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

Oliver Kirchkamp† J. Philipp Reiß‡

3 April 2019

We present results from a series of experiments that allow us to measure overbidding and, in particular, underbidding in first-price auctions. We investigate the extent to which the amount of underbidding depends on the seemingly innocuous 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 described by Bayesian Nash equilibrium bids, a larger fraction seems either to use constant markdown bids or 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 a feature of bidding behaviour in first-price auction experiments with private values that has received little attention: bidding less than the risk-neutral Bayesian Nash equilibrium (RNBNE) for low valuations. We refer to this as underbidding. This feature may have gone unnoticed because underbidding is difficult to observe in standard experimental setups; in addition, underbidding is hard to reconcile with several established theories.

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

†Universität 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, Schlossbezirk 12, 76131 Karlsruhe, +49-721-60848681, Fax +49-721-60848682, philipp.reiss@kit.edu

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

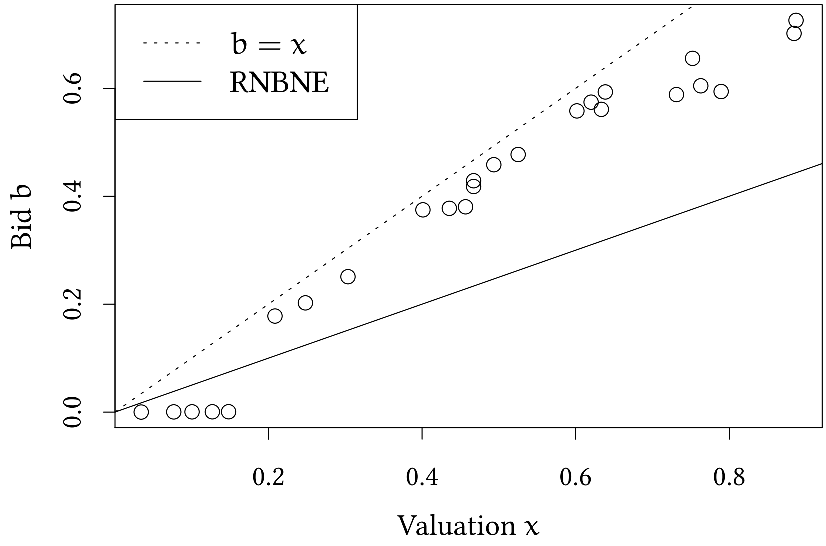

(page 83, Figure 8, series 4, exp. 3, n=4, subject 2)

In this paper, we present a method that makes it possible to observe overbidding and underbidding in first-price auctions. We find that the amount of underbidding depends on the seemingly innocent parameters of the experimental setup such as the range of valuations and the restriction that permits only positive bids. 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 bidders. For each participant, valuations are drawn from a uniform distribution. Figure 1 depicts normalised bidding data for a specific subject.[1]

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

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 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.[2]

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 appears to be an

[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 Figure 1 the valuations are normalised to $[0, 1]$ and bids are normalised correspondingly.

[2] Another explanation for bids of zero or close to zero is that bidders, who cannot opt out of participating in the auction, may want to document their unwillingness to participate in the auction, because the probability of winning with low valuations is deemed negligible.

innocent simplification. In this paper, we show that this simplification implies strong behavioural effects. In addition, bids for small valuations seldom win and might be perceived as less important for bidders. Hence, it might be harder to observe them precisely. In this paper, we use a variant of the strategy method that makes it easier to observe these bids.

We present a simple theory of heterogeneous bidders—who differ in the degree of rationality—that is supported by our data. The previous literature demonstrates heterogeneity in subjects' behaviour in many experimental settings, including auction experiments. Thus, approximating bidding behaviour by a single bidding function with a fixed functional form that describes all bids reasonably well may be too demanding. Instead, we propose three types of bidders that can be linked to different degrees of rationality. We know from experiments with other games that decision makers apply different levels of reasoning when 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. Apart from rationality, another important factor that might influence bids is the attitude towards risk. We will allow for different degrees of risk aversion in the theoretical model. However, the experiment does not focus on risk. Participants will be sampled from the same pool so that their risk attitude should be on average the same in all treatments.

Section 2.1 reviews the model of an extreme case: a rational bidder who assumes that the opponent is also rational. In Section 2.2 we study the opposite extreme: a bidder who is restricted to using absolute markdown bids and assumes that the opponent obeys the same restriction. The idea of boundedly rational bidders is not new. Kagel et al. (1987) 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, on the aggregate level, CRRA, which, in the context studied by Chen and Plott (1998), implies non-linear bidding functions, 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, and if some of them follow simple rules, while other use more complex, non-linear rules, then the aggregate behaviour might be better described by a non-linear rule. In our experiment, we want to examine whether some bidders systematically do something else. As a natural next step, we consider a bidder who is rational but assumes a mix of rational and restricted competitors in Section 2.3.

Section 3 describes the experiment, and Section 4 presents the results. Section 5 concludes. Anticipating our result, we classify bidders into groups, similar to those outlined above. If the experimental setup restricts underbidding, bidding seems to follow a (risk averse) Bayesian Nash equilibrium. Once underbidding is unrestricted in the experiment, some decision makers still use Bayesian Nash equilibrium bids, 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 contexts that differ in the composition of rational and boundedly rational bidders within the population. 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, which 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. 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 $r = 1$, and a risk-averse individual is described by $r < 1$. We confine our attention to the case of $r \in (0, 1]$. The derivation of the symmetric increasing bidding function $\gamma(x)$ is standard, see e.g., Krishna (2010). The unique equilibrium bidding function is given by

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

If valuations are drawn from the interval $[\underline{\omega}, \overline{\omega}]$ instead of $[0, 1]$, the equilibrium bid is

$$\gamma^*(x) - \underline{\omega} = \frac{x - \underline{\omega}}{1+r}\,. \tag{2}$$

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 absolute markdown bids

In the Bayesian Nash equilibrium of the first-price auction, bidders 'shade their bids' proportionally to $x - \underline{\omega}$. However, in a post-experimental questionnaire of another first-price auction experiment, some participants explained that they shade their bids not by a relative but, instead, by a constant amount.[3] More broadly, there may be various reasons for shading by a constant amount:

- It may be cognitively too difficult to work out the exact form of equation (2) 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 in the event of 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 a principal first has to define a bidding rule (before the individual valuation is revealed) and then the agent who follows this rule learns the valuation, it might be simpler for the principal to require a fixed amount that the agent is supposed to gain from each trade.

[3] Kirchkamp et al. (2009).

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

$$\bar{\gamma}_i(x) = x - \alpha_i \tag{3}$$

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 in behavioural bidding as described previously but choose their amounts of autonomous bid shading α_i simultaneously prior to 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, the Bayesian Nash equilibrium with rational bidders as reported in Section 2.1 emerges because α_i is then conditioned on the valuation x_i such that it replicates equation (1). 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

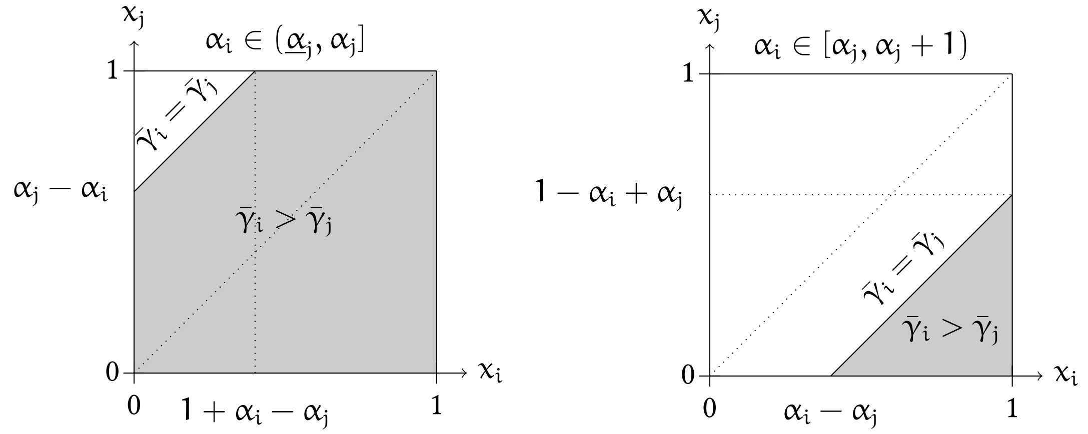

Bidder i wins the auction with the higher bid, $\gamma_i > \gamma_j$, for all bidders' values $(x_i, x_j) \in [0,1]^2$ such that bid shading implies $x_i - \alpha_i \geq x_j - \alpha_j$. For any α_j, bidder i finds it optimal to respond with an amount of bid shading α_i^* such that $\alpha_i^* \in [\underline{\alpha}_i, \alpha_j + 1]$, 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:[4]

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

With $u(x) = x^r$ and uniformly distributed values, bidder i's optimal amount of autonomous bid shading α_i^* is given by:[5]

$$\alpha_i^*(\alpha_j) = \begin{cases} \frac{r(\alpha_j+1)}{2+r} & \text{if } 0 \le \alpha_j \le \frac{r}{2} \\ \frac{(\alpha_j-1)(1+r)}{2+r} + \frac{1}{2+r}\sqrt{2r\,(2+r) + (\alpha_j - 1)^2} & \text{if } \alpha_j \ge \frac{r}{2} \end{cases} \tag{4}$$

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

$$\bar{\gamma}^*(x) = x - \frac{r}{2}. \tag{5}$$

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 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 heterogeneous 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 be playing against another markdown bidder with probability one.

