\rho}\right\},$$

where $\rho \equiv (1+\alpha)/(1+\beta)$.[4]

[3] Filiz-Ozbay and Ozbay (2007) report on a set of experiments that also test the regret model and show that in a setting with four human bidders bidding in a one-shot auction, bidders react to anticipated regret in a way consistent with the Engelbrecht-Wiggans (1989) model when they are asked to report their bid functions prior to learning the actual value.

[4] See Appendix A1.1 for the derivation of this result.

Corollary 1. *The best reply bid, when it is an interior solution, is a multiple of the bidder's value independent of the opponent's bid.*

Corollary 2. *If bidders are homogeneous in α and β (and this is common knowledge) and each bidder adopts the following strategy, then a Nash equilibrium results:*

$$b^*(v) = \frac{(1+\beta)v(N-1)}{(1+\beta)(N-1)+(1+\alpha)} = \frac{v(N-1)}{(N-1)+\rho},$$

where $\rho \equiv (1+\alpha)/(1+\beta)$.

Corollary 3. *For homogeneous (risk-neutral) bidders who are oblivious to regret (and have full information), the optimal bid is $b^*(v) = \min\{A, v(N-1)/N\}$, and a Nash equilibrium results if each bidder adopts the strategy $b^*(v) = v(N-1)/N$.*

This example has several interesting properties. In particular, note that the best reply can be written in terms of a single unknown parameter ρ; although there are two types of regret, only the relative weight—appropriately defined—really matters.[5] Note also that, by Corollary 1, the best reply is a fixed multiple of the bidder's value and independent of A whenever the best reply is within the range of the opponent's bids. This has several practical implications. For one, the example bidder's best reply is relatively insensitive to what multiple of their values we program our computerized opponents to bid; we can reasonably expect that the results of our experiment are not specific to the particular strategy that we chose for the competitors. Also, the fact that the best reply strategy has a very simple form should make it much easier for human subjects to converge to it during the course of the experiment.

For our example bidder, it is clear how the bidder's best reply changes as the saliency of either type of regret changes. In particular, as α increases—e.g., as winner's regret becomes more salient—the best reply decreases. Similarly, as β to increases—e.g., as loser's regret becomes more salient—the best reply increases.

In fact, regret has similar effects much more generally. In particular, now allow the bidder's utility to be an arbitrary function $u(v, b, \underline{v}-; \underline{B}-)$ of the bidder's type v, the bidder's bid b, the competitors' types $\underline{v}-$, and the competitors' bidding strategies $\underline{B}-$. This is much more general than the example considered above. For one, the bidder's value need not be privately known; indeed, the bidder's utility could depend on things like the relative amounts of profit made by each of the other opponents (see, for example, Bolton and Ockenfels 2000). Nor need the bidders' types be independent; in fact, the individual bidders could have multidimensional types. For this general setting in Engelbrecht-Wiggans and Katok (2007) we show that, roughly speaking, the more sensitive a bidder is to loser's regret, the higher that bidder should bid; and the more sensitive a bidder is to winner's regret, the lower that bidder should bid. This is what our experiment will test.

[5] Note that the best reply is then the same as the best reply, in the absence of regret, for a constant relative risk-aversion bidder who has risk parameter ρ. See, for example, Cox et al. (1988).

3. Design of the Experiment

In our experiments, one human bidder bids against two computerized opponents, so $N = 3$. The computerized bidders' values are integers uniformly distributed from 1 to 100, and independent of one another; this is public knowledge. For three bidders with independent, uniformly distributed values, the risk-neutral Nash equilibrium would be to bid 2/3 of one's value, and this is the strategy used by our computerized bidders. However, the subjects are told simply that the computerized opponents' strategy is one that would maximize a computerized opponent's expected profits under the assumption that all of its opponents follow the identical strategy (see the appendix for complete instructions). We opted for explaining the computerized rivals' behavior in this way for several reasons. First, we wanted to preserve the auction frame, and provide our participants with the same kind of information that participants in all-human experiments are likely to have. Generally, auction experiment participants know the distribution from which values are drawn, but do not know the actual bidding strategy of the opponents or the true distribution of the opponent's bids. Therefore, we told our participants the distribution of the automated rivals' values, but not that they were programmed to bid 2/3 of this value or that the opponents' bids are uniformly distributed from 0 to 66.67 (all bids were transmitted with two decimal places).

The human bidders cycle through the values of 50, 60, 70, 80, and 90, and each value is repeated for 20 consecutive decisions before going to the next value; thus, each session consists of 100 bidding decisions. All subjects cycled through the values in the same, increasing order, but different subjects started at different points in the cycle. We chose to give human bidders large values instead of generating their values from the entire 0 to 100 range for two practical reasons, both related to the fact that bidders with smaller values are unlikely to win. First, the outcome of the auction—who wins and the winner's price—tends to be driven by what bidders do when they have larger values. So, by focusing on the larger values, we focus on the effect of regret on the outcome of the auction.

Second, bidders who are unlikely to win may well not take as much care in bidding as those who are more likely to win; bidders with smaller values may produce "noisier" data than those with larger values. So, by only considering larger values, we reduce the amount of noise in the data.[6]

Each bidding decision was used in 10 independent auctions, with the computerized rivals' values (and bids) changing in each of the 10 auctions, while the human bidder's value and bid remained the same. Thus, each session consisted of 1,000 auctions. The purpose of this design is to create an environment in which participants are able to understand the effect that their decisions have on auction outcomes by experiencing a large number of auctions. There is evidence that this design speeds up learning (see Bolton and Katok 2008). In the same spirit of improving learning, we also displayed to the participants, as part of the bid confirmation screen, the probability of winning for the bid they were about to enter. There is some evidence that subjects in first-price auctions misperceive the probability of winning, and showing them this information improves the quality of their decisions (see, for example, Armantier and Treich 2007, Dorsey and Razzolini 2003).

