Heterogeneous bids in auctions with rational and boundedly rational bidders—Theory and Experiment*

Oliver Kirchkamp† J. Philipp Reiß‡

August 5, 2014

Abstract

We present results from a series of experiments that allow us to measure overbidding and, in particular, underbidding in first-price auctions. We investigate how the amount of underbidding depends on seemingly innocent parameters of the experimental setup.

To structure our data we present and test a theory that introduces constant markdown bidders into a population of fully rational bidders. While a fraction of bidders in the experiment can be well described by Bayesian Nash equilibrium bids, a larger fraction seems to either use constant markdown bids or seems to rationally optimise against a population with fully rational and boundedly rational markdown bidders.

Keywords: Experiments, Auction, Bounded rationality, Overbidding, Underbidding, Markdown bidding

(JEL C92, D44)

1 Introduction

In this paper we study one feature of bidding behaviour in first-price auction experiments with private values that was paid little attention to for a long time: underbidding for low valuations. One reason for taking insufficient notice of this feature might be that underbidding is difficult to observe with standard experimental setups, another reason might be that underbidding is hard to reconcile with several established theories.

In this paper we present a method that allows to observe overbidding and underbidding in first-price auctions. We find that the amount of underbidding depends on seemingly innocent parameters of the experimental setup. To organise the data we introduce (boundedly rational) constant markdown bidders into a population of fully rational bidders. In our analysis we consider a heterogeneous population of rational bidders and (boundedly rational) markdown bidders.

To set the stage for our paper, let us review a seminal series of first-price auction experiments presented by Cox et al. (1983, 1985, 1988). Figure 1 shows bidding data from one of their experiments. Participants repeatedly play a first-price auction with a fixed number of

*We thank audiences at Amsterdam, Atlanta, Bonn, Cologne, Düsseldorf, Essen, Graz, and OSU for useful comments and suggestions. Financial support from the Deutsche Forschungsgemeinschaft through SFB 504 is

Figure 1 An example from an experiment by Cox et al. (1988)

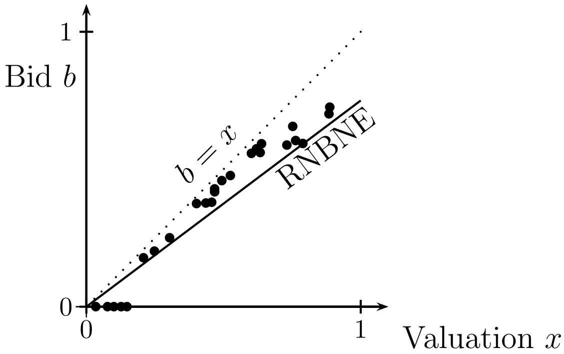

(page 83, figure 8, series 4, exp. 3, n=4, subject 2)

bidders. For each participant valuations are drawn from a uniform distribution. The figure depicts normalised[1] bidding data for a specific subject.

The solid line indicates the risk neutral Bayesian Nash equilibrium (RNBNE) bidding function. As commonly found, most bids exceed the risk neutral equilibrium bidding function. This is what we refer to as overbidding and what was replicated in many first-price auction experiments.

Interestingly, closer examination of the bidding data in figure 1 reveals that for low valuations many bids are below, not above, the equilibrium bid. This characteristic of bidding data is not pathological: Cox et al. (1988) find in some cases negative intercepts when approximating bids by linear bidding functions in some cases; further, Ivanova-Stenzel and Sonsino (2004) report that 7.4% of the bids in their first-price auction experiments are below the lowest possible valuation. If bidders attach any utility to money, these bids cannot be part of an equilibrium.

Despite these findings, underbidding does not receive much attention in the experimental literature. One reason might be that underbidding is often ruled out implicitly or explicitly through the design of the experiment. Choosing zero as the smallest possible valuation looks like an innocent simplification. In this paper we will show that this simplification implies strong behavioral effects. Another reason might be that bids for small valuations are not easy to observe precisely. In this paper we will use a variant of the strategy method that makes it easier to observe these bids.

In this paper we present a simple theory of heterogenous bidders, differing in the degree of rationality, that is supported by our data. Previous literature demonstrates heterogeneity in subjects' behavior in many experimental settings, including auction experiments. Thus, approximating bidding behavior by a single bidding function with a fixed functional form that describes all bids reasonably well might be too demanding. Instead, we propose three different types of bidders which can be linked to different degrees of rationality. We know from experiments with other games that decision makers apply different levels of reasoning when

gratefully acknowledged.

†Universitẗ Jena, School of Economics, Carl-Zeiß-Str. 3, 07743 Jena, Germany, +49-3641-943240, Fax +49-3641-943242, oliver@kirchkamp.de

‡Karlsruhe Institute of Technology (KIT), Institute of Economics, philipp.reiss@kit.edu

[1]In this experiment the smallest possible valuation was \$0 or \$0.10 and the largest possible valuation ranges from \$4.90-\$36.10. In the figure the valuations are normalised to $[0, 1]$ and bids are normalised correspondingly.

choosing a strategy in a game (see Bosch-Domenech et al., 2002). In the context of first-price auctions we should mention Crawford and Iriberri (2007) who analyse bidding under different degrees of rationality. In our paper we suggest a specific starting point for such a sequence of different levels of rationality: absolute markdown bids. For a population-mix model in the double auction, see Saran (2011).

Section 2.1 reviews the model of an extreme case: a rational bidder who assumes that the opponent is rational, too. In section 2.2 we study the opposite extreme: a bidder who is restricted to use absolute markdown bids and assumes that the opponent obeys to the same restriction. The idea of boundedly rational bidders is not new. Kagel et al. (1987) have used a more flexible form of markdown bids in the context of affiliated private value auctions and found some explanatory power. In contrast, we introduce an equilibrium foundation for absolute markdown bids that also accommodates heterogeneous bidding behaviour with perfectly rational bidders along with boundedly rational markdown bidders. Chen and Plott (1998) also compare several variants of markdown bids with Bayesian Nash equilibrium bids when bidders exhibit constant relative risk aversion (CRRA). They find that CRRA provides a more accurate model than their variants of markdown bids. We do not deny such a possibility. If all bidders must fit a single type of bidding function, then CRRA might be a good choice. In our experiment we want to examine whether some bidders systematically do something else. As a natural next step, we will consider a bidder who is rational but assumes to meet a mix of rational and restricted opponents in section 2.3.

Section 3 then describes the experiment and section 4 presents the results. Section 5 concludes. Anticipating our result we will be able to classify bidders into these three groups that we outlined above. If the experimental setup rules out underbidding then these groups are indistinguishable. Bids can be well explained by established theories, e.g. CRRA. Once underbidding is possible in the experiment, some decision makers continue to bid in line with CRRA, but the majority of decision makers follow a quite different bidding pattern.

2 The theoretical framework

In this section, we derive optimal bidding functions for three different contexts that differ in the population's composition of rational and boundedly rational bidders. We concentrate on a first-price sealed-bid auction with private valuations and two bidders. First, we report the well-known Bayesian Nash equilibrium with rational bidders where the rationality of bidders is common knowledge. Second, we introduce boundedly rational bidders that we refer to as markdown bidders. We derive the optimal bid function for the situation where the bounded rationality of bidders is common knowledge. Third, we derive the optimal bid function of rational bidders where it is common knowledge that there are boundedly rational bidders alongside rational bidders in the underlying population of bidders.

2.1 Bayesian Nash Equilibrium bids

Deriving the Bayesian Nash Equilibrium for the first-price sealed-bid auction is standard and is repeated here to introduce the notation. Consider the case where valuations are distributed uniformly over $[0, 1]$ and bidders have constant relative risk aversion (CRRA), i.e., utility is given by $u(x) = x^r$ where r is a parameter of risk tolerance. A risk neutral individual is described by $r = 1$, a risk averse individual has $r < 1$. We confine our attention to the case of

$r \in (0,1]$. Let us assume that bidders use a symmetric increasing bidding function $\gamma(x)$ with inverse $\gamma^{-1}(\cdot)$. In equilibrium all bidders bid according to γ. Bidder 1 has valuation x and bids b. Hence, bidder 1 wins the auction if the valuation of bidder 2 is smaller than $z = \gamma^{-1}(b)$ which occurs with probability $G(z) \equiv z$. Bidder 1 chooses z to maximise $\mathrm{EU} = G(z) \cdot u(x - \gamma(z))$ which yields the following first-order condition:

$$(x - \gamma(z))^r - r \cdot z \cdot (x - \gamma(z))^{r-1} \gamma'(z) = 0 \tag{1}$$

In the symmetric equilibrium we have $z = x$:

$$(x - \gamma(x))^{r-1} \cdot (\gamma(x) + x \cdot (r\gamma'(x) - 1)) = 0 \tag{2}$$

It is easy to see that in equilibrium $\gamma(0) = 0$ which yields the unique solution

$$\gamma^*(x) = \frac{x}{1+r} \,. \tag{3}$$

The second derivative $\partial^2 \mathrm{EU}/\partial z^2 = -(rx/(1+r))^{r-1}$ is negative, so we have indeed found a maximum. If valuations are drawn from the interval $[\underline{\omega}, \overline{\omega}]$ instead of $[0,1]$, one finds similarly that the equilibrium bid is

$$\gamma^*(x) - \underline{\omega} = \frac{x - \underline{\omega}}{1+r} \,. \tag{4}$$

As is well-known, the more risk averse a bidder is (the smaller the value of r), the larger is γ^*. Further, for finitely risk-averse bidders $\gamma^*(x) < x$, so that bidders "shade their bids" by a fraction of $x - \underline{\omega}$ that depends on risk tolerance r.

2.2 Equilibrium with markdown bids

In the Bayesian Nash equilibrium of the first-price auction bidders 'shade their bids' proportional to $x - \underline{\omega}$. However, in a post-experimental questionnaire of another first-price auction experiment[2] some participants explained that they shade their bids not by a relative but, instead, by a constant amount. More broadly, shading by a constant amount may have various reasons:

- It may be cognitively too difficult to work out the exact form of equation (4) or to intuitively behave in full accordance with it. However, participants quickly understand that the bid must be somewhat lower than the valuation to have the opportunity of gaining a positive payoff, hence they may resort to finding a suitable constant by trial and error.

- Shading by a constant amount could be due to satisficing behaviour. A bidder who wants to gain a pre-determined amount if winning the auction must bid the own valuation minus this amount.

