

Rent Seeking with Regretful Agents: Theory and Experiment

Kyle Hyndman

Maastricht University & Southern Methodist University
k.hyndman@maastrichtuniversity.nl, <http://www.personeel.unimaas.nl/k-hyndman>

Erkut Y. Ozbay

University of Maryland, ozbay@umd.edu, <http://econweb.umd.edu/~ozbay/>

Pacharasut Sujarittanonta

University of Maryland

We investigate both theoretically and experimentally the role that information disclosure has on behavior in all pay environments in which all agents must exert costly effort, but only the winner is rewarded. Through the lens of all pay auctions, we show that bidders who have regret concerns when they lose an auction may alter their bids. Whether or not information about the winning bid is disclosed, all losing bidders regret not bidding 0. However, when feedback on the winning bid is provided, those bidders who lost at an affordable price may regret not bidding more. We show that the former effect causes bidders to lower their bids, while the latter effect causes bidders to raise their bids. Our experimental results indicate overbidding whether or not the winning bid is disclosed to losing bidders. However, disclosing the winning bid leads to even more aggressive bidding, increases revenues, decreases the frequency of dropouts and also leads to more efficient allocations. Thus our results show that information disclosure is a powerful tool that mechanism designers may wish to exploit to extract more rents in all pay environments.

Key words: all-pay auctions; regret; experiment; rent-seeking

1. Introduction

All-pay environments are those settings in which many agents engage in some costly activity but only one “winner” is awarded with a prize. Numerous practical examples are present both in the literature and in everyday life. For example, employees in many industries must exert costly effort in order to win a bonus or a job promotion. Special interest groups hire lobbyists in order to influence legislation. Businesses devote considerable resources to developing and patenting new products or processes.¹

Competitions to promote technological development are similar in flavor and have a long history, dating back hundreds of years. They have been used to improve navigation (British government,

¹ For a discussion of these and other examples, see Tullock (1980), Baye et al. (1993), Che and Gale (1998) and Müller and Schotter (2010), among others.

1700s), advance chemical engineering processes (French Academy, 1800s), to encourage advances in aviation (Orteig Prize) as well as unsolved mathematics problems (e.g., Fermat’s Last Theorem).² Two prominent recent examples were the so-called Ansari X Prize and the Netflix Prize. In October 2004, the \$10 million Ansari X Prize was awarded after SpaceShipOne made two successful trips 100KM above the Earth’s surface in a two-week period. In total, 26 teams devoted more than \$100 million to compete for the prize.³ Starting in October 2006, Netflix offered a \$1 million prize for the person or group who could develop an algorithm to improve upon their own movie recommendation system by 10%. After nearly 3 years which saw 51,051 participants make valid submissions, the prize was awarded in September 2009.⁴ It seems clear that these competitions fit the basic description of an all-pay environment: participants provide their solutions by exerting costly effort, while only the best is awarded.

Though less ambitious in terms of scale to the competitive prizes described above, online community knowledge markets have recently become much more common and can be seen through the lens of an all-pay environment (DiPalantino and Vojnovic 2009). In these markets one seeks solutions to a problem via open calls to large-scale communities (known as crowdsourcing) and community-members submit solutions, with the best solution, sometimes, receiving a reward. Amazon’s Mechanical Turk, Yahoo! Answers, www.innocentive.com, www.taskcn.com, are some of the widely used examples of crowdsourcing where the users submit their answers to the proposed questions (see e.g. Yang et al. (2009)).

Because of their wide usage in the real world, contests and tournaments have received considerable attention both among theorists and experimentalists. The most well-known theoretical analysis is due to Lazear and Rosen (1981) who discuss rank-order contests in relation to labor contracts. Somewhat later, Moldovanu and Sela (2001) characterized the optimal allocation of prizes in contests. If costs are linear, then a single prize is optimal, while if costs are quadratic, it may be optimal to award one large and one smaller prize.

Many of the experimental tests of contests have been concerned with comparing effort under various incentive mechanisms such as a piece rate, proportional prizes or more traditional “winner-take-all” contests (see, e.g., Bull et al. (1987), Schotter and Weigelt (1992), van Dijk et al. (2001),

² Further details can also be found at <http://www.xprize.org/x-prizes/incentivized-competition-heritage>. Boudreau et al. (Forthcoming) contains more historical examples. They also report on a study of over 9,500 software development contests. Their main purpose is to understand the effect of adding “competitors”. While adding competitors reduces incentives to exert effort, they show that it also increases the likelihood of finding “extreme” solutions, especially for highly uncertain problems.

³ For more information, see <http://space.xprize.org/ansari-x-prize>.

⁴ For more information, see www.netflixprize.com.

Orrison et al. (2004) and Cason et al. (2010), among others). Guided by key features of R&D contests and political campaigns, Sheremeta (2010a,b) has studied behavior in multi-stage contests where, in the first stage (e.g., primaries), players compete to make it to the next stage of the contest (e.g., presidential campaign). Focusing on the results about contests, one fairly robust conclusion — and one that is consistent with our results — is that average effort is generally at least as high as theory predicts, and often higher, but also that effort is highly variable. Indeed, this result seems consistent with conventional wisdom, which suggests that in many all-pay environments the amount of effort exerted is bimodal, with some exerting very little effort and others, seemingly, exerting too much effort. A recent and convincing experimental study of this phenomenon is Müller and Schotter (2010). They show that “low-ability” subjects have a tendency to dropout and not exert any effort, while “high-ability” subjects exert significantly more effort than predicted by theory.

In this paper, we take this conventional wisdom for granted and ask whether there are any environmental changes that one can make in order to influence both the frequency of dropouts as well as the amount of effort that is exerted. We do so both theoretically and experimentally. In particular, we study whether changing the feedback provided to agents affects behavior.

Because of its analytical convenience, and the fact that it has been well-studied, both theoretically and experimentally, we adopt the framework of all-pay auctions. The findings of the all-pay auctions are quite robust in the experimental literature: subjects tend to overbid (see, e.g., Gneezy and Smorodinsky (1999); Lugovskyy et al. (Forthcoming)). Moreover, when the bidders’ valuations are private information, it is only the bidders with higher valuations who tend to overbid (see Noussair and Silver (2006)). The results from all-pay auctions are also consistent with the aforementioned results on contests. To explain the deviations from theory, authors have turned to risk aversion (Fibich et al. 2006) and logit equilibrium (Anderson et al. 1998). However, these explanations are silent when the information regarding the bids is disclosed at the end of the auction. We argue that this effect is due, in part, to the anticipation of regret.⁵

Naturally, a lobbyist knows his political contribution; a fan knows when he started waiting for ticket to a popular concert; a bidder knows how much he bids in an auction. However, the contributions, the waiting times or other bids may or may not be learned.⁶ With this information,

⁵ As argued in Filiz-Ozbay and Ozbay (2007), Engelbrecht-Wiggans and Katok (2007, 2008), Filiz-Ozbay and Ozbay (2010), the anticipation of this regret (loser regret) predicts overbidding in first- and third-price auction in line with the findings of the experimental literature (e.g., Kagel and Levin (1993)).