We varied feedback information across treatments. In two treatments, we compute the amount of loser's regret, and report both the amount of loser's regret and the winning price. In two treatments, we compute the amount of winner's regret, and report both the amount of winner's regret and the second highest bid (see the instructions in the appendix for the exact wording we used to define the amounts of winner's and loser's regrets).

In the core of our experiment, we cross these two information conditions for a 2×2 full-factorial design. Table 1 summarizes the four treatments, their labels, and the sample sizes.

As an additional manipulation to check the effect of showing subjects the probability of winning, we replicated the *Loser's Regret* treatment and the *Both* treatment without showing participants their probability of winning. Those two additional treatments are labeled "*Loser's Regret* (*no prob*)" and "*Both* (*no prob*)" and have sample sizes of 24 and 25, respectively.

Each bidder participated in a single treatment only. Each session lasted for approximately 45 minutes and average earnings, including a $5 participation fee, were $18 (standard deviation of about $1.50).

[6] We should also add that focusing on larger values should not bias the analysis. Specifically, for any fixed information condition, the theory predicts that subjects overbid the risk-neutral Nash equilibrium by the same percentage regardless of their value; the ratio of bid/value should be independent of value. We report the results in terms of the bid/value ratio. Therefore, our choice of value parameters does not bias the theory's predictions.

Table 1 **Summary of the Experimental Design**

Loser's regret information	Winner's regret information: No	Winner's regret information: Yes
No	*None* treatment (20)	*Winner's Regret* treatment (20)
Yes	*Loser's Regret* treatment (20)	*Both* treatment (20)

Note. Treatment labels indicate what information is given to participants (sample sizes are in parentheses).

All sessions were conducted at the Laboratory for Economic Management and Auctions at Penn State University Smeal College of Business during the summer and fall of 2004. Participants were Penn State students, mostly undergraduates, from a variety of majors, recruited through a web-based recruitment system, with earning cash being the only incentive offered. The auction software we used was web-based and was built using PHP and mySQL.

4. Results

4.1. Experimental Hypothesis

The core study directly tests the effect of providing winner's regret- and loser's regret-related information. Arguably, the more apparent (or the more easily observed) a particular form of regret is, the more salient that type of regret will be to the decision maker in practice. For example, if losing bidders do not see how much the winner paid, it may be that only sufficiently astute losers would infer the possibility of having missed an opportunity to win at a favorable price and the corresponding loser's regret. Furthermore, even if bidders are aware of loser's regret, they may well underestimate the amount of regret if they do not know the winner's price. As Engelbrecht-Wiggans (1989) suggests, rather than compute the expected difference between their value and the winner's price in those cases that this is positive, bidders may simply compute the difference between their value and the expected amount paid by the winner. This systematically underestimates the actual regret. Therefore, not showing losers the winner's price may result in loser's underestimating the importance of loser's regret to them.

We will compare bids in treatments with winner's regret and/or loser's regret information to bids in treatments in which the winner's regret and the loser's regret information (as well as the amount of the highest and the second highest bid) are not displayed. This design tests the regret model in conjunction with the auxiliary hypothesis that providing participants with specific information about regret increases the awareness of the feeling of regret relative to not providing it, and this increased awareness

Figure 1 Average Bid/Value in the Four Treatments Over the 20 Decisions Made with the Same Value

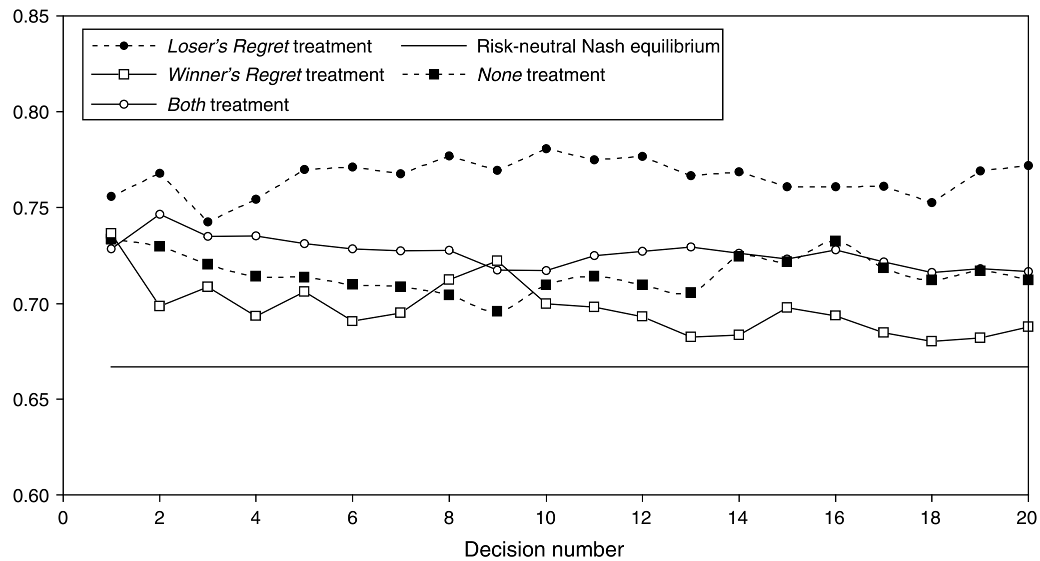

increases the intensity of regret, and thus changes the parameters α and β. Therefore, if we find that the data shifts in the direction consistent with the regret model, we will be able to conclude that we found evidence in support of the regret model jointly with the auxiliary hypothesis. Data inconsistent with the regret model, on the other hand, would imply that we reject either the regret model or the auxiliary hypothesis (or both). Thus, our design provides a tough test of the regret model.

