Decision Analysis
Vol. 6, No. 2, June 2009, pp. 75–86
issn 1545-8490 | eissn 1545-8504 | 09 | 0602 | 0075
doi 10.1287/deca.1090.0135

A Direct Test of Risk Aversion and Regret in First Price Sealed-Bid Auctions

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

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

Why do bidders tend to bid higher than the risk-neutral Nash equilibrium in sealed-bid first price auction experiments? The effect of risk aversion has long been offered as a possible explanation. More recently, several studies proposed regret as another explanation, citing strong experimental evidence. But which effect is more important? We design an experiment to separate the effects of risk aversion from those of regret. We find overwhelming evidence in support of the regret model, and virtually no support for the constant relative risk aversion model.

Key words: auctions; competitive bidding; regret; risk aversion; experimental economics
History: Received on September 25, 2008. Accepted on January 14, 2009, after 1 revision. Published online in *Articles in Advance* March 11, 2009.

1. Introduction

Risk aversion in general, and constant relative risk aversion (CRRA) in particular, has been offered as an explanation for the commonly observed "overbidding" relative to the risk-neutral Nash equilibrium (RNNE) in sealed-bid first price (SBFP) auction experiments (see, for example, Cox et al. 1988). If bidders are risk averse rather than risk neutral, the resulting models fit the actual SBFP auction data reasonably well. Risk aversion, however, has been reported to be inconsistent with observed behavior in a number of other auction and auction-like settings. Examples include the experiments of Kagel and Levin (1993), Cason (1995), Isaac and James (2000). In short, risk aversion may have very little effect in auction experiments even though the observed behavior in SBFP auctions is consistent with risk aversion—the risk aversion model may fit the SBFP auction data for the wrong reason.

Several recent studies report data for which risk aversion by itself does not provide a sufficient explanation. Ivanova-Stenzel and Salmon (2008) report that when bidders can select whether to enter an SBFP or an ascending auction, more bidders enter the ascending auction than risk aversion by itself implies. Ivanova-Stenzel and Salmon (2004) also find that subjects are willing to pay more than risk aversion would imply for the right to participate in an ascending auction than in an SBFP auction. Ivanova-Stenzel and Salmon (2008, p. 17) note that both sets of results imply that "people have a nonpecuniary preference for the ascending auction" and proceed to argue that this preference can be explained by "minimizing the possibility of regret."

The basic idea of regret in auctions, first proposed by Engelbrecht-Wiggans (1989), which builds on earlier work on regret by Bell (1982) and Bell (1983) as well as related work on disappointment (Bell 1985), is that bidders are motivated not only by expected profits, but also by two types of regret. Winner regret occurs when the winner pays more than the second highest bid, and loser regret occurs when a loser misses a profitable trade opportunity by bidding too low. Engelbrecht-Wiggans (1989) argues that the equilibrium bids will be above or below the RNNE bids depending on whether the bidders are more sensitive to loser regret or to winner regret. In Engelbrecht-Wiggans and Katok (2007), we extend the model to allow the decision maker to have a more general utility function; for example, social preferences may matter in addition to risk aversion and regret. Again, optimal bids increase as the sensitivity to loser regret increases, and decrease as the sensitivity to winner regret increases. We also show that regret can explain not only the observed overbidding in SBFP auctions, but also the behavior observed by, for example, Kagel and Levin (1993), Cason (1995), and Isaac and James (2000)—behavior that could not be explained by risk aversion. Given how well regret organizes a variety of data, it should be considered as a promising explanation for bidding behavior.

The effect of regret on bidding in SBFP auctions has been examined in the laboratory by several different researchers. These experiments manipulate feedback information, making one or the other type of regret more or less salient. Regret theory predicts that average bids should shift down when winner regret is emphasized or up when loser regret is emphasized. In Engelbrecht-Wiggans and Katok (2008) we report on an experiment in which human bidders compete against computerized opponents in a large number of auctions.[1] We find that manipulating the saliency of the winner and the loser regret has the effect predicted by the theory, and this effect persists over time. Filiz-Ozbay and Ozbay (2007) focus on bidders' reaction to anticipated regret and find an effect for loser regret, but not for winner regret. Their experiment does not control for potential effects that may result from interpersonal interactions, such as of collusion or inequality aversion (for example, those of Isaac and Walker 1985, Dufwenberg and Gneezy 2002, Ockenfels and Selten 2005, and Neugebauer and Selten 2006) or address how reaction to regret evolves over time.

At this point, we have presented two possible explanations for observed behavior. Risk aversion fails to explain certain behavior that can be explained by regret. This, however, does not rule out the possibility that risk aversion is still a significant part of the explanation; it does not answer the question of how important is risk aversion relative to regret in explaining observed bidding behavior in SBFP auctions.[2] This paper directly addresses that question.