⁶ It has been shown that the information regarding the other bids alters the bidding behavior (see, e.g., Isaac and Walker (1985); Dufwenberg and Gneezy (2002), Ockenfels and Selten (2005)). For example, in first-price sealed bid auction, informing the winning bid to all bidders increases the bids, and hence the revenue of the auctioneer is boosted (see, e.g., Filiz-Ozbay and Ozbay (2007) and Engelbrecht-Wiggans and Katok (2007, 2008)).

a loser can figure out how far he was away from being a winner. If he realizes that he lost at an affordable level, he may experience regret for not being sufficiently aggressive. Anticipating that he may later come to regret his decision, he may exert more effort to win the prize. Hence we will observe workaholics, excessive lobbying or people forming lines days before a new product launch or concert tickets go on sale.

Of course, in an all-pay environment, since all agents are required to pay, there is another source of regret for the losing agent. Any losing agent regrets that he exerted a positive amount rather than dropping out. We call this type of loser regret as “all-pay loser regret”. Since the agents know their own effort, it is enough to learn that he lost for the activation of the all-pay loser regret.

In order to control for any additional effect other than regret, we conduct an experiment of an all-pay auction consisting of two treatments, with and without feedback on the highest bid. This allows us to test our regret hypothesis, while leaving risk preferences unaffected across treatments.⁷

Our paper presents four main experimental results. First, overbidding is present and persistent in both treatments, though it is significantly more prevalent in the full feedback treatment in which the winning bid was disclosed to the loser of the auction. Second, in both treatments, there is a bifurcation in behavior. Those who make positive bids tend to overbid, while those with lower valuations appear to drop out and bid zero. Though this bifurcation is present in both treatments, it is stronger in our partial feedback treatment. In particular, bidders in the partial feedback treatment are more likely to submit a bid of zero and do so for higher valuations. Third, conditional on bidders placing a strictly positive bid, bidders bid more aggressively in the full feedback treatment. Surprisingly, this only appears to be so for intermediate bidder valuations. For extremely low and extremely high bidder valuations, there is no difference in the bid functions. Finally, while in the partial feedback treatment, there is some convergence to the risk neutral Nash equilibrium, bids are still substantially higher in the full feedback treatment. Moreover, our results also suggest that subjects adjust their bids from round to round in a manner consistent with regret. That is, controlling for changes in one’s value, subjects raise their bid if it is revealed that they lost at an affordable price and they often lower their bids if they find out they lose (possibly also learning about the winning bid).

⁷ The LC-1 and LC-1 (control) treatments of Müller and Schotter (2010) also provide suggestive experimental evidence on the role of feedback, though the discussion is limited, and is not attributed to regret. Their LC-1 treatment, which is virtually isomorphic to the all-pay auction experiments we conduct, provides feedback on the ability of the subject who exerted the highest effort, while their LC-1 (control) treatment did not. Müller and Schotter (2010) found, like us, that the frequency of dropouts declines when feedback is provided. Although the evidence is suggestive, we must be careful about drawing conclusions from this comparison because, in addition to changing the feedback provided to losing subjects, both the matching protocol and the strategy space were changed between the two treatments.

Taken as a whole, our results suggest that feedback about the winning bid, which we argue activates a kind of loser regret, causes bidders to be less likely to drop out of the auction and causes them to bid more aggressively. Both of these features lead to higher revenue for the auctioneer and higher auction efficiency. Thus, in all-pay auctions, or other similar environments, it may be beneficial (from the perspective of the mechanism designer) to provide feedback on the effort exerted by the winners.

Our paper also contributes to the literature on gender differences in behavior. The conventional wisdom, as expressed in the survey of Croson and Gneezy (2009) is that women are generally more risk averse and less inclined to enter into competitive environments. In our setting, this should imply that (i) women are more likely to drop out (*i.e.*, place a bid of 0) and (ii) conditional upon placing a positive bid, they should bid more aggressively than men. While our results confirm this latter prediction, we actually find that women are *less* likely to drop out than men. Both of these findings help explain why women earn significantly less than men in our experiments.

The paper proceeds as follows: Section 2 provides the equilibrium analysis for the bidders with regret concerns. Section 3 describes our experimental design. Section 4 presents our experimental results. Section 5 concludes.

2. Model

There is one item for sale. For notational simplicity, consider two symmetric bidders, indexed by $i = 1, 2$. Bidder i 's valuation, x_i , is independently distributed on the interval $[0, 1]$ according to the distribution function F with associated density function f . Each bidder submits a bid b_i . The bidder who submits the highest bid wins the item. Each bidder pays his bid regardless of winning or losing.

In our full feedback treatment in which the winning bid is disclosed to the loser, bidder i 's payoff is

$$U_F(x_i, b_i) = \begin{cases} x_i - b_i, & \text{if win} \\ -b_i - \beta b_i - \gamma(x_i - b_w), & \text{if lose and } b_w < x_i, \\ -b_i - \beta b_i, & \text{if lose and } x_i < b_w \end{cases}$$

where b_w is the winning bid, β is the all-pay loser regret coefficient and γ is the loser regret coefficient. Assume that β and γ are non-negative. The all-pay loser regret coefficient indicates how strongly the loser feels regret for paying without winning the item. The loser regret coefficient indicates how strongly the negative emotion a loser feels for losing at an affordable price.

Let $B_F(x)$ be a bid function in a symmetric equilibrium. Bidder i 's expected payoff is

$$u_F(x_i, b_i) = x_i F(B_F^{-1}(b_i)) - b_i - \beta b_i (1 - F(B_F^{-1}(b_i))) - \gamma \int_{B_F^{-1}(b_i)}^{B_F^{-1}(x_i)} (x_i - B_F(s)) dF(s).$$

The first term is expected payoff conditioned on winning. The second term is a monetary payment made to the seller. The third term is a negative expected payoff from all-pay loser regret. The last term is a negative expected payoff from loser regret conditioned on losing at affordable price—the opponent's bid is between his bid and his valuation.

Differentiating the expected payoff with respect to b_i and rearranging yields

$$\frac{(1 + \gamma)xf(x)}{1 + \beta(1 - F(x))} = B'_F(x) + \frac{(\gamma - \beta)B_F(x)f(x)}{1 + \beta(1 - F(x))}.$$

With integrating factor $e^{\int \frac{(\gamma - \beta)f(x)}{1 + \beta(1 - F(x))} dx}$, we can solve for the equilibrium bidding function as follows:

$$B_F(x) = \frac{1}{e^{\int \frac{(\gamma - \beta)f(x)}{1 + \beta(1 - F(x))} dx}} \int_0^x \frac{(1 + \gamma)xf(x)}{1 + \beta(1 - F(x))} e^{\int \frac{(\gamma - \beta)f(x)}{1 + \beta(1 - F(x))} dx} dx.$$

In our partial feedback treatment, the final prices were not reported. Therefore, bidders do not feel loser regret. Each bidder makes a payment to the seller equal to his bid. Bidder i 's payoff is

$$U_P(x_i, b_i) = \begin{cases} x_i - b_i, & \text{if win} \\ -b_i - \beta b_i, & \text{if lose} \end{cases}.$$