Our model implies two hypotheses, each pertinent to a specific type of regret:

Hypothesis 1 (The Effect of Winner's Regret). *If bidders suffer from winner's regret and adding this information makes winner's regret more salient, then the average bids should decrease.*

Specifically, bids in the *Both* treatment should be lower than in the *Loser's Regret* treatment, and bids in the *Winner's Regret* treatment should be lower than bids in the *None* treatment.

Hypothesis 2 (The Effect of Loser's Regret). *If bidders suffer from loser's regret and removing this information makes loser's regret less salient, then the average bids should decrease.*

Specifically, bids in the *Both* treatment should be higher than bids in the *Winner's Regret* treatment, and bids in the *Loser's Regret* treatment should be higher than bids in the *None* treatment.

These two hypotheses imply that of the four treatments, the bids in the *Loser's Regret* treatment should be the highest, the bids in the *Winner's Regret* treatment should be the lowest, and the order of the *Both* treatment and the *None* treatment would be determined by the relative strength of the two types of regret, and thus cannot be predicted a priori.

4.2. Descriptive Statistics

Figure 1 shows the average bid/value over the 20 decisions participants made with the same value. Each point on the graph represents the average bid/value for 20 participants and five values. The figure provides a sense of how bid levels vary in response to information, as well as over time.

The average bid/value by treatment is as follows: *Both* treatment, 0.7263; *Winner's Regret* treatment, 0.6973; *Loser's Regret* treatment, 0.7660, *None* treatment, 0.7154. Of course, these averages do not take into account the dynamics that may be present in the data: for example, bids in the *Winner's Regret* treatment decrease over time (ordinary least squares p-value $= 0.0013$), whereas bids in the other three treatments do not. To make an adjustment for this learning, we do the analysis based on all decisions as well as on the last half of each set of decisions with the same value. The average bid/value is above the risk neutral Nash equilibrium in three of the four treatments (all but the *Winner's Regret* treatment), and overbidding in those three treatments persists even when we consider only the last half of the decisions. We make the comparisons between the treatments using a one-sided t-test, and the null hypothesis as implied by Hypotheses 1 and 2, and summarize these results along with theoretical predictions in Table 2.

Table 2 shows that we find support for both hypotheses. Each hypothesis includes two predictions, and in both cases one of the predictions is strongly supported and other prediction is in the direction implied by the hypothesis. Additionally, the prediction that the *Loser's Regret* treatment results in the highest average bids and the *Winner's Regret* treatment results in the lowest average bids is also strongly supported (p-value $= 0.0002$).

Table 2 **Summary of the Theoretical Predictions and Experimental Results**

		Both treatment	*Winner's Regret* treatment	*Loser's Regret* treatment	*None* treatment	*p*-value decisions 1–20	*p*-value decisions 11–20
Decisions 1–20							
Average bid/value (std. dev.)		0.7263 (0.0529)	0.6973 (0.0652)	0.7660 (0.0479)	0.7154 (0.0686)		
Decisions 11–20							
Average bid/value (std. dev.)		0.7231 (0.0583)	0.6882 (0.0683)	0.7664 (0.0536)	0.7168 (0.0663)		
Hypothesis	Prediction						
1: Winner's regret	*Both* treatment < *Loser's Regret* treatment	X		X		0.0103	0.0105
	Winner's Regret treatment < *None* treatment		X		X	0.1992	0.0945
2: Loser's regret	*Both* treatment > *Winner's Regret* treatment	X	X			0.0703	0.0479
	Loser's Regret treatment > *None* treatment			X	X	0.0051	0.0066

Notes. Average bid/value (standard deviations in parentheses are displayed in row 1 (for all rounds) and in row 2 (for the last 10 rounds with each value)). The four bottom rows correspond to a comparison of two treatments, as indicated in the second column. The *p*-values (one sided) in the last two columns refer to results of a *t*-test comparing the average bid/value in the two treatments. The unit of observation is the average bid/value for an individual subject for all rounds (column 7) and for last 10 rounds with each value (column 8). There are 20 subjects in each treatment.

When we analyze the decisions in the second half of each decision block (decisions 11–20) all the differences become significant (the second Hypothesis 1 comparison at the 10% level and all the rest at the 5% level). Note, in particular, that the treatment effects persist over time; subjects do not learn to ignore regret—to converge to the risk-neutral Nash equilibrium—despite the uncommonly large amounts of feedback provided by our experiments.

The fact that only the *Winner's Regret* treatment bids decrease significantly over time may have a simple explanation. When we ask our executive MBA students why they bid the way they do in informal class experiments, many express the concern that they can't make a profit unless they win, some explicitly mention the possibility that too low of a bid may result in a missed opportunity to win at a favorable price, but very few (if any) anticipate the possibility of winner's regret. So, imagine that subjects tend to anticipate loser's regret, but tend not to consider winner's regret until they have actually experienced it several times. Then, winner's regret information would lower bids over time, whereas loser's regret would have a smaller effect in the other direction. Bids in the *Winner's Regret* treatment and the *Loser's Regret* treatment do go in the expected directions, but only the winner's regret bids change enough for the change to be statistically significant. The two types of information work in opposite directions, but if winner's regret information has a stronger effect over time, then the bids in the *Both* treatment should decrease over time, but not by as much as in the *Winner's Regret* treatment; they do indeed decrease, but not significantly, whereas they do decrease significantly in the *Winner's Regret* treatment.