- Shading by a constant amount can also be interpreted as a simple rule given to a bidding agent. If first a principal has to define a bidding rule (before the individual valuation is revealed) and then the agent who follows this rule learns the valuation, then it might be simpler for the principal to require a fixed amount that the agent is supposed to gain from each trade.

[2] Kirchkamp et al. (2009).

To formalise this type of bidding behaviour, we utilise a notion of behavioural bidding that assumes behavioural player i to subtract a fixed amount from the valuation as follows:

$$\bar{\gamma}_i(x) = x - \alpha_i \tag{5}$$

where the parameter of autonomous bid shading, $\alpha_i \geq 0$, is individual-specific. To capture that real-life participants in auctions do not arbitrarily select the size of autonomous bid shading, we move beyond a purely behavioural approach with α_i given by a draw from some distribution. Instead, we endogenise α_i by assuming that bidders are boundedly rational and maximise their expected utility by choosing the parameter of constant bid shading independently.

In the following we derive the Bayesian Nash equilibrium of the auction game where both bidders engage into behavioural bidding as described before but choose their amounts of autonomous bid shading α_i simultaneously before learning their valuations. After learning the valuation, they bid according to the implied bidding rule. The assumption that bidders cannot update their bidding rules in response to learning their valuation is essential for autonomous bid shading. If bidder i could change the bidding rule by selecting a different α_i after observing valuation x_i, then the Bayesian Nash equilibrium with rational bidders as reported in section 2.1 emerges since then α_i is conditioned on the valuation x_i such that it replicates equation (3). As a result the equilibrium value of α with autonomous bid shading is optimal in expectation, although it is not the best response after learning the realisation of a particular valuation.

Figure 2 Space of bidders' values

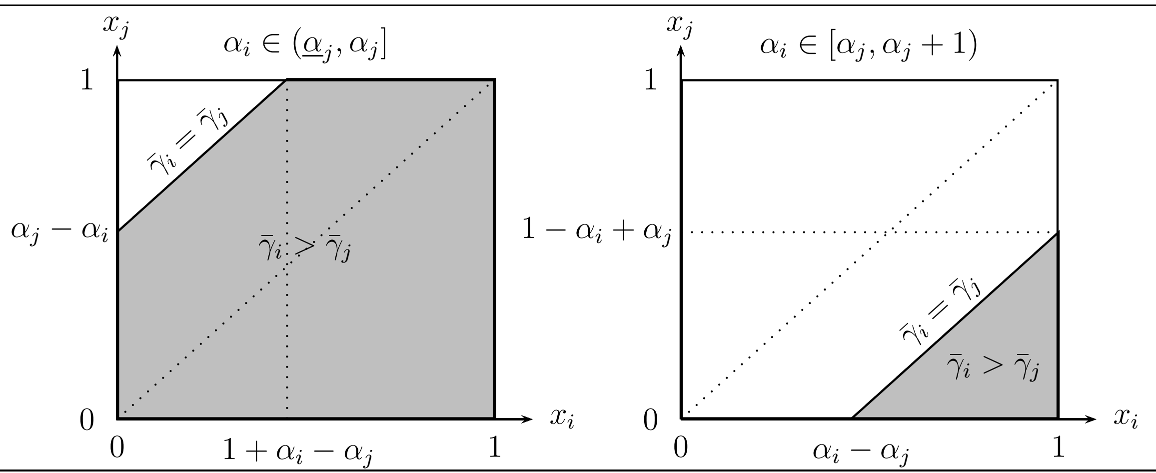

Bidder i wins the auction with the higher bid, $\gamma_i > \gamma_j$, for all bidders' values $(x_i, x_j) \in [0,1]^2$ such that bid shading implies $x_i - \alpha_i \geq x_j - \alpha_j$. For any α_j, bidder i finds it optimal to respond with an amount of bid shading α_i^* such that $\alpha_i^* \in [\underline{\alpha}_i, \alpha_j + 1]$ where $\underline{\alpha}_i = \max\{0, \alpha_j - 1\}$. With bid shading beyond $\alpha_j + 1$ there is no realisation of values (x_i, x_j) that allows bidder i to win the auction, hence any choice of $\alpha_i > \alpha_j + 1$ is strictly dominated by, e.g., $\alpha_i = \alpha_j$. Similarly, any amount of bid shading smaller than $\alpha_j - 1$ allows bidder i to win the auction for any realisation of values so that shading bids by $\alpha_i = \alpha_j - 1$ strictly dominates any smaller amount of bid shading.

Figure 2 indicates the set of bidders' values that lead bidder i to win the auction with autonomous bid shading of (α_i, α_j) as grey-shaded regions; the left panel assumes that bidder i shades bids less than bidder j, while the right panel assumes the opposite. Bidder i wins the auction if bidder j submits the smaller bid, i.e. if $x_j < x_i + \alpha_j - \alpha_i$. It follows that the expected utility of bidder i is given by:[3]

$$\mathrm{EU}_i(\alpha_i) = \begin{cases} [1 - \int_0^{1+\alpha_i-\alpha_j} \int_{x_i-\alpha_i+\alpha_j}^{1} f(x_i)\, f(x_j)\, dx_j\, dx_i]\, u(\alpha_i) & \text{if } \underline{\alpha}_i \le \alpha_i \le \alpha_j \\ [\int_{\alpha_i-\alpha_j}^{1} \int_0^{1+\alpha_i-\alpha_j} f(x_i)\, f(x_j)\, dx_j\, dx_i]\, u(\alpha_i) & \text{if } \alpha_j \le \alpha_i \le \alpha_j + 1 \end{cases}$$

With $u(x) = x^r$ and uniformly distributed values, bidder i's optimal amount of autonomous bid shading α_i^* is given by:[4]

$$\alpha_i^*(\alpha_j) = \begin{cases} \frac{r(\alpha_j+1)}{2+r} & \text{if } 0 \le \alpha_j \le \frac{r}{2} \\ \frac{(\alpha_j-1)(1+r)}{2+r} + \frac{1}{2+r}\sqrt{2r\,(2+r) + (\alpha_j - 1)^2} & \text{if } \alpha_j \ge \frac{r}{2} \end{cases} \tag{6}$$

The best-response function $\alpha_i^*(\alpha_j)$ is continuous and strictly increasing in the other bidder's amount of bid shading α_j. Solving for the unique Bayesian Nash equilibrium yields the equilibrium value of bid shading

$$\alpha^* = \frac{r}{2}$$

and the equilibrium bid function follows as

$$\bar{\gamma}^*(x) = x - \frac{r}{2}. \tag{7}$$

With risk neutrality, $r = 1$, bidders shade their bids by 1/2; for an increasing degree of risk aversion, i.e. for decreasing r, the equilibrium amount of constant bid shading decreases.

2.3 Equilibrium with rational bidders alongside markdown bidders

In the previous two subsections we have considered the two polar cases of a homogeneous population with either rational agents or with boundedly rational agents only. In real-life or in an experiment the population might be heterogeneous in terms of cognitive abilities—some players might be more rational or less cognitively limited than other players. For recent evidence that heterogenous cognitive abilities and beliefs about cognitive heterogeneity of players can influence behaviour in games see Blume and Gneezy (2010). To address the possibility of heterogeneous levels of rationality, we assume that the underlying population of potential bidders is composed of rational players alongside boundedly rational players.

Specifically, let $\rho \in [0, 1)$ be the share of all perfectly rational bidders in the population of potential opponents while the remaining population with share $1 - \rho$ consists of markdown bidders. This population composition is common knowledge among rational bidders only, while boundedly rational markdown bidders are assumed to believe to play against another markdown bidder with probability one. Let $\theta_j \in \{\mathrm{R}, \overline{\mathrm{R}}\}$ denote the rationality type of player j that can be either fully rational, $\theta_j = \mathrm{R}$, or boundedly rational in the sense of markdown bidding, $\theta_j = \overline{\mathrm{R}}$. Then a fully rational bidder's prior of competing with another fully rational bidder is ρ and that of facing a markdown bidder is $1 - \rho$.

Assume that there is an equilibrium such that the fully rational type bids according to $\gamma(x)$ and the boundedly rational type bids according to $\bar{\gamma}(x)$ where both equilibrium bid

[3]The expected utility for any $\alpha_i < \alpha_j - 1$ or any $\alpha_i > \alpha_j + 1$ is given by $u(\alpha_i)$ or 0, respectively.

[4]See appendix A for the detailed derivation.

functions are strictly increasing. The expected utility of a fully rational bidder i facing the competitor j, who is randomly drawn from the population of bidders, is (assuming that bidder j bids according to the proposed equilibrium) given by:

$$\begin{aligned}\mathrm{EU}_i = \rho \cdot \Pr\{b_i = \max\{b_i, \gamma(x_j)\} \mid \theta_j = \mathrm{R}\} \cdot u(x_i - b_i) \\ + (1-\rho) \cdot \Pr\{b_i = \max\{b_i, \bar{\gamma}(x_j)\} \mid \theta_j = \overline{\mathrm{R}}\} \cdot u(x_i - b_i).\end{aligned}$$

Since the assumed equilibrium bid function of the fully rational type is strictly increasing, there exists the inverse $\chi(b) := \gamma^{-1}(b)$ that maps a fully rational player's bid b to the corresponding value x. The probability that player i outbids another fully rational bidder follows as $F(\chi(b_i))$. Let $G(b)$ denote the cumulative distribution function of *bids* submitted by the boundedly rational type so that the probability of player i outbidding this type is $G(b_i)$. By markdown bidding as described by (7) together with the distribution of values, $F(x)$, we have $G(b) = b + r/2$ for $b \in [-r/2, (2-r)/2]$. Therefore, the maximization problem of fully rational bidder i that competes with bid b_i against an equilibrium bidder of unknown rationality type is given by

$$\max_{b_i} \quad \mathrm{EU}_i = [\rho\, F(\chi(b_i)) + (1-\rho)\, G(b_i)] \cdot (x_i - b_i)^r.$$

The first-order condition follows as

$$[\rho F'(\chi(b_i))\chi'(b_i) + (1-\rho)G'(b_i)]\,(x_i - b_i)^r - r\,[\rho F(\chi(b_i)) + (1-\rho)G(b_i)]\,(x_i - b_i)^{r-1} = 0.$$

For fully rational bidder i it cannot be beneficial to deviate from the equilibrium strategy in equilibrium, hence, $x_i = \chi(b_i)$. Using this property and substituting for probability densities leads to the following differential equation whose solution (with an appropriate initial value to be determined below) is the inverse of the equilibrium bid function of the fully rational type, $\chi(b)$,

$$\rho\,[\chi(b_i) - b_i]\,\chi'(b_i) = [(1+r)\rho - 1]\,\chi(b_i) + (1+r)(1-\rho)b_i + (1-\rho)\frac{r^2}{2} \tag{8}$$

In equilibrium, a rational bidder with the smallest possible value of 0 never wins against another rational bidder but only against boundedly rational bidders. With the distribution of bids submitted by markdown bidders, $G(b)$, the optimal bid of rational bidder i with x_i follows as[5]

$$\gamma(0) = -\frac{r^2}{2(1+r)}$$

The initial condition follows as $\chi(-r^2/(2(1+r))) = 0$. Since differential equation (8) is non-linear and non-autonomous an explicit solution is not known in general. Figure 3 shows the (inverted) numerical solution for different attitudes towards risk r and various population mixes ρ.