Our experimental design has two novel aspects. First, bidders bid in a large number of auctions—each bidder makes 100 bidding decisions. This allows us to determine that the effects of regret persist—even increase—with experience despite the fact that reacting to regret hurts the bidder's expected profit; this argues against the possibility that bidders actually want to maximize expected profit and simply don't realize—as they may well not in a one-shot setting—that reacting to regret adversely affects expected profit. Second, and most importantly, each of our human bidders bids against computerized rivals rather than other human bidders, so they are facing a decision analysis problem. This decision analytic focus plays three critical roles in our design. (1) It controls for interpersonal interactions. Other explanations that have been offered for overbidding that rely on interpersonal interactions, including collusion (Isaac and Walker 1985), spite (Morgan et al. 2003), and inequality aversion (Ockenfels and Selten 2005). By using computerized rivals, we can rule these out as possible explanations for our experimental results. (2) It allows us to vary the amount (rather then the type) of feedback by showing participants how their bid would fare in several identical independent auctions. (3) It allows us to vary the payment rule. Specifically, subjects can be paid based on how their particular bid fared in one auction, or how it fared, on average, in many auctions. This allows us to directly manipulate the amount of risk of the auction.[3]

By varying the amount of feedback and the payment rule in addition to varying the type of feedback, we can cleanly and directly separate the effects of

[1] The present manuscript reports on a larger data set, and the data that we analyzed earlier in Engelbrecht-Wiggans and Katok (2008) is a part of the data set on which we report presently. There are, however, two additional information conditions that are new to the present manuscript and have not been previously analyzed anywhere.

[2] We are not questioning the importance of risk aversion as a concept, as there is plenty of evidence that in many settings risk aversion organizes behavioral data well. We are, however, questioning whether risk aversion is needed to explain SBFP bidding behavior.

[3] See Katok and Salmon (2008) for a laboratory experiment that measures a similar effect of variance reduction on decisions in a simpler setting of selecting between pairs of lotteries, and on measuring certainty equivalence.

regret versus risk aversion. If risk aversion matters, then bidders should overbid less when their payment is based on the average outcome of many auctions rather than just one auction (we show this formally in the next section). However, the amount (and even type) of feedback provided should have no effect on a risk-averse bidder. In contrast, changing the type of feedback has already been shown to affect regret. We discover that changing the type of feedback has a significant effect, but varying the payment rule does not, which suggests that regret plays a much larger role in bidders' bidding decisions than does risk aversion.

2. The Theory for the Direct Test of the CRRA Model

Imagine a human subject bidding against computerized opponents in a sealed-bid first price auction. The human subject has a privately known value v for the good that is being auctioned and must decide on a bid b. This bid will be applied to k independent replications of the auction—v remains fixed but the computerized bids change. The subject is paid the average of his winnings in the k auctions.

Intuitively, as k increases, the variance in the human subject's payoff decreases and a risk-averse bidder's best bid decreases. More precisely, a risk-averse bidder should bid higher than a risk-neutral bidder when $k=1$. As k increases, the payoff resulting from any specific bid converges (in distribution) to the expected payoff from that bid. In the limit as k goes to infinity, there is no uncertainty, and therefore, any bidder whose utility function increases as payoff increases should bid the same as a risk-neutral bidder. So, at least in the limit of infinite k, one might expect that a risk-averse bidder's best bid decreases as k increases.

Arguing that best bids decrease as k goes from 1 to some finite $k>1$ is more difficult. Indeed, best bids do not always decrease as k increases.[4] However, we now show that best bids indeed do increase as k increases for the most commonly assumed forms of risk-averse utility functions. Specifically, let $u(x)$ denote my utility for \$$x$. Then my expected utility from bidding b in k independent auctions when my value is v is

$$U(b;v,k) = \sum_{i=0}^{k} u\left(\frac{(v-b)i}{k}\right)\frac{k!}{i!(k-i)!}F(b)^i(1-F(b))^{k-1},$$

where $F(b)$ denotes the probability that all the computerized bids in an auction are at most b.

A necessary condition for a best bid is that $dU/db=0$. So, we examine dU/db for $k>1$ and evaluate it at the bid that is optimal when $k=1$. Assume that u is a strictly increasing, differentiable function of x. Also, without loss of generality, assume $u(x)=0$. Then (see Appendix A.1 for an outline of the derivation) $dU(b;v,k)/db$ may be written as

$$\sum_{i=0}^{k-1}\left\{\left[\left(\frac{u(x(i+1)/k)-u(xi/k)}{u(x)}\right)k\right]u(x)F'(b) - \left(\frac{u'(x(i+1)/k)}{u'(x)}\right)u'(x)F(b)\right\} \times \frac{(k-1)!}{i!((k-1)-i)!}F(b)^i(1-F(b))^{(k-1)-i},$$

where $x\equiv v-b$. Let b^* denote the best bid for $k=1$. Then, $u(x^*)F'(b^*)-u(x^*)F'(b^*)=0$, where $x^*\equiv v-b^*$, and dominance arguments imply $x^*\geq 0$. So, for $b=b^*$ and $x=x^*$, the ith term of the above sum will have the same sign as

$$\Delta(x^*;i,k)\equiv\left[u\left(\frac{x^*(i+1)}{k}\right)-u\left(\frac{x^*i}{k}\right)\right]k\Big/u(x^*) - u'\left(\frac{x^*(i+1)}{k}\right)\Big/u'(x^*).$$

Note that if this expression is nonpositive for all i, and is negative for at least one i, then $dU(b;v,k)/db|_{b=b^*}<0$.