The payoff is similar to that in Treatment F with the loser regret coefficient, γ , equal to zero. Let $B_P(x)$ be a bid function in a symmetric equilibrium:

$$B_P(x) = \frac{1}{1 + \beta(1 - F(x))} \int_0^x x dF(x).$$


In the experiment, the bidders' valuations were uniformly distributed. Hence, the symmetric equilibrium bid function in an all-pay auction with feedback becomes

$$B_F(x, \beta, \gamma) = \begin{cases} \frac{(1 + \gamma)(1 + e^{x\gamma}(x\gamma - 1))}{\gamma^2 e^{x\gamma}}, & \text{if } \beta = 0 \\ \frac{(1 + \beta)[-x\beta + (1 + \beta)(\log(1 + \beta) - \log(1 + \beta - x\beta))]}{\beta^2}, & \text{if } \beta = \gamma \\ \frac{(1 + 2\beta)(x\beta + (-1 - \beta + \beta x)\log(1 + \beta) + (1 + \beta - \beta x)\log(1 + \beta - \beta x))}{\beta^2}, & \text{if } 2\beta = \gamma \\ \frac{(1 + \beta - x\beta)^{\gamma/\beta - 1}(1 + \gamma)((1 + \beta)^{2 - \gamma/\beta} + (1 + \beta - x\beta)^{1 - \gamma/\beta}(x\gamma - 1 - (1 + x)\beta))}{(\gamma - 2\beta)(\gamma - \beta)}, & \text{o.w.} \end{cases}$$

and the equilibrium bid function with partial feedback becomes

$$B_P(x, \beta) = \frac{x^2}{2 + 2\beta(1 - x)}.$$

For illustrative purposes, Figure 1 shows comparison of bid functions with various parameters when valuations are uniformly distributed on the interval $[0, 1]$. As in first-price sealed bid auction, a larger loser regret coefficient leads to more aggressive bidding. Without loser regret, the bid function is decreasing in all-pay loser regret.

Figure 1 BID FUNCTIONS WITH VARIOUS PARAMETERS (β, γ) 

3. Experimental Design

The experiments were conducted at the Experimental Economics Lab at the University of Maryland. All subjects were undergraduate students at the University of Maryland, College Park, at the time of the experiments. We conducted two treatments: partial feedback (P) and full feedback (F), and four sessions per treatment. No subject participated in more than one session. For the full feedback treatment one session had 14 subjects and three had 16 subjects, while for the partial feedback treatment two sessions had 16 subjects, one had 14 subjects and one had 12 subjects. Overall, therefore, we had 62 subjects in the full feedback treatment and 58 subjects in the partial feedback treatment and, conservatively, four independent observations per treatment. Each session took approximately one hour. At the beginning of the experiment, subjects were seated at isolated cubicles with a computer and given a set of experimental instructions (see Appendix). After reading the instructions, subjects were allowed to ask questions regarding the auction process. Once all questions were answered, the experiment started. The experiment was programmed in z-Tree (Fischbacher 2007).

In each session, subjects participated in 21 auctions. The first auction was a practice auction and, as such, was not used in the calculation of earnings. In each auction, subjects were randomly matched in groups of two, and players within each group competed for a fictitious item. Subjects were randomly rematched at the beginning of each auction and bidders' identities were anonymous.

At the beginning of each auction, subjects were given their private valuations for the item on their computer screens. All valuations were independently and uniformly distributed between 0 and 100, rounded to the nearest cent. Bidders had an unlimited amount of time to submit their bid, but were encouraged to do so within 30 seconds. In both treatments, after both bids were submitted, players learned their payoff and whether or not they won the fictitious item. Additionally, in the full feedback treatment, the price paid by the winner was also revealed to both bidders.

Valuations, bids and payoffs were denoted in Experimental Currency Units (ECU). For the purposes of calculating subjects' final payment, one of the 20 auctions was randomly selected by the computer software. The number of ECU earned in this randomly chosen auction was converted to US dollars at the rate of 20¢ per ECU. This amount was added to the \$20 participation fee. Cash payments were made at the conclusion of the experiment. On average, subjects earned \$20.36 (inclusive of the participation fee) with a minimum payment of \$6.20 and a maximum payment of \$33.04.

4. Experimental Results

As noted above, for each treatment, we conducted four sessions, with each session consisting of 20 periods. Therefore, our dataset contains 1,240 auctions in the full feedback treatment and 1,160 auctions in the partial feedback treatment. We begin by taking a very conservative approach and assuming that the unit of independent observation is the session average, giving us four independent observations per treatment. Table 1 provides a number of summary statistics, and a preliminary comparison of the full feedback and partial feedback treatments. Panel (a) takes session averages over all 20 periods, while panel (b) takes session averages over only the last 5 periods. In both cases, the conclusions are the same.

For the moment, focus on panel (a). The first row looks at the average revenue generated in each session. As can be seen, revenue is about 28.7% higher in the full feedback treatment. Moreover, a ranksum test rejects the null hypothesis (at the 2% level) that revenue is the same across treatments. The second row looks at the difference between actual revenue and expected revenue (assuming risk neutral bidders). As can be seen, in both treatments, over-bidding is prevalent; however, it is more-so in the full feedback treatment. Again, the ranksum test rejects the null hypothesis that the revenue difference is the same. The third row looks at the frequency with which the high-valuation bidder won each auction. As can be seen, efficiency is approximately 72.6% in the partial feedback treatment, and 92.3% in the full feedback treatment, and the difference between treatments is statistically significant. Finally, the fourth row looks at what we call the “realized gain from trade,”

Table 1 Summary statistics (based on session averages)

(a) ALL 20 PERIODS						
	Partial Feedback		Full Feedback		Ranksum test [†]	<i>p</i> -value [†]
	Mean	Std. Dev.	Mean	Std. Dev.		
Revenue	50.83	1.98	65.40	4.21	2.31	0.021
Revenue difference	17.00	3.03	30.52	4.84	2.31	0.021
Efficiency	0.726	0.027	0.823	0.054	2.02	0.043
Realized gain from trade	0.889	0.017	0.939	0.021	2.31	0.021

(b) LAST 5 PERIODS						
	Partial Feedback		Full Feedback		Ranksum test [†]	<i>p</i> -value [†]
	Mean	Std. Dev.	Mean	Std. Dev.		
Revenue	42.03	5.97	59.90	5.83	2.31	0.021
Revenue difference	9.26	6.01	25.39	6.52	2.02	0.043
Efficiency	0.691	0.078	0.871	0.048	2.32	0.020
Realized gain from trade	0.868	0.063	0.948	0.014	2.02	0.043

[†]These tests are based on session averages. For each treatment, we have 4 sessions, for a total of 8 independent observations. Efficiency is the frequency that high valuation bidder won the auction. Realized gain from trade is defined as the winner's value divided by the highest value.

which we define to be the average ratio of the winning bid to the highest valuation for each auction. As with the other measures, the realized gains from trade are significantly higher in the full feedback treatment. Given that we have taken a very conservative approach and that we find such strong differences between the treatments, we already feel justified in saying that the feedback provided — and, therefore, regret — has a strong influence on behavior.