[5]The maximization problem of rational bidder i with value $x_i = 0$ is

$$\max_{b_i} \quad \rho \cdot 0 + (1-\rho) \cdot G(b_i) \cdot (0 - b_i)^r$$

where $G(b_i) = b_i + r/2$ for $b_i \in [-r/2, (2-r)/2]$ and the first-order condition follows as $(-b)^r - r(b + r/2)(-b)^{r-1} = 0$ and is necessary and sufficient for a unique maximum.

Figure 3 Equilibrium bid function of rational bidders $\gamma(x)$ for risk aversion $r \in \{1, \frac{2}{3}, \frac{1}{3}\}$ and population mix $\rho \in \{.1, .2, \ldots, .9\}$.

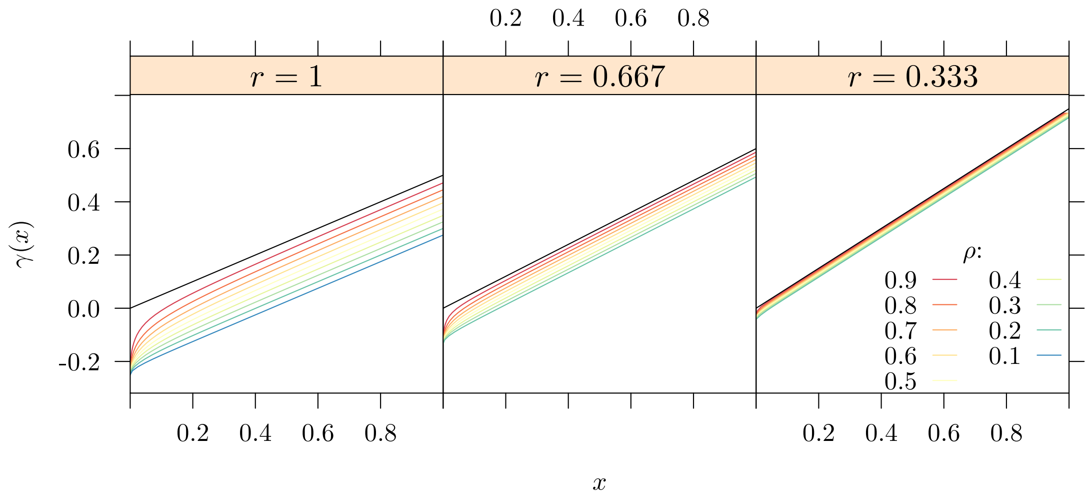

Coloured lines show equilibrium bids in the mixed population for different values of ρ. The black line shows the equilibum bid for a population with only rational bidders.

3 Experimental setup

The purpose of the experiment to twofold: We want to examine to what extent the existing experimental evidence on first-price auctions is an artefact of the design and we want to find out how far absolute markdown bids are consistent with actual behaviour.

Chen and Plott (1998) study a situation where Bayesian Nash equilibrium bids are not linear. In this setup they do not find much evidence of markdown bids. To give markdown bids a good chance to be observable we use here a situation where Bayesian Nash equilibrium bids are linear and clearly distinguishable from markdown bids. Of course, our design does not allow us to assess the capability of markdown bids to explain bidding behaviour in all conceivable auctions. However, it allows us to establish whether markdown bids are an element contributing to actual bidding behaviour.

Comparing equations (4) and (7) shows that absolute markdown bids differ from Bayesian Nash equilibrium bids—in particular for low valuations: Absolute markdown bids can be smaller than the smallest valuation while Bayesian Nash equilibrium bids cannot. We exploit this difference to distinguish absolute markdown bids from Bayesian Nash equilibrium bids. This has two implications for the experiment:

First, we must observe bids also for low valuations in a reliable way. To allow bidders to gain as much experience as possible for low valuations we use a setup with two bidders only. Furthermore, we use the strategy method and play five independent auctions in each round which increases the chance of feedback with low valuations. The idea of this setup is similar to that in Kirchkamp et al. (2009) and Kirchkamp and Reiß (2011).

Second, we must provide the realistic possibility for bidders to submit bids that are lower than the lowest valuation. This might be difficult if the lower bound of valuations is equal to zero as in many experimental studies. To this end we want to understand how seemingly

Table 1 Treatments

Treatment	$[\underline{\omega}, \overline{\omega}]$	restriction of bids	auction
-25	$[-25, 25]$	—	first-price
0	$[0, 50]$	—	first-price
0+	$[0, 50]$	only positive bids	first-price
25	$[25, 75]$	—	first-price
50	$[50, 100]$	—	first-price
50+	$[50, 100]$	only positive bids	first-price
50+II	$[50, 100]$	only positive bids	second-price

innocent changes in the parameters of the experiment affect the choice between Bayesian Nash equilibrium bids and absolute markdown bids. In our experiment we vary the range of valuations and the restriction to submit only positive bids.

In section 2 we determined equilibrium bids and absolute markdown bids for valuations which are distributed uniformly over an interval $[0, 1]$. These bids can be easily generalised to valuations which follow a uniform distribution over any interval $[\underline{\omega}, \overline{\omega}]$. Table 1 lists these intervals that we study in our experiments and provides treatment names. With our experimental design, we test the following hypotheses:

Hypothesis 1 (pure Bayesian Nash equilibrium bidding) *If all bidders use Bayesian Nash equilibrium bids we should not find any significant amount of underbidding. Also if bidders are risk averse, or if regret or spite plays a substantial role, we should not find underbidding.*

Hypothesis 2 (partial markdown bidding) *If some bidders use markdown bids or if some bidders believe that there are absolute markdown bidders with positive probability we should find underbidding for small and overbidding for large valuations in all treatments where underbidding is possible (i.e. the* -25*,* 25*,* 50*, and* 50+ *treatment).*

Even with markdown bids we should find no underbidding in the 0+ treatment since there it is not possible to submit negative bids. The 0 treatment where negative bids are allowed is an intermediate case. Some participants might be tempted to assume that bids should not be smaller than zero, others might not.

Hypothesis 3 (suppression of markdown bidding) *We should find more absolute markdown bids in the 0 treatment than in the 0+ treatment.*

To check whether the restriction to positive bids has any confounding effects even with an interval where the restriction should not matter we compare the 50 to the 50+ treatment leading to hypothesis 4.

Hypothesis 4 (strong negative bids exclusion effect) *We should find more absolute markdown bids in the 50 treatment than in the* 50+ *treatment.*

While the -25 treatment is theoretically equivalent to the 25 and 50 treatment, the -25 treatment involves negative and positive valuations at the same time. This might be perceived as more difficult and, thus, may give an additional incentive to use (simpler) absolute markdown bids.

Figure 4 A typical input screen in the experiment (translated into English)

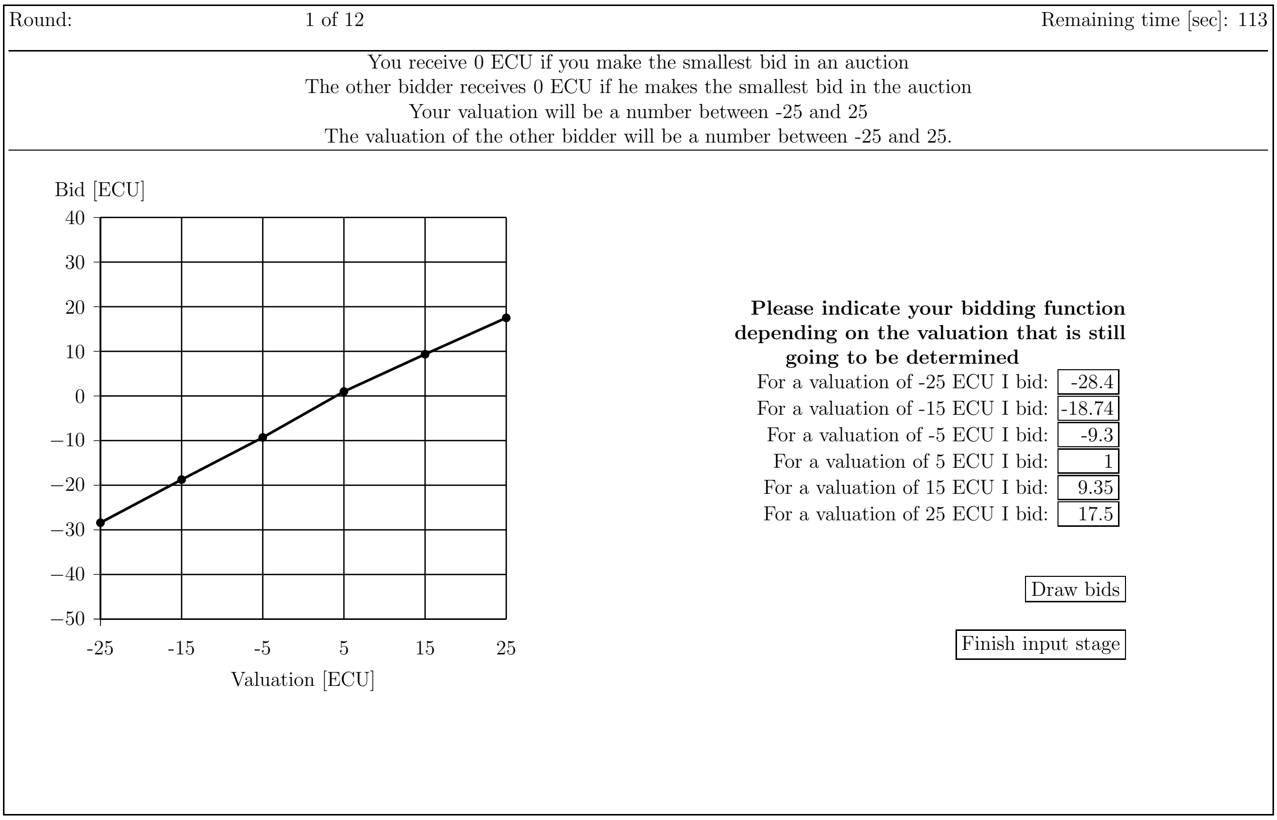

Hypothesis 5 (complexity favours markdown bidding) *We should find more absolute markdown bids in the* -25 *treatment than in the* 25 *or* 50 *or* $50+$ *treatments due to increased difficulty.*

While for first-price auctions absolute markdown bids differ substantially from Bayesian Nash equilibrium bids there is no such difference for second-price auctions. Underbidding for small valuations can be the result of absolute markdown bids in first-price auctions, but it should disappear (even with absolute markdown bids) in second-price auctions (treatment $50 + \mathbb{I}$).

