Overbidding in First Price Private Value Auctions Revisited: Implications of a Multi–Unit Auctions Experiment *

VERONIKA GRIMM

grimm@merlin.fae.ua.es

Departamento de Fundamentos del Análisis Económico

University of Alicante

Campus San Vicente

03071 Alicante

Spain

DIRK ENGELMANN

dirk.engelmann@rhul.ac.uk

Department of Economics

Royal Holloway, University of London

Egham, Surrey

TW20 0EX

United Kingdom

December 21, 2004

*We thank Bob Sherman as well as seminar participants at Nottingham for helpful comments and suggestions. Financial Support by the *Deutsche Forschungsgemeinschaft*, through SFB 373 ("Quantifikation und Simulation Ökonomischer Prozesse"), Humboldt-Universität zu Berlin, and through grant number EN459/1-1 is gratefully acknowledged. Part of this research was conducted while Dirk Engelmann visited the Institute for Empirical Research in Economics, University of Zurich. The hospitality of this institution is gratefully acknowledged.

1 Introduction

One of the to date most intense debates in experimental economics has evolved from a series of papers by Cox, Roberson and Smith (1982) and Cox, Smith and Walker (1983a, 1983b, 1985, 1988) on bidding behavior in single–unit first–price sealed–bid auctions. In several laboratory experiments they observe persistent overbidding of the risk neutral Nash equilibrium (RNNE) strategies, which they argue to be due to risk aversion of the bidders. They show that data of various experiments fit a model of bidders that exhibit constant relative risk aversion (CRRA) and demonstrate that the data yield rather similar estimates of the bidders' average degree of CRRA.

Their conclusion has been criticized by Harrison (1989) who argues that due to the low cost of deviation from RNNE behavior in the experimental settings of Cox, Smith, and Walker, the results cannot be considered significant evidence for risk aversion of the bidders (the so–called "flat maximum critique").

In the subsequent debate, several authors came up with evidence against the CRRA hypothesis and suggested different possible explanations for the observed behavior. In the present paper, we investigate the consistency of the different hypotheses with data obtained from a multi-unit discriminatory auction experiment. Before we do so, let us give an overview over the debate that has been going on up to now.

Kagel and Roth (1992, p.1379) state that " [...] risk aversion cannot be the only factor and may well not be the most important factor behind bidding above the risk neutral Nash equilibrium found so often in first–price private value auctions." They provide evidence for other possible explanations:

First, they note that overbidding of the (dominant) optimal strategy is also observed in second–price sealed–bid auctions, where it cannot be explained by risk aversion.[1] Second, they mention that Cox, Smith and Walker themselves (1985) find evidence against the CRRA hypothesis in an experiment where they pay the subjects one time in money and the second time in lottery tickets. In the second series of experiments the overbidding of RNNE theoretically should disappear, but it does not, which leads Cox, Smith and Walker to reject the empirical adequacy of the lottery technique, rather than revising their hypothesis. Third, they refer to an experiment on multiple

[1] See studies by Kagel, Harstad and Levin (1987) and Kagel and Levin (1990). The same evidence is found in a multi–unit setting by Engelmann and Grimm (2004).

unit discriminatory auctions (Cox, Smith and Walker, 1984) where bids are found to be significantly lower than the RNNE prediction.

Friedman (1992), on the other hand, notes that asymmetric costs of deviation from RNNE would be needed in order to explain the observed "misbehavior" as a consequence of payoff function flatness. Since, however, in most of the relevant experimental studies the loss function is almost symmetric, Harrison's argument cannot be sufficient to explain the observed deviations.

Goeree, Holt and Palfrey (2000) take this point into account and compare behavior of subjects in two different first–price sealed–bid auctions that have the same equilibria but differ with respect to the curvature of the loss function. Inspired by the long lasting debate, they compare several competing explanations for overbidding of RNNE observed also in their data:

(1) Constant relative risk aversion.

(2) Misperception of the probability distribution over outcomes (rank dependent utility).

In their estimations, Goeree, Holt and Palfrey find that risk aversion and misperception of probabilities both yield a good fit of their data, whereas the joy of winning hypothesis is still reasonable but does significantly worse. Their estimated degree of CRRA, moreover, coincides with many other studies in the literature.

In this paper, we contrast those findings with data from multi–unit auction experiments where two bidders compete for two units of a homogenous good. In our analysis we focus on a discriminatory auction and use results from a Vickrey and a uniform–price auction as benchmarks.[2]

In the discriminatory auction data, we observe a high degree of bid spreading, which can be explained neither by risk aversion,[3] nor by misperception of probabilities. A myopic joy of winning seems to fit these data better. Moreover, it is consistent with some subjects' statements in the post–experimental questionnaire: that they used the first bid to ensure getting a unit and the second one for making money. This leads us to a last point that is in sharp contrast to the risk aversion hypothesis: the majority of lower bids (58%, without any discernible time trend) are below the RNNE prediction.

[2]We present a detailed analysis of the other auction formats in Engelmann and Grimm (2004)

[3]Decreasing absolute risk aversion can yield unequal bids, but would imply bids above the RNNE for *both* units which we do not observe.

The paper is organized as follows: In section 2 we introduce the model, derive the RNNE of the game and discuss the implications of the three different alternative hypothesis on equilibrium bidding behavior. The experimental design is presented in section 3. In section 4 we report the experimental results and, in section 5, we contrast them with the three different hypothesis that might explain deviations from RNNE behavior. Section 6 concludes.

2 Theoretical Background and Hypotheses

We investigate bidding behavior in independent private value discriminatory auctions (DA) with two bidders and two indivisible identical objects for sale. In this format, the two highest bids win a unit each and the respective prices equal these bids. Each bidder i, $i = 1, 2$, demands at most two units and places the same value v_i on each of the two units. The bidders' valuations are drawn independently from the uniform distribution on the interval $[0, V]$.

2.1 Risk Neutrality

We start our theoretical analysis by deriving the Risk Neutral Nash equilibrium (RNNE) of the auction. An important observation in order to derive the optimal strategy is that with flat demand a bidder places the same bid on both units.[4] To see this, suppose the other bidder placed two different bids. Then, in order to win one unit, a bidder has to overbid only the other bidder's lower bid and in order to get two units both his bids have to exceed the other bidder's higher bid. Therefore, a bid on the first unit solves the optimal trade–off between the probability of winning (against the other bidder's lower bid) and profit in this case. Now observe that the probability of winning the second unit is even lower (one has to overbid the other bidder's higher bid) and therefore, the optimal trade–off for the second unit cannot be solved at a lower bid.[5] Thus, both bids will be equal since by definition the bid for the second unit cannot be higher than the bid for the first unit. If the other bidder chooses identical bids, the argument is even more obvious, since the trade–off is the same for both units.