Now consider the specific case of CRRA. In this case, $u(x)=x^r$ for $x\geq 0$, with $0<r\leq 1$. It is straightforward to verify (see Appendix A.2) that $\Delta(x^*;i,k)=0$ for all $x>0$ when $i=0$, and $\Delta(x^*;i,k)<0$ when

[4] For example, let $u(x)=x$ for $0\leq x\leq 0.1$, and $u(x)=0.1+0.85(x-0.1)$ for $0.1\leq x$; this is a piecewise linear function with a slope of 1.0 for small x and a slope of 0.85 for larger x. Imagine that my value is $v=0.64$ and I'm bidding against two opponents whose bids are independent Uniform(0, 2/3). Then, it is straightforward to calculate that my best bid is 0.4384 for $k=1$, but is 0.44 (putting $(v-b^*)/2$ right at the kink) for $k=2$; as k increases, so too does my best bid (for this specific v, and for other v's that are large enough). The RNNE bid is $(2/3)(0.64)\approx 0.42667$.

$i > 0$. Therefore, when $k > 1$, $dU(b; v, k)/db|_{b=b^*} < 0$. In words, a CRRA bidder's expected utility decreases as the bidder increases his bid past the amount that was best when $k = 1$; roughly speaking, a CRRA bidder should bid lower when $k > 1$ than when $k = 1$.[5] Furthermore, once we have an estimate of the parameter r, we can predict by how much bids should decrease due to CRRA as k increases.

3. Experimental Design and Research Hypotheses

In our experiments we manipulate two factors: the first we will call the *payment* and the second we will call the *feedback*. We manipulate the payment in three different ways and the feedback in four different ways for the total of 12 treatments. Each treatment included 20 participants, with 240 participants in all. The factors we manipulate are as follows.

Each bidding decision affects k independent auctions, and after each decision participants are given feedback for the k auctions and are paid based on either all k auctions or just one of the k auctions. We have three payment conditions:

(1) In the $k = 1$ condition, each bidding decision affects a single auction; participants are paid for this auction and observe feedback only from this auction.

(2) In the $k = 10$ condition, each bidding decision affects 10 independent Auctions; participants are paid for the average of the 10 auctions and observe feedback for all 10 auctions.[6]

(3) Because participants in the $k = 10$ condition see 10 times as many auction outcomes as participants in the $k = 1$ condition, we also conducted a set of $k = 1^*$ treatments in which upon receiving feedback from the one auction that affects the payment, participants see a screen with the outcome of nine additional random auctions that do not have any effect on the payment, but do provide information about the distribution of the computerized opponents' bids.

We manipulate the saliency of the two types of regret by varying the feedback provided after each decision. We either show the winner the second highest bid or not (*winner regret*) *and either show the losers the winning bid or not* (*loser regret*). *This manipulation results in four feedback conditions*:

(1) *Winner Regret*: winners know the second highest bid (as well as the winning bid, because the winning bid is their own). Losers receive no feedback information except that their bid did not win.

(2) *Loser Regret*: winners do not know the second highest bid but losers know the winning bid.

(3) *Both Regrets*: winners know the second highest bid and losers know the winning bid.

(4) *No Regret*: winners do not know the second highest bid and losers do not know the winning bid.

Each human bidder bids against two computerized opponents and has the same value v for 20 consecutive decisions. Subjects cycle through five different values (50, 60, 70, 80, and 90) for a total of 100 bidding decisions (affecting a total of $100k$ auctions) per session.[7] We tell the subjects in the written instructions that the computerized bidders' values are integers uniformly distributed from 1 to 100, and that the computerized bidders are programmed to bid so as to maximize their expected profits under the assumption that all of their opponents follow the identical strategy (see the example instructions in Appendix B). Consistent with this explanation, we programmed the computerized bidders to bid 2/3 of their values (rounded to two decimals), which is the RNNE strategy for three computerized bidders bidding against each other.[8]

[5] Similarly, Appendix A.3 shows that if $k > 1$ and $u(x) = 1 - e^{-ax}$ with $a > 0$, then $dU(b; v, k)/db|_{b=b^*} < 0$. Therefore, a constant absolute risk-averse (CARA) bidder's expected utility increases as the bidder increases his bid past the amount that was best when $k = 1$; roughly speaking, a CARA bidder should bid lower when $k > 1$ than when $k = 1$. A similar result holds for the quadratic risk-averse utility function $u(x) = ax - x^2$ with $a > 2$.

[6] The data in this condition are reported in Engelbrecht-Wiggans and Katok (2008). The data in the other two conditions are new to the present manuscript.

[7] All subjects cycled through the values in the same increasing order, but different subjects started at different points in the cycle.

[8] To keep the information consistent with standard laboratory auction experiments, we tell our subjects the distribution of the opponents' values but not their bids. To avoid leading the subjects in their bidding decisions, we do not tell the subjects that the computerized bidders' strategy is to bid 2/3 of their values. To help subjects better infer the actual distribution of the computerized bidders' bids, we compute and display to them the actual probability their bid wins, which they see prior to finalizing their bidding

Each bidder participated in only a single treatment. Each session lasted for approximately 45 minutes, and the average earnings (including a \$5 participation fee) were \$16. All sessions were conducted at the Laboratory for Economic Management and Auctions at Penn State Smeal College of Business. 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.

The regret hypothesis we test is as follows (see Table 2 in §5 for the formal statements of the four specific predictions of the regret theory):

Hypothesis 1 (Regret). *For* all *payment conditions, providing information about the winning bid to losers increases bids, and providing information about the second highest bid to winners decreases bids.*

Our design also provides a direct test of risk aversion as the explanation for overbidding in SBFP auctions because risk aversion implies a shift (lower bids in $k = 10$ condition than in $k = 1^*$ condition), and this shift should be independent of feedback information. The risk aversion hypothesis we test is as follows (formal statements in Table 2 in §5).