Hypothesis 6 (no underbidding in second-price auctions) *There should be no significant amount of underbidding in the* $50 + \mathbb{I}$ *treatment.*

All experiments were conducted in the experimental laboratory of the SFB 504 in Mannheim. In total 304 subjects participated in these experiments. A detailed list of the treatments is given in appendix B, instructions are provided in appendix C. The experiments were computerised and we used the software package z-Tree (Fischbacher (2007)).

A typical input screen used in the experiments is shown in Figure 4 (translated into English). In each round participants enter bids for six valuations which are equally spaced between $\underline{\omega}$ and $\overline{\omega}$. Bids for all other valuations are interpolated linearly. Upon determination of bidding functions by all participants we draw five random and independent valuations for each participant. Each of these five random draws corresponds to an auction for which the winner is determined and the gain of each player is calculated. The sum of the gains obtained in these five auctions determines the total gain from this round.

Figure 5 A typical feedback screen in the experiment (translated into English)

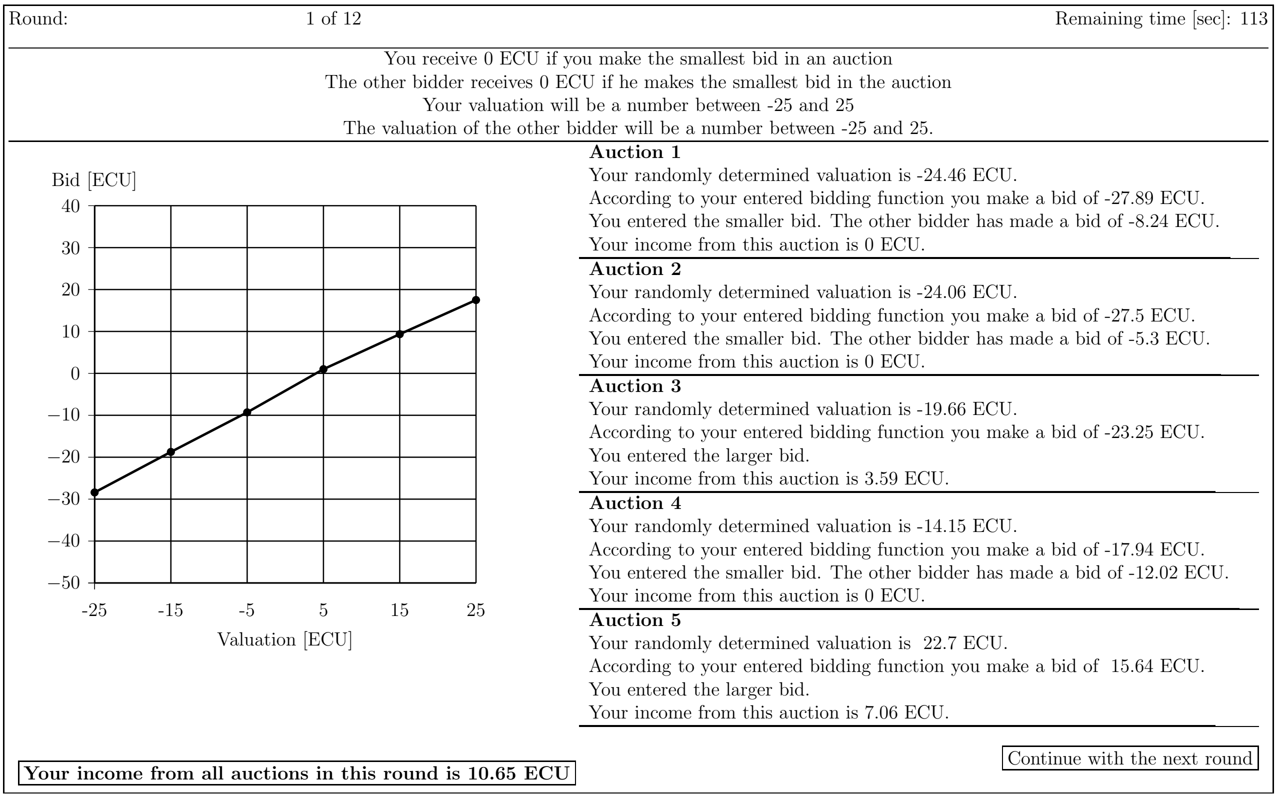

A typical feedback screen is shown in figure 5. Participants play 12 rounds. Each round consists of a bid input stage and a feedback stage. At the end of these 12 rounds participants complete a short questionnaire and are paid in cash according to their gains throughout in the experiment.

The strategy method has been used before in other auction experiments by Selten and Buchta (1999), Güth et al. (2003), Pezanis-Christou and Sadrieh (2003), Kirchkamp et al. (2009), and Kirchkamp and Reiß (2011). From our own experience with this method we know that bids that are observed with the strategy method are very similar to bids observed with alternative methods.

In the context of this paper we should note that the three benchmark solutions we described in sections 2.1, 2.2, and 2.3 can be represented as three different bidding functions which are all (almost) straight lines in the experimental interface. It is, however, up to the participants, whether they choose any of these three lines or any other curve.

4 Results

4.1 Convergence of bidding behaviour

Before we look at details of bidding behaviour we have to check whether behaviour stabilises over the course of the experiment. To do this we rely on three indicators of stability. First we count how often participants change support points of their bidding function. In each period and for each participant this can be a number between zero and six. It is zero if the participant continues to use the bidding function from the last period, and it is six if all bids are changed.

Figure 6 Convergence of bids over time

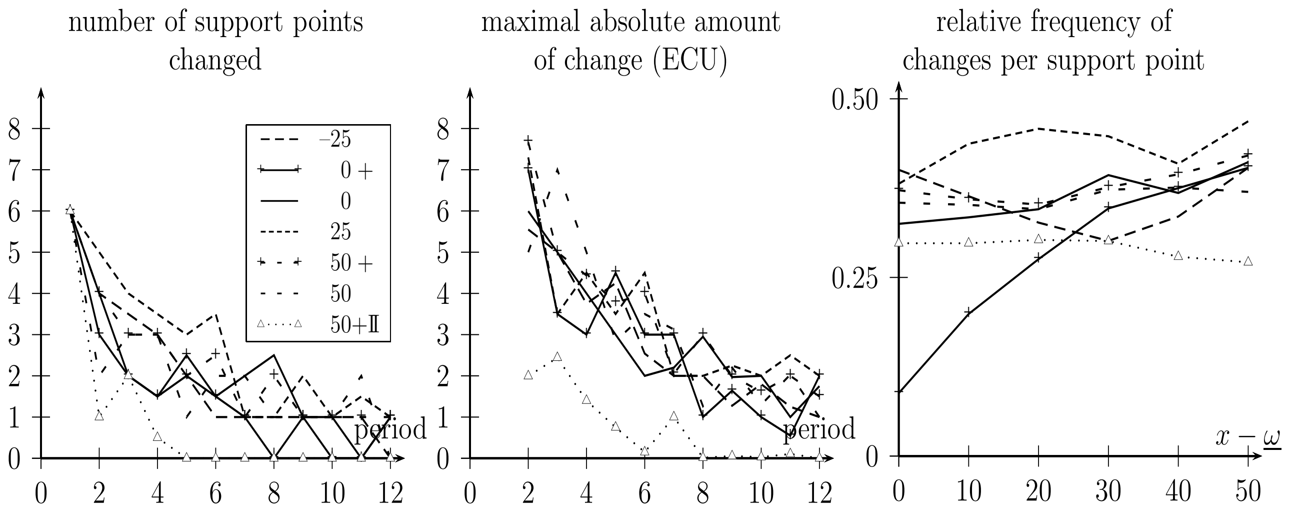

The figure shows how many of the six support points (hypothetical bids) of a bid function a median bidder changes and the maximal absolute amount of this change.

The development of the median of this distribution is reported in the left graph of figure 6. By definition all six support points are new in the first period, thus, period 1 must start with 6 changes for all treatments. We see that, after some adjustments during the first few periods, participants apply a stable bidding function. During the second half of the experiment the median bidder does not change more than one or two support points in each period.

Second, the graph in the middle panel of figure 6 shows the absolute amount of these changes over time. For each participant and each period we determine the largest absolute change in the six hypothetical bids from one period to the next. The median of this distribution is shown in the graph. We see that these changes are small compared to the range of the valuation.

Third, the right panel in the figure shows that changes are distributed fairly evenly over valuations for most treatments. The exception is the 0+ treatment where bidders are, indeed, restricted in their changes for small valuations.

We conclude here that bidding behaviour is stable in the second half of the experiment.

4.2 Visual inspection of aggregate bids

For a first impression of bidding behaviour, figure 7 shows the median and interquartile range[6] amount of overbidding $b(x) - \gamma^*(x)$ as a function of the valuation x. As above $\gamma^*(x)$ is the Bayesian Nash equilibrium bidding function with risk neutrality ($r = 1$) as given by equation (4). Let us briefly inspect the individual treatments:

Second-price auction: In the second-price treatment bidders have a weakly dominant bidding strategy. Many participants follow this strategy. Overbidding is zero for the 25% quantile and for the median bid. The 75% quantile is, for all valuations, larger than 0, i.e. there are some bidders which bid more than the weakly dominant bidding strategy. This is consistent with

[6]Median and quartiles are taken over all bidders and all periods (after period 6) in a given treatment.