[4]See Lebrun and Tremblay (2003) for a formal proof of this fact for much more general demand functions.

[5]"First unit" ("second unit") always refers to the unit on which the bidder places the higher (lower) bid.

Suppose that there exists a symmetric and increasing equilibrium and denote by $b(\cdot)$ and $b^{-1}(\cdot)$ the equilibrium strategy and its inverse function, respectively. Given that the other bidder bids $b(\cdot)$, a bidder with value v on each unit bids

$$\arg\max_{\beta} F(b^{-1}(\beta))[v-\beta], \tag{1}$$

where $F(\cdot)$ is the distribution function of the bidders' values. In the case of uniformly distributed valuations on $[0, V]$ it holds that $F(b^{-1}(\beta)) = \frac{b^{-1}(\beta)}{V}$ and the equilibrium bid functions are

$$b_1(v) = b_2(v) = \frac{1}{2}v, \tag{2}$$

where $b_1(v)$ ($b_2(v)$) is the bid on the first (second) unit.

2.2 Risk Aversion

The most prominent explanation of overbidding in single–unit first price auctions is risk aversion. Thus, in this section we consider the effect of risk aversion on the optimal strategies in our setting.

First, note that a standard result from single–unit auction theory is that (symmetric) risk aversion of any type increases bids above the RNNE level. Thus, independent of the type of risk aversion we assume, we should expect subjects to bid more than half their valuation (the RNNE bid) on any of the two units.[6]

Now consider the case that bidders exhibit constant absolute risk aversion (CARA). Then, a bidder's bids would still be equal, although at a higher level. The reason is the same as under risk neutrality. In order to be able to employ the above argument, it is, however, important to note that under CARA a bidder's wealth does not affect his degree of risk aversion. Then, if a bidder faces the same probability of winning for his first and his second bid, the optimal tradeoff between a higher probability of winning and a higher profit in case of winning is solved by the same bid.[7]

If absolute risk aversion is increasing in wealth[8] optimal bids on the two units will still be equal. A bidder facing two equal bids would have an

[6]See Krishna (2002), Maskin and Riley (1984).

[7]This argument applies if the other bidder places identical bids, which hence holds in equilibrium. Otherwise, we get, as in the case of risk neutrality, that the second–unit bid should be higher than the first–unit bid, which naturally cannot hold.

[8]This is not very plausible in many situations.

incentive to bid higher on the second unit he bids for, because given he obtained the first one, he will be "more risk averse". However, since bidding higher on the second unit is not possible (the bid would turn into the first unit bid), he will bid the same on both units. Note also that increasing absolute risk aversion would make the bids increase over time (i. e. from one auction to the next as long as he makes a profit on the first), depending on the wealth already accumulated by the bidders.

Now consider a bidder with decreasing absolute risk aversion (e. g. constant relative risk aversion (CRRA), which is most often assumed in the literature we referred to). Under decreasing absolute risk aversion, a bidder who has already won one unit will exhibit a lower degree of risk aversion due to his higher wealth. Therefore, if the other player would place equal bids, a bidder would like to bid lower on the second unit. We should therefore expect to observe bid spreading to some extent.

Solving for the equilibrium of the discriminatory auction with bidders that have a CRRA utility function seems to be untractable. Thus, we try to shed light on the behavior of risk averse agents by the following considerations: We simplify the problem by assuming that a bidder decides about his bids sequentially. That is, he first chooses a first–unit bid, ignoring that he will also place a second–unit bid and then he decides on the optimal second–unit bid, conditional on having won with his first–unit bid.

The coefficient of relative risk aversion usually estimated is about 0.5 [see Goeree, Holt, and Palfrey (2000)]. This corresponds to a utility function $U(x) = 2x^{\frac{1}{2}}$. Now consider the case that bidder 2 bids according to the simple linear bid functions $d_1(v) = k_1 v$ and $d_2(v) = k_2 v$ with $k_2 \leq k_1$. For ease of notation assume here that $V = 1$, hence that valuations are uniformly distributed on [0,1]. Then the distribution functions of bidder 2's bids are $F_1(z) = \frac{z}{k_1}$, and $F_2(z) = \frac{z}{k_2}$, for $z < k_1$ and $z < k_2$, respectively, and 1 otherwise. Consider first the case that bidder 1's bids are smaller than k_2. His second unit bid is only relevant if he already wins with his first unit bid. Bidder 1's first–unit bid has to maximize the utility that can be obtained by winning the first unit.

$$
\begin{aligned}
U\left(b_1, v\right) &= 2\left(v - b_1\right)^{\frac{1}{2}} P\left(b_1 > k_2 v_2\right) = \frac{2b_1}{k_2}\left(v - b_1\right)^{\frac{1}{2}}, \\
U'\left(b_1, v\right) &= 0 \Leftrightarrow b_1 = \frac{2}{3} v.
\end{aligned}
$$

If $\frac{2}{3}v > k_2$, then $b_1 = k_2$. Hence, as long as bidder 2's second–unit bid is a simple linear function of his valuation, bidder 1's first unit bid is independent of the precise form of bidder 2's bidding function (except for large valuations, because the bid can then be capped at k_2 which becomes relevant if the other bidder bids relatively low). Conditional on bidder 1's first unit bid being higher than bidder 2's second–unit bid, bidder 2's first–unit bid is uniformly distributed on the interval $[0, \frac{k_1}{k_2}b_1]$, so for any $b_2 < b_1$ we get for the conditional probability $P(b_2 > k_1 v_2) = \frac{b_2}{k}$ with $k := \frac{k_1}{k_2}b_1$.

Given the first–unit bid, the second–unit bid should maximize

$$\begin{aligned} U(b_2, v) &= 2(2v - b_1 - b_2)^{\frac{1}{2}} P(b_2 > k_1 v_2) + 2(v - b_1)^{\frac{1}{2}}(1 - P(b_2 > k_1 v_2)) \\ &= 2\left(\frac{4}{3}v - b_2\right)^{\frac{1}{2}} \frac{b_2}{k} + 2\left(\frac{1}{3}v\right)^{\frac{1}{2}}\left(1 - \frac{b_2}{k}\right), \\ U'(b_2, v) = 0 \qquad &\Leftrightarrow \qquad b_2 = \left(\frac{22}{27} - \frac{2\sqrt{13}}{27}\right) v \approx 0.5477v, \end{aligned}$$

discarding the second solution which implies $b_2 > v$.