The derivation of the equilibrium condition is given in Appendix B. Figure 3 shows the equilibrium bids for different attitudes towards risk r and various population mixes ρ.

[4] 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.

[5] See appendix A for the detailed derivation.

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 $\rho \in \{.1, .2, \ldots, .9\}$.

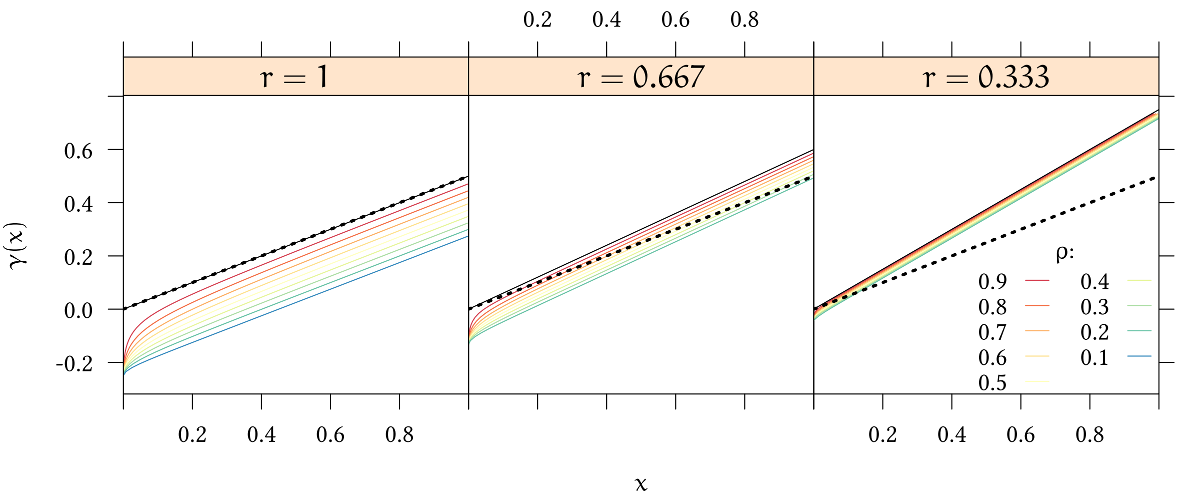

Coloured lines show equilibrium bids in the mixed population for different values of ρ. The black line shows the equilibrium bid for a population with only rational (though risk averse) bidders. The dotted line denotes the risk-neutral equilibrium bid (RNBNE) for a population of only rational bidders.

3. Experimental setup

The purpose of the experiment is twofold: We want to examine the extent to which the existing experimental evidence on first-price auctions is an artefact[6] of the design and we want to find out the extent to which 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 chance to be observable, we use 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 potential 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 (2) and (5) 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 also observe bids 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.[7] 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

[6] More specifically, the shape of bidding functions for low valuations that is biased upwards in experiments not admitting substantial underbidding with low valuations, e.g., by ruling out negative bids while using a lowest possible value of zero or close to it.

[7] A smaller number of bidders increases the probability to win the auction. An increase in the probability to win the auction increases the number of learning opportunities for bidders.

Table 1 Treatments

Treatment	$[\underline{\omega}, \overline{\omega}]$	$\underline{b}$	auction type	indep. observations	participants
-25	[-25,25]	-125	1st	4	32
0	[0,50]	-100	1st	6	40
0+	[0,50]	0	1st	4	32
25	[25,75]	-75	1st	3	26
50	[50,100]	-50	1st	4	30
50+	[50,100]	0	1st	8	86
50II+	[50,100]	0	2nd	6	58

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 $\overline{\omega} + 100$.

Kirchkamp et al. (2009) and Kirchkamp and Reiß (2011).

Second, we must provide a 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 it is in many experimental studies. To this end, we want to understand how seemingly 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 that permits only positive bids.

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

Hypothesis 1 (pure Bayesian Nash equilibrium bidding) *If all bidders use Bayesian Nash equilibrium bids, we should not see much underbidding. Additionally, 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 absolute 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 absolute markdown bids we should find no underbidding in the $0+$ treatment, because 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, and 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 being restricted 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 (negative bids admission priming effect) *We should find more absolute markdown bids in the 50 treatment than in the* $50+$ *treatment.*

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

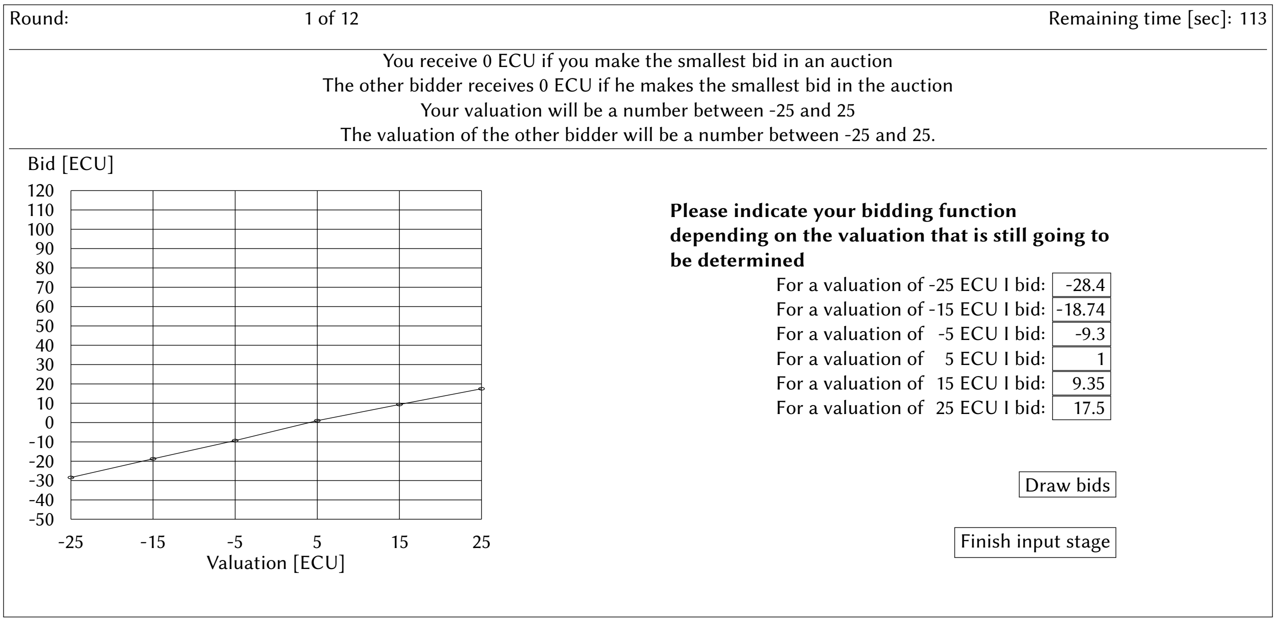

Whereas the -25 treatment is just a transformation of the 25 and 50 treatment, the -25 treatment involves negative and positive valuations at the same time. This might be perceived as more difficult and, thus, may give an additional incentive to use (simpler) absolute markdown bids.

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. An overview of the treatments is shown in Table 1, and instructions are provided in Appendix D. The experiments were computerised with z-Tree (Fischbacher (2007)).

A typical input screen used in the experiments is shown in Figure 4 (translated into English). In each round, participants are matched randomly to pairs, and they simultanously enter bids for six valuations that are equally spaced between $\underline{\omega}$ and $\overline{\omega}$. Bids for all other valuations are interpolated linearly. The bidding function is shown as a graph in the left part of the screen.[8] 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.

[8] The bid axis actually shown to subjects ranged from -50 to 120 in all treatments allowing for negative bids and from 0 to 120 in all treatments with non-negative bids. In Figures 4 and 5 we use a condensed range to improve the diagrams' exposition.