Figure 7 Overbidding for different treatments

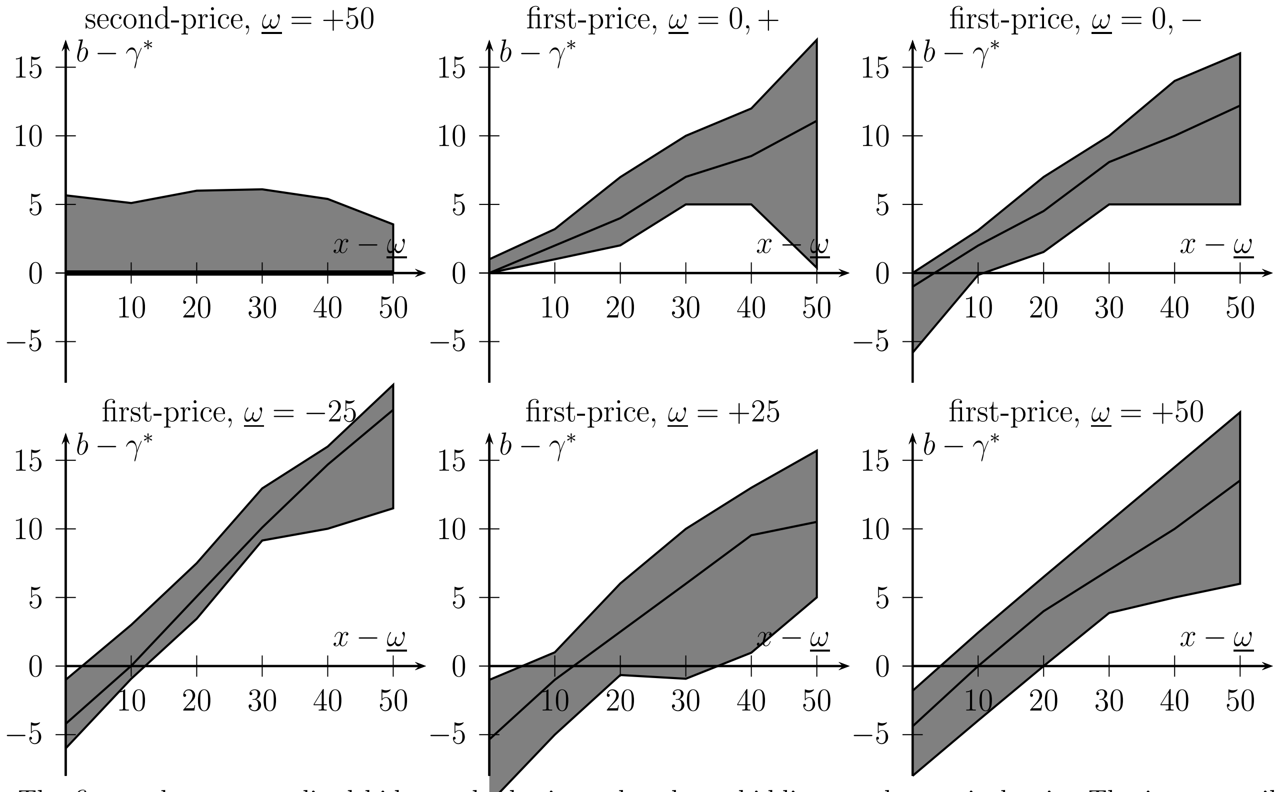

The figure shows normalised bids on the horizontal and overbidding on the vertical axis. The interquartile range of overbidding is shown as a grey area. The median amount of overbidding is shown as a black line. The first 6 periods from each session are discarded.

the experimental literature. Already Kagel et al. (1987) find a small amount of overbidding in second-price auctions. Kagel and Levin (1993) confirm that only a small fraction of bidders bid less than the equilibrium strategy while a substantial fraction bid more in second-price auctions.

First price auction, 0+: The traditional first-price treatment prevailing in the experimental literature is characterised by $\underline{\omega} = 0$ and '+', where the sign indicates the restriction to positive bids. The lowest possible valuation is 0, and bids are constrained to be positive. As we should expect we find overbidding in this treatment. Median overbidding and 75% quantile overbidding increase with the valuation. Except for the highest valuation also the 25% quantile increases with the valuation. What we see at the right end of the 0+ graph is a decrease in the amount of overbidding for the 25% quantile. The value of the bid is still increasing for these players, although the slope of the bidding function is now smaller than one. This finding is consistent with risk-aversion and confirms results from several previous experiments, starting with Cox et al. (1982).

First price auction, all other treatments: All the other treatments allow for bids that are smaller than the smallest possible valuation. Similar to the 0+ treatment, we find overbidding for high valuations. In contrast to the 0+ treatment, we find underbidding for low valuations.

Table 2 Overbidding for different treatments

treatment		n	$b(\underline{\omega}) - \gamma^*(\underline{\omega})$ mean	t	$P_{>t}$	P_{bin}	$b(\overline{\omega}) - \gamma^*(\overline{\omega})$ mean	t	$P_{>t}$	P_{bin}
-25	$-$	4	-4.653	-7.703	.0023	.0625	16.229	20.753	.0001	.0625
0	$-$	6	-3.453	-2.003	.0508	.0156	10.586	5.798	.0011	.0156
0	$+$	4	.841	2.009	.9309	1	9.397	14.154	.0004	.0625
25	$-$	3	-4.953	-5.574	.0154	.125	8.461	4.176	.0264	.125
50	$-$	4	-7.53	-4.643	.0094	.0625	6.639	2.806	.0338	.0625
50	$+$	14	-2.597	-1.704	.0561	.212	17.272	7.169	.0000	.0001
all firstprice		29	-4.856	-7.466	.0000	.0000	10.389	13.882	.0000	.0000
secondprice		6	3.549	3.98	.9947	.9844	27.485	36.084	.0000	.0156

The table compares overbidding for the lowest and highest valuation, $\underline{\omega}$ and $\overline{\omega}$. For each treatment n is the number of independent observations. We test whether $b(\underline{\omega}) < \gamma^*(\underline{\omega})$ (underbidding for low valuations) and whether $b(\overline{\omega}) > \gamma^*(\overline{\omega})$ (overbidding for high valuations). Mean deviations from the risk neutral Bayesian Nash equilibrium bid are shown together with results of a parametric t-test ($P_{>t}$) and a non-parametric binomial test (P_{bin}). As in many other experiments we find a significant amount of overbidding for the highest possible valuation $\overline{\omega}$ in all treatments. However, we also find underbidding for the smallest possible valuation $\underline{\omega}$ in all first-price treatments where underbidding is possible, i.e. always, except in the 0+ treatment.

4.3 Results of statistical tests of aggregate behaviour

Table 2 shows mean overbidding for the highest and lowest valuation, $\underline{\omega}$ and $\overline{\omega}$. Overbidding for high valuations is consistent with e.g. CRRA and several other theories of bidding behaviour in auctions that we mentioned above. Underbidding for low valuations is harder to explain. However, underbidding for low valuations is consistent with the absolute markdown bids presented in section 2. The table compares overbidding for the lowest and highest valuation, $\underline{\omega}$ and $\overline{\omega}$. For each treatment n is the number of independent observations. We test whether $b(\underline{\omega}) < \gamma^*(\underline{\omega})$ (underbidding for low valuation) and whether $b(\overline{\omega}) > \gamma^*(\overline{\omega})$ (overbidding for high valuation). Mean bids are shown together with results of a parametric t-test ($P_{>t}$) and a non-parametric binomial test (P_{bin}). To test whether the t-test on the level of the independent observations is appropriate we have applied a Shapiro-Wilk test on means over the independent observations. For none of the treatments the test rejects the assumption of normality.

Hypotheses 1 and 2: Not surprisingly, and as in many other experiments with first-price auctions, we find a significant amount of overbidding for the highest possible valuation $\overline{\omega}$ in all treatments. More interestingly, we find underbidding for the smallest possible valuation $\underline{\omega}$ in all first-price treatments where underbidding is possible, i.e. always, except for the 0+ treatment. We find, thus, no support for hypothesis 1 (pure Bayesian Nash equilibrium bidding), but we can confirm hypothesis 2 (partial markdown bidding).

Hypothesis 3: To see if bidding behaviour is affected by the exclusion of negative bids, we compare differences in bids for the smallest and the highest valuation $b(\overline{\omega}) - b(\underline{\omega})$ observed in the 0 treatment to those observed in the 0+ treatment. If the possibility to submit negative bids does not affect behaviour, observed differences in bids should equal zero. The first line in table 3 shows the results. In the 0 treatment the difference in mean bids exceeds that in

Table 3 Slope of bidding function $b(\overline{\omega}) - b(\underline{\omega})$

groups of treatments	independent observations	difference in means	t	$P_{>t}$	z	$P_{>z}$
$(0) - (0+)$	10	5.482	2.469	0.018	1.492	0.068
$(50) - (50+)$	12	-2.958	-1.000	0.831	-1.019	0.846
$(-25) - (25, 50, 50+)$	19	5.059	3.657	0.001	2.500	0.006

The figure compares average slopes of the bidding functions measured as $b(\overline{\omega}) - b(\underline{\omega})$. The column "difference in means" shows the difference between the average slope in the first treatment group and the average slope in the second treatment group. The columns t and $P_{>t}$ show the result of a parametric t-test, the columns z and $P_{>z}$ show the result of a non-parametric Mann-Whitney test.

the 0+ treatment by 5.482 ECU. This difference is significant under a parametric test, thus, we can confirm hypothesis 3 (suppression of markdown bidding).

Hypothesis 4: The next line of table 3 shows the difference in slopes of the bidding function between the 50 and 50+ treatment. This difference is not significant. Furthermore, not even the sign of the effect is the same as what we expect under a strong treatment effect of explicitly allowing negative bids. Excluding negative bids (and mentioning this fact in the instructions to the experiment) does not per-se trigger behavioural changes towards absolute markdown bids, so that we cannot confirm hypothesis 4 (strong negative bids exclusion effect).

Hypothesis 5: The third line of table 3 shows the difference in slopes of the bidding function between the -25 treatment and the 25, 50, and 50+ treatment. According to hypothesis 5 (complexity favours markdown bidding) we should expect a steeper slope of the bidding function under the -25 treatment. This is confirmed by a parametric and a non-parametric test.

Hypothesis 6: Here we have to go back to table 2. The last line shows the difference between bids in the experiment and equilibrium bids. We see that this difference is positive, not negative. Hence, we do not observe a significant amount of underbidding which supports hypothesis 6 (no underbidding in second-price auctions).