We can use the average bid/value in the four information conditions to estimate three of the four regret-related parameters, α (winner's regret weight when winner's regret information is not given), α^W (winner's regret weight when winner's regret information is given), β (loser's regret weight when loser's regret information is not given), and β^L (loser's regret weight when loser's regret information is given). Because there are four parameters to be estimated and only three independent relationships, we can provide an example of possible regret weights by arbitrarily setting $\alpha = 0$ (to establish the scale), and estimate $\alpha^W = 0.2349$, $\beta = 0.4224$, and $\beta^L = 0.6371$ using data from the *Loser's Regret*, *Winner's Regret*, and *Both* treatments. The above estimates are consistent with the data from the *None* treatment. See Appendix A1.2 for the details of the calculations.

4.3. Probability of Winning

To check the effect of showing the participants their probability of winning, we conducted additional versions of the *Loser's Regret* treatment and *Both* treatment without displaying the probability of winning to the subjects. Figure 2 shows the average bid/value over the 20 decisions.

The average bid/value in the *Both* (*no prob*) treatment is 0.7257 (not statistically different from the average of 0.7263 in the *Both* treatment, p-value $=$ 0.4871). And the average bid/value in the *Loser's Regret* (*no prob*) treatment is 0.7828 (not statistically different from the average of 0.7660 in the *Loser's*

Figure 2 Average Bid/Value in the *Loser's Regret* and *Both* Treatments, with and without Winning Probability Information, Over the 20 Decisions Made with the Same Value

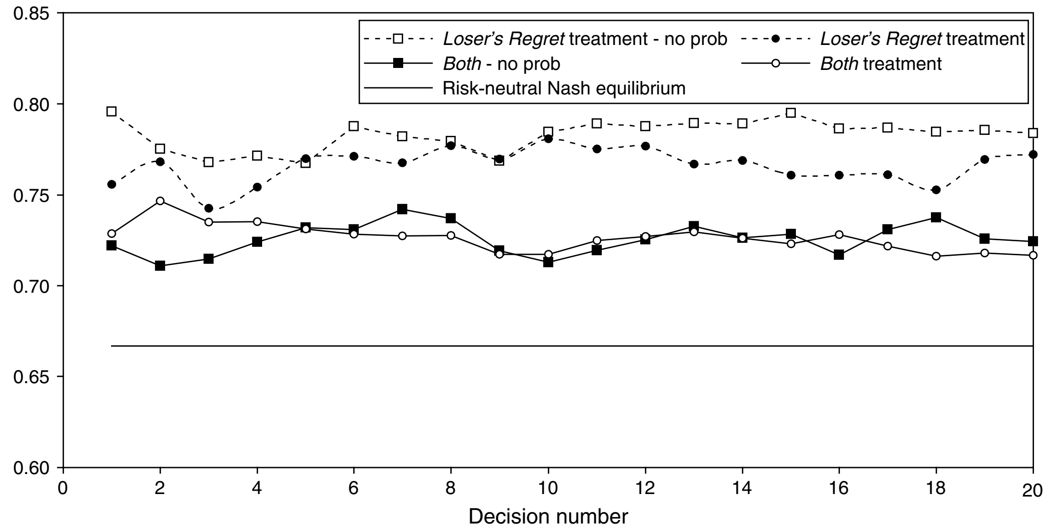

Regret treatment, p-value = 0.1842). Therefore we conclude that, in our setting, showing subjects the probability of winning does not affect behavior. This finding is not particularly surprising in view of the fact that bidders observe the outcome of 10 auctions for every decision they make and, therefore, are likely to be able to obtain the sense of the probability of winning based on their experience. The difference between the *Both* (*no prob*) treatment and the *Loser's Regret* (*no prob*) treatment continues to be significant (p-value = 0.0063), providing additional evidence in support of Hypothesis 1.

4.4. Learning

Before concluding the discussion of our results, we would like to discuss the *learning conjecture* as an alternative explanation for the shifts observed in our data.[7] The four information conditions we analyze in §4.2 provide subjects with a very different amount of feedback information, and the learning conjecture is that our results can be explained by the fact that participants learn faster in conditions with better feedback.

The learning conjecture implies that bids in the *Both* treatment should be lower than bids in the other three treatments. In fact, the observed bids in the *Both* treatment are significantly higher than bids in the *Winner's Regret* treatment. This is consistent with the regret theory but inconsistent with the learning conjecture. The observed bids are significantly lower than bids in the *Loser's Regret* treatment, which is consistent with both, the regret theory and the learning conjecture. However, the observed bids are not significantly different from bids in the *None* treatment, which is inconsistent with the learning conjecture (the regret theory makes no prediction about this comparison).

The learning conjecture also implies that bids in the *None* treatment should be higher than bids in the other three treatments. In fact, the observed bids in the *None* treatment are weakly higher than bids in the *Winner's Regret* treatment, which is consistent with both the regret theory and the learning conjecture. However, the observed bids are significantly lower than bids in the *Loser's Regret* treatment, which is consistent with the regret theory, but not with the learning conjecture.

Regret theory and the learning conjecture make opposite predictions in two cases. These give us a way to separate the two explanations, and we find that the data are consistent with the regret theory, but not with the learning conjecture. The lack of difference between the *None* treatment and the *Both* treatment provides further evidence against the learning conjecture.

