How to Boost Revenues in First-Price Auctions? The Magic of Disclosing Only Winning Bids from Past Auctions

Philippe Jehiel* Peter Katuščák† Fabio Michelucci‡

December 29, 2015

Abstract

A long-term auctioneer who repeatedly sells identical or similar items might disclose selective information about past bidding. We present an experimental design to evaluate the revenue implications of two common policies: disclosure of all past bids versus winning bids only. The analysis of our data points out that disclosing winning bids dominates in terms of revenue generation. We propose that when presented historical winning bids some of the bidders mistakenly best-respond to that distribution, failing to realize that winning bids are not representative of all bids. In the steady state, this *selection bias* results in higher auction revenues relative to when all bids are presented. On the theory side, the findings challenge the predictive power of Bayesian Nash Equilibrium based on rational bidders (that would yield revenue equivalence). On the market design side, they underline the role of historical market information as a key design choice.

PRELIMINARY AND INCOMPLETE DRAFT. DO NOT CITE OR DISTRIBUTE.

Keywords: auctions, feedback, selection bias, market design.

JEL Classification: C91, C92, D44

We are adding Section 5.3 (Estimation of the Fraction of Naive Bidders).

*Paris School of Economics and Department of Economics and ELSE, University College London, Campus Jourdan, 48 Boulevard Jourdan, Building E, 2nd floor, office 25, 75014 Paris, jehiel@enpc.fr.

†University of Economics in Prague, Peter.Katuscak@gmail.com.

‡CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Science, Politických vězňů 7, 111 21 Prague, Czech Republic, Fabio.Michelucci@cerge-ei.cz.

1 Introduction

Consider an auction house or a long-term auctioneer who repeatedly sells identical or similar items. Typically, such an auction house has gathered a lot of relevant information about bidding data from past auctions and can disclose some (or potentially all) of this information to current bidders. A familiar example of this is eBay in which bidders can routinely access bidding data from previous auctions of identical or similar items. A natural question is about the revenue implication of different disclosure policies.

This paper considers an experimental design with symmetric two-bidder first-price auctions in which prior to making their bidding choice bidders are informed of their valuation for the good for sale (private value setting). We consider two disclosure policies. Under one policy, bidders have access to historical information (HI) on all bids from previous auctions over a certain time period, while under the other policy, bidders have access to HI that describes only the winning bids over the same time period.

We have in mind applications in which bidders participate in just one auction so that there is no dynamic consideration in the underlying strategic interaction, and where it is reasonable to consider that bidders' beliefs regarding opponents behavior are based on historical information. Because of the focus above, even though in our experiment each individual subject participated in several auctions (to save on the number of required subjects), a given subject was never informed of the outcome of the auction in which she/he participated until the very end of the experiment.[1]

If bidders were fully rational, the type of disclosure policy would make no difference (assuming the dataset is arbitrarily large). This is because fully rational bidders can get an unbiased estimate of the distribution of all bids from the distribution of winning bids with arbitrarily large datasets (recall symmetry among bidders). Thus, assuming bidders are fully rational, the steady state of a dynamic process in which bidders would assess the distribution of opponent's bid based on HI and best-respond to it would correspond to the Bayes Nash Equilibrium of the underlying Bayesian game under either type of HI, and this would hold true irrespective of the risk attitude of bidders.

Our main finding is that the choice of disclosure policy has a massive effect on the distribution of bids, even after that distribution has stabilized (i.e., reached steady state). Bids are 6-7% higher when only winning bids are shown as opposed to when all bids are shown. As a result, revenue is also significantly higher (again, 6-7%) when only winning bids are shown.

Such a difference cannot be explained by risk attitudes, as argued above. It cannot either be explained by theories related to ex post regret that would be differently affected according to whether winning bids or all bids are displayed as in Filiz-Ozbay and Ozbay (2007) essentially

[1]This focus on one-shot bidding is similar to designs of Neugebauer and Perote (2008) and Füllbrunn and Neugebauer (2013). Unlike our paper, however, they do not provide any HI in the course of bidding.

because here individual subjects do not know the bids chosen by their opponent in their own past auctions, and thus have no way to directly feel the regret of having lost the auction when they could have profitably won it with a higher bid.[2] Another explanation that could come to mind is that the more aggressive bidding in the winning bid treatment could be due to subjects imitating or anchoring their decision on previous winning bids. However, in our setting there is no clean way for a subject to imitate because she/he is not informed about the valuations behind the winning bids. Furthermore, anchoring a bidding decision on previous winning bids may be infeasible when a subject receives a low valuation.[3]

We propose to explain the difference in the two treatments by the neglect of selection bias in the treatment in which only winning bids are disclosed. Specifically, the distribution of winning bids is not representative of the distribution of all bids, as winning bids are a biased sample of all bids (this bias is clearly toward higher bids given that the distribution of winning bids first-order stochastically dominates the distribution of all bids). Yet, we postulate that some share -that we estimate with our experimental data- erroneously consider that the shown distribution of winning bids is representative of all bids and accordingly best-respond as if it were the distribution of all bids. In this way, bidders in the winning bid treatment bid more aggressively when they are subject to selection bias than in the all bid treatment, and the overall distribution of bids is shifted toward higher bids in the winning bid treatment as compared with the all bid treatment.