4.4 Individual bids

The quartiles of bidding behaviour, as depicted in the left panel of figure 7, suggest heterogeneity among bidders. To better understand individual bidding behaviour we estimate for each bidder a linear bidding function:

$$b(x) = \underline{\omega} - \alpha + \beta \cdot (x - \underline{\omega}) + u \tag{9}$$

The regression specification normalises valuations and bids such that the point $(\underline{\omega}, \underline{\omega})$ transforms to the origin $(0, 0)$ in valuation-bid space. We normalise to facilitate the comparison of estimated intercept α (as the markdown amount) across treatments with different valuation domains. Estimated intercepts and slopes can be interpreted as if the valuation domain is $[0, 50]$ for any treatment. Again we discarded the first six periods of the experiment. Outliers have been eliminated using Hadi's method. The fit of the estimations of equation (9) is very good, e.g., the median R^2 is 0.9918.

Figure 8 Individual bidding

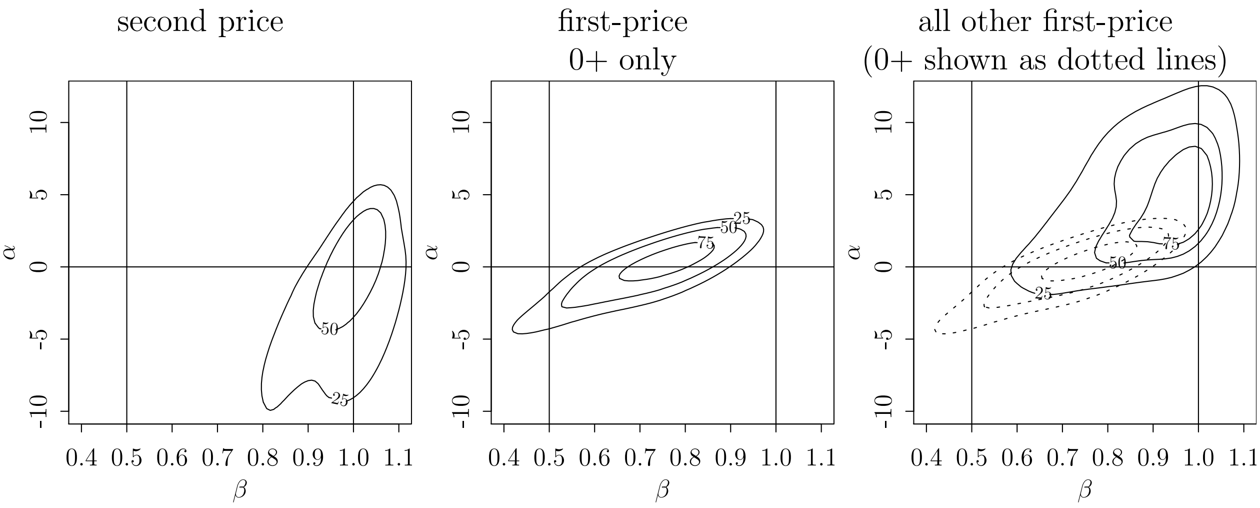

Each graph shows contour lines of the kernel density estimate of the distribution of individual bidding functions (see equation 9) with the first six periods of the experiment discarded. Numbers next to the contour lines are estimated percentiles.

Since a scatterplot of all individually estimated coefficients is rather confusing, figure 8 shows the contour lines of the estimated joint distribution of α and β. We aggregate the data in three graphs:

Second price auction: The left panel of figure 8 shows the distribution of estimated bidding functions for second price auctions. In the weakly dominant equilibrium we have $\alpha = 0$ and $\beta = 1$. Indeed, the distribution of estimated values is nicely centered around this value.

First price auction, 0+**:** The middle panel depicts coefficient estimates for the 0+ treatment. The risk neutral Bayesian Nash equilibrium predicts $\alpha = 0$ and $\beta = 1/2$. Risk averse equilibrium predicts $\alpha = 0$ and $\beta > 1/2$. As the figure illustrates, the distribution of coefficients is concentrated around $\alpha = 0$ and its support includes values for β between 1/2 and 1. In this treatment risk averse Bayesian Nash equilibrium and other theories that we mentioned above explain the data quite well.

First price auction, all other treatments: The right panel of figure 8 illustrates coefficient estimates for all the other treatments. To facilitate the comparison we also indicate the distribution estimated for the 0+ treatment as dotted lines. To further facilitate the discussion, figure 9 shows the same graph with additional lines and labels. First consider the right panel of figure 8. It is easy to see that the possibility to make bids smaller than the lower bound of the valuation domain, $\underline{\omega}$, changes bidding behaviour substantially:

- Now a large group of bidders is characterised by $\beta \approx 1$ and a substantial markdown amounts $\alpha > 1$. These bidders are, in line with section 2.2, much better described by absolute markdown bids instead of following the Bayesian Nash equilibrium bidding.

Figure 9 Categorisation of bidders

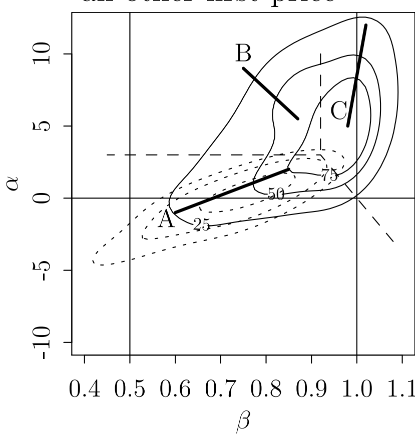

relative frequency of bidders [%]

treatment	A	B	C
−25	31.2	12.5	56.2
0+	93.8	3.1	3.1
0	65	20	15
25	37.4	26.3	36.3
50	30.1	28.6	41.3
50+	25.6	32.6	41.9
all except 0+	32.8	28.2	39

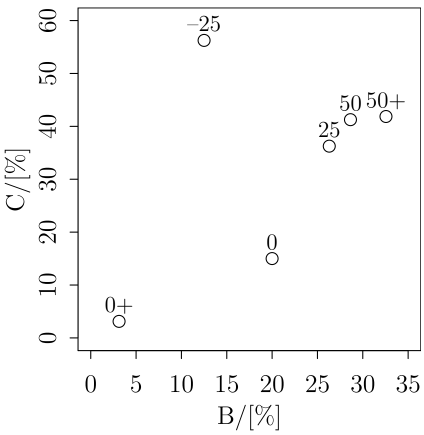

- There is still some overlap with the distribution estimated for the 0+ treatment. This region could characterise bidders that are not affected by our treatment conditions and always bid according to the risk averse Bayesian Nash equilibrium (see section 2.1), or are driven by motives like spite or regret.

- Finally, there is a group of bidders with $\beta < 1$ but still a positive markdown amount α. In the context of section 2.3 we interpret these bidders as rational bidders who realise that not all bidders are perfectly rational. Alternatively they might be viewed as markdown bidders.

4.5 Categorising individual bidders

When interpreting the contour lines of figure 8 as contour lines of a mountain, the mountain in the graph on the right has three ridges. In the left panel of figure 9 we indicate these ridges approximately with bold lines and letters A, B, and C. These lines are merely intended to clarify the discussion. One could certainly draw them in a slightly different manner.

Ridge A runs parallel to the distribution of bidding functions in the 0+ treatment. We categorise bidders close to this ridge as rational bidders who believe that they are in a rational world.

Ridge C is close to $\beta = 1$. All bidding functions on this ridge have an $-\alpha < 0$. We categorise bidders close to this ridge as bidders who use absolute markdown bids and who expect their opponents to do the same.

Ridge B is in between (and less pronounced than A or C). As above, we interpret these bidders as rational bidders who realise that not all bidders are perfectly rational.

We now categorise bidders into different groups. The dashed lines indicate the borders of the regions that we use. Clearly, there is some degree of arbitrariness in choosing threshold values for α and β. We have tried to find a region for the Bayesian Nash (A) bidders which includes almost all of the 0+ bidders. A suitable threshold seemed to be $\alpha = 3$. We have also tried to make region (C) symmetric to $\beta = 1$. This led to a threshold of $\beta = 0.92$.

The table in the center of figure 9 shows the result of this categorisation. The graph on the right hand side of the same figure depicts the sample shares of bidder categories B and C for all treatments. For the 0+ treatment the shares of B and C are negligible. However, for the 0 treatment shares of B and C increase substantially. We attribute this to the possibility to submit negative bids so that markdown bidding is feasible as opposed to the 0+ treatment. While for the remaining four treatments most players are either Bs or Cs, there is still a fraction of about 25% of Bayesian Nash players.

We are now ready to examine our hypotheses at the individual level. Again, we reject hypothesis 1 (pure Bayesian Nash equilibrium bidding) and support hypothesis 2 (partial markdown bidding), except for the 0+ treatment. We also find clear support for hypothesis 3 (suppression of markdown bidding), since there are more Cs in the 0 treatment than in the 0+ treatment. In line with the aggregate data, we cannot support hypothesis 4 (strong negative bids exclusion effect); there is essentially no difference between the 50 and the 50+ treatment. Hypothesis 5 (complexity favours markdown bidding) is supported—the largest share of Cs is found in the −25 treatment.

5 Concluding remarks

Many first-price auction experiments find that subjects bid more than the risk neutral equilibrium bid, they 'overbid'. We can confirm this finding. However, the approaches that have been used so far to explain overbidding are not in line with our second finding: underbidding for small valuations.

The idea we are proposing here, namely that some bidders use absolute markdown bids, is independent of the representation of payoffs as lottery tickets or as money and consistent with the traditional experimental evidence. We have seen in section 2 that optimal absolute markdown bids imply underbidding for small valuations and that the presence of a small proportion of bidders with absolute markdown bids is sufficient to make rational bidders behave as if they were constrained in a similar way.

In our experiment we find that a fraction of bidders can still be described well with traditional models, e.g. risk averse Bayesian Nash equilibrium. In many treatments this is a small fraction. We have seen that a substantially larger fraction of bidders follows absolute markdown bids. A third group, finally, behaves like optimisers against such a mixed population.

We can, hence, not conclude that *all* bidders can be described better with absolute markdown bids. We also cannot say much about the importance of markdown bids in other, perhaps more complicated contexts. We have, however, seen that in our experiment absolute markdown bids become more prominent when the environment becomes more complicated (as in the −25 treatment). Hence, when we find already a fair amount of markdown bids in a very simple auction, we might suspect that markdown bids play an even more important role in more complex auctions.

References

Andreas Blume and Uri Gneezy. Cognitive forward induction and coordination without common knowledge: An experimental study. *Games and Economic Behavior*, 68(2):488–511, 2010.