5. Discussion and Conclusions

This paper presents a laboratory test of the Engelbrecht-Wiggans (1989) regret theory in auctions. We manipulate the saliency of regret by varying the feedback information provided at the end of the auction. The theory predicts two shifts: (1) When "money left on the table" (winner's) regret is made more salient through announcing that the second highest bid will be revealed, the average bids should decrease, and (2) when missed opportunities to win at a favorable price (loser's) regret is made more salient through announcing that the winning bid will be revealed, the average bids should increase. We observed both of these shifts in our data.

[7] We thank an anonymous referee for pointing out this issue to us.

The design of our experiment differs from standard auction experiments, and understanding the effect of these differences provides additional insight into bidding behavior in first-price auctions. The three major differences are that (1) human bidders compete against computerized rather than human rivals, (2) each decision affects 10 independent auctions, and (3) each value is repeated for 20 consecutive decisions. We use a subject pool that consists of undergraduate students, which is the norm in experimental laboratory economics studies of auctions. The participants in our study typically have little, if any, experience with auctions.

There is recent evidence that correctly understanding gains from sealed-bid auctions is a difficult task for inexperienced bidders (see, for example, Engelbrecht-Wiggans and Katok 2005, Engelbrecht-Wiggans et al. 2007, Harrison and List 2005). Because bidding in auctions has proven to be a difficult task, one of the challenges in designing the experiment is to come up with a setting in which participants are able to fully understand the game. The goal of our design was to provide hands-on learning experience to the bidders as a part of the experimental session and improve their understanding of bidding in auctions.

One pilot study looked at the effect of repeating values. In this treatment, each decision affected only a single auction, and both the winner's regret and the loser's regret information was shown after each auction. We found that the lack of repetition results in higher average bids[8] and in a gradual decrease of bids over time.[9]

In another pilot study, we compared the bidding data from two-person auctions reported by Ockenfels and Selten (2005) to the bidding data in a setting identical to theirs in every way except that one human bids against one computerized rival (the computerized rival was programmed to place bids identical to the bids placed by the corresponding human bidder in the Ockenfels and Selten 2005 experiment). The bidding behavior against computerized opponents was not significantly different from that against human opponents.

In Engelbrecht-Wiggans and Katok (2008), we compared the behavior in treatments in which each decision is used in a single auction to the behavior in treatments in which each decision is used in 10 independent auctions. One human competes against two computerized rivals and each value is repeated 20 times in all treatments. We found that under some feedback conditions, bidding in 10 auctions increases average bids, and in others it decreases average bids. The critical feedback seems to be the winning price—when the winning price is revealed following an auction or a set of 10 auctions, the average bid is higher when it affects 10 auctions than when it affects one. The shift is reversed when the winning price is not revealed. The loser's regret information did not cause a significant shift in that setting. However, just as in the present study, winner's regret information caused average bids to decrease even when each bid affects just one auction.

Regret is also consistent with several other auction-like settings discussed in the literature and known to be inconsistent with risk aversion. In the third-price (Kagel and Levin 1993) and the random-price (Cason 1995) auctions, the winner's bid does not directly affect the price, so winner's regret comes from winning at an unfavorable price. Loser's regret remains the same as in first-price auctions. Just as in first-price auctions, being more sensitive to the winner's regret decreases bids and being more sensitive to the loser's regret increases them (see Engelbrecht-Wiggans and Katok 2007 for a formal derivation). As we have shown, laboratory participants put more weight on the loser's regret than on the winner's regret, and these weights imply bids above the risk-neutral Nash equilibrium in the Cason (1995) random-price auction, which is consistent with the data.

In the third-price auction (Kagel and Levin 1993), expected losses conditional on winning are more than twice as large in auctions with five bidders than they are in auctions with 10 bidders and, consequently, we can expect the winner's regret to be more salient, and depress bids more, in auctions with five bidders than in auctions with 10 bidders. This may qualitatively explain the Kagel and Levin (1993) data: the stronger winner's regret drives bids below the risk-neutral Nash equilibrium level in auctions with five bidders, whereas the weaker winner's regret in auctions with 10 bidders leaves bids above the risk-neutral Nash equilibrium level.

Our results have two practical implications. First, the U.S. Government's policy of disclosing losing bids after the auction may be having a significant effect on its revenue from such auctions. Second, aversion to regret may explain the "overbidding" relative to the risk-neutral Nash equilibrium so commonly observed in the laboratory.

Acknowledgments

The authors gratefully acknowledge support from the National Science Foundation, and thank Axel Ockenfels and

[8] In a treatment in which values were repeated 20 times each, average bids were 37.70, 43.86, 46.41, 51.45, and 52.72 for values of 50, 60, 70, 80, and 90 respectively. Corresponding average bids in the treatment in which values were $\sim U(45, 94)$ and presented without repetition were 40.64, 46.91, 53.16, 56.36, and 56.96 for ranges 45–54, 55–64, 65–74, 75–84, and 85–94, respectively. All differences are highly significant, with $p < 0.05$.

[9] The average bid/value decreases over time in the treatment without repetition ($p < 0.05$), but not in the corresponding treatment with repetition.

the Deutsche Forschungsgemeinschaft for financial support through the Leibniz-Program. The second author gratefully acknowledges the support from the Smeal College of Business and the Center for Supply Chain Research (CSCR) at Penn State's Smeal College of Business.