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

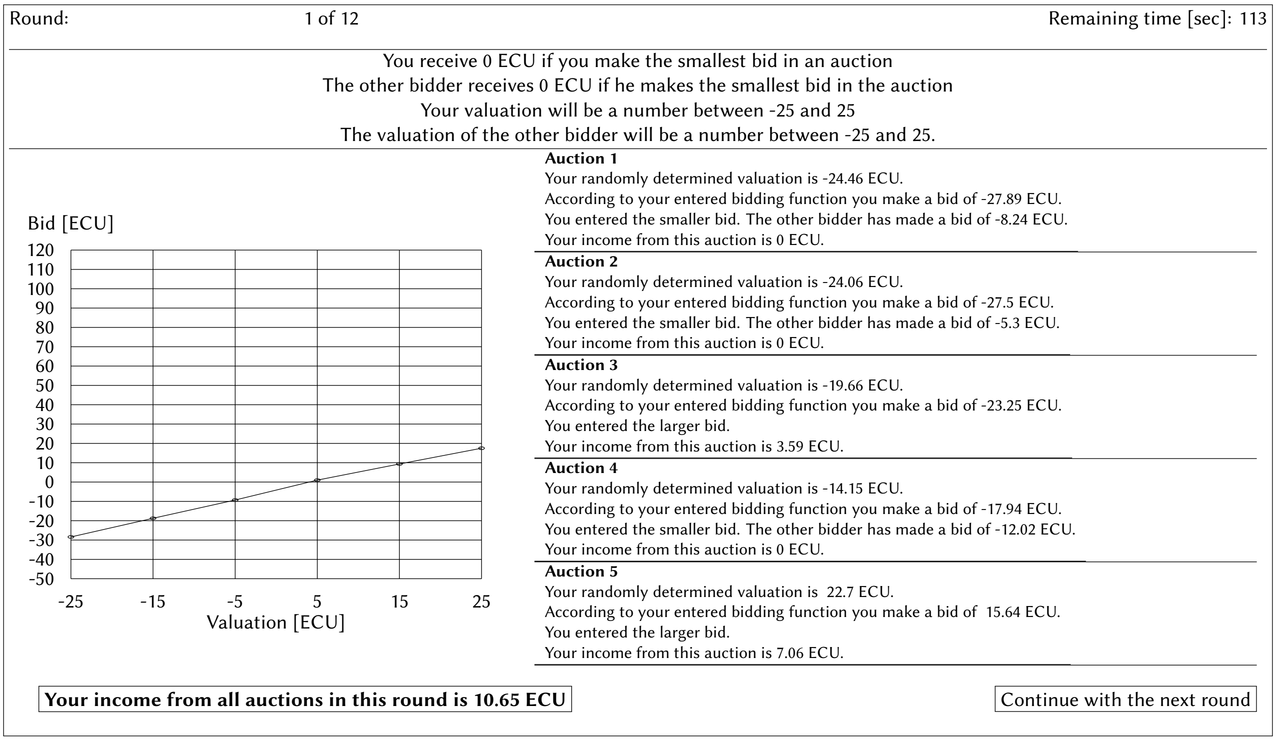

A typical feedback screen is shown in Figure 5. Participants play 12 rounds. Each round consists of a bid input stage and a feedback stage. At the end of these 12 rounds, participants complete a short questionnaire and are paid in cash according to their gains throughout the experiment.

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, the decision of the participants to 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 check whether behaviour stabilises over the course of the experiment. We rely on two indicators of stability. First, we count how often participants change the 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. The result is shown in the graph on the left in 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. After some adjustments during the first few periods, participants apply a more stable bidding function, adjusting fewer and fewer support points in each

Figure 6 Convergence of bids over time

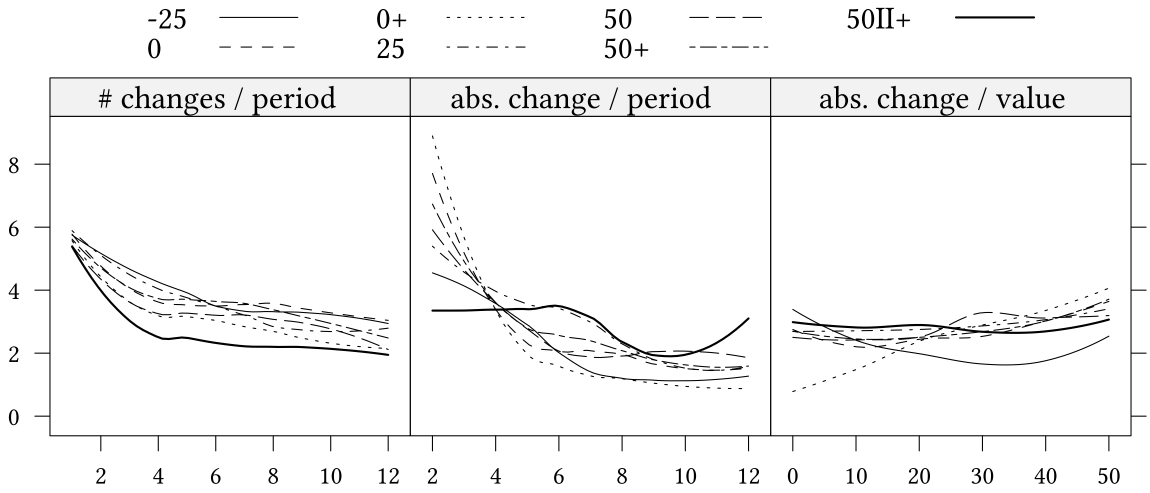

The panel on the left shows how many of the six support points (hypothetical bids) of a bid function bidders change for different periods in the game. The middle panel shows the absolute amount of this change for each period. The right panel shows the absolute amount of this change for the different valuations. All panels show loess splines (using the default parameters for loess).

period.

Second, the graph in the middle panel of Figure 6 shows the average absolute amount of these changes over time. 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 restricted in their changes for small valuations.

We conclude 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 amount of overbidding $b - \text{RNBNE}$ as a function of the normalised valuation $x - \underline{\omega}$.[9] RNBNE is the Bayesian Nash equilibrium bidding function with risk neutrality ($r = 1$) as given by equation (2). Let us briefly inspect the individual treatments:

Second-price auction, 50II+**:** 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 who bid more than the weakly dominant bidding strategy. This is consistent with 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[10].

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

[10] Garratt et al. (2012) report overbidding along with underbidding in second-price auctions for bidders with extensive eBay experience.

Figure 7 Overbidding for different treatments

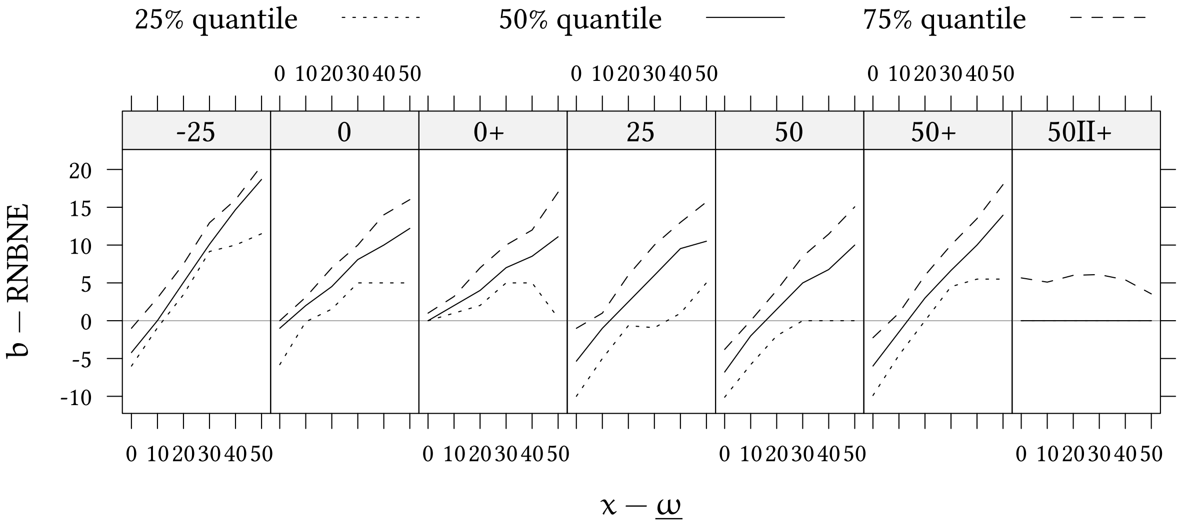

The figure shows normalised valuations ($x - \underline{\omega}$) on the horizontal and overbidding (b − RNBNE) on the vertical axis. The median amount of overbidding is shown as a black line. The interquartile range of overbidding is indicated with a dashed and a dotted line. The first 6 periods of each session are discarded.

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, the 25% quantile also 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 of 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.

4.3. Inference from aggregate bids

Hypotheses 1 and 2: To study Hypotheses 1 (pure Bayesian Nash equilibrium bidding) and 2 (partial markdown bidding), we use a Bayesian binomial model as an analogue to the binomial test in the frequentist approach. We treat each session as one independent observation and observe for each session whether for a given valuation the average bid in this session is higher or lower than the RNBNE. Let q denote the probability that the average bid in a session for a given valuation is above the RNBNE. For an (uninformed) uniform prior $q \sim \text{Beta}(1, 1)$, Figure 8 shows the posterior odds $o_{q>1/2} \equiv P_{q>1/2}/(1 - P_{q>1/2})$, i.e., the odds that bids above the RNBNE (overbidding) are more likely than bids below the RNBNE (underbidding). For high valuations ($x = \underline{\omega} + 50$), the odds of average overbidding are above 10:1 for all treatments. In the 0+ and 50II+ treatments for small valuations ($x = \underline{\omega}$) as well, the odds favour overbidding, although less strongly.