Hence, if a bidder with a degree of constant relative risk aversion of 0.5 bids against a bidder whose bids are given by any simple linear bid functions $b_1^2 = k_1 v$ and $b_2^2 = k_2 v$ his optimal bids would then be $b_1 = \frac{2}{3}v$ and $b_2 \approx 0.55v$ where b_1 and b_2 are capped at k_2 for large v (b_2 is really capped by k_1, but since it is also constrained to be no larger than b_1, it is in fact capped at k_2.)

The resulting bid spread is quite substantial (about 24% of the RNNE equilibrium bid), but the average equilibrium spread for two risk averse bidders would be smaller. First, since the other bidder's maximal second unit bid is smaller than his maximal first–unit bid, the first–unit bids are capped at the maximum second–unit bid, which lowers the bid–spread. Second, notice that simultaneous maximization would imply that at least for low values of k_2, b_1 and b_2 would be larger than k_2, enabling the bidder to win *both units* with a high probability. This would clearly reduce the average bid spread.

Hence, for reasonable degrees of risk–aversion we should expect bid–spreads clearly lower than 25% of the RNNE bid for low valuations and lower to no bid spreads for high valuations. Note, however, that this should be coupled with substantial overbidding on both units.[9]

[9]Furthermore, according to Rabin (2000) and Rabin and Thaler (2001) risk aversion

Summarizing, the important observations are (1) that all bids placed by any risk averse bidder should be above the RNNE bids, (2) under risk neutrality, constant, or increasing absolute risk aversion bids on both units should be equal (and in the latter case they should increase over time, depending on the wealth already accumulated by the bidder), and (3) under decreasing absolute risk aversion (e. g. CRRA) bids might be different across units and should decrease over time (i. e. over the course of several auctions) depending on the wealth already accumulated by the bidder.

2.3 Misperception of Probability Outcomes

Some skepticism concerning the CRRA hypothesis may arise from the fact that experimental evidence from lottery choice experiments often suggests that subjects do not even behave consistent with expected utility theory. Thus, several authors have proposed (and tested) models of probability misperception to explain upwards deviations from the RNNE bids in first price auctions. Goeree, Holt, and Palfrey (2000) propose a model of rank dependent utility, where bidders maximize expected utility, but misperceive probabilities. They estimate the parameters of a "S"–shaped probability weighting function proposed by Prelec (1998),

$$w(p) = \exp\left(-\beta\left(-\ln(p)\right)^{\alpha}\right) \tag{3}$$

Subjects behaving according to this function will overestimate the probabilities close to 0, but will underweight probabilities close to 1. This can explain why subjects are willing to bet on gains if the probability of winning is low (this would imply lower bids on the second unit) while they shy away from doing so when in fact the probability of winning is high (which would lead to higher bids on the first unit).[10]

on such small stakes cannot be reconciled with the maximization of the expected utility of wealth. (According to Cox and Sadiraj, 2001, however, it is consistent with the maximization of the expected utility of income). The fact that the small gains from winning the first unit should cause substantially smaller second–unit bids appears to be a good illustration that small stakes risk–aversion is not very plausible in the first place. While we observe even larger bid–spreads, it appears counterintuitive that they result from a dramatic decrease in risk aversion due to a such a small income gain.

[10]The parameters that Goeree, Holt, and Palfrey estimate do actually not correspond to an S–shaped function, but closely to a quadratic function. This corresponds to risk aversion. According to the authors, this does not come as a surprise, because single unit

Misperception of probabilities, however, does not destroy equal bidding as an equilibrium. If the other bidder places two equal bids, the probability to win the first and the second unit if one places two identical bids are the same. Hence while misperception of probabilities can lead to higher or lower bids, it would have the same effect on both bids and hence as long as the other bidder places two equal bids, the best reply always consists of two equal bids.

If, however, the other bidder places two different bids, this might imply bid spreading. If the probability weighting function is S–shaped, and in particular the density is convex, this would mean that if the probability to win is very small, then the perceived $H(b)/h(b)$ [where H (h) is the distribution (density) of the other's bid] is larger, which implies a lower optimal bid. On the other hand, if the probability to win is very large, then the perceived $H(b)/h(b)$ is smaller, implying a larger optimal bid. This can lead to the optimal first and second unit bids being different if the other bidder places different bids. Since the probability to win with the first unit is higher than with the second unit, the distorted perception of the probabilities would bias the second–unit bid down relative to the first–unit bid. Unless the perception of probabilities is dramatically distorted, the effect would, however, not be very large, even if the other bidder's bids are very different. Hence a possible bidspread in equilibrium would be small. In particular, they would be limited for large valuations because the first–unit bid need never be higher than the maximum of the other bidder's second unit bid. Furthermore, the argument above has additional implications. If the valuation is high, and hence the probability to win either of the units is high, both bids should be higher than the RNNE. On the other hand, if the valuation is low, and hence the probability to win either of the units is low, both bids should be below the RNNE.[11] Finally, if bidders are risk neutral and learn over time, bids should converge to RNNE bids.

We summarize that under misperception of probabilities (1) bidding the same on both units is still an equilibrium, (2) depending on the shape of the

auctions cannot discriminate between nonlinear utility and nonlinear probability weighting.

[11] If the probability weighting function is not S–shaped, but for example as estimated by Goeree, Holt, and Palfrey, quadratic, bids could be above the RNNE throughout. A quadratic probability weighting function would, however, not be distinguishable from risk aversion and hence the problems that occur for risk aversion as explanation would apply as well.

probability weighting function and the distribution of values, equilibria with moderate bid spreading might exist, (3) whether bids are above or below RNNE bids depends on the shape of the probability weighting function and (4) if bidders are risk neutral and learn over time, bids should converge to RNNE bids.

2.4 Joy of Winning

Cox, Smith, and Walker (1983b, 1988) and Goeree, Holt, and Palfrey (2000) suggest as an alternative explanation for overbidding in first–price auctions a model, where bidders receive a utility from the event of winning the auction. A pure joy of winning model (without incorporating risk aversion) explains overbidding in single–unit first–price auctions, although (according to Goeree, Holt, and Palfrey) not as good as the previous two explanations. In a multiple unit setting, joy of winning has further implications on the structure of bids which allows us to distinguish it better from the previous two alternatives.