Hypothesis 2 (Constant Relative Risk Aversion). *For* all *feedback information conditions, average bids will be lower in the $k = 10$ treatments than in the $k = 1^*$ treatments.*

The third hypothesis has to do with the effect of the amount of information. The critical point to note is that the four regret feedback information conditions provide (by design) different amounts of information to winners and losers. To see this, note that in the Winner Regret condition participants learn new information when they win (the second highest bid) but not when they lose. When a single decision results in feedback for 10 auctions, the probability that any given bid provides winner information for at least one auction is much higher than when a single decision results in feedback for a single auction. Therefore, it is possible that at least in some of the feedback conditions in the $k = 1$ payment treatments subjects may bid higher to improve their chances of winning and to acquire additional information only available to winners (for example, the amount of the winning bid). We test for this as follows:

Hypothesis 3 (Information). *For* all *feedback information conditions, average bids will be lower in the $k = 1^*$ treatments than in the $k = 1$ treatments.*

4. CRRA Predictions

Another test for the effect of risk aversion is to compare the actual bidding behavior in the $k = 10$ condition with that implied by the CRRA model. To make this test requires a prediction of what the $k = 10$ bids would be if bidders were affected by CRRA. Calculating this prediction for $k = 10$ involves two steps: (1) using the $k = 1^*$ data to estimate the distribution of the CRRA parameter r in the subject pool under the assumption that subjects are maximizing CRRA utility, and then (2) using this distribution to simulate the utility-maximizing bids for a players with those CRRA utility functions. We elaborate on each in turn.

In our experiment, the subject competes against $n - 1$ computerized bidders who bid uniformly on $[0, 66.67]$. The best response for the CRRA bidder who has a value v and a CRRA parameter of r is to bid

$$b = \min\left\{v\frac{n-1}{n-1+r}, 66.67\right\}, \tag{1}$$

where b is the optimal bid, v is value, and r is the true CRRA parameter. This is similar to other researchers' models in that it predicts bids being proportional to value for all small enough bids.

The standard method for estimating rs is first to fit a regression model for each subject:

$$b_t = \alpha + \beta v_t + \text{error}_t, \tag{2}$$

where b_t and v_t are the observed bid and the value in auction t. Because the prediction that bids are proportion to value applies only to small enough bids, it is common practice (Kagel 1995) to use only data for

decisions. If human subjects do not infer the probability of winning correctly on their own (see Dorsey and Razzolini 2003), providing this information helps stimulate learning as well as allow them to infer the true distribution of the computerized bids. Having this information in all 12 treatments ensures that it cannot account for any of the treatment effects.

which the observed bids fall below the highest possible bid of a risk-neutral bidder (in our case, 66.67) and exclude data from subjects whose α estimate is significantly different from 0. The resulting β is an estimate of $(n-1)/(n-1+r)$. Setting this estimate equal to $(n-1)/(n-1+r)$ and solving for r provides the estimate of r.

We departed from the above-described procedure in how we estimated $(n-1)/(n-1+r)$. Instead of using the β obtained from the linear regression, to we used the average $\overline{b/v} \equiv (\sum_t b_t/v_t)/T$ for each subject (where T is the number of decisions by the subject).

We believe that this method provides a better estimate than the regression. Specifically, we question the assumption, implicit in the linear regression, that the errors are i.i.d. normal with mean zero and variance independent of v; without this assumption, there is no assurance that β will be an unbiased estimate of $(n-1)/(n-1+r)$. For $\overline{b/v}$ to be an unbiased estimate of $(n-1)/(n-1+r)$, we only need to assume that errors have the mean of zero. We actually used all 100 samples for each subject in computing the estimates, and we verified that censoring samples that are at or above 66 makes virtually no difference to any of our conclusions.[9]

Given an estimate for the parameter r of a subject's CRRA, we then compute the utility-maximizing bid for a player with that CRRA utility function.[10] Specifically, for each value v of 50, 60, 70, 80, and 90, we (numerically) calculate the expected utility for the $k=10$ condition for each possible bid from 1 to 100 and select the bid that would maximize the expected utility. Then, we calculate the average b/v ratio for each r (i.e., for each treatment). Finally, the figures reported in the last row of Table 1 are averages of the (CRRA estimated best bid)/(value) for 20 subjects.

[9] Using regression model (2) and β to estimate $(n-1)/(n-1+r)$ actually noticeably lowers estimated optimal bids, further strengthening our conclusions; but as we explained above, we do not think this is a good method.

[10] We do not attempt to estimate the standard error of the risk aversion parameter, and in that our approach is simplistic. Although there may be more sophisticated methods for estimating distributions of the risk aversion parameter, most are not without their own problems and they are beyond the scope of our paper. We thank an anonymous referee for pointing this out and acknowledge that this as a limitation.