Figure 8 Odds of overbidding for different treatments

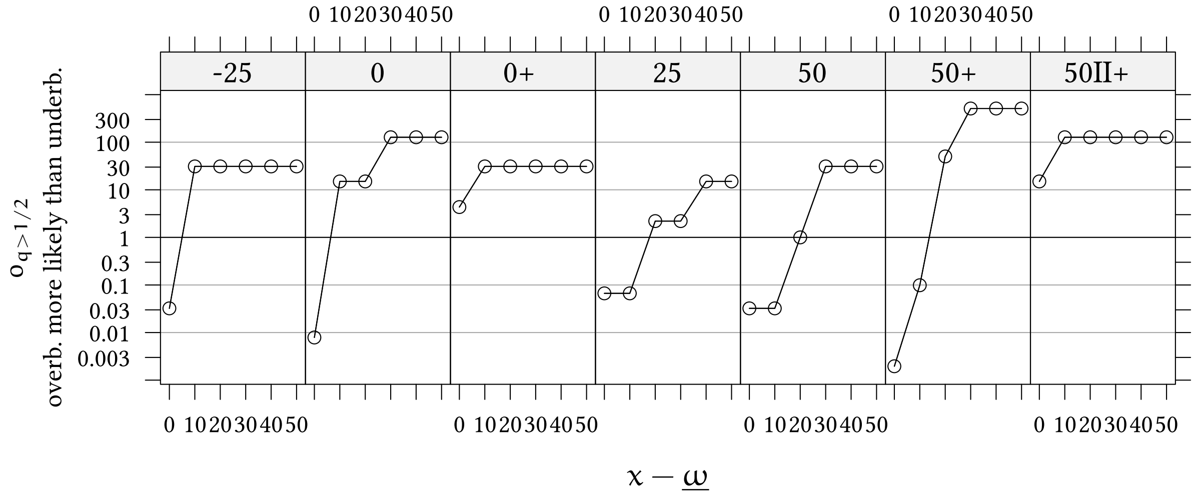

The figure shows normalised valuations ($x - \underline{\omega}$) on the horizontal axis. Let q be the probability that the average bid in a given session and for a given valuation is above the RNBNE. For an (uninformed) uniform prior $q \sim \text{Beta}(1, 1)$, the vertical axis displays (on a log scale) the posterior odds $o_{q>1/2} \equiv P_{q>1/2}/(1 - P_{q>1/2})$, i.e., the odds that bids above the RNBNE (overbidding) are more likely than bids below the RNBNE (underbidding). The first 6 periods of each session are discarded.

For each of the other treatments, however, the odds are strongly in favour of underbidding for small valuations (10:1 or more, i.e., for each treatment, we have strong evidence following the terminology of Jeffreys, 1961). Aggregating the evidence from the single treatments yields the posterior odds in favour of underbidding (for $x = \underline{\omega}$) to be more than 3×10^7:1; i.e., we find decisive evidence for underbidding for small valuations. This is consistent with the absolute markdown bids presented in Section 2.

Not surprisingly, as in many other experiments with first-price auctions, we find strong support for overbidding for the highest possible valuation $\overline{\omega}$ in all treatments. More interestingly, we find strong support for underbidding for the smallest possible valuation $\underline{\omega}$ in all first-price treatments where underbidding is possible, that is except for the $0+$ treatment. Thus, we find no support for Hypothesis 1 (pure Bayesian Nash equilibrium bidding), but we can confirm Hypothesis 2 (partial markdown bidding).

Sometimes the unwillingness to participate in the auction is put forward to explain bids of zero or close to zero for the results in the experiments of, e.g., Cox et al. (1988) as shown in Figure 1. Then, the bidding behavior for low valuations in the $0+$ treatment should be similar to that in all other treatments; i.e., for low valuations, there should be bids around the lowest possible valuation. This is, however, inconsistent with the underbidding observed in all other treatments, so that this explanation for zero bids with low valuations is rejected by the data.

Hypotheses 3–6: With absolute markdown bids, the slope of the bidding function should be one. With Bayesian Nash equilibrium bids, the slope of the bidding function should be smaller than one. To discuss Hypotheses 3 to 6, we compare the slopes of the aggregate bidding function in the different treatments. We take steeper slopes of the aggregate bidding function in a given condition as an indication of the presence of absolute markdown bids. In Sections 4.4 and 4.5, we

Table 2 Slope of bidding function β_3 (Equation (7))

$\underline{\omega}$	C	$\hat{\beta}_3$	95% -	C.I.	$\Pr(\beta_3 \leq 0)$	$\Pr(\beta_3 \geq 0)$
0	$\underline{b} \neq 0$	0.1000	0.0796	0.1204	0.0000	1.0000
50	$\underline{b} \neq 0$	-0.0616	-0.0854	-0.0376	1.0000	0.0000
all except 0	$\underline{\omega} = -25$	0.1840	0.1649	0.2031	0.0000	1.0000

Further details on convergence are given in Appendix C.1.

then consider individual bidding functions.

Here, we estimate a Bayesian hierarchical model where b_{igt} is the bid of individual i in matching group g at period t and ν_{igt} is the valuation of this bidder. We include random effects ϵ_i^I and ϵ_g^G for the individual bidder i and for the matching group g. The different levels of random effects are denoted $j \in \{U, I, G\}$ for the residual, the individual and the matching group, respectively. C is a dummy that is equal to one when the condition under C in Table 2 is fulfilled and zero otherwise. We use (almost) uninformative priors. The first 6 periods of each session are discarded.

$$(b_{igt} - \underline{\omega}) \sim N(\mu_{ig}, 1/\tau_U) \tag{6}$$

$$\text{with } \mu_{ig} = \beta_0 + \beta_1(\nu_{igt} - \underline{\omega}) + \beta_2 C + \beta_3 (\nu_{igt} - \underline{\omega})\, C + \epsilon_i^I + \epsilon_g^G \tag{7}$$

$$\epsilon_i^I \sim N(0, 1/\tau_I) \quad \text{and} \quad \epsilon_g^G \sim N(0, 1/\tau_G) \tag{8}$$

$$\text{with priors } \beta_j \sim N(0, 10^4) \qquad \text{for } j \in \{0, .., 4\} \tag{9}$$

$$\tau_j \sim \Gamma({m_j}^2/{s_j}^2, m_j/{s_j}^2) \text{ with } m_j \sim \text{Exp}(1), s_j \sim \text{Exp}(1) \tag{10}$$

Hypothesis 3 (suppression of markdown bidding): The first line in Table 2 shows the comparison of the 0 with the 0+ treatment; i.e., we consider only observations with $\underline{\omega} = 0$, and we have $C = (\underline{b} \neq 0)$. In the 0 treatment, the slope of the bidding function is larger than the slope in the 0+ treatment by $\beta_3 = 0.10$. In 100.00% of our 80000 samples, the posterior β_3 was positive. We can, therefore, confirm Hypothesis 3.

Hypothesis 4 (negative bids admission priming effect): The next line of Table 2 shows the difference in slopes of the bidding function between the 50 and 50+ treatment. According to the hypothesis, we should find a steeper slope in the 50 treatment, i.e., we should expect a positive β_3. The difference in slope between the two treatments is $\beta_3 = -0.06$. In 0.00% of our 80000 samples, the posterior β_3 was positive; hence, Hypothesis 4 is not confirmed. We find no priming effect when allowing for negative bids per se. To summarise: A restriction to positive bids does have an effect in the context of Hypothesis 3, i.e. in a range where some bidders otherwise make negative bids. Perhaps not surprisingly, the same restriction to positive bids does not have an effect in the context of Hypothesis 4, i.e. in a range where, even for a markdown bidder, the restriction should not have an effect.

Hypothesis 5 (complexity favours markdown bidding): The third line of Table 2 shows the difference in the slope 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. Indeed, our estimate for

Figure 9 Individual bidding

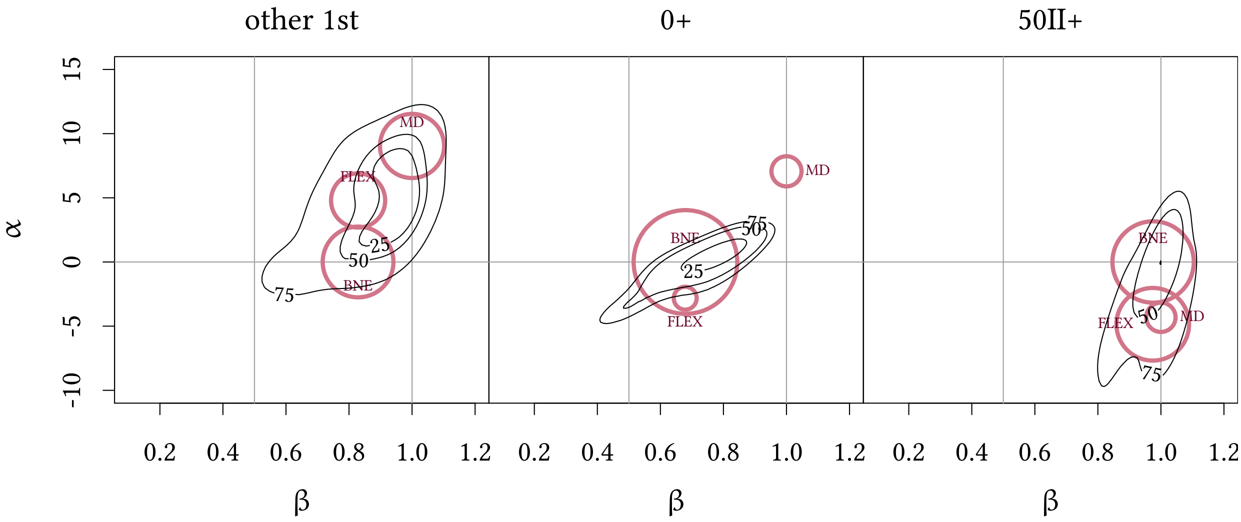

Each graph shows in black contour lines of the kernel density estimate of the distribution of individual bidding functions (see equation 11), with the first six periods of the experiment discarded. Numbers next to the contour lines are estimated percentiles. The centre points of the red circles (areas proportional to frequency) and their labels denote the centres of distributions of strategies from the mixture model we study in Section 4.5.