Suppose that the additional utility from winning the auction is proportional to the observed valuation, so that a bidder with valuation v who is bidding (b_1, b_2) has expected utility

$$U(b_1, b_2, v) = H_2(b_1)(vw - b_1) + H_1(b_2)(v - b_2), \tag{4}$$

where $H_1(\cdot)$ $(H_2(\cdot))$ denotes the distribution of the other bidder's higher (lower) bid, and $w > 1$ models the joy of winning.[12] For w big enough it can be shown that bidders always bid higher on the first unit than on the second one, and also that the second unit bid is above RNNE.

Moreover, joy of winning as modeled above could also explain overbidding in second–price auctions (as observed in Kagel, Harstad, and Levin (1987), Kagel and Levin (1990), and Engelmann and Grimm (2004)), which the alternative models suggested above can not.[13]

Summarizing, joy of winning would imply (1) extreme bid spreading if the parameter w is big, (2) higher than RNNE bids on both units, and (3) no adjustments of bids over time since joy of winning as introduced here is

[12] Note that in this formulation winning a second unit does not yield additional joy.

[13] In this case, however, we would require winning to also yield joy if the monetary gain is negative, which might appear less plausible.

a myopic concept in the sense that there is always a joy of winning at least one unit in each auction.[14]

2.5 Hypotheses from the Theory

Table 1 summarizes the predictions that follow from the alternative explanations:

	RNNE	CARA	CRRA	Probability Misperception	Joy of Winning + risk neutrality
first unit bid	$\frac{1}{2}v$	$> \frac{1}{2}v$	$> \frac{1}{2}v$	(?)	$> \frac{1}{2}v$
second unit bid	$\frac{1}{2}v$	$> \frac{1}{2}v$	$> \frac{1}{2}v$	(?)	$> \frac{1}{2}v$
bid spreads	no	no	moderate	no/ moderate for certain distribution and weighting functions	possibly large
bids over time	const.	const.	decreasing	converging to RNNE	const.

Table 1: Hypothesis from the different theories

[14]In particular, since subjects get the result they aim for, namely almost always a positive profit and occasionally a large profit, reinforcement learning would not lead to a decrease in bid spreading in spite of it generating sub–optimal profits.

3 Experimental Design

In each auction two units of a homogeneous object were auctioned off among two bidders with flat demand for two units. The bidders' private valuations for both units were drawn independently in each auction from the same uniform distribution on $[0, 100]$ experimental currency units (ECU).[15] The bidders were undergraduate students from Humboldt University Berlin, the University of Zürich, and the ETH Zürich. Pairs of bidders were randomly formed and each of the nine pairs played ten auctions.

Subjects were placed at isolated computer terminals, so that they could not determine whom they formed a pair with. Then the instructions (see appendix A for the translation) were read aloud. Before the start of a sequence of ten auctions, subjects played three dry runs, where they knew that their partner was simulated by a pre–programmed strategy. This strategy and the valuations of the subjects in the three dry runs were chosen in such a way that it was likely that each subject was exposed to winning 0 units in one auction, 1 unit in another and 2 units in the third. The pre–programmed strategy did not reflect any characteristics of the equilibrium and the subjects were explicitly advised that they should not see this strategy as an example of a good or a bad strategy (because they only observed the bids, they could not really copy the programmed strategy in any case).

The auctions were run in a straightforward way, i. e. both bidders simultaneously placed two bids. Subjects were informed that the order of the bids was irrelevant. After each auction bidders were informed about all four bids, as well as the resulting allocation, their own gains or losses and their aggregate profits.

The experimental software was developed in zTree (Fischbacher, 1999). The sessions lasted for about 30 minutes. At the end of each session, experimental currency units were exchanged in real currency at a rate of DM 0.04 (Berlin) or CHF 0.04 (Zürich) per ECU. In addition subjects received DM 5 (Berlin) or CHF10 (Zürich) as show–up fee.[16] Average total payoffs were 270 ECU. This resulted in average earnings (including show–up fees) of DM

[15] Valuations were in fact drawn from the set of integers in [0,100] and also bids were restricted to integers. This does not, however, influence the predictions.

[16] In order to relate the earnings, the exchange rates are 1 CHF = 0.65 Euro and 1 DM = 0.51 Euro. Cost of living is higher in Zurich, which justified the higher returns. The higher show–up fee in Zurich is based on a longer average commute to the laboratory than in Berlin.

14.79 (about EURO 7.56) in Berlin and CHF 21.68 (about EURO 14.09) in Zuerich.

4 The Data

In this section, we first summarize the results from the experiment before in section 5 we contrast the data with the hypotheses derived in section 2. Throughout our discussion of the experimental results, we use non–parametric Mann–Whitney tests for comparisons between treatments. These are always based on aggregate data per pair. The aggregate is computed over all periods. For comparisons between the first five and the second five auctions, as well as for comparisons with equilibrium predictions, we use non–parametric Wilcoxon signed–rank tests, because the data are paired. Again the tests are based on aggregate data per pair.