Table 1 Averages and Standard Deviations (in Parentheses) of b/v for the 12 Treatments, and the CRRA Prediction for the $k=10$ Condition Based on the $k=1^*$ Data

	Feedback			
Payment	Both Regrets (both)	Loser Regret (LR)	No Regret (none)	Winner Regret (WR)
$k=1$	0.7039 (0.0385)	0.7603 (0.0813)	0.764 (0.0925)	0.7375 (0.0122)
$k=1^*$	0.7183 (0.0634)	0.7578 (0.0604)	0.7268 (0.0331)	0.6751 (0.0601)
$k=10$	0.7261 (0.0529)	0.7660 (0.048)	0.7154 (0.0687)	0.6973 (0.0652)
$k=10$ CRRA prediction	0.6988 (0.0368)	0.7228 (0.0386)	0.7021 (0.0197)	0.6721 (0.0161)

5. Results

On average, subjects earned \$10.76 in the $k=1$ condition, \$11.06 in the $k=1^*$ conditions, and \$11.41 in the $k=10$ condition. The expected profit-maximizing strategy for a subject bidding against our computerized bidders is for the subject to bid 2/3 of her value. Given the value draws in our experiments, the average expected profits of the human bidders from always following this strategy were \$13.67 in the $k=1$ and $k=1^*$ conditions and \$12.83 in the $k=10$ conditions (we kept the value draws constant between treatments in the same payment condition and in the $k=1$ and $k=1^*$ conditions). Thus, the subjects earned 27% less in the $k=1$ and $k=1^*$ conditions and 12% less in the $k=10$ condition then they could have done had they played the best reply to uniformly distributed bids.

Table 1 shows the averages and standard deviations (in parentheses) of bid/value (b/v) for the 12 treatments, as well as the CRRA predictions discussed in §4.

Table 2 summarizes statistical tests for the three hypotheses. The p-values are from the Wilcoxon test, and p-values below 0.05 are in bold.

5.1. Regret Hypotheses

We generally find support for both regret hypotheses. One part of the regret hypothesis is that bids in the Loser Regret feedback condition are higher than in the No Regret feedback condition. We find support for this hypothesis in the $k=1^*$ and $k=10$ payment conditions, but not in the $k=1$ payment condition.

Table 2 Hypotheses Testing

	Hypotheses statement	$k=1$	$k=1^*$	$k=10$
Winner Regret	H_o: Both $\geq$ LR H_a: Both $<$ LR	**0.004**	**0.021**	**0.000**
	H_o: WR $\geq$ None H_a: WR $<$ None	**0.027**	**0.003**	0.212
Loser Regret	H_o: Both $\leq$ WR H_a: Both $>$ WR	0.999	**0.013**	**0.038**
	H_o: LR $\leq$ None H_a: LR $>$ None	0.656	**0.021**	**0.011**

	Hypotheses statement	Both Regrets	Loser Regret	No Regret	Winner Regret
Risk aversion: Shift	H_o: $(k=1^*) \leq (k=10)$ H_a: $(k=1^*) > (k=10)$	0.831	0.841	0.127	0.788
Risk aversion: Prediction	H_o: $(k=10) =$ CRRA H_a: $(k=10) \neq$ CRRA	**0.033**	**0.002**	0.204	**0.032**
Information	H_o: $(k=1^*) \geq (k=1)$ H_a: $(k=1^*) < (k=1)$	0.940	0.399	**0.014**	**0.001**

Notes. The p-values are from the Wilcoxon test. p-values below 0.05 are in bold. H_o, null hypothesis; H_a, alternative hypothesis; LR, Loser Regret; WR, Winner Regret.

We also find support (in the $k=1^*$ and $k=10$ conditions) for our second hypothesis dealing with loser regret: bids in the Both Regrets feedback condition are higher than in the Winner Regret feedback condition in those two payment conditions, but not in the $k=1$ payment condition.

One part of the winner regret hypothesis is that bids in the Winner Regret feedback condition are lower than in the No Regret feedback condition. Based on data aggregated over all 100 decisions, we do find support for this hypothesis in our $k=1$ and $k=1^*$ payment conditions, but not in our $k=10$ payment condition (although the direction of the shift is consistent with winner regret even there, and, as subsequent analysis will show, the effect becomes significant toward the end of the session). Additionally, we find support in all three payment conditions for our second winner regret hypothesis, that bids in the Both Regrets feedback condition are lower than in the Loser Regret feedback condition.[11]

[11] The regret model also makes a prediction about a direct comparison between the Loser Regret and the Winner Regret treatments: H_o: LR $\leq$ WR; H_a: LR $>$ WR. This comparison is also consistent with the regret model in all three payment conditions: Wilcoxon test p-values are 0.0739 for $k=1$, 0.0001 for $k=1^*$, and 0.0006 for $k=10$.

Does regret persist with experience? Our experimental design gives us the ability to answer this question. For each of the four regret hypotheses in Table 2, regret implies that the b/v ratio in one feedback condition is strictly smaller than in another feedback condition. To analyze the regret effect over time, we first computed the average of the b/v ratios across all subjects in each decision period, separately for each treatment. We then calculated the difference between the larger and the smaller b/v ratio, as implied by the corresponding alternative hypothesis in Table 2, for each of the four hypotheses in each of the 100 decision periods. And finally, for each of the four hypotheses in Table 2, we estimated the following linear regression model:

$$\left(\overline{\frac{b^{\text{larger}}}{v}}\right)_t - \left(\overline{\frac{b^{\text{smaller}}}{v}}\right)_t = \beta_0 + \beta^t \times t + \varepsilon_t. \qquad (3)$$

We are interested in the estimated β_0s because these estimates tell us the magnitude of the regret effect at the beginning of the session (a measure of anticipated regret, if you will). We can also estimate the magnitude of the regret effect at the end of the session by fitting $(\overline{b^{\text{larger}}/v})_t - (\overline{b^{\text{smaller}}/v})_t = \beta_{100} + \beta^{(100-t)} \times (100 - t) + \varepsilon_t$. Of course $\beta^{(100-t)} = -\beta^t$, but β_{100}s provide estimates of the regret effect at the end of the 100-decision session. In Figure 1, we plot $(\overline{b^{\text{larger}}/v})_t - (\overline{b^{\text{smaller}}/v})_t$ (on the y-axis) over time (on the x-axis) for all four regret hypotheses stated in Table 2, and also report corresponding estimates of β_0, β_{100}, and β^t.