β_3 is 0.18. Furthermore, β_3 is positive in 100.00% of our 80000 samples of the posterior. We can, hence, confirm Hypothesis 5.

Does this mean that a sufficiently complex design eventually leads all bidders to follow a markdown rule? Perhaps not. Chen and Plott (1998) study a more complex environment which, in equilibrium, implies non-linear bids. Still, aggregate bids in their experiment are better described by a non-linear function. Hence, at least some bidders in Chen and Plott (1998)'s experiment must use non-linear bids.

Hypothesis 6 (no underbidding in second-price auctions): Here we have to go back to Figure 8. The line "50II+" shows the difference between bids in the experiment and equilibrium bids. For all values, the probability of overbidding is clearly larger than 1/2, which supports Hypothesis 6.

4.4. Individual bids

The quartiles of bidding behaviour in Figure 7, suggest heterogeneity among bidders. To better understand individual bidding behaviour we estimate for each bidder a linear bidding function:

$$b(x) = \underline{\omega} - \alpha + \beta \cdot (x - \underline{\omega}) + u \tag{11}$$

The regression specification normalises valuations and bids such that the point $(\underline{\omega}, \underline{\omega})$ transforms to the origin $(0, 0)$ in the 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 was $[0, 50]$ for any treatment. Again, we discard the first six periods of the experiment. The fit of the estimations of equation (11) is very good; e.g., the median R^2 is 0.9767.

Figure 9 shows the contour lines of the estimated joint distribution of α and β. We aggregate the data in three graphs:

First-price auction, treatments other than $0+$: The left-hand panel in Figure 9 illustrates coefficient estimates for first-price auction treatments other than $0+$. Here, bids smaller than the lower bound of the valuation domain $\underline{\omega}$ are possible.

- A large group of bidders is characterised by $\beta \approx 1$ and substantial markdown amounts $\alpha > 0$. In line with Section 2.2, these bidders are better described by absolute markdown bids instead of Bayesian Nash equilibrium bids ($\beta < 1, \alpha = 0$). This is evidence against Hypothesis 1 and in favour of Hypothesis 2.
- At the (vertically) lower end of the distribution, there are still some bidders with $\alpha \approx 0$. These could be 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 can interpret these bidders as rational bidders who realise that not all bidders are perfectly rational. They might also be viewed as markdown bidders.

First-price auction, $0+$: The panel in the middle of Figure 9 depicts coefficient estimates for the $0+$ treatment. The risk-neutral Bayesian Nash equilibrium predicts $\alpha = 0$ and $\beta = 1/2$. The risk-averse equilibrium predicts $\alpha = 0$ and $\beta > 1/2$. As Figure 9 illustrates, the distribution of coefficients is concentrated around $\alpha \approx 0$ and $\beta \in [1/2, 1]$. In this treatment, risk-averse Bayesian Nash equilibrium and other theories that we mentioned above explain the data quite well. Compared with the other first-price treatments in the left panel, we find fewer bidders with $\alpha > 0$, i.e., evidence in support of Hypothesis 3.

Second-price auction, $50\mathbb{I}+$: The rightmost panel in Figure 9 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 centred around this value.

4.5. Categorising individual bidders

To more formally categorise individuals according to their bidding behaviour, we estimate the following Bayesian mixture model. b_{igt} is the bid of bidder i in matching group g and in period t. v_{igt} is the valuation of this bidder. $c_i \in \{\text{RAND}, \text{MD}, \text{FLEX}, \text{BNE}\}$ is the subpopulation to which bidder i belongs. These subpopulations are distinguished by different restrictions (18) on coefficients $\beta_{m,i,c}$. Which restriction is in place is determined by the category c_i, which follows a categorical distribution and can take one of the following four values:

RAND: $b_{igt} = \beta_{0,i,\text{RAND}} + u$. This is a random bidder that bids a constant plus some noise.

MD: $b_{igt} = \beta_{0,i,\text{MD}} + v_{igt} + u$. This is the markdown bidder, which reduces the own valuation by a fixed amount (see Equation (5)).

BNE: $b_{igt} = \beta_{1,i,\text{BNE}} v_{igt} + u$. This bidder follows a (possibly risk-averse) Bayesian Nash equilibrium (see Equation (1)).

FLEX: $b_{igt} = \beta_{0,i,\text{FLEX}} + \beta_{1,i,\text{FLEX}} v_{igt} + u$. This is a linear approximation, capturing all bidders that cannot be better explained by (possibly risk-averse) Bayesian Nash equilibrium bidding (BNE), by markdown bidding (MD), or by random bidding (RAND). For first-price auctions, it describes the rational bidder who plays a best reply against a population mix of other markdown and other rational bidders (see Figure 3).

Obviously, the above list is not meant to be exhaustive. There might be other bidders which follow more complicated rules which are approximated by the above rules. As before, we allow for random effects for the individual bidder $\epsilon^I_{i,c}$ and for the matching group $\epsilon^G_{g,c}$. The different levels of random effects are denoted $j \in \{U, I, G\}$ for the residual, the individual and the matching group, respectively. We describe bids as follows:

$$\tau_j \sim \Gamma(m_j'^2/s_j'^2, m_j'/s_j'^2) \text{ with } m_j' \sim \text{Exp}(1), s_j' \sim \text{Exp}(1) \tag{12}$$

$$\epsilon^I_{i,c} \sim N(0, 1/\tau_I) \text{ and } \epsilon^G_{g,c} \sim N(0, 1/\tau_G) \tag{13}$$

$$b_{igt} - \underline{\omega} \sim N(\beta_{0,i,c_i} + \beta_{1,i,c_i}(v_{igt} - \underline{\omega}) + \epsilon^I_{i,c_i} + \epsilon^G_{g,c_i}, 1/\tau_U) \tag{14}$$

For each bidder i and for each of the four categories c_i, we need two coefficients $\beta_{k,i,c}$ with $k \in \{0, 1\}$ for Equation 14:

$$\mu_{k,c} \sim N(0, 10^4) \tag{15}$$

$$t_{k,c} \sim \Gamma(m^2_{k,c}/s^2_{k,c}, m_{k,c}/s^2_{k,c}) \text{ with } m_{k,c} \sim \text{Exp}(1), s_{k,c} \sim \text{Exp}(1) \tag{16}$$

$$\beta_{k,i,c} \sim N(\mu_{k,c}, 1/t_{k,c}) \tag{17}$$

$$\beta_{1,i,\text{RAND}} = 0, \beta_{1,i,\text{MD}} = 1, \beta_{0,i,\text{BNE}} = 0 \tag{18}$$

With Cat denoting the categorical distribution, δ is the vector of probabilities of the different categories $c \in \{\text{RAND}, \text{MD}, \text{FLEX}, \text{BNE}\}$.

$$\delta_c = \gamma_c / \textstyle\sum_c \gamma_c \text{ with } \gamma_c \sim \text{Exp}(1) \tag{19}$$

$$c_i \sim \text{Cat}(\delta) \tag{20}$$

Figure 10 shows relative frequencies of the categories in the different treatments and Figure 11 shows means of the estimated $\beta_{k,i,c}$ for each treatment. Table 4 in Appendix C.2 provides information on δ_c. Table 5 in Appendix C.2 provides statistics on the convergence of δ_c. We are now ready to examine our hypotheses at the individual level:

Again, we find no support for Hypothesis 1 (pure Bayesian Nash equilibrium bidding). Over all treatments with the first-price auction, more than a median of 54.4% of bidders are not classified as Bayesian Nash equilibrium bidders.

We do find support for Hypothesis 2 (partial markdown bidding). Over all treatments of the first-price auctions, a median of 30.3% of bidders are classified as markdown bidders with the smallest amount of markdown bidders in treatment 0+.

The data also support Hypothesis 3 (suppression of markdown bidding). We have a median of 22% markdown bidders in the 0 treatment but only 5% in 0+.

We find no support for Hypothesis 4 (negative bids admission priming effect): There are actually fewer markdown bidders (median 27.3%) in the 50 than in the 50+ treatment (median 40.7%).