Appendix 1

A1.1. Proof of Proposition 1

The $N-1$ competitors' bids are independent and distributed uniformly on $[0, A]$. Observe that for $b \le 0$, $\Pi(b;\, v) = \int_{z:\, 0\le z\le v} \beta(v-z)\, dF(z)$, which is independent of b. Therefore, $b=0$ is an expected utility-maximizing bid if $b\le 0$. Similarly, for any bid $b \ge A$,

$$\begin{aligned}\Pi(b;\, v) &= (v-b) - \int_{z:\, z\le A} \alpha(b-z)\, dF(z)\\ &= (v-b) - \int_{z:\, z\le A} \alpha(b-z)(N-1)(z/A)^{N-2}(1/A)\, dz\\ &= (v-b) - \int_{z:\, z\le A} \alpha b(N-1)(z/A)^{N-2}(1/A)\, dz\\ &\quad + \int_{z:\, z\le A} \alpha(N-1)(z/A)^{N-1}\, dz\\ &= v - b(1+\alpha) + \alpha A(N-1).\end{aligned}$$

This is a strictly decreasing function of b, and $b=A$ is the (unique) expected utility-maximizing bid if $b\ge A$. As a consequence of these two observations, all that remains to be done is to determine the expected utility-maximizing bid whenever it is in the interval $[0, A]$.

For any bid $b\in[0, A]$,

$$\begin{aligned}\Pi(b;\, v) &\equiv (v-b)F(b) - \int_{z:\, z\le b} \alpha(b-z)\, dF(z)\\ &\quad - \int_{z:\, b\le z\le v} \beta(v-z)\, dF(z)\\ &= (v-b)(b/A)^{N-1} - \int_{z:z\le b} \alpha(b-z)(N-1)(z/A)^{N-2}(1/A)\, dz\\ &\quad - \int_{z:\, b\le z\le v} \beta(v-z)(N-1)(z/A)^{N-2}(1/A)\, dz\\ &= (v-b)(b/A)^{N-1} - \int_{z:z\le b} \alpha b(N-1)(z/A)^{N-2}(1/A)\, dz\\ &\quad + \int_{z:\, z\le b} \alpha(N-1)(z/A)^{N-1}\, dz\\ &\quad - \int_{z:\, b\le z\le v} \beta(v-z)(N-1)(z/A)^{N-2}(1/A)\, dz\\ &= (v-b)(b/A)^{N-1} - \alpha A(b/A)^N\\ &\quad + \int_{z:\, z\le b} \alpha(N-1)(z/A)^{N-1}\, dz\\ &\quad - \int_{z:\, b\le z\le v} \beta(v-z)(N-1)(z/A)^{N-2}(1/A)\, dz.\end{aligned}$$

Then, $d\Pi(b;\, v)/db = [(v-b)(N-1) - b - \alpha Nb + \alpha(N-1)b + \beta(v-b)(N-1)][(b/A)^{N-2}/A]$. The first-order condition $d\Pi(b;\, v)/db = 0$ implies that $b(v) = (1+\beta)v(N-1)/((1+\beta)(N-1)+(1+\alpha))$ whenever $0\le b\le A$. It is straight forward (though quite messy) to verify that the second-order condition $d^2\Pi(b;\, v)db^2 \le 0$ is satisfied when $b(v) = (1+\beta)v(N-1)/((1+\beta)(N-1)+(1+\alpha))$. In short, $b(v) = (1+\beta)v(N-1)/((1+\beta)(N-1)+(1+\alpha))$ is an expected utility-maximizing bid if $0\le b\le A$. Therefore,

$$\begin{aligned}b^*(v) &= \min\left\{A, \frac{(1+\beta)v(N-1)}{(1+\beta)(N-1)+(1+\alpha)}\right\}\\ &= \min\left\{A, \frac{v(N-1)}{(N-1)+\rho}\right\}, \quad \text{where } \rho \equiv (1+\alpha)/(1+\beta). \quad \square\end{aligned}$$

A1.2. Derivation of Regret Weights

It follows from Proposition 1 that given the parameters of our experiments $b^*(v)/v = (N-1)/(N-1+\rho)$. Because $\rho = (1+\alpha)/(1+\beta)$ and $N=2$, this can be rewritten as $(1+\alpha)/(1+\beta) = 2((v/b^*(v))-1)$. Let $(b/v)^{Treatment}$ represent the average (b/v) in a given treatment, as summarized in the first row of Table 2. The four treatments yield the following four equations:

$$\frac{1+\alpha}{1+\beta} = 2\left(\frac{1}{(b/v)^{None}} - 1\right) = 2\left(\frac{1}{0.7154} - 1\right) = 0.7956$$

$$\frac{1+\alpha}{1+\beta^L} = 2\left(\frac{1}{(b/v)^{LR}} - 1\right) = 2\left(\frac{1}{0.7660} - 1\right) = 0.6108$$

$$\frac{1+\alpha^W}{1+\beta} = 2\left(\frac{1}{(b/v)^{WR}} - 1\right) = 2\left(\frac{1}{0.6973} - 1\right) = 0.8681$$

$$\frac{1+\alpha^W}{1+\beta^L} = 2\left(\frac{1}{(b/v)^{Both}} - 1\right) = 2\left(\frac{1}{0.7261} - 1\right) = 0.7543.$$

Because

$$\frac{(1+\alpha)}{(1+\beta)} = \frac{(1+\alpha)}{(1+\beta^L)}\frac{(1+\alpha^W)}{(1+\beta)} \Big/ \frac{(1+\alpha^W)}{(1+\beta^L)},$$

we have only three independent equations and the four equations together are inconsistent (because parameters are estimated using data). If we arbitrarily set $\alpha=0$, to establish scale, and disregard the *None* treatment, then the reader can verify that $\alpha^W = 0.2348$, $\beta = 0.4224$, and $\beta^L = 0.6371$ satisfy the relationships for the other three treatments. These values also imply $(b/v)^{None} = 0.7399$, which is slightly higher than the actual ratio for that treatment of 0.7154, but is well within a 90% confidence interval for it.