We conclude that generally the magnitude of the regret effect persists with experience. In all nine instances (highlighted in bold in the top panel of Table 2) in which the data are consistent with the regret model, on aggregate, it continues to be consistent, both at the beginning of the sessions and at the end of the sessions. Additionally, the one Winner Regret (H_o: WR $\geq$ None) comparison in the $k=10$ payment condition, that was in the direction consistent with the regret hypothesis but not statistically significant using aggregate data, becomes consistent with the regret model in later periods. The two Loser Regret comparisons that were not consistent with the regret model based on aggregate data in the $k=1$ condition are not consistent with it both at the beginning and at the end of the session (these two treatments will be the subject of the discussion in §5.3, about the

Figure 1 Magnitude of the Regret Effect over Time

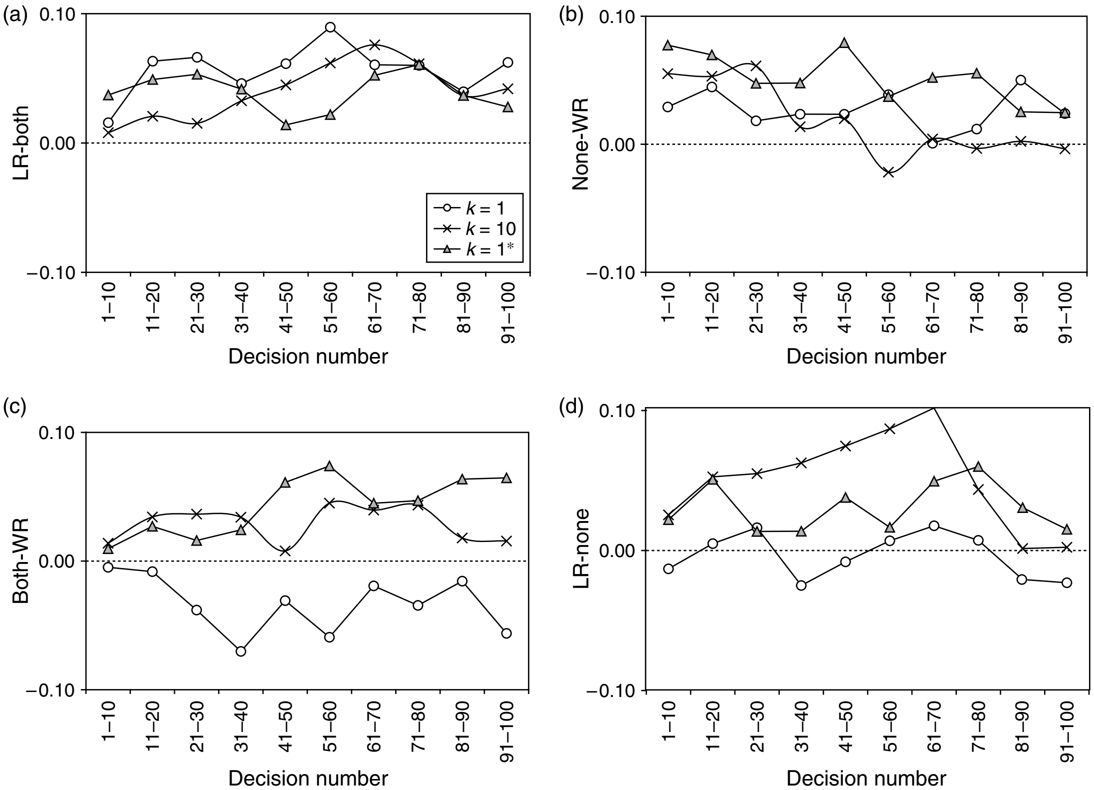

Notes. Panel (a) $k=1$: $\beta_0 = 0.047^{**}$, $\beta_{100} = 0.066^{**}$, $\beta^t = 0.0002^{*}$; $k=10$: $\beta_0 = 0.014^{**}$, $\beta_{100} = 0.065^{*}$, $\beta^t = 0.0005^{**}$; $k=1^*$: $\beta_0 = 0.042^{*}$, $\beta_{100} = 0.037^{*}$, $\beta^t = -0.0002^{**}$. Panel (b) $k=1$: $\beta_0 = 0.023^{**}$, $\beta_{100} = 0.030^{**}$, $\beta^t = 0.0001$; $k=10$: $\beta_0 = -0.020^{**}$, $\beta_{100} = 0.057^{**}$, $\beta^t = 0.0008^{**}$; $k=1^*$: $\beta_0 = 0.027^{**}$, $\beta_{100} = 0.076^{**}$, $\beta^t = 0.0005^{**}$. Panel (c) $k=1$: $\beta_0 = -0.021^{**}$, $\beta_{100} = -0.046^{**}$, $\beta^t = -0.0002^{*}$; $k=10$: $\beta_0 = 0.030^{**}$, $\beta_{100} = 0.028^{*}$, $\beta^t = -0.0001^{**}$; $k=1^*$: $\beta_0 = 0.012^{**}$, $\beta_{100} = 0.075^{**}$, $\beta^t = 0.0006^{**}$. Panel (d) $k=1$: $\beta_0 = 0.003$, $\beta_{100} = -0.010^{*}$, $\beta^t = -0.0001$; $k=10$: $\beta_0 = 0.064^{**}$, $\beta_{100} = 0.037^{**}$, $\beta^t = -0.0003^{**}$; $k=1^*$: $\beta_0 = 0.027^{**}$, $\beta_{100} = 0.035^{**}$, $\beta^t = 0.0001^{**}$.