We find some support for Hypothesis 5 (complexity favours markdown bidding): The largest share of markdown bidders (median 46.3%) is found in the −25 treatment. However, the effect is

Figure 10 Categorisation of bidders (Equation 19)

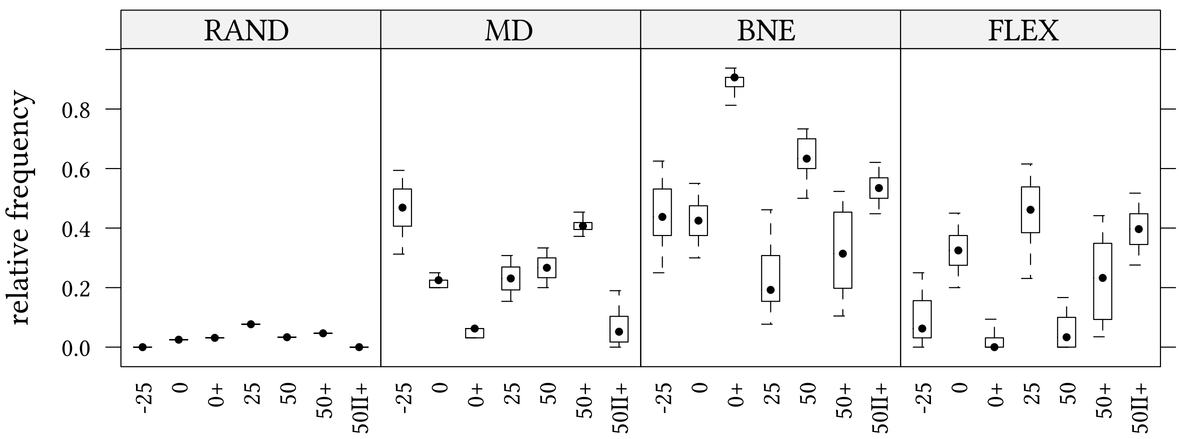

The boxplot shows relative frequencies of the different types c. Whiskers range from 5% to 95% quantiles. Boxes denote 25% and 75% quantiles. Medians are shown as dots.

quite small. The share of markdown bidders in the 50+ treatment is only slightly smaller (median 40.7%).

We can also support Hypothesis 6 (no underbidding in second-price auctions). Although some bidders (median 6.8%) are classified as using a markdown-strategy (MD), these bidders actually increase—not decrease—their bids on average; to see this, consider the right panel of Figure 11 which indicates the means of the estimated regression coefficients in Equation (14) separately for each category of bidders by treatment. There, FLEX-bidders in the second-price auction show a strictly positive intercept (vertical axis) indicating overbidding by 4.3 ECU on average.

Finally, compare the starkly different behaviour of FLEX-bidders in treatment 0+ (in this treatment a median share of 2.5% are classified as FLEX) to that bidding behaviour found for all other first-price treatments (there a median share of 23.1% are classified as FLEX). As we can see in the right panel of Figure 11, bidding functions estimated for the FLEX-bidders in treatment 0+ with the mixture model have a positive average intercept. In contrast, in any other first-price treatment, FLEX-bidders are characterised by negative intercepts on average indicating underbidding, consistent with the best-response against another bidder drawn from a mixed population with rational and markdown bidders.

5. Concluding remarks

Many first-price auction experiments find that subjects bid more than the risk-neutral equilibrium bid; i.e., 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, theoretically, optimal absolute

Figure 11 Averages of β_0 and β_1 from Equation (14).

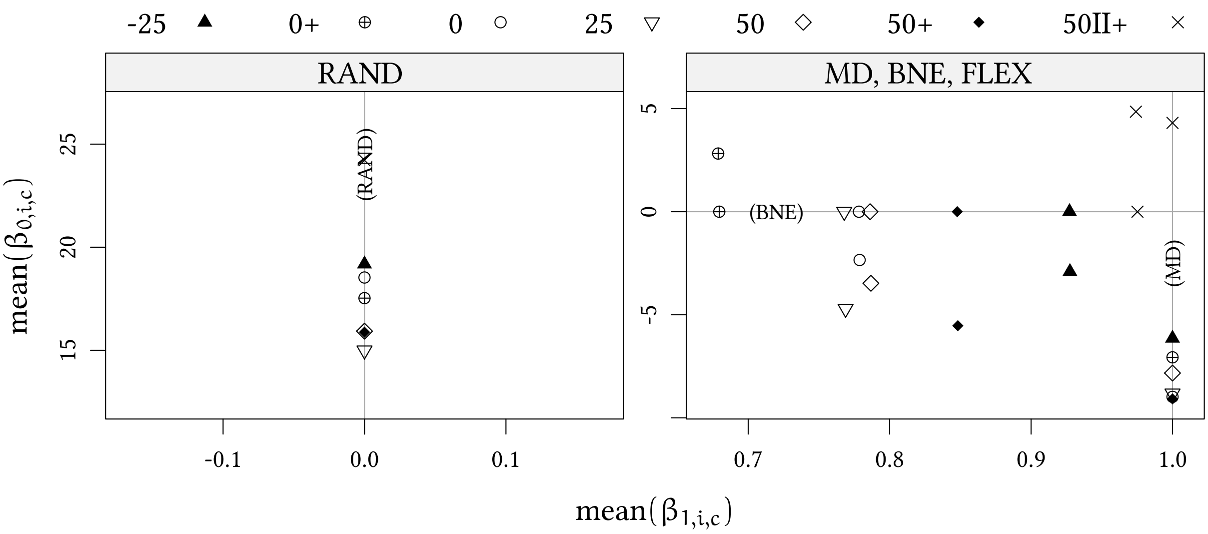

The Figure shows the averages of the coefficients estimated with the mixture model that classifies bidders into the bidding categories RAND, MD, BNE, and FLEX (see Equation (14) for the specification of the bidding model). In line with Equation (18), the case RAND (left panel) must be on the vertical line $\beta_{1,i,\text{RAND}} = 0$; the case MD must be on the vertical line $\beta_{1,i,\text{MD}} = 1$ (right panel); the case BNE must be on the horizontal line $\beta_{0,i,\text{BNE}} = 0$ (right panel); all other points in the right panel refer to the case FLEX.

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.

We found that there are very different types of bidding behaviour that we estimated with a mixture model basing on our theoretical framework. On average, 45.6% of bidders in our experiment were classified as Bayesian Nash bidders, most of which could be described by a risk averse bidding function. A substantial fraction of bidders—over all treatments 30.3%—seems to follow absolute markdown bids. A third group (20.4%) behaves like optimisers against such a mixed population. The remaining 3.7% were classified as random.

These types of bidders seem not to be exogeneously fixed before bidders enter the auction, at least they are not fixed for all bidders. Instead, the auction environment itself seems to influence their type. For example, in the traditional treatment, 0+, there are 89.4% of Bayesian Nash bidders along with 5% absolute markdown bidders, while for the most complex treatment, −25, these shares of bidder types amount to 44.1% and 46.3%, respectively. As in Masiliunas, Mengel, and Reiss (2014), we find more often less sophisticated strategies in more complex situations.

In our experiment, the manipulation of the range of valuations allows us to shed light on the richness of the interaction of the auction environment and bidding behaviour. For example, in the most complex treatment, −25, bidders are, for the most part, classified as (quite unsophisticated) markdown bidders, but in the least complex treatments (0, 25, 50), the majority of bidders is in line with rational bidding, in the sense of either best-responding to homogeneous BNE beliefs or to a heterogeneous population. This finding points to a challenge of auction design: Auction design has to incorporate heterogeneity; however, at the same time, the auction design actually determines heterogeneity.

For the designer of a mechanism it might be reassuring that in a context which is traditionally used in experiments, many bids follow Bayesian Nash equilibria. This would be good news for two reasons: First, it is reassuring to see that standard theory describes bids well. Second, if most bidders follow equilibrium bids, then it is more likely to reach an efficient allocation. Unfortunately, this property seems to depend on what looks like a minor detail: the range of valuations. We did not change much. We only shifted the range of valuations by a small amount. As a result fewer bidders can be described by equilibrium bids. This would not be a problem for efficiency if all bidders moved to markdown bids in the same way. Unfortunately, this is not the case. We find more heterogeneity than in the standard case which implies less efficiency.

In a broader context, our results also bear direct consequences for structural estimations with auction data. The application range of structural estimation is vast and it includes auctions of spectrum licenses in telecommunication (e.g. Hong and Shum, 2003 and Fox and Bajari, 2013) and procurement auctions (e.g. Decarolis, 2018). Hortaçsu and McAdams (2018) address many more applications and review the status quo of structural estimation from auction data. For structural estimations, the first-price auction case, the case we study, is of central importance, since "Recent advances in the empirical study of multiple-object auctions build on methods developed to estimate bidder values in single-object auctions, ..." (Hortaçsu and McAdams, 2018, p. 159). The estimation's key building block is the structural model. For first-price private-value auctions, this is the BNE equilibrium bidding function as stated in the review's equations (1) and (2).[11] Most importantly, the BNE equilibrium bidding function emerging under full rationality is used as identifying assumption.