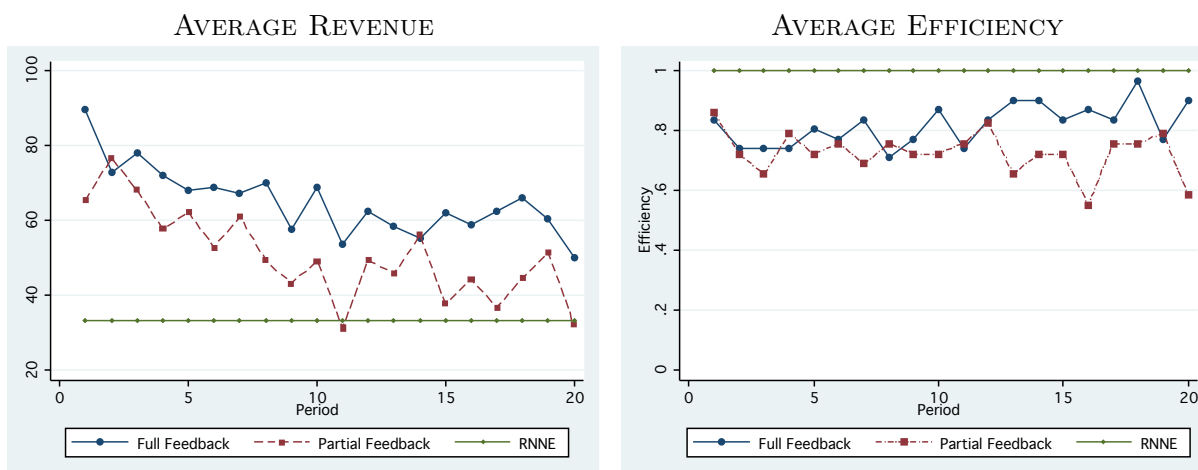
As can be seen, whether we look at all 20 periods or only the last 5 periods, the full feedback treatment generates significantly higher revenues, more over-bidding and greater auction efficiency. That being said, both revenues and overbidding are lower over the last 5 periods, which is consistent with learning. Somewhat surprisingly, however, we see that efficiency is actually somewhat lower in both treatments over the final 5 periods.

Since revenue in an all-pay auction is the sum of all bids, it is therefore an indicative measure of the amount of overbidding. Figure 2 shows average revenue by period. As can be seen, revenues in the full feedback treatment are higher than the partial feedback treatment in all but two periods. In both treatments, we observe a clear downward trend in revenues towards the risk neutral equilibrium expected revenue of 33.33 ECU. Regardless, as Table 1(b) and Figure 2 suggest, overbidding remained prevalent throughout the experiment — especially in the full feedback treatment, revenues were still well above the risk neutral equilibrium prediction over the last 5 periods.

4.1. Bidding Behavior

In this section we turn to a deeper analysis of bidder behavior. The fact that the bias was higher in the full feedback treatment than in the partial feedback treatment, while there was no difference

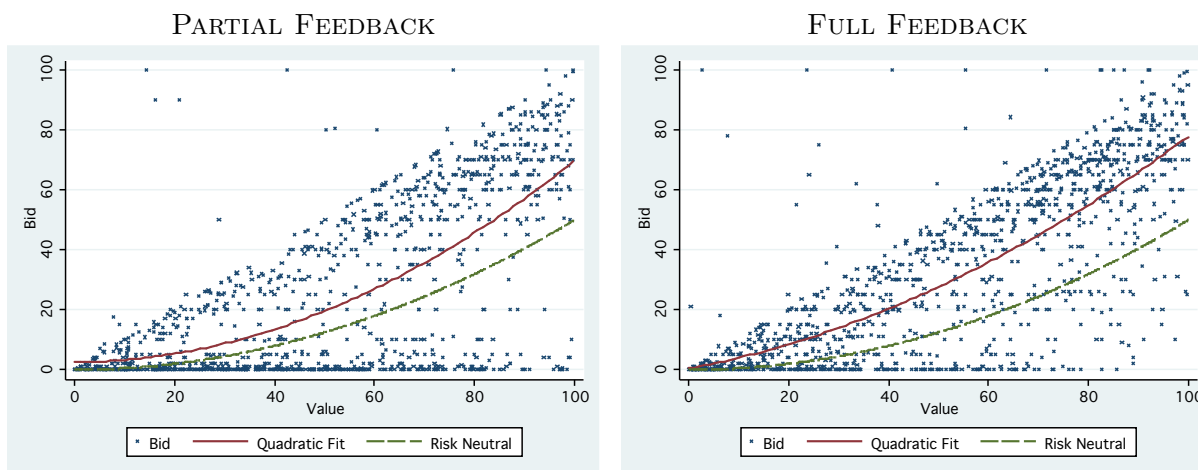
Figure 2 AVERAGE REVENUE AND EFFICIENCY BY PERIOD



in absolute bias between the two treatments, suggests that behavior was more bifurcated in the partial feedback treatment with a substantial amount of both over- and under-bidding.

We begin in Figure 3, which contains a scatterplot of all bids for both treatments. The solid line represents the best-fitting quadratic function, while the dashed line represents the risk neutral equilibrium bid function. As can be seen, and consistent with our earlier results, overbidding is present in both treatments but appears to be more pronounced under full feedback. It also looks as if there was a greater number of zero bids in the partial feedback treatment.

Figure 3 SCATTERPLOT OF BIDS



Presently, we will provide a more rigorous estimation of the bid functions for each treatment. However, before we do that, we wish to make more precise some of the apparent features from Figure 3. Specifically, in Table 2, we report the average frequency of dropouts, the average bid

as well as the average bid conditional on placing a strictly positive bid. We report the session averages, averaging over all realized values as well as in increments of 20 points. As can be seen, approximately 24.4% of bids were zero in the partial feedback treatment, while only 12.5% of bids were zero in the full feedback treatment, and this difference is significant at the 5% level. In both treatments, the frequency of dropouts declines as valuations increase; however, in every interval, there are significantly more dropouts in the partial feedback treatment.

Table 2 ANALYSIS OF BIDS

	Range of Values	Partial Feedback		Full Feedback		Ranksum test	p -value
		Mean	Std. Dev.	Mean	Std. Dev.		
Frequency of dropouts	0 – 100	0.244	0.055	0.125	0.020	2.31	0.021
	0 – 20	0.478	0.054	0.346	0.099	2.02	0.043
	20 – 40	0.325	0.088	0.160	0.035	2.31	0.021
	40 – 60	0.263	0.100	0.118	0.063	1.73	0.083
	60 – 80	0.101	0.060	0.025	0.018	2.02	0.043
	80 – 100	0.068	0.063	0.000	0.000	2.46	0.014
Average bid	0 – 100	25.416	0.99	32.702	2.10	2.31	0.021
	0 – 20	3.964	1.47	3.732	1.44	0.29	0.773
	20 – 40	9.162	1.93	13.831	1.96	2.02	0.043
	40 – 60	18.185	3.78	27.757	4.83	2.31	0.021
	60 – 80	37.334	3.37	46.299	4.77	2.31	0.021
	80 – 100	56.407	6.78	65.655	3.10	2.02	0.043
Average Bid ($b > 0$)	0 – 100	33.687	1.70	37.406	3.05	2.02	0.043
	0 – 20	7.656	3.09	5.710	1.99	0.87	0.387
	20 – 40	13.664	2.68	16.527	2.75	1.44	0.149
	40 – 60	24.602	3.56	31.505	5.05	1.73	0.083
	60 – 80	41.513	1.71	47.422	4.12	2.31	0.021
	80 – 100	60.399	4.53	65.655	3.10	1.44	0.149

Next look at the average bids in the two treatments. Bids are significantly higher in the full feedback treatment overall. One difference, however, is that average bids are no different between treatments when the valuations are small (i.e., in the interval $[0, 20]$); in all other cases, subjects bid significantly more aggressively in the full feedback treatment. Finally, look at the average bid conditional upon placing a strictly positive bid. While the difference between the two treatments is significant overall, when we look at 20-point intervals, it is only when bids are between 40 and 80 that there is a significant treatment effect.