4.1 A First Look at the Data

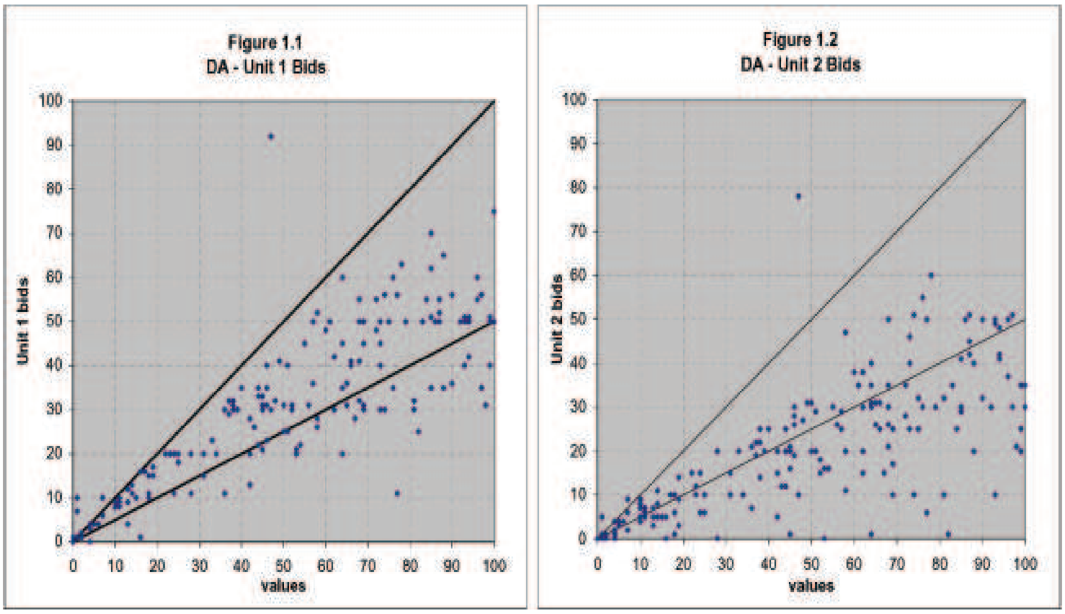

Figure 1: Scatter Diagrams

The scatter diagrams in Figure 1 provide a first impression of the behavior of the bidders. "unit1 bids" refers to the (weakly) higher, and "unit2 bids" to the (weakly) lower bid of a bidder. According to the RNNE prediction, in a discriminatory auction the bidders should place equal bids ($b_1 = b_2 = \frac{1}{2}v$) on both units. However, as the scatter diagrams show, subjects placed substantially different bids on unit 1 and unit 2. The first unit bids seem to be well above the RNNE prediction, whereas the second unit bids are mostly below that level. According to Wilcoxon signed–rank tests, first–unit bids were significantly higher ($p = 0.021$) than the RNNE bid (average difference 5.48 ECU). The average second–unit bid is 3.73 ECU smaller than the RNNE equilibrium bid ($p = 0.139$).

As can also be seen in Figure 1, except for one subject in one auction, we observed overbidding of the valuation only for very small valuations and to a very small degree. It seems that it is obvious to bidders in DA that overbidding is dominated.

4.2 Estimation of Bid Functions for the First and the Second Unit

Our initial observations are supported by estimating first–unit (b_1) and second–unit (b_2) bid functions that are linear in the valuation, i. e.

$$b_i = \alpha_i + \beta_i v. \tag{5}$$

Over all subjects, in a regression of the higher bid (with robust standard errors taking the dependence of observations within each pair into account) the coefficient for the valuation is $\beta_1 = 0.516$ (see Table 2), which is close to the equilibrium value of 0.5, while it is substantially smaller in a regression of the lower bid ($\beta_2 = 0.379$). Combined with estimated constants of $\alpha_1 = 4.706$ and $\alpha_2 = 2.25$ this is consistent with first–unit bids substantially above the RNNE and second–unit bids below the RNNE. In bid functions estimated for individual subjects, β_1 is within 10% deviation of the equilibrium prediction only for 7 out of 18 subjects. For β_2, this is the case for only 5 subjects (see Table 2).

4.3 Bid Spreading

The above results suggest that bids on the first and the second unit were rather different, contrary to the RNNE prediction. Table 3 contrasts the observed bid spreading with the bidspreads observed by Engelmann and

bidder	α_1	β_1	α_2	β_2
1	6.608	0.533	4.601	0.468
2	0.512	0.612	-1.072	0.493
3	0.113	0.602	2.261	0.244
4	14.481	0.299	5.354	0.343
5	8.881	0.479	3.704	0.459
6	6.630	0.744	5.638	0.667
7	14.534	0.227	16.453	-0.071
8	10.647	0.462	10.693	0.220
9	5.434	0.573	2.973	0.401
10	7.205	0.532	1.715	0.519
11	6.829	0.593	0.353	0.306
12	2.090	0.749	2.777	0.408
13	4.753	0.328	4.858	0.252
14	6.165	0.355	2.068	0.334
15	3.572	0.549	4.463	0.326
16	1.781	0.537	0.163	0.511
17	-1.587	0.406	-1.679	0.157
18	3.002	0.449	-1.438	0.354
all	4.706	0.516	2.250	0.379

Table 2: Parameter estimates for the bidding functions.

Grimm (2004) in two other multiple unit sealed auction formats: the Vickrey–Auction (VA), where it is a bidder's dominant strategy to bid his true value on both units (i. e. we expect no bid spreading) and the Uniform–Price Sealed–Bid Auction (UPS) (where we expect up to 100% bid spreading).

We observe that in the discriminatory auction in only 12% of cases the bids were exactly equal and in only 15% (including the 12% equal bids) the difference was smaller than 10% of the risk–neutral equilibrium bid (i. e. 5% of the valuation, see Table 3). More than half of these nearly equal bids (12 out of 21) were submitted by only two subjects (8 by subject 16 and 4 by subject 13, see Table 2 for their estimated bid functions). 49% of the bid spreads were larger than or equal to 40% of the equilibrium bid. The aggregate bid spread is 37%. This corresponds, for example, to bids of 21 and 30 for a valuation of 50 where the risk–neutral equilibrium bids would be 25.

maxbid-minbid	UPS	VA	DA
$= 0$	18%	49%	12%
$< 10\%$ RNNE	34%	62%	15%
$\geq 40\%$ RNNE	33%	14%	49%

Table 3: Share of bid pairs that are exactly equal, where the difference is smaller than 10, or larger than 40 percent of the RNNE bids.

According to Kolmogorov–Smirnov tests, the hypothesis that both, the higher and the lower bids in DA (relative to RNNE bids) are drawn from the same distribution, can be rejected at the 5%–level for 12 out of 18 bidders.

In comparison, in the Vickrey auction (VA) the aggregate bid spread is 13% (see Engelmann and Grimm (2004)) and the hypothesis that both, the higher and the lower bid are drawn from the same distribution can be rejected for only 4 out of 20 bidders at the 5% level. Hence, bid spreading (relative to equilibrium bids) was clearly more prominent in DA than in VA, which is also confirmed by a Mann–Whitney test ($p = 0.0025$). Recall that in both auctions RNNE bids on both units are equal.

In UPS, the aggregate bid spread is 41% (see Engelmann and Grimm (2004)) and the hypothesis that both, the higher and the lower bid are drawn from the same distribution can be rejected for 13 out of 20 bidders at the 5% level. Hence bid–spreading was of the same order in UPS as in DA. Bid spreading (relative to equilibrium bids) was indeed indistinguishable from that in DA (Mann–Whitney test, $p = 0.807$). This is surprising, since in UPS extreme bid spreading is predicted by equilibrium analysis, whereas in DA it is not.

To summarize, bid spreading in DA is much larger than in VA, although it should be zero in both auction formats, and is similar to that in UPS, where it is predicted to be large.

4.4 Time Trends

A linear regression of the bidspread yields, over all subjects and periods (with robust standard errors) a negative coefficient (-0.05) for period, which is, however, not significantly smaller than 0 ($p = 0.83$). Hence, on average the bidspread decreased over time, but the effect is very small and insignificant. Indeed, the aggregate bid spread is 38% in periods 1 to 5 and 36% in periods

6 to 10. Moreover, the aggregate bid spread increased from the first to the second half of the experiment in five pairs, but decreased in only four.