In our statistical work, we estimate the distribution of risk aversion parameters that would allow to best-fit the data (in a CRRA specification similar to the one used in Cox, Smith and Walker (1982), Cox, Smith and Walker (1992)) in the treatment in which all past bids are disclosed assuming in each period bidders best-respond to the distribution of past bids shown to them. Imposing that the distribution of risk aversion parameters is the same both in the treatment in which only winning bids are shown and in the treatment in which all bids are shown (this is just an assumption about the homogeneity of the population of subjects in both treatments), we estimate how the share of naive subjects should be chosen to best fit the data in the winning bid treatment. We find that the share of naive subjects is x% + add elaborations on how this share depends on risk aversion.

Our findings have potentially profound implications. On the theory side, they challenge the predictive power of Bayes Nash Equilibrium based on rational bidders and suggest in the context of the winning bid disclosure rule how steady state should be modified to account for the selection bias that some share of bidders fail to appreciate. On the market design side, they open new ground for HI to have a key design role as a virtually costless and effective instrument to raise expected revenues. In particular, our findings may shed light as to why

[2]By the same token, it cannot be explained by learning theories based on individual experience as in Ockenfels and Selten (2005) (see below for a description of this approach).

[3]It is not even clear what a subject could use as anchor as what he sees is a distribution rather than a single observation.

it is a relatively common practice to disclose only winning bids in the context of a number of auctions (for example, Priceline.com -a site matching Hotels with potential customers, and allowing them to bid a price- provides samples of winning bids to bidders willing to secure a Hotel room, see `http://www.priceline.com/promo/hotel_winning_bids.html`), since such a disclosure rule is shown to generate higher revenues in the lab.[4]

The rest of the paper is structured as follows. Section 2 surveys the related theoretical and experimental work and positions our paper relative to this literature. Section 3 presents the theoretical model, which serves as the basis for our analysis of the experimental data. Section 4 presents details of our experimental design. Section 5 presents empirical findings from the experiment. Finally, Section 6 concludes and suggests avenues for further research. The Appendix contains experimental instructions and a demographic questionnaire we used at the end of the experiment.

2 Related Literature

We start with some theoretical contributions that also point out that manipulating HI from past auctions can affect bidders' beliefs and ultimately their behavior. Then we move to some theoretical and experimental literature that has looked at the effect of feedback about the outcome of past auctions the bidder has directly participated in. There the theoretical focus is not on beliefs. On the one hand, there is a learning approach based on impulses and, on the other, a modification of standard preferences (allowing for regret) without learning motivation. Finally, we mention some interesting but less related papers that deal with information disclosure in strategic repeated bidding effects (absent by design in our work).[5]

A first related theoretical approach is based on the application of analogy-based expectation equilibrium (ABEE) (Jehiel 2005) to auctions (Jehiel 2011).[6] A shared feature is that in the underlying learning model that gives rise to the proposed steady state, bidders fictitiously best respond to the information they are given. But, the manipulation of HI differs. In our setting, bidders receive HI purely on the auction type they participate in. However, they are not offered the whole dataset-only the winning bids' selection. In Jehiel (2011), datasets are complete, but there is a pooling of data regarding HI of different types of auctions and/or different types of bidders. In our setting, there is one auction format and bidders are ex ante symmetric, thereby ensuring that the manipulation in ABEE is not at work in the present setting.

Esponda (2008) applies the concept of self-confirming equilibrium (SCE) (Battigalli 1987,

[4]Of course, there could be other reasons for not disclosing losing bids for example related to anonymity considerations, but our findings suggest another argument more directly beneficial to sellers.

[5]It should be mentioned that we do not refer to the vast literature documenting overbidding in experimental FPA, nor to the various explanations (risk aversion, joy of winning, level k quantal response equilibrium) that have been proposed to give some account to the observation.

[6]See also Jehiel and Koessler (2008) for an application of this idea to general Bayesian games and Huck, Jehiel and Rutter (2011) for a related experimental investigation.

Fudenberg and Levine 1993, Dekel, Fudenberg and Levine 2004) to first-price auctions in a context where bidders plays repeatedly and bidders learn from their own experience either being informed of the winning bid or of all other bids (as well as being informed of their own payoff). In a SCE, beliefs are only required to be consistent with the feedback that players obtain. Hence, by manipulating the richness of the feedback one might reach SCE for which beliefs are not correct and that differ from the Bayesian Nash equilibrium (BNE) in affiliated value settings. However, this is not the case for the private value setting we use in which BNE is the only possible steady state whatever HI in Esponda' setting.