These results suggest that feedback increases bids in two ways. First, providing feedback on the winning bid leads to significantly fewer dropouts. Second, conditional on bidding a positive amount, providing feedback on the winning bid increases bids, though only for intermediate valuations.

That the effect is only present for intermediate valuations would appear to make sense since this is the region where one's chances of winning are most uncertain and, therefore, where regret is most likely to be activated. When one's valuation is high, the bidder can be reasonably confident that she has the highest valuation and is likely to win (hence regret is inactive), while when one's valuation is low, it is very unlikely that the bidder will lose the auction *at an affordable price* (hence regret is inactive).

Let us now return to a more rigorous estimation of the bid functions. While Figure 3 and Table 2 already provide a lot of information, they do not provide the complete story. In Table 3, we report results from parametric estimates of the bid function assuming, as is true for the risk neutral bid function, that it has a quadratic structure. In particular, we estimate:

$$b_{it} = \beta_0 + \beta_1 v_{it} + \beta_2 v_{it}^2 + \mathbb{I}[P]_{it} \cdot (\beta_3 + \beta_4 v_{it} + \beta_5 v_{it}^2) + \mu_i + \epsilon_{it},$$

where $\mathbb{I}[P]_{it}$ is an indicator variable taking value 1 if the observation is from the partial feedback treatment and 0 otherwise. We report results for both a random effects OLS and a random effects Tobit, which controls for censoring at 0 and 100. We also report the results of a random effects logit model which looks at how the probability of bidding 0 varies with one's valuation.

Table 3 RANDOM EFFECTS REGRESSIONS: BID FUNCTIONS

	ALL BIDS			POSITIVE BIDS	
	Tobit	OLS	Logit	Tobit	OLS
cons	-6.829*** [2.426]	0.845 [2.014]	-0.553 [0.390]	1.923 [2.265]	1.886 [2.251]
value	0.524*** [0.0751]	0.305*** [0.0627]	-0.0697*** [0.00619]	0.414*** [0.0699]	0.416*** [0.0696]
value ² × 10 ⁻³	3.12*** [0.72]	4.62*** [0.612]		3.24*** [0.655]	3.20*** [0.652]
$\mathbb{I}[P]$	-5.217 [3.612]	1.314 [2.926]	0.778 [0.549]	0.0699 [3.433]	0.0935 [3.412]
value × $\mathbb{I}[P]$	-0.168 [0.113]	-0.272*** [0.0913]	0.0135* [0.00766]	-0.177 [0.108]	-0.179* [0.107]
value ² × 10 ⁻³ × $\mathbb{I}[P]$	1.33 [1.07]	1.76** [0.888]		1.26 [0.999]	1.29 [0.994]
N	2400	2400	2400	1967	1967
LL / R ²	-8841	0.556	-696.2	-8214	0.539

Standard errors in brackets.

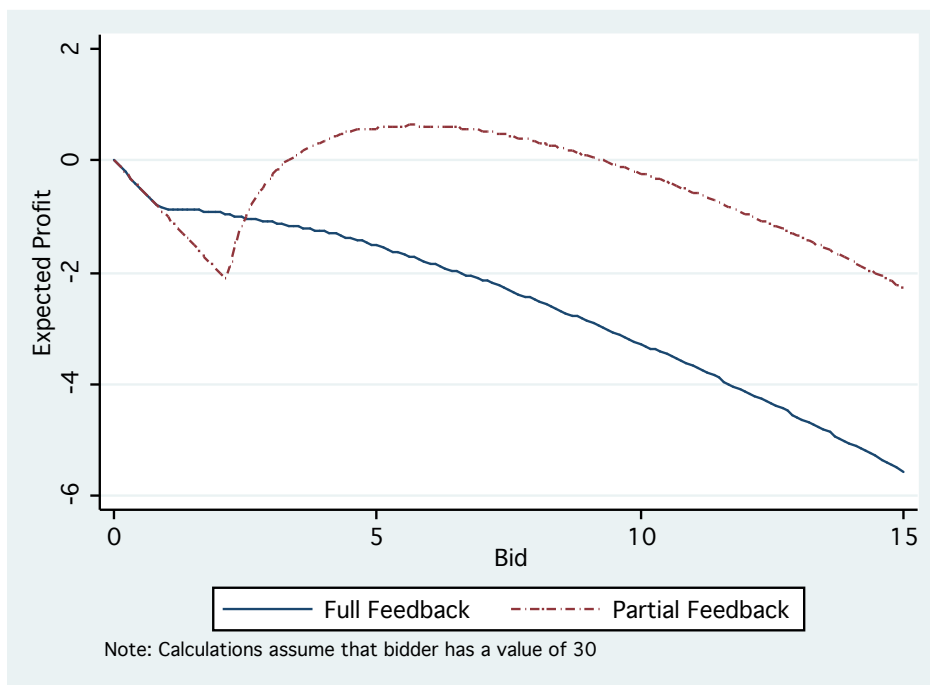
*** significant at 1%; ** significant at 5%; * significant at 10%.

Notice that the coefficient on any one treatment interaction (i.e., β_3 , β_4 or β_5) is not often significant. However, if we conduct the test that our three treatment interactions are *jointly* zero (i.e., $H_0 : \beta_3 = \beta_4 = \beta_5 = 0$) then in all cases we can easily reject this hypothesis. When we condition

on all bids, for both our OLS and Tobit, we have $p \ll 0.01$, while when we condition only on positive bids, we have that $p = 0.03$. Thus, as suggested by Table 2 and Figure 3, the treatment effect is diminished when we look at positive bids. Finally, observe that, consistent with Table 2, the logit estimation indicates that subjects in the partial feedback treatment are more likely to place a zero bid for higher values than are subjects in the full feedback treatment.

As we have seen, there are clear deviations from the risk neutral Nash equilibrium, with substantial overbidding in both treatments (more so in the full feedback treatment). A natural question to ask is then, given the estimated behavior in each treatment, what would have been a best response? In Figure 4 we give a flavor of this for the specific example of a bidder with a value of 30 in each of the treatments. Specifically, for each treatment, we use the estimated bid function from the OLS regression in Table 3 and use this to calculate the expected payoff to a bidder with value 30 from making a bid of b .

Figure 4 EXAMPLE: THE BEST RESPONSE TO THE EMPIRICAL BID FUNCTIONS



From this figure, we observe that the best response of a bidder in the full feedback treatment with a value of 30 is actually to bid 0, while for the same bidder in the partial feedback treatment, the best response is a bid of approximately 5.4. Thus if bidders were best-responding to the average behavior of subjects, we should actually see a *higher* frequency of dropouts in the full feedback treatment. However, as was shown in Table 2, we observe the exact opposite pattern. Thus, feedback

has an even stronger effect because more aggressive bidding (at relatively higher values) should go hand-in-hand with more frequent dropouts (for relatively lower values), but we actually observe fewer dropouts.