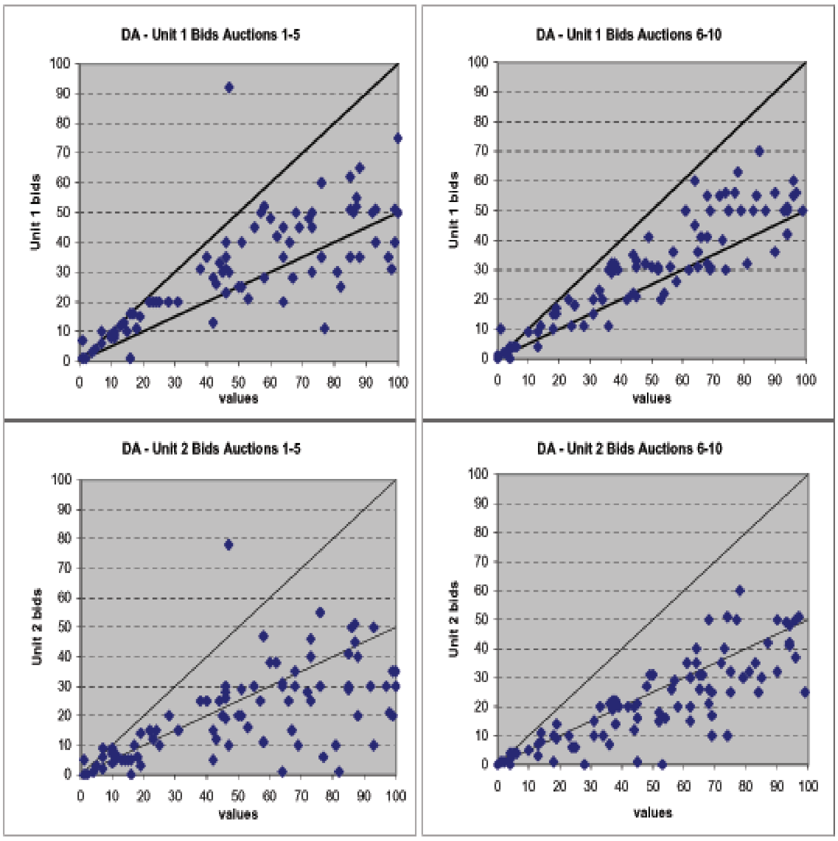

Figure 2: Scatter Diagrams for the first and the second five periods

The first– and second–unit bids by themselves do not exhibit any clear time trends either. Indeed, aggregating over all pairs and either all bids in the first five periods or all bids in the second five periods (see, as an illustration, Figure 2), the first–unit bids amount to 1.22 times the RNNE bid in both the early and the late periods, while the second–unit bids amount to 0.83 times the RNNE bids in periods 1–5, and to 0.86 times the RNNE bids in periods 6–10. The pattern is also highly heterogenous across pairs. First–unit bids relative to the RNNE increase in five pairs from the first five to the last five auctions, and decrease in four. With respect to second–unit bids, the result is just the opposite.[17] In particular, there is no discernible trend towards

[17]Note that these results are unlikely to follow from different draws of valuations in the

lower bids that would be implied by decreasing absolute risk aversion.

5 Comparative Performance of the Suggested Explanations

In this section we discuss the performance of the different theories with respect to organizing our data.

5.1 Risk Aversion — Cannot Be All that Matters

While in single–unit first–price auction experiments risk aversion seems to explain the observed behavior considerably well, it cannot be a satisfactory explanation of our multiple–unit auction data. Several of the observed and significant patterns are not consistent with risk aversion:

(1) **Low second–unit bids.** Bids on the second unit are lower than the RNNE bid. This is clearly inconsistent with the risk aversion hypothesis. Under any kind of risk aversion, the other bidder's first–unit bid is higher than the RNNE bid. Thus, it is even harder for a bidder to obtain the second unit. The lower probability of winning the second unit (due to the high first–unit bid of the opponent) together with risk aversion should yield second–unit bids that are considerably higher than the RNNE bids.

(2) **Extreme bidspreads.** We observe extreme bidspreads in the discriminatory auction. Recall that bidspreads in DA are of the same order as in UPS (where bidspreading should occur in equilibrium) and significantly higher than in VA (where bids should be equal in equilibrium). While the bidspreads are significantly smaller in VA than in DA, they are still present and in this auction format they cannot possibly be explained by risk aversion. Therefore, although mild bidspreading in DA could be explained by decreasing absolute risk aversion, there must still be another motivation for the observed behavior.

(3) **No significant time trends.** As shown in section 2, only decreasing absolute risk aversion could possibly be consistent with a positive bid

different auctions, because the aggregate valuations across all pairs and all of either the first or the last five auctions increases by only about 2% from the former to the latter.

spread. However, this would also imply that bids should decrease over time, depending on the wealth accumulated by a bidder. Since we do not observe this, decreasing absolute risk aversion would have to be highly myopic, i. e. bidders would be required to consider the utility of income and this for each auction separately. To see most clearly that decreasing absolute risk aversion does not work as an explanation, consider the frequent case (32% of the bid pairs) that a bidder places a first–unit bid above the RNNE and a second–unit bid below the RNNE. Apart from the fact that this would require risk aversion to decrease so dramatically that it is actually turned into risk seeking, it further implies, that in all future auctions, both bids should be below the RNNE (as long as the bidder is successful with at least one bid in the current auction), which we clearly do not observe.

To summarize, risk aversion is only a viable explanation for the observed bidspreads if absolute risk aversion is decreasing and stronger than usually estimated. On the other hand, this would imply that both bids are substantially higher than the RNNE and decrease over time, neither of which we observe.

5.2 Misperception of Probabilities

While misperception of probabilities can explain overbidding in single–unit first–price auctions (Goeree, Holt, and Palfrey estimate a concave probability weighting function), it also fails to explain our multi–unit auction data because it is not consistent with the following aspects:

(1) **Bid spreads.** Misperception of probabilities does not eliminate equal bidding as equilibrium. Moreover, while misperception of probabilities might also be consistent with mild bid–spreading, the distortion would have to be dramatic to explain the large spreads that we observe. Furthermore, this would have additional implications not consistent with our data.

(2) **No learning.** In case bidders misperceive probabilities, they should notice that they do so during the course of the experiment. Therefore, over time one should expect bids to get closer to the RNNE prediction. This, however, cannot be observed in our data. The problem could be that subjects played too few rounds in order to be able to learn.