Appendix 2. Instructions for the *Both* Treatment

Overview

You are about to participate in an experiment in the economics of decision making. If you follow these instructions carefully and make good decisions you will earn a considerable amount of money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the monitor will answer it. We ask that you not talk with one another for the duration of the experiment.

In each round of today's session you will be competing with two other bidders to purchase a unit of a fictitious asset. You will be bidding in an auction against two computerized competitors. The computerized competitors have been programmed to bid in a way that would maximized their expected earnings when they bid against likewise programmed competitors. You will make a total of 100 bidding decisions.

On your desks you should have a check-out form, a pen, and two copies of the consent form.

How You Make Money

In the beginning of each bidding decision, you will learn your resale value for a fictitious asset. The resale values for your two computerized opponents have already been predetermined for all auctions in today's session, and they are integers from 1 to 100, with each integer being equally likely. Their resale values in one round have no correlation with their resale values in any other round or with the resale values of any of the other bidders (in other words, all resale values have been drawn independently). The bids of the computerized bidders have also been determined, and they cannot be affected by your decisions today.

Your own value for the asset will be 90 in 20 bidding decisions, 80 in 20 bidding decisions, 70 in 20 bidding decisions, 60 in 20 bidding decisions, and 50 in 20 bidding decisions. You will have the same value in 20 consecutive bidding decisions and then the value will change (and will then stay at this new value for the next 20 consecutive auctions, etc.). The order of your resale values has been determined randomly.

You make one bidding decision for a block of 10 consecutive auctions. In each of those 10 auctions, your competitors will have different values and place different bids, although your own bid and value will remain the same.

You make money by winning the auction at a favorable price. If you win an auction at a price that is below your resale value, then your profit is:

Your resale value − Auction Price.

For example, if your resale value is 60 and you win the auction at a price of 45, then your profit in this auction is $60 - 45 = 15$. Note, if you win the auction at an unfavorable price (at a price that is above your resale value), you will lose money. Because you will know your resale value prior to bidding you can avoid the possibility of losing any money in an auction by not bidding at unfavorable prices. If you do not win the auction, your profit for the round is 0.

The Mechanics of the Auction

You bid in the auction by clicking the "Bid" button and then typing your bid into a box on your screen. On the next screen you will see a message asking you to confirm your bid. The confirmation screen also displays the following information:

- Your value: this is a reminder of your value from the previous screen
- Your bid: this is the bid you have just entered
- Your profit if your bid wins: this is always your value − your bid
- Profit if you lose: 0
- Your probability of winning: this is the percentage of times the bid you just entered would win in this auction). Note: this information is helpful in deciding on the bid amount.
- Your expected profit: this would be your average profit if you made this same bid in this same auction situation many times. (Mathematically, it is your profit if your bid wins multiplied by your probability of winning.)

If you wish to confirm your bid, click the "Confirm" button, and if you wish to change your bid, click the "Cancel" button. You can change your bid as many times as you wish. Your bid will be entered after you have clicked the "Confirm" button.

Your two computerized opponents have been programmed to bid in the beginning of each round, before you have entered your bid. Please note that just as you are not aware of the bid amounts your computerized opponents have placed, neither are they aware of your bid amount at the time their bids are placed.

The bidder who places the highest bid wins the auction and pays the amount they bid. The winner earns Resale value − Purchase Price. The other two bidders who did not win the auction earn zero.

Example

Suppose your resale value is 80, and you place the bid of 65. On the confirmation screen you will see the following information:

Your bid: 65

Expected profit if you win: 15

Profit if you lose: 0

Winning probability: 0.95 Note: this means that 95% of the time a bid of 65 will win

Expected Profit: 14.25 Note: $0.95 \times 15 = 14.25$

Suppose the two bids your computerized opponents placed are 47 and 51. In this case, because your bid of 65 is higher than the other two bids, you win the auction, and earn $80 - 65 = 15$. The two computerized bidders earn 0.

Now suppose that instead, the two bids placed by the computerized bidders were 47 and 66. In this case, the bidder who bid 66 wins the auction and pays 66. You do not win the auction, and earn 0.

Summary Information You Will See at the End of Each Auction

After each bidding decision (at the end of each block of 10 auctions, after you have confirmed your own bid) you will see the following information:

- Your own resale value
- Your own bid amount

For each of the 10 auctions:

- The selling price
- The second highest bid amount
- Your profit and whether or not you won

In addition, the computer will calculate and display for you, in each of the 10 auctions:

- Money left on the table, which is always 0 if you DO NOT win the auction and is your bid—the second highest bid when you do win the auction.
- Missed opportunity to win, which is always 0 when you DO win as well as when your resale value is below the highest bid amount the auction, and otherwise it is Your Resale Value − Winning Bid Amount.

You will also see the average selling price, the average second highest, the number of times you won, the total profit, the total money left on the table, and the total missed opportunity to win for ALL 10 auctions.

How the Session Will Progress

The session will include 1,000 auctions in blocks of 10. You will make 100 bidding decisions, and each decision will be used in 10 consecutive auctions. You will have the same resale value for each 20 consecutive decisions (200 consecutive auctions).

Your earnings from all auctions will contribute to your total earnings from the session. Remember that you will be bidding against two computerized competitors in all 1,000 auctions, and the resale values of your competitors will be integers from 1 to 100, each integer equally likely. The resale values of your competitors will change in each auction (even when your own resale value stays the same).