4.2. Learning and Convergence

We consider two measures of bid deviation from the equilibrium prediction: (1) *bias*: the difference between the actual bid and the risk neutral equilibrium bid, and (2) *absolute bias*: the absolute value of bias. An upward bias implies that subjects were, on average, overbidding, while changes in the absolute bias can be useful in determining whether or not there is convergence to the equilibrium.

Figure 5 AVERAGE ABSOLUTE DIFFERENCE AND BIAS BY PERIOD

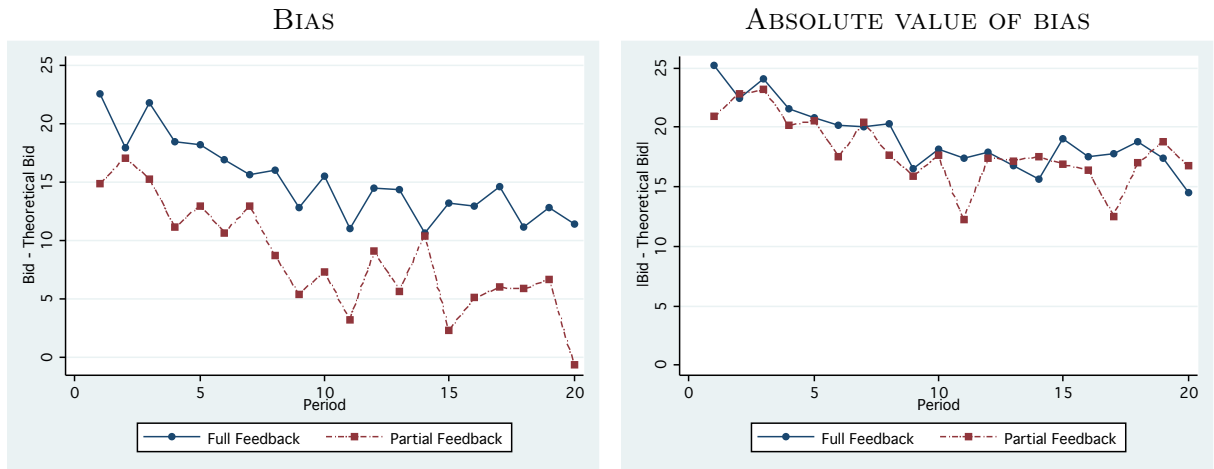


Figure 5 shows the average bias and average absolute bias for both treatments by period. As can be seen, by looking at the left-hand panel, there appears to be a persistent pattern of over-bidding, though the trend is downward in both treatments, it appears stronger in the partial feedback treatment. Notice also that the bias is always higher in the full feedback treatment than in the partial feedback treatment; this is consistent with our earlier findings, reported in Table 1, and indeed goes further by showing that the difference between treatments persists over all periods.

It is interesting to note that while the bias is higher in the full feedback treatment than in the partial feedback treatment, there is almost no difference between the two treatments in terms of absolute bias. This result confirms our earlier observation in Table 2 that there was more underbidding (e.g., dropouts) in the partial feedback treatment. Finally, note that there appears to be a downward trend, for both treatments, in the absolute bias. This indicates that some learning is taking place; however, it is difficult to say that subjects are converging to the equilibrium since the absolute bias is still quite high (and only minimally decreasing) after about period 15. An

analysis of the bias and absolute bias along the lines of Noussair and Silver (2006) suggests that the biases are persistent in both treatments and stronger in the full feedback treatment.⁸

Observe that regret can be both anticipated and result from experience. Therefore, along the lines of the learning direction theory of Selten and Stoecker (1986), we analyze how bids change given the outcome of the round and on the feedback available. Specifically, in Table 4 we report the results of random effects regressions where we regress Δb_{it} on $\Delta \text{value}_{i,t}$, `period` and feedback variables. The feedback variables are `lostt-1`, which takes value 1 if the bidder lost in the previous period and 0 otherwise; `lostt-1 × bidt-1`, which captures all-pay loser regret and is active in both treatments; and `lost afford.t-1 × (win bid - value)t-1`, which captures loser regret and is only active in the full-feedback treatment.

Table 4 RANDOM EFFECTS REGRESSION RESULTS: CHANGES IN THE BID FUNCTION WITH TIME AND FEEDBACK EFFECTS

	FF	PF
<code>cons</code>	-5.841*** [0.981]	-6.142*** [0.937]
Δvalue	0.744*** [0.029]	0.599*** [0.041]
<code>period</code>	0.028 [0.053]	-0.147** [0.064]
<code>lost_{t-1}</code>	10.37*** [1.751]	17.00*** [1.777]
<code>lost_{t-1} × bid_{t-1}</code>	-0.199*** [0.045]	-0.306*** [0.066]
<code>lost afford._{t-1} × (win bid - value)_{t-1}</code>	0.489*** [0.059]	
<code>N</code>	1178	1102
R^2	0.772	0.606

Clustered standard errors (subject level) in brackets.

Highlighted rows indicate a significant difference in the relevant coefficient between the full feedback and partial feedback treatments at the 5% level or better.

*** significant at 1%; ** significant at 5%; * significant at 10%.

As should be the case, the coefficient on Δvalue is significantly positive in both treatments. Furthermore, consistent with our previous results, the coefficient on `period` is negative and significant for the partial feedback treatment. That is, there was a general trend towards lower bids. Next, when it comes to the feedback variables, we see that the coefficient on `lostt-1` is positive and significant, while the coefficient on `lostt-1 × bidt-1` is negative and significant. The estimates imply that for losing bids smaller than about 50, the subsequent bid will increase, while for losing

⁸ These results are available upon request.

bids higher than about 50, the subsequent bid will decrease. This is consistent with the pattern we saw of substantial underbidding (via dropouts) for low values and overbidding for high values. It also suggests that all-pay loser regret is most active for high losing bids. Finally, in the full feedback treatment we observe that bidders who lost at an affordable price are likely to increase their bid in the next round, and the magnitude is increasing in the difference between the winning bid and the subject's value. Thus, loser regret appears to have a strong influence on behavior — in particular, leading to the more persistent overbidding that we have seen in the full feedback treatment.

4.3. Gender Effects

As we have noted in the introduction, work environments (in which workers must all exert costly effort, while only a few are rewarded with a bonus or promotion) share many similarities with the all-pay auction environment considered here. Previous research (see, e.g., Gneezy et al. (2003) and Croson and Gneezy (2009), among others) has shown that men out-perform women in competitive environments. Given this, a reasonable initial hypothesis is that women should earn less than men in our all-pay auction experiment.

It is also a fairly robust finding that women are more risk averse than men (Croson and Gneezy 2009). In first-price auctions, this means that one would expect women to bid more aggressively than men, which is precisely what Chen et al. (2009) has reported.⁹ Ham and Kagel (2006) found that women were more likely to go bankrupt and bid more aggressively than men in a two-stage private values auction. Casari et al. (2007) found that women started out bidding more aggressively than men in a common value auction, but that women also appeared to learn faster than men, noting that by the end of the auction, the differences between men and women largely disappear.