Thus, although it may well be that subjects misperceive probabilities, this can definitely not be the only driving force behind the observed behavior. The data are not consistent with any of the unambiguous predictions implied by this model, that is, small bidspreads and convergence over time.

Still, among the models discussed in section 2, misperception of probabilities is the only one that could possibly explain lower than RNNE bids (which we observed on the second unit). However, for low valuations we should then also observe first–unit bids below the RNNE, which we clearly do not.

Finally, in the postexperimental questionaire some subjects state that they placed "a high secure bid and a lower bid that could yield a higher profit". This suggests that they *willingly* bid rather low on the second unit and did not misperceive the probability of winning.

5.3 Joy of Winning

Joy of winning does not perfectly explain the data, but does considerably better than the alternatives discussed above. As already mentioned, some statements in the postexperimental questionnaires suggest that bidders wanted to secure one unit by placing a high bid on the first one, while they aimed at realizing a high profit by placing a low bid on the second. This seems to describe a (highly myopic) joy of being successful in each single auction, which is consistent with the "Joy of Winning" models that have been discussed in section 2. This explanation is consistent with the following main aspects of out data:

(1) **Bid spreading.** Those models indeed could predict extreme bid spreading as observed in the data, if the additional utility received from winning at least one unit is sufficiently high.

(2) **Overbidding in Vickrey and uniform–price auctions.** Joy of winning is the only among the discussed models that could also explain the often observed overbidding of the valuation in auction formats where the equilibrium first unit bid equals the valuation (VA and UPS). This is sometimes interpreted as a bidding error. It does, however, not disappear even if it is explained to bidders why they should not overbid (see Kagel and Levin, 2001), and it is persistent across almost all experiments on those auctions.

Already Kagel and Roth (1992) made the point that bidders also overbid in auctions where risk aversion plays no role and concluded that there must be something different from risk aversion driving this behavior. Given that bidders run the risk of paying a price higher than their valuation, the joy seems to be present even if winning could imply monetary losses.[18]

(3) **No learning** Finally, the model could also explain why subjects did not revise their behavior in the course of the auction. Actually, they rather get reinforced by frequently winning one unit and occasionally making a large profit on the second one.

The only feature of the data that could not be explained by a "Joy of Winning plus Risk Neutrality" hypothesis (but neither by the other theories discussed above in isolation) is that second unit bids are frequently below the RNNE bid. Hence, a combination with either risk seeking behavior or misperception of probabilities (in the sense that the probability of receiving the second unit is overestimated) would get further in explaining our results. However, in order to fit our data as equilibrium, either of these effects would have to be very strong, because due to first–unit bids well above RNNE bids, the probability of getting a second unit is rather low which should imply rather high second–unit bids.

The data of auction experiments in general strongly suggest that bidders aim at being "successful" on each single occasion. Our data strengthen this point, since they exclude alternative explanations that usually cannot be well distinguished on the basis of data from single–unit auction experiments. A possible explanation for the observed behavior might combine a joy of winning (high first–unit bid) with a joy of gambling (low second–unit bids). Hence our bidders would act like people who buy insurance (against the risk of having zero profit) while at the same time buying lottery tickets.

[18] Possibly bidders have a distorted view of the game in the sense that they realize that overbidding increases both the probability of winning a unit and the probability of making a loss, without realizing that the additional units are *exactly* won when they result in monetary losses. If this is the case then the joy of winning would not have to be so strong as to compensate actual monetary losses (which seems unlikely in the first place) but only bias the perceived trade–off between higher chances to win a unit and higher risks of a monetary loss in favor of the first.

6 Overbidding in First–Price Auctions Revisited

At a first glance the observed bidding behavior in our multi–unit auction experiments looks more or less consistent with the well known phenomenon of overbidding in first–price single–unit auctions. However, as we have shown, it can be explained neither by risk aversion nor by misperception of probabilities, which are the two most prominent hypotheses in the literature on overbidding in single–unit first–price auctions. The fact that in our data the majority of second–unit bids is below the RNNE is clearly inconsistent with risk aversion. Furthermore, the observed bidspreads are of a magnitude that is inconsistent with misperception of probabilities and reasonable degrees of risk aversion. Since an explanatory model should be consistent across different auction formats (e. g. not only be valid for single–unit auctions), the data from our multi–unit auction experiments raise doubts about the explanatory adequacy of risk aversion and misperception of probabilities for overbidding in single–unit auctions.

A further insight from our multi–unit auctions for the interpretation of behavior in single–unit auctions follows from the comparison between risk aversion and misperception of probabilities. In contrast to our auctions, in a single–unit setting the two explanations are usually not distinguishable: Goeree, Holt and Palfrey (2000) study single–unit auctions with asymmetric loss functions, which also allows them to compare these two hypotheses. They find that they perform equally well and better than a joy of winning model. In our setting, however, this order seems to be reversed and, moreover, the first two models perform quite differently.

The behavioral pattern observed in our experiment seems to be caused by a myopic “joy of winning”, which leads subjects to increase the probability of acquiring at least one unit in each auction at the expense of expected profits. This has a lower distorting effect in the other auction mechanisms we mentioned in section 4 (UPS and VA), since the probability of acquiring at least one unit (without making losses) is maximized by bidding the valuation on the first unit, consistent with equilibrium behavior. However, in those auction formats some bidders even risk a loss in order to further increase the probability of winning one unit. Again, a joy of winning hypothesis is the only one that could explain the observed behavior.[19]

For a complete explanation of our data, several reasons have to come together. Our analysis suggests that in any such combination joy of winning would play a prominent role, because no possible combination of the other models can explain the observed behavior, in particular the bidspreading, in a satisfactory way. On the other hand, the observed underbidding on the second unit, which looks like risk–seeking behavior, implies that joy of winning alone cannot provide a complete explanation. The data could be explained by combining joy of winning with either risk prone behavior or with joy of gambling, which could be due to the low stakes in an experiment, or with misperception of probabilities.

A somewhat puzzling observation is that the bidders in all auction experiments appear to be driven by a highly myopic (per auction) desire to win. While such a myopic joy of winning does not appear so surprising in a single–unit auction since it would just suggest that a bidder likes to win as many auctions as possible, it is interesting that in our auction it appears to apply just to win one unit per auction. Hence bidders appear to want to win something in each auction, but not necessarily all the available units.[20]

To conclude, reevaluating the hypotheses that have been suggested for the explanation of the common behavioral pattern in single–unit auctions has cast significant doubt on the performance of the usual suspects. Further experimental research on multi–unit auctions may substantially improve our understanding of the behavior in single–unit auctions, because hypotheses that imply only subtle differences in single–unit auctions can have substantially different implications in multi–unit auctions, making the latter a more powerful tool to discriminate among them.