Antoni Bosch-Domenech, Jose G. Montalvo, Rosemarie Nagel, and Albert Satorra. One, two, (three), infinity,... and lab beauty-contest experiments. *The American Economic Review*, 92(5):1687–1701, 2002.

Kay-Yut Chen and Charles R. Plott. Nonlinear behavior in sealed bid first price auctions. *Games and Economic Behavior*, 25:34–78, 1998.

James C. Cox, B. Roberson, and Vernon L. Smith. Theory and behavior of single object auctions. In Vernon L. Smith, editor, *Research in experimental economics*. JAI Press, Greenwich, Conn., 1982.

James C. Cox, Vernon L. Smith, and James M. Walker. Test of a heterogeneous bidder's theory of first price auctions. *Economic Letters*, 12(3-4):207–212, 1983.

James C. Cox, Vernon L. Smith, and James M. Walker. Experimental development of sealed-bid auction theory: Calibrating controls for risk aversion. *American Economic Review*, 75 (2):160–165, 1985.

James C. Cox, Vernon L. Smith, and James M. Walker. Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, 1:61–99, 1988.

Vincent P. Crawford and Nagore Iriberri. Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770, November 2007.

Urs Fischbacher. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007. URL `http://ideas.repec.org/a/kap/expeco/v10y2007i2p171-178.html`.

Werner Güth, Radosveta Ivanova-Stenzel, Manfred Königstein, and Martin Strobel. Learning to bid—an experimental study of bid function adjustments in auctions and fair division games. *The Economic Journal*, 113(487):477–494, 2003.

Radosveta Ivanova-Stenzel and Doron Sonsino. Comparative study of one bid versus two bid auctions. *Journal of Economic Behavior and Organization*, 54(4):561–583, 2004.

John H. Kagel and Dan Levin. Independent private value auctions: Bidder behavior in first-, second- and third-price auctions with varying numbers of bidders. *The Economic Journal*, 103:868–879, 1993.

John H. Kagel, R. M. Harstad, and Dan Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55:1275–1304, 1987.

Oliver Kirchkamp and J. Philipp Reiß. Out-of-equilibrium bids in first-price auctions: Wrong expectations or wrong bids. *The Economic Journal*, 121(157):1361–1397, 2011.

Oliver Kirchkamp, Eva Poen, and J. Philipp Reiß. Outside options: Another reason to choose the first-price auction. *European Economic Review*, 53(2):153–169, 2009.

Paul Pezanis-Christou and Abdolkarim Sadrieh. Elicited bid functions in (a)symmetric first-price auctions. Discussion Paper 2003-58, CentER, Tilburg University, 2003.

Rene Saran. Bilateral trading with naive traders. *Games and Economic Behavior*, 72(2): 544–557, 2011.

Reinhard Selten and Joachim Buchta. Experimental sealed bid first price auctions with directly observed bid functions. In D. Bodescu, I. Erev, and R. Zwick, editors, *Games and Human Behaviour*, pages 79–102. Lawrence Erlbaum Aussociates Inc., Mahwah (NJ), 1999.

A Derivation of the best-response function $\alpha_i^*(\alpha_j)$ with markdown bids as given by (6)

With the uniform distribution, $f(x_i) = f(x_j) = 1$ for $x_i, x_j \in [0,1]$, and CRRA utility, the expected utility function simplifies to

$$\mathrm{EU}_i(\alpha_i) = \begin{cases} [1 - \frac{1}{2}(\alpha_i - \alpha_j + 1)^2]\,\alpha_i^r & \text{if } \max\{0, \alpha_j - 1\} \le \alpha_i \le \alpha_j, \\ \frac{1}{2}(\alpha_j - \alpha_i + 1)^2\,\alpha_i^r & \text{if } \alpha_j \le \alpha_i \le \alpha_j + 1. \end{cases}$$

To ease the exposition, define auxiliary functions $h_-(z)$ and $h_+(z)$ corresponding to the two cases of the expected utility function as follows:

$$h_-(z) = [1 - \frac{1}{2}(z - \alpha_j + 1)^2]\,z^r,$$
$$h_+(z) = \frac{1}{2}(\alpha_j - z + 1)^2\,z^r.$$

Case I: $h_-(z)$ is maximised on the interval $[\underline{z}, \alpha_j]$ at $z_-^* = \min\{z_1,\, \alpha_j\}$ where z_1 is defined further below. The first derivative of $h_-(z)$ is

$$h_-'(z) = \frac{r\,z^{r-1}}{2}\left[2 - (z - \alpha_j + 1)^2 - \frac{2z}{r}(z - \alpha_j + 1)\right]$$

and exhibits two non-zero roots $z_1, z_2 \neq 0$ such that

$$z_{1,2} = \frac{(1+r)(\alpha_j - 1)}{2+r} \pm \frac{1}{2+r}\sqrt{2r(2+r) + (\alpha_j - 1)^2}.$$

It is straightforward to show that $\underline{z} < z_1 < \alpha_j - 1 + \sqrt{2}$. There exists $z' \in (\underline{z}, z_1)$ such that $h_-'(z') > 0$, further, $h_-'(z_1) = 0$, and $h_-'(\alpha_j - 1 + \sqrt{2}) < 0$. Since $z_2 < \underline{z}$, z_1 is the unique root of $h_-'(z)$ for $z > \underline{z}$ so that, with continuous differentiability of $h_-(z)$ for $z > \underline{z}$ and continuity of $h_-(z)$ for $z \geq \underline{z}$, $h_-(z_1)$ is the maximum on interval $[\underline{z}, \alpha_j - 1 + \sqrt{2}]$. Therefore, $z_-^* = z_1$ is the maximiser on interval $[\underline{z}, \alpha_j]$ for $z_1 \le \alpha_j$ and $z_-^* = \alpha_j$ emerges as the boundary solution for $z_1 > \alpha_j$.

Case II: $h_+(z)$ is maximised on $[\alpha_j,\, \alpha_j + 1]$ at $z_+^* = \max\{z_4, \alpha_j\}$ where z_4 is defined further below. The first derivative of $h_+(z)$ is

$$h_+'(z) = \frac{1}{2}\,z^{r-1}\,(\alpha_j - z + 1)\,[r\,(\alpha_j - z + 1) - 2z]$$

and exhibits two non-zero roots: $z_3 = \alpha_j + 1$ and $z_4 = r(\alpha_j + 1)/(2 + r)$. Since $h_+(z_3) = 0$ while $h_+(z) > 0$ for $z \in [\alpha_j,\, \alpha_j + 1)$, z_3 identifies a minimum. It is obvious that $0 < z_4 < \alpha_j + 1$. There exists $z' \in (0, z_4)$ such that $h_+'(z') > 0$, further, $h_+'(z_4) = 0$, and there exists $z'' \in (z_4, \alpha_j + 1)$ such that $h_+'(z'') < 0$. Since z_4 is the unique root of $h_+'(z)$ for $0 < z < \alpha_j + 1$, with continuous differentiability of $h_+(z)$ for $0 < z < \alpha_j + 1$ and continuity of $h_+(z)$ for $z \geq 0$, $h_+(z_4)$ is the maximum on interval $[0, \alpha_j + 1]$. For $\alpha_j \le r/2$, $z_4 \geq \alpha_j$, hence, $z_+^* = z_4$ is the maximiser of $h_+(z)$ on interval $[\alpha_j, \alpha_j + 1]$ for $\alpha_j \le r/2$. Further, $z_+^* = \alpha_j$ for $\alpha_j > r/2$ since then $z_4 < \alpha_j$.

By $h_-(\alpha_j) = h_+(\alpha_j)$ and, for $\alpha_j > 0$, $h_-'(\alpha_j) = h_+'(\alpha_j) = \alpha_j^{r-1}(r - 2\alpha_j)/2$, expected utility $\mathrm{EU}_i(\alpha_i)$ is maximised (i) by z_1 for $\alpha_j > r/2$, (ii) by z_1 and z_4 for $\alpha_j = r/2$ (implying $z_1 = z_4$), and (iii) by z_4 for $0 < \alpha_j < r/2$. The comparison of $h_-(z_-^* = 0) = 0$ and $h_+(z_+^* = z_4) > 0$ implies that $\mathrm{EU}_i(\alpha_i)$ is maximised by z_4 for $\alpha_j = 0$. The best-response function $\alpha_i^*(\alpha_j)$ as given by (6) follows immediately.

B List of independent observations

session and matching group	group	$\underline{\omega}$	$\underline{b}$	second-price	participants
A1	-25	-125	0	8	
A2	-25	-125	0	8	
B1	-25	-125	0	8	
B2	-25	-125	0	8	
C1	0	-100	0	8	
C2	0	-100	0	6	
D1	0	-100	0	6	
D2	0	-100	0	6	
E1	0	-100	0	8	
E2	0	-100	0	6	
F1	0	0	0	8	
F2	0	0	0	8	
G1	0	0	0	8	
G2	0	0	0	8	
H1	25	-75	0	8	
H2	25	-75	0	8	
I1	25	-75	0	10	
J1	50	-50	0	8	
J2	50	-50	0	6	
K1	50	-50	0	8	
K2	50	-50	0	8	
L1	50	0	0	14	
M1	50	0	0	14	
N1	50	0	0	8	
N2	50	0	0	10	
O1	50	0	0	10	
O2	50	0	0	10	
P1	50	0	0	10	
P2	50	0	0	10	
Q1	50	0	1	10	
Q2	50	0	1	10	
R1	50	0	1	10	
R2	50	0	1	10	
S1	50	0	1	10	
S2	50	0	1	8	

The parameter $\underline{b}$ is the smallest possible bid. In the +treatments $\underline{b} = 0$, otherwise $\underline{b} = \underline{\omega} - 100$. The highest bid that participants could enter was always $\overline{\omega} + 100$.

C Conducting the experiment and instructions

Participants were recruited by email and could register for the experiment on the internet. At the beginning of the experiment participants drew balls from an urn to determine their allocation to seats. Being seated participants then obtained written instructions in German. These instructions very slightly depending on the treatment. In the following we give a translation of the instructions.

After answering control questions on the screen subjects entered the treatment described in the instructions. After completing the treatment they answered a short questionnaire on the screen and where then paid in cash. The experiment was done with the help of z-Tree (Fischbacher (2007)).

C.1 General information

You are participating in a scientific experiment that is sponsored by the Deutsche Forschungsgemeinschaft (German Research Foundation). If you read the following instructions carefully then you can—depending on your decision—gain a considerable amount of money. It is, hence, very important that you read the instructions carefully.