In all-pay auctions, Fibich et al. (2006) have shown that risk averse bidders with low valuations bid less than risk neutral bidders, while bidders with high valuations bid higher than risk neutral bidders. As noted above, our general results are consistent with this finding. With regards to gender differences, we would, therefore, expect women to drop out more frequently than men, and (especially for higher values) women to bid more aggressively than men. A higher frequency of dropouts would also be consistent with results summarized in the survey by Croson and Gneezy (2009); namely, that women are less likely to enter into competitive games like tournaments.

Before presenting results on gender difference, first note that in the full feedback treatment, 27 out of 62 subjects were female, while in the partial feedback treatment, 21 out of 58 subjects were female. Turn now to Table 5, which contains results on the frequency with which men and women

⁹ The also reported that there is no difference in bidding behavior of men and women in second-price auctions, where it is a dominant strategy to bid one's value independent of one's risk preferences.

submit zero bids, the difference between the actual and theoretical bid, the difference between the actual and theoretical bid (conditional on placing a strictly positive bid) and also the average profits of men and women. Panel (a) contains the results for the full feedback treatment, while panel (b) contains the results for the partial feedback treatment, and panel (c) provides the p -values of ranksum tests for gender-specific treatment effects on each of our four variables of interest. All hypothesis tests are based on session averages (of which we have four for each treatment), after separating males and females.

Table 5 GENDER EFFECTS

(a) FULL FEEDBACK TREATMENT						
	Males		Females		Ranksum test	p -value
	Mean	Std. Dev.	Mean	Std. Dev.		
Actual – Theo. Bid	9.11	2.80	23.33	2.94	2.31	0.021
Frequency $b = 0$	0.14	0.02	0.10	0.03	1.73	0.083
Actual – Theo. Bid ($b > 0$)	11.42	2.89	26.46	3.96	2.31	0.021
Profit	2.27	2.92	-3.39	2.51	2.02	0.043

(b) PARTIAL FEEDBACK TREATMENT						
	Males		Females		Ranksum test	p -value
	Mean	Std. Dev.	Mean	Std. Dev.		
Actual – Theo. Bid	4.18	4.07	15.96	5.39	2.31	0.021
Frequency $b = 0$	0.26	0.09	0.23	0.15	0.00	1.000
Actual – Theo. Bid ($b > 0$)	8.67	3.95	22.84	2.26	2.31	0.021
Profit	6.19	2.44	2.28	3.09	1.73	0.083

(c) GENDER TREATMENT EFFECTS		
	Males	Females
	p -value	p -value
Actual – Theo. Bid	0.083	0.043
Frequency $b = 0$	0.043	0.248
Actual – Theo. Bid ($b > 0$)	0.564	0.149
Profit	0.083	0.043

As can be seen, women tend to drop out less frequently than men, although the difference is only statistically significant in the full feedback treatment. Furthermore, in both treatments, women tend to overbid significantly more than men — a result that holds true whether we consider all bids or just those bids that are strictly positive. Given these results, it is, therefore, no surprise that men earn significantly more than women. Observe that these results are only partially consistent with findings in the literature on gender differences. While we observe that women overbid more than men (which is consistent with women being more risk averse than men), we find that, if

anything, women drop out less frequently than men (which is *inconsistent* with women being more risk averse than men). Although we do not report it, there is also some evidence, which goes against the results of Casari et al. (2007), that women learn more slowly than men.

Finally, consider panel (c) which looks at treatment effects at the gender level. First consider men. It appears that men overbid significantly more in the full feedback treatment, but that this result is due to the fact that men are less likely to drop out under full feedback. The fact that overbidding increases, then leads to a significant decline in profits going from partial to full feedback. Turn now to women. While there is no significant treatment effect for either the frequency of drop outs or the extend of overbidding (conditional on placing a positive bid), these two effects, taken together, do lead to significantly more overbidding in the full feedback treatment and, therefore, to significantly lower profits.

5. Conclusion

Recent developments in the in the online community knowledge markets has increased the appeal of the all-pay auctions. In this paper, we investigated the effect of revealing information regarding the highest (winning) bid in all-pay auction. Consistent with the existing literature, our results showed that bidders who submit strictly positive bids tend to over-bid, while those with low valuations tend to submit zero bids. Moreover, the revenue of the seller is higher than the prediction of the standard theory with risk neutral bidders.

Our experimental results indicates the joint importance of the regret and risk aversion hypotheses. When the information regarding the winning bid is concealed, risk aversion is responsible for overbidding. Informing both winners and losers about the value of the winning bid activates loser regret on top of risk aversion. Hence, we observe even more overbidding and fewer dropouts.

Dropouts of the low value bidders may be due to the fact that they are more likely to lose. Consequently, they bid 0 in an attempt to avoid losing money. Hence, they put more weight on the all-pay loser regret. With revelation of the winning bid, they realize that they would regret if they drop out and lose at an affordable price. Our experimental results suggest that the all-pay loser regret carries less weight than standard loser regret, which makes subjects less likely to submit zero bids and, conditional on submitting a positive bid, makes them more likely to overbid (especially for intermediate valuations where there is more uncertainty about one's chances of winning).

Our experiment also uncovered some interesting gender effects that are partially at odds with the extant literature. It a robust finding that women are more risk averse than men. In our setting, this implies that women should be more likely to place a bid of zero and, conditional upon placing a

positive bid, they should bid more aggressively than men. While we confirm this latter prediction, we find that, if anything, women are actually less likely to drop out of the auction by placing a bid of zero.

One implication of our experiment is that information disclosure is a powerful tool that auctioneers and mechanism designers may wish to exploit. While it is very easy to disclose the winning bid in an all-pay auction, in other similar settings information disclosure may be more difficult to the extent that “winning” is determined by partially unobservable factors such as effort or ability. When such disclosures are difficult, it may still be possible to provide signals that are correlated with effort. For example, when announcing bonuses, a law/consulting firm, seeking to activate loser regret, may also wish to disclose the billable hours of employees. Therefore, it may be fruitful to conduct a deeper examination of the different kinds of possible disclosures, and their implications on behavior, in real-world settings other than all-pay auctions.

Acknowledgments

We thank the University of Maryland, Department of Economics for its generous financial support for this research. We have also benefited from helpful discussions with Philipp Reiss.

Appendix. Instructions

Welcome to the auction experiment. In this experiment, you will participate in auctions as a bidder. The precise rules and procedures that govern the operation of these auctions will be explained to you below. Various research foundations have provided funds for this research. The instructions are simple, and if you follow them carefully and make good decisions you can finish the experiment with a considerable amount of money, which will be paid to you in cash privately at the end of the experiment. The experiment will last about one hour.

The type of currency used in this experiment is Experimental Currency Units (ECU). At the end of the experiment all your earnings will be calculated and converted to US Dollars ($5 \text{ ECU} = \$1$). The more ECU you earn, the more US Dollars you earn. The experiment will be broken up into a series of 20 rounds in which you will be bidding in a series of auctions.

Auction Process. In every auction, there is a fictitious good that is sold. When the first auction starts, you will observe your own valuation of the fictitious good of that round. Your valuation is chosen randomly and is equally likely to be any number between 0 and 100, rounded to the nearest cent. You will receive a new randomly chosen value in each round, and your value will typically be different from the one of every other player. Other bidders' values are independent of your value, that is, each other bidder's value is also chosen randomly and is equally likely to be any number between 0 and 100, rounded to the nearest cent, regardless of what your number happens to be. Each bidder will know only his or her own valuation.