[19] An interesting conclusion from this observation is that auction formats where everyone obtains something in equilibrium are likely to raise rather low revenues. The effect is the stronger, the more the auction permits the bidders to ensure their opponents winning a unit, i. e. open auctions, where a bidder can do this by dropping out immediately. Sealed bid formats, in contrast, maintain a certain extent of uncertainty about winning a unit and therefore trigger more aggressive behavior. This is consistent with findings in Engelmann and Grimm (2004), where sealed bid formats yield significantly higher revenues than open auctions.

[20] Joy of winning might in some cases be driven by an aversion against zero–profits. While such a myopic zero–profit aversion might explain our data for DA, it is not a viable explanation for the overbidding in UPS and VA, since a dislike for zero–profits is hardly strong enough to risk negative profits.

References

Cox, J. C. and Sadiraj, V. 2001. Risk Aversion and Expected–Utility Theory: Coherence for Small– and Large–Stakes Games, Working Paper, University of Arizona.

Cox, J. C., Roberson, B., and Smith, V. L. 1982. Theory and Behavior of Single Object Auctions, in: V. L. Smith, ed., Research in Experimental Economics, Vol. 2, Greenwich, CT: JAI Press, pp. 1–44.

Cox, J. C., Smith, V. L., and Walker, J. M. 1983a. Tests of a Heterogenous Bidder's Theory of First Price Auctions, Economics Letters 12),pp. 207–212.

Cox, J. C., Smith, V. L., and Walker, J. M. 1983b. A Test that Discriminates Between Two Models of the Dutch–First Auction Non–isomorphism, Journal of Economic Behavior and Organization, 4,pp. 205–219.

Cox, J. C., Smith, V. L., and Walker, J. M. 1984. Theory and Behavior of Multiple Unit Discriminative Auctions, Journal of Finance 39, pp. 983–1010.

Cox, J. C., Smith, V. L., and Walker, J. M. 1985. Experimental Development of Sealed–Bid Auction Theory: Calibrating Controls for Risk Aversion, American Economic Review 75 (Papers and Proceedings), pp. 160–165.

Cox, J. C., Smith, V. L., and Walker, J. M. 1988. Theory and Individual Behavior of First–Price Auctions, Journal of Risk and Uncertainty 1, pp. 61–99.

Engelmann, D. and Grimm, V. 2004. Bidding Behavior in Multi–Unit Auctions — An Experimental Investigation and some Theoretical Insights, Working Paper, University of Alicante.

Fischbacher, U. 1999. Z–Tree: Zurich Toolbox for Readymade Economic Experiments, Working paper No. 21, Institute for Empirical Research in Economics, University of Zurich.

Friedman, D. 1992. Theory and Misbehavior of First–Price Auctions: Comment, American Economic Review 82, pp. 1374–1378.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. 2000. Quantal Response Equilibrium and Overbidding in Private Value Auctions, Working Paper, California Institute of Technology.

Harrison, G. W. 1989. Theory and Misbehavior of First–Price Auctions, American Economic Review 79, pp. 749–762.

Kagel, J. H., Harstad, R. M., and Levin, D. 1987. Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study, Econometrica 55, pp. 1275–1304.

Kagel, J. H. and Levin, D. 1990. Independent Private Value Auctions: Bidder Behavior in First–, Second– and Third–Price Auctions with Varying Numbers of Bidders, mimeo, University of Houston.

Kagel, J. H. and Levin, D. 2001. Behavior in Multi–Unit Demand Auctions: Experiments with Uniform Price and Dynamic Auctions, Econometrica 69, pp. 413 – 454.

Kagel, J. H. and Roth, A. E. 1992. Theory and Misbehavior of First–Price Auctions: Comment, American Economic Review 82, pp. 1379–1391.

Krishna, V. 2002. Auction Theory. Academic Press.

Lebrun, B. and Tremblay, M.–C. 2003. Multi–Unit Pay–Your–Bid Auction with One–Dimensional Multi–Unit Demands, International Economic Review 44, pp. 1135–1172.

Maskin, E. and Riley, J. 1984. Optimal Auctions with Risk Averse Buyers, Econometrica 52, pp. 1473–1518.

Prelec, D. 1998. The Probability Weighting Function, Econometrica 66, pp. 497–527.

Rabin, M. 2000. Risk Aversion and Expected–Utility Theory: A Calibration Theorem, Econometrica 68, pp. 1281–1292.

Rabin, M. and Thaler, R. H. 2001. Anomalies: Risk Aversion, Journal of Economic Perspectives 15, pp. 219–232.

A Instructions (Original Instructions Were in German)

Please read these instructions carefully. If there is anything you do not understand, please raise your hand. We will then answer your questions privately. The instructions are identical for all participants.

In the course of the experiment you will participate in 10 auctions. In each auction you and another bidder will bid for two units of a fictitious good. This other bidder will be the same in each auction. Each unit that you acquire will be sold to the experimenters for your private resale value v. Before each auction this value **per unit**, v, will be randomly drawn independently for each bidder from the interval $0 \leq v \leq 100$ ECU (Experimental Currency Unit). Any number between 0 and 100 is equally probable. The private resale values of different bidders are independent. **In each auction any unit that you acquire will have the same value for you. This value will be drawn anew before each auction.**

Before each auction you will be informed about your resale value **per unit**, v. Each participant will be informed only about his or her own resale value, but not about the other bidder's resale value.

Subsequently, you have to make your bids b_1 and b_2. You enter your bids in the designated fields (one each for the first and the second unit) and click the field *OK*.

The two highest bids win the units. Hence you will win one unit if one of your bids is among the highest two units and you obtain both units if both your bids are higher than those of the other bidder. If because of identical bids the highest bids are not uniquely determined, then the buyers will be chosen randomly.

If you win a unit then you pay the amount you have bid for this unit. Your profit per unit that you obtain amounts thus to your resale value minus the bid you have won the unit for. If you do not win any unit then you will not obtain anything and also not pay anything, hence your profit is 0.

Note that you can make losses as well. It is always possible, however, to bid in such a way that you can prevent losses for sure.

You will make your decision via the computer terminal. You will not get to know the names and code numbers of the other participants. Thus all decisions remain confidential.

One ECU corresponds to 0,04 DM. You will obtain an initial endowment of 5 DM. If you make losses in an auction these will be deducted from your previous gains (or from your initial endowment). You will receive your final profit in cash at the end of the experiment. The other participants will not get to know your profits.

If there is something you have not understood, please raise your hand. We will then answer your questions privately.