role of information). The estimates of the slope coefficients β^t are positive in seven of the 12 comparisons, negative in four, and not significantly different from 0 in two. Such distribution is no different from random (sign test p-value $= 0.54$). Thus, there is no evidence that the regret effect disappears with experience.

5.2. Risk Aversion

We find virtually no support for CRRA. In all four feedback conditions we cannot reject the null hypothesis that average bids in the $k = 1^*$ payment condition are equal to the average bids in the $k = 10$ payment condition. In three of the four feedback conditions, average bids when $k = 10$ are higher than when $k = 1^*$ (not statistically significantly so), which is the opposite of the risk aversion prediction. Additionally, we compare the $k = 1^*$ and $k = 10$ bids using an aggregate matched-pair t-test treating the average b/v in each of the four information conditions as a single observation, and the p-value of this test is 0.4032. We can also reject, in three of the four feedback conditions, the null hypothesis that the average $k = 10$ bids are equal to the CRRA prediction.[12] Again, the only feedback condition in which the data are consistent with risk aversion is the No Regret condition. Additionally, we compare the $k = 10$ data and CRRA prediction for $k = 10$ based on $k = 1^*$ data using an aggregate matched-pair t-test treating the average b/v in each of the four

[12] The standard errors in our CRRA prediction assume that the risk aversion coefficient r is estimated without an error, thus the CRRA standard errors in Table 2 may be lower than they should be and the actual p-values may be higher than reported in Table 2.

information conditions as a single observation, and the p-value of this test is 0.0213.

5.3. Information

Our findings about the information hypothesis are more nuanced. We find support for the information hypothesis in feedback conditions in which the losers do not see the winning bid (No Regret and Winner Regret) but not in the feedback conditions in which losers do see the winning bid (Both Regrets and Loser Regret). This result indicates that the feedback conditions in which the selling price is not revealed to losers are especially challenging for subjects to contend with in our settings; one critical piece of information subjects need to learn is the distribution of the bids of the computerized rivals—what is the highest the computerized rivals ever bid? No information about the computerized rivals' bids is ever reported directly in the No Regret condition; winners can infer that their own bid is an upper bound on the rivals' bids, but losers can't even do that. In the Winner Regret condition, winners are told the highest rival bid, but losers are not. In both conditions, the only way to collect information about how computerized rivals bid is by observing winning auctions; losing auctions provide no information at all. In the $k = 1^*$ treatments, each decision results in being able to observe 10 auctions, and each decision is 10 times as likely to produce at least one win than is the $k = 1$ treatment. Therefore, it is possible that in the No Regret and Winner Regret conditions, bidders bid higher in the $k = 1$ treatments to win auctions, which is the only way to obtain some information about how high the computerized bidders bid, and the decrease in average bids in the $k = 1^*$ treatments is due to the fact that, in that setting, lower bids are more likely to generate at least some winning auctions. Consequently, the lower average bids in the $k = 1^*$ condition than in the $k = 1$ condition in the No Regret and Winner Regret treatments may be due to this desire to learn about the computerized rivals.

The above explanation does account for our failure to find evidence for the loser regret hypotheses in the $k = 1$ payment condition. The alternative hypotheses state that the average Winner Regret bids (the first) and the average No Regret bids (the second) should be lower than average bids in corresponding treatments with the winning price revealed. If bidders bid high in the Winner Regret and No Regret treatments to win auctions and learn about the distribution of the rivals' bids, that would explain the lack of a shift predicted by the regret model. The regret effect pushes bids down, but the desire to acquire additional information from winning auctions pushes the bids back up. The fact that average bids in the $k = 1^*$ treatments are lower is consistent with the above logic.

6. Summary

We investigate causes for bidding above the risk-neutral Nash equilibrium in sealed-bid first price auctions with independently known private values, and present experimental evidence to evaluate two explanations: the long-standing risk aversion explanation and the regret explanation. To provide a direct test of the CRRA model, we consider applying a bidder's bid to k independent, stochastically identical auctions and then paying the bidder the average outcome. To keep the amount of feedback information constant for the $k = 1$ and $k = 10$ payment conditions, we show bidders in the $k = 1$ payment condition outcomes of 10 auctions, but pay them only for one. We call this condition $k = 1^*$. We show that the best bid of a CRRA bidder should decrease as k increases. Indeed, by using the bidding data from the $k = 1^*$ case, we can predict how much the bids should decrease. We find virtually no support for the risk aversion model. In most cases, bids in the $k = 10$ condition are slightly higher than bids in the $k = 1^*$ condition, and are significantly higher than the CRRA prediction.