At first glance, empirical support for such an assumption seems to be provided by Bajari and Hortaçsu (2005) who estimate the distribution of bidders' valuations from experimental first-price auction data to find that the risk-averse model 'is able to generate reasonable estimates of bidder valuations' (p.703). The dataset underlying the evaluation is experimental auction data that is based on uniformly distributed valuations from the interval $[0, 30]$ (see Dyer et al., 1989, for a more detailed description of the experimental design). As we have shown in this paper, experimental designs with a lowest valuation of zero or close to zero and the restriction to non-zero bids suppress underbidding for low valuations. In the field, however, firms' valuations of a spectrum license or firms' costs to provide a highway mile of tarred road admit underbidding, because valuations close to zero may not be relevant; in the case of procurement auctions, where the analogue of underbidding for low valuations is overbidding for high cost levels, overbidding high cost is always possible. Thus, the relevant experimental data for an evaluation is experimental auction data that allow for underbidding the lowest possible valuations, as, e.g., in our treatments 50 or $50+$. Ignoring systematic underbidding for low valuations in structural estimations might yield biased distributions of bidders' valuations and other biased estimated parameters, e.g., the parameter of risk-aversion with CRRA preferences, and derived results, e.g., on efficiency and prices. It seems that the use of the markdown equilibrium as a basis for structural estimation, or perhaps even more promising, the mixed-population model, could help avoiding biased estimates.

[11] See Hortaçsu and McAdams (2018), p.159.

References

Patrick Bajari and Ali Hortaçsu. Are structural estimates of auction models reasonable? evidence from experimental data. Journal of Political Economy, 113(41):703–741, 2005. doi: 10.3386/w9889.

Andreas Blume and Uri Gneezy. Cognitive forward induction and coordination without common knowledge: An experimental study. Games and Economic Behavior, 68(2):488–511, March 2010. URL http://ideas.repec.org/a/eee/gamebe/v68y2010i2p488-511.html.

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, pages 1–43. 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.

Francesco Decarolis. Comparing public procurement auctions. International Economic Review, 59(2):391–419, 2018. doi: 10.1111/iere.12274. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12274.

Douglas Dyer, John H. Kagel, and Dan Levin. Resolving uncertainty about the number of bidders in independent private-value auctions: An experimental analysis. RAND Journal of Economics, 20(2):268–279, Summer 1989. URL https://ideas.repec.org/a/rje/randje/v20y1989isummerp268-279.html.

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.

Jeremy T. Fox and Patrick Bajari. Measuring the efficiency of an FCC spectrum auction. American Economic Journal: Microeconomics, 5(1):100–146, February 2013. URL https://ideas.repec.org/a/aea/aejmic/v5y2013i1p100-146.html.

Rodney Garratt, Mark Walker, and John Wooders. Behavior in second-price auctions by highly experienced eBay buyers and sellers. Experimental Economics, 15(1):44–57, March 2012. URL http://ideas.repec.org/a/kap/expeco/v15y2012i1p44-57.html.

Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. Statistical Science, 7(4):457–472, 11 1992. doi: 10.1214/ss/1177011136. URL http://dx.doi.org/10.1214/ss/1177011136.

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.

Han Hong and Matthew Shum. Econometric models of asymmetric ascending auctions. Journal of Econometrics, 112(2):327–358, February 2003. URL https://ideas.repec.org/a/eee/econom/v112y2003i2p327-358.html.

Ali Hortaçsu and David McAdams. Empirical work on auctions of multiple objects. Journal of Economic Literature, 56(1):157–84, March 2018. doi: 10.1257/jel.20160961. URL http://www.aeaweb.org/articles?id=10.1257/jel.20160961.

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.

Harold Jeffreys. Theory of probability. Third edition. Clarendon Press, Oxford, 1961.

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.

Vijay Krishna. Auction Theory. Academic Press/Elsevier, 2010. ISBN 9780123745071.

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

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 (4)

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 $h'_-(\alpha_j - 1 + \sqrt{2}) < 0$. Because $z_2 < \underline{z}$, 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 $z \ge \underline{z}$, $h_-(z_1)$ 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 $z_4 = r(\alpha_j + 1)/(2+r)$. Because $h_+(z_3) = 0$ and $h_+(z) > 0$ 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 $h'_+(z'') < 0$. Because 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 $z \ge 0$, $h_+(z_4)$ is the maximum on interval $[0, \alpha_j + 1]$. For $\alpha_j \le r/2$, $z_4 \ge \alpha_j$; 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 $\alpha_j > r/2$ because 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 $EU_i(\alpha_i)$ 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 $EU_i(\alpha_i)$ is maximised by z_4 for $\alpha_j = 0$. The best-response function $\alpha_i^*(\alpha_j)$ as given by (4) follows immediately.

B. Rational bidders alongside markdown bidders

In Section 2.3 we considered the situation of a heterogeneous population. A share $\rho \in [0, 1)$ of the population is perfectly rational. The rest, a share $1 - \rho$, consists of markdown bidders. In this appendix we derive the equilibrium conditions for this case.

Let $\theta_j \in \{R, \overline{R}\}$ denote the rationality type of player j that can be either fully rational, $\theta_j = R$, or boundedly rational in the sense of markdown bidding, $\theta_j = \overline{R}$. Then, a fully rational bidder's prior probability of competing with another fully rational bidder is ρ and that of facing a markdown bidder is $1 - \rho$. Assume 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 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:

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

Because 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 (5) 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

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

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 a 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 yields 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 $\chi(b)$:

$$\rho\,[\chi(b_i) - b_i]\,\chi'(b_i) = [(1 + r)\rho - 1]\,\chi(b_i) + (1 + r)(1 - \rho)b_i + (1 - \rho)\frac{r^2}{2} \tag{21}$$

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 $x_i = 0$ follows as[12]

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

The initial condition follows as $\chi(-r^2/(2(1+r))) = 0$. Because differential equation (21) is non-linear and non-autonomous, an explicit solution is not known in general. We can, however, derive a good numerical approximation. Figure 3 shows the equilibrium bids for different attitudes towards risk r and various population mixes ρ.

For the special case of a population with equilibrium markdown bidders only, i.e., $\rho = 0$, equation (21) simplifies to

$$\chi(b_i) = (1+r)b_i + \frac{r^2}{2} \tag{22}$$

implying for a rational bidder who optimises against an equilibrium markdown bidder with $\alpha^* = r/2$ to bid according to

$$b(x) = \frac{x_i}{1+r} - \frac{r^2}{2(1+r)} \,. \tag{23}$$

With non-equilibrium markdown bidders who shade valuations by arbitrary $\alpha > 0$, so that $G(b) = b + \alpha$ for $b \in [-\alpha, 1-\alpha]$, a rational bidder's best-response is

$$b(x) = \frac{x_i}{1+r} - \frac{r\,\alpha}{1+r} \tag{24}$$

conditional on non-extreme markdowns of $\alpha \leq r$, which ensures that a rational bidder with the highest possible valuation does not overbid the highest possible bid of a markdown bidder of $1-\alpha$; otherwise, for $\alpha > r$, there emerges a flat portion of a rational bidder's best-response, where $b(x) = 1 - \alpha$ for all $x \geq 1 + r - \alpha$. For extreme markdowns, $\alpha \geq 1 + r$, the rational bidder – whatever the valuation – would prefer to always win the auction. In this case the best response would be a flat bid of $1 - \alpha$ (equal to the highest possible bid of a markdown bidder) for all $x \geq 0$.

C. Further estimation results

C.1. Estimating Equations (6)–(10)

To calculate the posterior distribution for Equations (6)–(10), we use JAGS 4.0.0 with 8 different chains, each with 5000 burnin steps. For each chain, we take 10000 actual samples. The posterior distributions for each case are therefore based on 80000 actual samples.[13] Table 3 provides the effective sample size and potential scale reduction factor (Gelman and Rubin, 1992) for β_3 from Equation (7).

[12] 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.

[13] This took 5 minutes with 8 parallel threads on an i7-2600 CPU @ 3.40GHz.

Table 3 Convergence statistics for β_3 from Equation (7).

ω	eff. size	psrf
0	79668	1.00000
50	80458	1.00000
all except 0	80000	0.99999

Table 4 Medians of the distribution for δ_c from Equation 19.

treatment	RAND	MD	BNE	FLEX
-25	2%	45.7%	39.5%	9.3%
	[0.1,10.4]	[24.3,65.4]	[19.1,65.3]	[0.4,33.9]
0	4%	22.5%	38.6%	33.3%
	[0.6,12.7]	[11.3,36.9]	[21.9,59.1]	[14.4,52.7]
0+	5.2%	6.9%	81.6%	4%
	[0.8,16.1]	[1.1,19.5]	[63.9,93]	[0.2,18.5]
25	8.9%	22.9%	23.1%	42.4%
	[2.1,22.3]	[9.1,41]	[11.3,50.7]	[15.1,63.7]
50	5.4%	28.3%	56.9%	6.5%
	[0.8,16.7]	[13.5,46.5]	[35.8,76.1]	[0.2,26]
50+	5.2%	40%	30.8%	22.7%
	[1.9,11.1]	[29.3,51.3]	[11.5,55]	[2.4,47]
50II+	1.2%	6.5%	50.2%	40.1%
	[0,6.1]	[0.3,25.3]	[34.6,65.2]	[21.6,57.8]

Numbers are given as percentages. 95% credible intervals are given in brackets.