At the beginning of the first auction, you will be randomly matched with another participant in this room; and at beginning of every auction, you will be randomly re-matched with a different participant. In each auction you will be bidding against the participant with whom you are matched.

During each auction, you may submit a bid for the fictitious good. After you and your opponent have submitted bids, the bidder who submitted the highest bid will be awarded the fictitious item. If there is a tie between the highest bids, the item is randomly.

In this auction, you pay the amount of your bid regardless of whether or not you receive the fictitious item. Your earnings in a round can be described as follows:

For example, you have a value of 82.55 and bid 60. If your bid (60) is higher than your opponent's bid, then you win the fictitious good and your earning in that round is 22.55 ECU ($82.55 - 60 = 22.55$). If your opponent's bid is higher, then you do not win the fictitious good and your earning in that round is -60 ECU.

When the first auction is completed, the second auction will start. At the beginning of the second auction, you will be randomly matched with another participant in this room and play with that person in this auction. Again, the computer will show you your new value for the good for this auction. It is again chosen randomly and is equally likely to be any number between 0 and 100, rounded to the nearest cent. Your opponent will also observe his or her own value of the fictitious item for this auction privately. The same auction rules as in the first auction will apply. This process will continue for 20 rounds.

Final Earning. There are 20 auctions in total. The computer will randomly select your earning in one of 20 auctions and convert this amount to US Dollars by using the rate of 0.2 dollar per ECU. This amount will be added (or subtract if negative) to \$20 participation fee. We will pay you this amount in cash at the end of the experiment in person.

Information Structure:

For treatment P: At the end of each auction, you will only learn whether you win or not. You will not learn any additional information regarding the bid of your opponent.

For treatment F: At the end of each auction, you will learn whether you win or not. Also, you will learn the bid of the winner.

References

- Anderson, S.P., J.K. Goeree, C.A. Holt. 1998. Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy* **106**(4) 828–853.
- Baye, M., D. Kovenock, C. de Vries. 1993. Rigging the lobbying process: an application of the all pay auction. *American Economic Review* **83** 289–294.
- Boudreau, K.J., N. Lacetera, K.R. Lakhani. Forthcoming. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science* .
- Bull, Clive, Andrew Schotter, Keith Weigelt. 1987. Tournaments and piece rates: An experimental study. *Journal of Political Economy* **95** 1–33.

- Casari, M., J. C. Ham, J. Kagel. 2007. Selection bias, demographic effects and ability effects in common value auction experiments. *American Economic Review* **97**(4) 1278–1304.
- Cason, Timothy N., William A. Masters, Roman M. Sheremeta. 2010. Entry into winner-take-all and proportional-prize contests: An experimental study. *Journal of Public Economics* **94** 604–611.
- Che, Y-K, I.L. Gale. 1998. Caps on political lobbying. *American Economic Review* **88**(3) 643–651.
- Chen, Y., P. Katuščák, E. Ozdenoren. 2009. Why can't a woman bid more like a man? Working Paper.
- Croson, R., U. Gneezy. 2009. Gender differences in preferences. *Journal of Economic Literature* **47**(2) 1–27.
- DiPalantino, D., M. Vojnovic. 2009. Crowdsourcing and all-pay auctions. *Proceedings of the Tenth ACM Conference on Electronic Commerce*.
- Dufwenberg, M., U. Gneezy. 2002. Information disclosure in auctions: an experiment. *Journal of Economic Behavior & Organization* **48**(4) 431–444.
- Engelbrecht-Wiggans, R., E. Katok. 2008. Regret and feedback information in first-price sealed-bid auctions. *Management Science* **54**(4) 808–819.
- Engelbrecht-Wiggans, Richard, Elena Katok. 2007. Regret in auctions: Theory and evidence. *Economic Theory* **33** 81–101.
- Fibich, G., A. Gavious, A. Sela. 2006. All-pay auctions with risk-averse players. *International Journal of Game Theory* **34** 583–599.
- Filiz-Ozbay, E., E. Y. Ozbay. 2007. Auctions with anticipated regret: Theory and experiment. *American Economic Review* **97**(4) 1407–1418.
- Filiz-Ozbay, E., E. Y. Ozbay. 2010. Anticipated loser regret in third price auctions. *Economics Letters* **107**(2) 217–219.
- Fischbacher, U. 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10**(2) 171–178.
- Gneezy, U., M. Niederle, A. Rustichini. 2003. Performance in competitive environments: Gender differences. *Quarterly Journal of Economics* **118**(3) 1049–1074.
- Gneezy, U., R. Smorodinsky. 1999. All-pay auctions — an experimental study. Discussion Paper 99-01, University of Haifa.
- Ham, J. C., J. Kagel. 2006. Gender effects in private value auctions. *Economics Letters* **92** 375–382.
- Isaac, M., J. Walker. 1985. Information and conspiracy in sealed bid auctions. *Journal of Economic Behavior & Organization* **6**(2) 139–159.
- Kagel, J., D. Levin. 1993. Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *Economic Journal* **103** 868–879.
- Lazear, Edward P., Sherwin Rosen. 1981. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* **89**(5) 841–864.

-
- Lugovskyy, V., D. Puzzello, S. Tucker. Forthcoming. An experimental investigation of overdissipation in the all pay auction. *European Economic Review* .
- Moldovanu, Benny, Aner Sela. 2001. The optimal allocation of prizes in contests. *American Economic Review* **91** 542–558.
- Müller, W., A. Schotter. 2010. Workaholics and dropouts in organizations. *Journal of the European Economic Association* **8**(4) 717–743.
- Noussair, C., J. Silver. 2006. Behavior in all-pay auctions with incomplete information. *Games and Economic Behavior* **55**(1) 189–206.
- Ockenfels, A., R. Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior* **51**(1) 155–170.
- Orrison, Alannah, Andrew Schotter, Keith Weigelt. 2004. Multiperson tournaments: An experimental examination. *Management Science* **50**(2) 268–279.
- Schotter, Andrew, Keith Weigelt. 1992. Asymmetric tournaments, equal opportunity laws, and affirmative action: Some experimental results. *Quarterly Journal of Economics* **107** 511–539.
- Selten, R., R. Stoecker. 1986. End behavior in sequences of finite prisoner’s dilemma supergames: A learning theory approach. *Journal of Economic Behavior & Organization* **7** 47–70.
- Sheremeta, R. 2010a. Expenditures and information disclosure in two-stage political contests. *Journal of Conflict Resolution* **54**(5) 771–798.
- Sheremeta, R. 2010b. Experimental comparison of multistage and one-stage contests. *Games and Economic Behavior* **68** 731–747.
- Tullock, G. 1980. Efficient rent-seeking. J. M. Buchanan, R. D. Tollison, G. Tullock, eds., *Toward a Theory of the Rent-Seeking Society*. Texas A & M University Press.
- van Dijk, Frans, Joep Sonnemans, Franc van Winden. 2001. Incentive systems in a real effort experiment. *European Economic Review* **45** 187–214.
- Yang, J., L. Adamic, M. S. Ackerman. 2009. Crowdsourcing and knowledge sharing: strategic user behavior on taskcn. *Proceedings of the Ninth ACM Conference on Electronic Commerce*.