We do, however, find overwhelming support for the regret model. Of the 12 predictions the regret model makes for our data, nine at the beginning of the session and 10 by the end of the session result in a statistically significant shift that is consistent with the model. Only two of the 12 comparisons are inconsistent with the regret model. Our findings about the information hypothesis provide an explanation for these deviations.

Acknowledgments

The authors gratefully acknowledge the support from the National Science Foundation and thank Axel Ockenfels and the Deutsche Forschungsgemeinschaft for financial support through the Leibniz-Program. The second author also thanks the Smeal College of Business for the support through the Summer Research Grants program.

Appendix A

A.1. Details in Deriving $dU(b; v, k)/db$:

$$\frac{dU(b;v,k)}{db} = \frac{d}{db}\sum_{i=0}^{k} u\left(\frac{(v-b)i}{k}\right)\frac{k!}{i!(k-i)!}F(b)^i(1-F(b))^{k-1}.$$

Differentiating inside the summation, defining $x \equiv v - b$, rearranging and dropping terms that are equal to zero, redefining the index so that each sum starts at $i = 0$, combining the three sums, and, finally, inserting a factor of $u(x)/u(x)$ in the first term and a factor of $u(x)/u(x)$ into the second term gives the expression presented in this paper.

A.2. Details for the Case of CRRA

Under CRRA, $u(x) = x^r$ where $0 < r \le 1$, and note that

$$u\left(\frac{xi}{k}\right) = \left(\frac{i}{k}\right)^r u(x) \quad \text{and} \quad u'\left(\frac{xi}{k}\right) = \left(\frac{k}{i}\right)\left(\frac{i}{k}\right)^r u'(x).$$

Therefore,

$$\begin{aligned}
&\Delta(x^*; i, k)\\
&\equiv \left[u\left(\frac{x^*(i+1)}{k}\right) - u\left(\frac{x^*i}{k}\right)\right]k \Big/ u(x^*) - u'\left(\frac{x^*(i+1)}{k}\right)\Big/ u'(x^*)\\
&= \left[\left(\frac{i+1}{k}\right)^r u(x^*) - \left(\frac{i}{k}\right)^r u(x^*)\right]k\Big/ u(x^*)\\
&\quad - \left(\frac{k}{i+1}\right)\left(\frac{i+1}{k}\right)^r u'(x^*)/u'(x^*)\\
&= \left(\frac{1}{k^{r-1}}\right)\left\{[(i+1)^r - i^r] - (i+1)^{r-1}\right\},
\end{aligned}$$

which is zero when $i = 0$, and negative for all $i > 0$. So, if $k > 1$, then $dU(b; v, k)/db|_{b=b^*} < 0$.

A.3. Details for the Case of CARA

In this case, without loss of generality, $u(x) = 1 - e^{-x}$ and $u'(x) = e^{-x}$. Therefore,

$$\begin{aligned}
&\Delta(x^*; i, k)\\
&\equiv \left[u\left(\frac{x^*(i+1)}{k}\right) - u\left(\frac{x^*i}{k}\right)\right]k \Big/ u(x^*) - u'\left(\frac{x^*(i+1)}{k}\right)\Big/ u'(x^*)\\
&= e^{-x^*i/k}\left(\frac{(1-e^{-x^*/k})k}{1-e^{-x^*}} - \frac{e^{-x^*/k}}{e^{-x^*}}\right)\\
&= \frac{y^{i-k}}{1-y^k}[(1-y)ky^k - y(1-y^k)]\\
&\qquad\qquad \text{if } y \equiv e^{-x^*/k} \text{ (where } x^* > 0 \text{ implies that } 0 < y < 1)\\
&= \frac{y^{i-k}}{1-y^k}(y(1-y)\big[(y^{k-1}-1) + (y^{k-1}-y)\\
&\qquad\qquad + (y^{k-1}-y^2) + \cdots + (y^{k-1}-y^{k-1})\big],
\end{aligned}$$

which is negative whenever $k > 1$. Therefore, if $k > 1$, then $dU(b; v, k)/db|_{b=b^*} < 0$.

Appendix B

Instructions for the Both Regrets treatment in the $k = 10$ condition are shown below. In the other treatments, references to the corresponding feedback information and/or $k = 10$ auctions were removed.

B.1. 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 maximize 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 checkout form, a pen, and two copies of the consent form.

B.2. 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, whereas 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 that 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.

B.3. 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 minus 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 that 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 minus purchase price. The other two bidders who did not win the auction earn zero.

Example 1. 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.

B.4. 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, in each of the 10 auctions, the computer will calculate and display the following for you:

- Money left on the table, which is always 0 if you do NOT win the auction, and is your bid minus 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 minus winning bid amount.

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

B.5. 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).

B.6. 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

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.

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, R., 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: An 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. 2007. Regret in auctions: Theory and evidence. *Econom. Theory* **33** 81–101.

Engelbrecht-Wiggans, R., E. Katok. 2008. Regret and feedback information in first-price sealed-bid auctions. *Management Sci.* **54**(4) 808–819.

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

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. 2004. Bidder preference among auction institutions. *Econom. Inquiry* **42** 223–236.

Ivanova-Stenzel, R., T. C. Salmon. 2008. Revenue equivalence revisited. *Games Econom. Behav.* **64**(1) 171–192.

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.

Katok, E., T. Salmon. 2008. Response of risky choice to variance reduction. Working paper, Florida State University, Tallahassee.

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

Neugebauer, T., R. Selten. 2006. Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets. *Games Econom. Behav.* **54** 183–204.

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