C.2. Estimates and Convergence for Equations (12)–(20)

For each treatment, we calculate posteriors separately, each time using JAGS with 8 different chains, each chain with 5000 burnin steps. For each chain, we take 10000 samples with a thinning parameter of 10. The posterior distributions for each treatment are therefore based on 8×10^5 actual samples.[14] Table 4 shows medians of the distribution for δ_c from Equation 19. Numbers are given as percentages. 95% credible intervals are given in brackets.

The left part of Table 5 shows the potential scale reduction factor for δ_c . The right part shows the effective sample size.

D. 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. 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

[14] This took 104 minutes with 8 parallel threads on an i7-2600 CPU @ 3.40GHz.

Table 5 psrf and effective sample size for δ_c from Equation 19.

	potential scale reduction factor				effective sample size			
treatment	RAND	MD	BNE	FLEX	RAND	MD	BNE	FLEX
-25	1.00	1.00	1.03	1.07	72681	40027	4712	2693
0	1.00	1.00	1.00	1.00	79159	38799	2593	2510
0+	1.00	1.00	1.01	1.01	78213	40623	16203	8860
25	1.00	1.00	1.00	1.00	79006	44241	2639	3357
50	1.00	1.00	1.00	1.01	80335	48610	2704	1040
50+	1.00	1.00	1.00	1.00	79004	49683	1729	1684
50II+	1.00	1.00	1.00	1.00	73288	12183	65774	20946

and, then, were paid in cash. The experiment was conducted with z-Tree, (Fischbacher (2007)).

D.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 experiment and all payments.

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

D.2. Information regarding the experiment

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

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

For the auctioned object you have a valuation in ECU. This valuation is between x and $x + 50$ ECU and is determined randomly in each round.[15] The range from x to $x + 50$ is shown to you at the beginning of the experiment on the screen and is the same in each round.[16] **From this range you obtain new and random valuations for the object in each round.** The other bidder in the auction also has a valuation for the object. The 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.

D.2.1. Experimental procedure

The experimental procedure is the same in each round and is 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:[17]

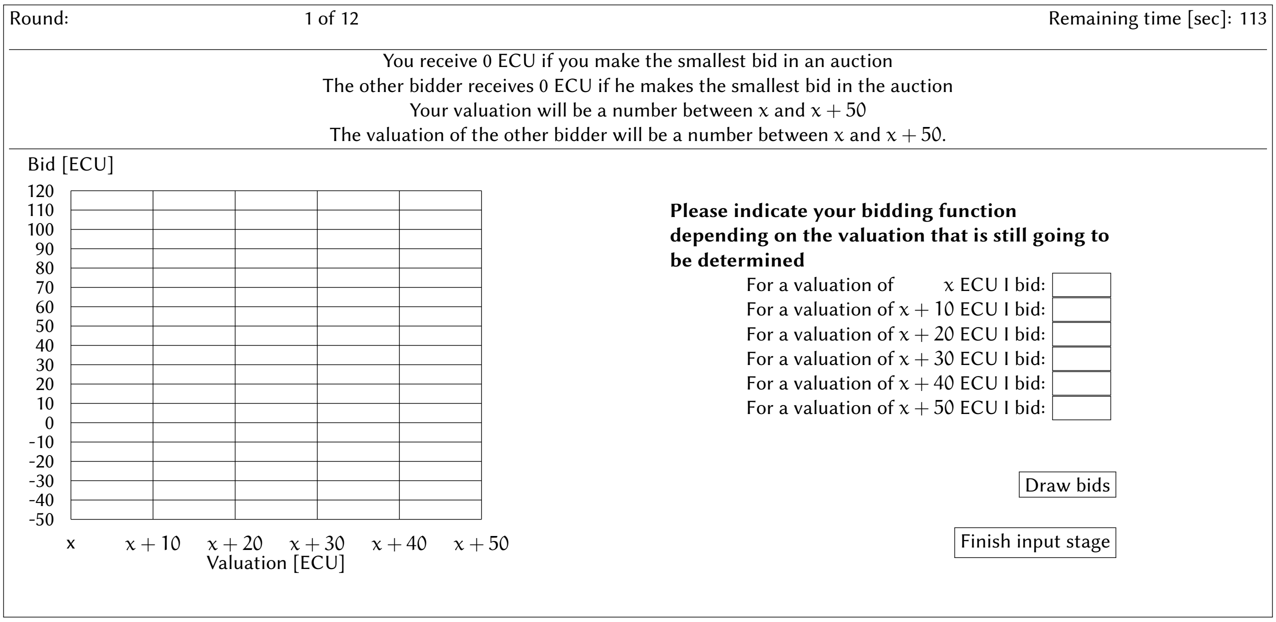

[15] In the 0+ and 50+ treatments the valuation would be announced precisely: “This valuation is between 0 and 50 ECU” in the 0+ treatment and “This valuation is between 50 and 100 ECU” in the 50+ treatment. Whenever x is mentioned in the remainder of the instructions, the same comment applies: In the 0+ and 50+ treatments the valuation is always announced precisely.

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

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

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 may 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 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 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:[18]

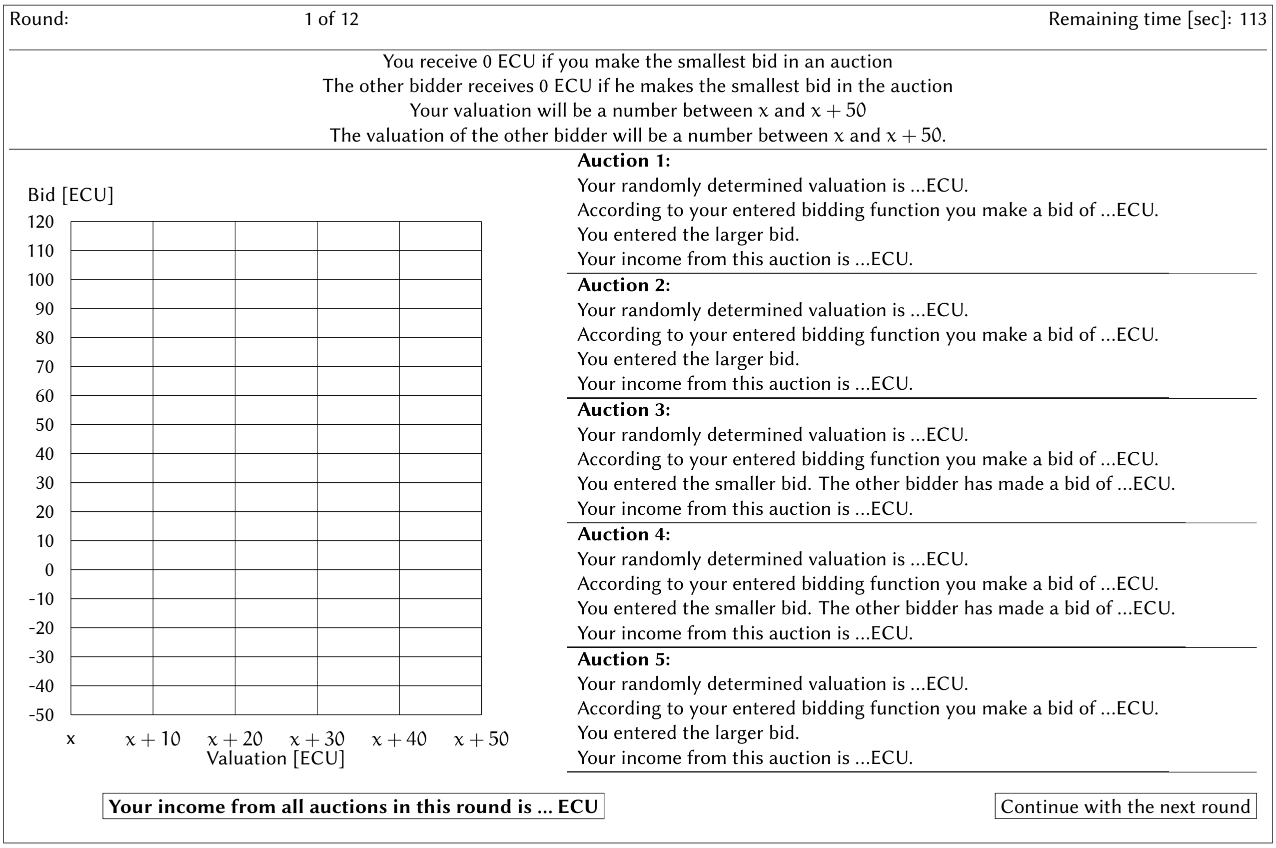

For each of your five valuations, the computer determines your bid according to the graph from stage 1. If a valuation is precisely at x, $x+10$, $x+20$, $x+30$, $x+40$, or $x+50$ the computer takes the bid that you gave for this valuation. If a valuation is between these points, your bid is determined

[18] 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.

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 obtains the object.

Your income from the auction:

For each of the five auctions the following holds:

- 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.[19]
- If the bid of bidder with the higher is a negative number, the amount is added to his ECU account.[20]
- 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 the input screen from stage 1 again in the next round.

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.

[19] 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.
[20] This item is not shown in the 0+ and 50+ treatments.