How You Will Be Paid

At the end of the session, the computer will calculate the total profit you earned in all auctions and will convert it to U.S. dollars at the rate of 1 cent per 10 tokens. Your dollar earnings will be added to your $5 participation fee and displayed on your computer screen. Please use this information to fill out the check-out form on your desk. All earnings will be paid in cash at the end of the session.

If you have any questions, please raise your hand and ask the monitor. If you understand these instructions and wish to continue to participate in this study, please sign one of the two copies of the consent forms on your desk and give it to the monitor before you start the session.

References

Allais, M. 1953. Le Comportement de l'Homme Rationnel Devant le Risqué; Critique Des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* **21** 503–546.

Armantier, O., N. Treich. 2007. Subjective probabilities in games: An application to the overbidding puzzle. CIRANO working paper, Université de Montréal, Quebec, Canada.

Becker, G. M., M. H. Degroot, J. Marschak. 1964. Measuring utility by a single-response sequential method. *Behavioral Sci.* **9** 226–232.

Bell, D. E. 1982. Regret in decision making under uncertainty. *Oper. Res.* **30**(5) 961–981.

Bell, D. E. 1983. Risk premium in decision regret. *Management Sci.* **29**(10) 1156–1166.

Bell, D. E. 1985. Disappointment in decision making under uncertainty. *Oper. Res.* **33**(1) 1–27.

Bolton, G., E. Katok. 2008. Learning-by-doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing Service Oper. Management*, ePub ahead of print January 25, http://msom.journal.informs.org/cgi/content/abstract/msom.1060.0190v1.

Bolton, G., A. Ockenfels. 2000. ERC: A theory of equity, reciprocity, and competition. *Amer. Econom. Rev.* **90**(1) 166–193.

Cason, T. N. 1995. An experimental investigation of the seller incentives in the EPA's emission trading auction. *Amer. Econom. Rev.* **85**(4) 905–922.

Cox, J. C., V. L. Smith, J. M. Walker. 1988. Theory and individual behavior of first price auctions. *J. Risk Uncertainty* **1** 61–99.

Dorsey, D., L. Razzolini. 2003. Explaining overbidding in first price auctions using controlled lotteries. *Experiment. Econom.* **6** 123–140.

Dufwenberg, M., U. Gneezy. 2002. Information disclosure in auctions: And experiment. *J. Econom. Behav. Organ.* **48** 431–444.

Engelbrecht-Wiggans, R. 1989. The effect of regret on optimal bidding in auctions. *Management Sci.* **35**(6) 685–692.

Engelbrecht-Wiggans, R., E. Katok. 2005. Experiments on auction valuation and endogenous entry, in experimental and behavior economics. J. Morgan, ed. *Advances in Applied Microeconomics*, Vol. 13. Chap. 7. Elsevier, Amsterdam, 169–193.

Engelbrecht-Wiggans, R., E. Katok. 2007. Regret in auctions: Theory and evidence. *Econom. Theory* **33** 81–101.

Engelbrecht-Wiggans, R., E. Katok. 2008. A direct test of risk aversion in first price sealed bid auctions. Working paper, Penn State University, University Park, PA. http://lema.smeal.psu.edu/katok/riskAversion_post.pdf.

Engelbrecht-Wiggans, R., E. Haruvy, E. Katok. 2007. A comparison of buyer-determined and price-based multiattribute mechanisms. *Marketing Sci.* **26**(5) 629–641.

Filiz-Ozbay, E., E. Y. Ozbay. 2007. Auctions with anticipated regret: Theory and experiment. *Amer. Econom. Rev.* **97**(4) 1407–1418.

Grether, D. M., R. C. Plott. 1979. Economic theory of choice and the preference reveral phenomenon. *Amer. Econom. Rev.* **69** 623–638.

Harrison, G. W., J. A. List. 2005. Naturally occurring markets and exogenous laboratory experiments: A case study of the winner's curse. Working paper, University of Chicago, Chicago.

Isaac, R. M., D. James. 2000. Just who are you calling risk averse? *J. Risk Uncertainty* **20** 177–187.

Isaac, R. M., J. M. Walker. 1985. Information and conspiracy in sealed bid auctions. *J. Econom. Behav. Organ.* **6** 139–159.

Ivanova-Stenzel, R., T. C. Salmon. 2008. Revenue equivalence revisited. *Games Econom. Behavior*. Forthcoming.

Kagel, J. H. 1995. Auctions: A survey of experimental research. J. H. Kagel, A. E. Roth, eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 501–585.

Kagel, J. H., D. Levin. 1993. Independent private value auctions: Bidder behavior in first-, second-, and third-price auctions with varying number of bidders. *Econom. J.* **103** 868–879.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decisions under risk. *Econometrica* **47** 263–291.

Loomes, G., R. Sugden. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Econom. J.* **92**(368) 805–824.

Luce, R. D., H. Raiffa. 1956. *Games and Decisions*. John Wiley & Sons, New York.

Morgan, J., K. Stiglitz, G. Reis. 2003. The spite motive and equilibrium behavior in auctions. *Contributions Econom. Anal. Policy*. **2**(1) article 5.

Ockenfels, A., R. Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games Econom. Behav.* **51** 155–179.

Rothkopf, M. H., S. Park. 2001. An elementary introduction to auctions. *Interfaces* **31**(6) 83–97.

Savage, L. J. 1951. The theory of statistical decisions. *J. Amer. Statist. Assoc.* **46** 55–67.