The instructions that you have received are only for your private information. **During the experiment no communication is permitted.** Whenever you have questions, please raise your hand. We then answer your question at your seat. Not following this rule leads to exclusion from the the experiment and all payments.

During the experiment we are not talking about Euro, but about ECU (Experimental Currency Unit). Your entire income will first be determined in ECU. The total amount of ECU that you have obtained during the experiment will be converted into Euro at the end and paid to you in **cash**. The conversion rate will be shown on your screen at the beginning of the experiment.

C.2 Information regarding the experiment

Today you are participating in an experiment on auctions. The experiment is divided into separate rounds. We will conduct **12 rounds**. In the following we explain what happens in each round.

In each round you bid for an object that is being auctioned. Together with you another participant is also bidding for the same object. Hence, in each round, there are **two bidders**. In each round you will be allocated randomly to another participant for the auction. *Your co-bidder in the auction changes in every round.* The bidder with the highest bid has obtained the object. If bids are the same the object will be allocated randomly.

For the auctioned object you have a valuation in ECU. This valuation lies between x and $x + 50$ ECU[7] and is determined randomly in each round. The range from x to $x + 50$ will be shown to you at the beginning of the experiment on the screen and is the same in each round.[8] **From this range you will obtain in each round new and random valuations for the object.** The other bidder in the auction also has a valuation for the object. The

[7]In the 0+ and 50+ treatments the valuation would be announced precisely: "This valuation lies between 0 and 50 ECU" in the 0+ treatment and "This valuation lies between 50 and 100 ECU" in the 50+ treatment. Whenever x is mentioned in the remainder of the instruction the same comment applies: In the 0+ and 50+ treatments the valuation is always announced precisely.

[8]This sentence was not shown in the 0+ and 50+ treatments, though in all treatments the range was shown on the screen.

valuation that the other bidder attributes to the object is determined by the same rules as your valuation and changes in each round, too. All possible valuations of the other bidder are also in the interval from x to $x+50$ from which also your valuations are drawn. All valuations between x and $x+50$ are equally probable. Your valuations and those of the other player are determined independently. You will be told your valuation in each round. *You will not know the valuation of the other bidder.*

C.2.1 Experimental procedure

The experimental procedure is the same in each round and will be described in the following. Each round in the experiment has two stages.

1. Stage

In the first stage of the experiment you see the following screen:[9]

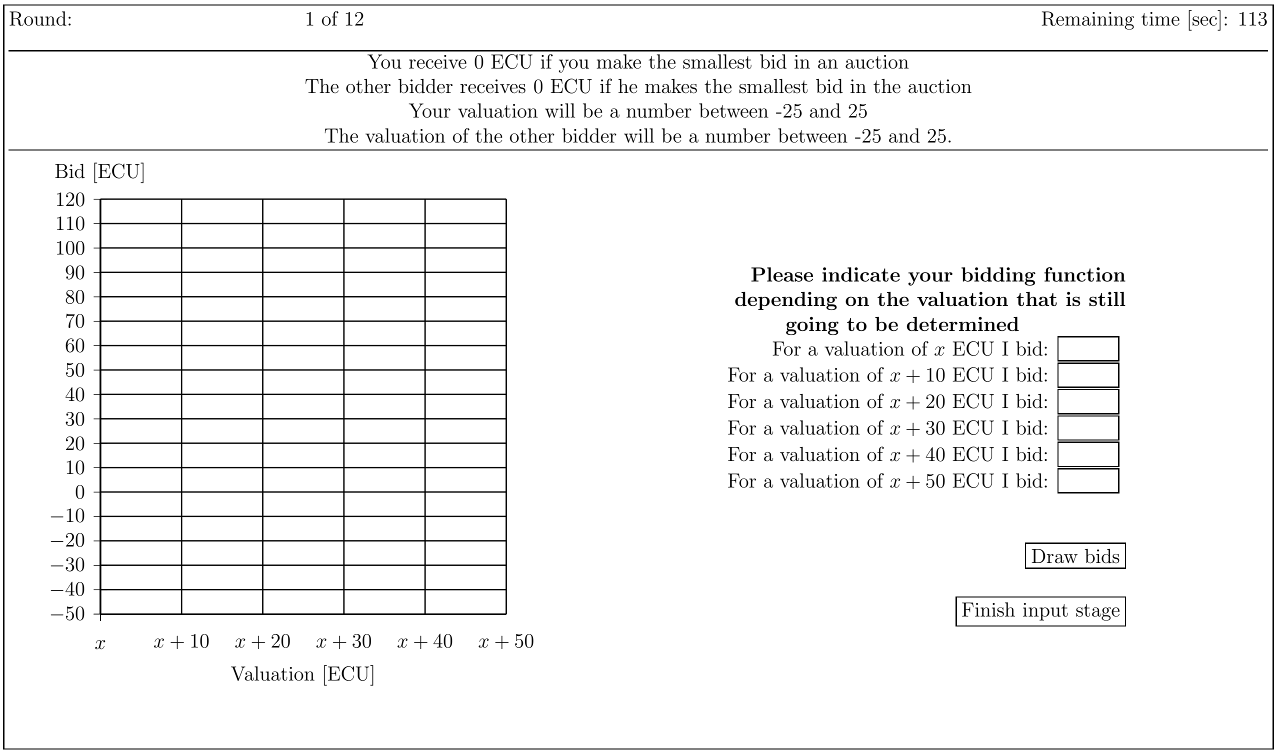

At that stage **you do not know your own valuation for the object in this round.** On the right side of the screen you are asked to enter a bid for **six hypothetical valuations** that you might have for the object. These six hypothetical valuations are x, $x+10$, $x+20$, $x+30$, $x+40$, and $x+50$ ECU. Your input into this table will be shown in the graph on the left side of the screen when you click on "draw bids". In the graph the hypothetical valuation is shown on the horizontal axis, the bids are shown on the vertical axis. Your input in the

[9]In the 0+ and 50+ treatments the interval was already shown exactly in the instructions and consistently also in the figures in the instructions. In the other treatments the interval x to $x+50$ was, as you see in the figure, described as x to $x+50$. From the first round of the experiment on the current numbers were given.

table is shown as six points in the diagram. **Neighbouring points are connected with a line automatically.** These lines determine your bid for all valuations *between* the six points for those you have made an input. For the other bidder the screen in the first stage looks the same and there are as well bids for six hypothetical valuations. The other bidder cannot see your input.

2. Stage

The actual auction takes place in the second stage of each round. In each round we will play not only a single auction but **five auctions**. This is done as follows: **Five times a random valuation is determined** that you have for the object. Similarly for the other bidder five random valuations are determined. You see the following screen:[10]

Round: 1 of 12 Remaining time [sec]: 113

You receive 0 ECU if you make the smallest bid in an auction
The other bidder receives 0 ECU if he makes the smallest bid in the auction
Your valuation will be a number between -25 and 25
The valuation of the other bidder will be a number between -25 and 25.

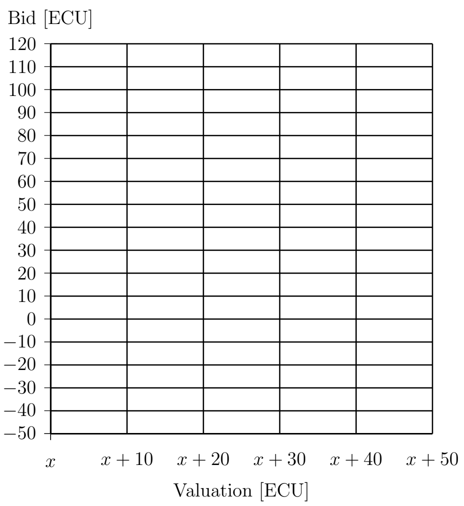

Auction 1:
Your randomly determined valuation is . . . ECU.
According to your entered bidding function you make a bid of . . . ECU.
You entered the larger bid.
Your income from this auction is . . . ECU.

Auction 2:
Your randomly determined valuation is . . . ECU.
According to your entered bidding function you make a bid of . . . ECU.
You entered the larger bid.
Your income from this auction is . . . ECU.

Auction 3:
Your randomly determined valuation is . . . ECU.
According to your entered bidding function you make a bid of . . . ECU.
You entered the smaller bid. The other bidder has made a bid of . . . ECU.
Your income from this auction is . . . ECU.

Auction 4:
Your randomly determined valuation is . . . ECU.
According to your entered bidding function you make a bid of . . . ECU.
You entered the smaller bid. The other bidder has made a bid of . . . ECU.
Your income from this auction is . . . ECU.

Auction 5:
Your randomly determined valuation is . . . ECU.
According to your entered bidding function you make a bid of . . . ECU.
You entered the larger bid.
Your income from this auction is . . . ECU.

Your income from all auctions in this round is . . . ECU

Continue with the next round

For each of your five valuations the computer determines your bid according to the graph from stage 1. If a valuation is precisely at x, $x + 10$, $x + 20$, $x + 30$, $x + 40$, or $x + 50$ the computer takes the bid that you gave for this valuation. If a valuation is between these points your bid is determined according to the joining line. In the same way the bids of the other bidder are determined for his five valuations. Your bid is compared with the one of the other bidder. The bidder with the higher bid has obtained the object.

Your income from the auction:

For each of the five auctions the following holds:

[10] In the instructions the following figure was shown. This figure does not show the bidding function in the graph and the specific bids, gains and losses that would be shown during the experiment.

- The bidder with the higher bid obtains the valuation he had for the object in this auction added to his account minus his bid for the object.
- If the bidder with the higher bid has a negative valuation for the object, the ECU account is reduced by this amount.[11]
- If the bid of bidder with the higher is a negative number, the amount is added to his ECU account.[12]
- The bidder with the smaller bid obtains **no income** from this auction.

You total income in a round is **the sum of the ECU income from those auctions in this round where you have made the higher bid.**

This ends one round of the experiment and you see in the next round again the input screen from stage 1.

At the end of the experiment your total ECU income from all rounds will be converted into Euro and paid to you in cash together with your Show-Up Fee of 3.00 Euro.

Please raise your hand if you have questions.

[11] This item is not shown in the 0+ and 50+ treatments.

Note that, in order to be able to use same instructions for all treatments we mention the possibility of negative valuations in all, except the 0+ and 50+ treatments, even if subjects learn later that their valuation is drawn from an interval that contains only positive numbers.

[12] This item not shown in the 0+ and 50+ treatments.