Our work focuses on HI about similar auctions played by other bidders in the past (our design does not provide any feedback about the outcome of past auctions so that a bidder does not know whether his bids from past auctions are winning bids or not). Some literature has instead looked at the effect of feedback about outcome of auctions the bidder has played.

Within this literature, one line of research treats feedback as an input for learning how to adjust one own'strategy. Ockenfels and Selten (2005) and Neugebauer and Selten (2006) consider an environment in which having lost, but learning that the winning bid was less than one's value, generates an impulse to bid higher in future auctions. On the other hand, having won but learning that the second highest bid was lower than one's own winning bid, generates an impulse to bid lower in future auctions. Equilibrium bidding is then determined by an "impulse balance" between the upward and downward impulses on bidding. The former impulse can be exacerbated relative to the latter by providing feedback on winning bids only, thereby leading to higher revenues.

One way of interpreting the origin of the bidding impulses described above is that they are triggered by an *ex post* experience of loser and winner regret, respectively. This brings us to a second related line of research. Engelbrecht-Wiggans (1989), Filiz-Ozbay and Ozbay (2007) propose that *ex ante* regret anticipation can extend even to a one-shot setting. The idea is that knowing that one will learn about the winning (second highest) bid could induce anticipation of loser (winner) regret and hence to bid more (less). Experimental evidence from repeated bidding (Isaac and Walker 1985, Ockenfels and Selten 2005, Neugebauer and Selten 2006, Engelbrecht-Wiggans and Katok 2008) robustly documents that providing feedback on winning bids in one's own past auctions results in higher bids and higher expected revenue compared to the case when bidders only learn whether they won or not.[7] In contrast, the evidence from one shot bidding experiments is mixed. Filiz-Ozbay and Ozbay (2007) find that bids and revenues are higher when feedback on the winning bid is given as opposed to only learning whether one won or not, but Katuščák, Michelucci and Zajíček (2015), in an experiment with a larger sample size and employing various experimental protocols do not

[7]There is weaker (mixed) evidence from studies employing repeated bidding that feedback on the second highest bid in one's own past auctions results in less aggressive bidding and lower expected revenue.

find any effect.[8]

While our design does not allow the bidders to know whether their bids are winning or losing bids, thereby making the above two approaches not directly applicable, it could be argued that the selection bias on which our theory builds could also be at work in contexts where this knowledge is accessible. More work would be required to disentangle these various effects in such contexts.

On a different dimension, our insights do not apply either to situations where the same set of bidders plays the same auction repeatedly. In those settings additional strategic considerations are present which may interplay with the choice of disclosure rule. Bergemann and Hörner (2010) and Cason, Kannan and Siebert (2011) consider repeated bidding by a fixed set of bidders, with each bidder having a value (or a cost, in procurement auctions) that remains constant in each auction repetition. In this environment, non-anonymous HI about previous repetitions creates a channel through which bidders' bids signal their values which has an impact on the rest of the interaction. In our setting, these considerations are absent since bidding groups are drawn from a large pool of bidders and HI is anonymous.

Another consideration that arises when bidders play repeatedly is tacit collusion. Dufwenberg and Gneezy (2002) consider a repeated common value auction in which each bidding group is randomly drawn from a fixed population.[9] If all bids from previous auctions are disclosed, bidders may attempt to bid low early on in order to induce other bidders to also bid low in future auctions. In our setting, collusion-building is unlikely to operate. First, given that the HI is anonymous, it is virtually impossible to identify whether a group of bidders is attempting to initiate collusion. Second, in a private value environment, a low bid might very well be competitive if the bidder's (unobserved) value is low, making it hard to detect deviations.

Finally, it is worth mentioning a paper by Sonsino and Ivanova-Stenzel (2006) also motivated by eBay, which tests whether subjects searching for more HI perform better.[10] This is a complementary question to the one addressed here in which the same HI is imposed on all bidders and the effect on revenues of varying HI is analyzed.

3 Theory

Consider the first-price auction in which two bidders simultaneously submit a bid, the winner is the bidder who submitted the highest bid (in case of ties, the winner is drawn at random) and the price paid is the bid submitted by the winning bidder. Bidders are ex ante alike in the sense that their distributions of valuations v are drawn from the same distribution with density

[8] All of these studies agree on the lack of any effect of providing feedback on the second highest bid.

[9] The common value is constant and publicly known.

[10] The authors provide evidence that bidders who use the option to retrieve information and who search with more effort on average obtain higher payoffs.

$g(\cdot)$ on $[0, 100]$. In the experiment, we consider the uniform distribution $g(v) = \frac{1}{100}$ for all $v \in [0, 100]$.

In order to accomodate the observed overbidding relative to the risk neutral benchmark, we allow the bidders to be risk averse, and in order to accomodate the apparent dispersion in bidding behaviors, we allow that different subjects be heterogeneous in their risk attitude. Specifically, we denote by $u(v - b)$ the utility derived by a bidder with valuation v winning the auction with a bid b, where a more concave u reflects a higher degree of risk aversion. Following Cox et al. (1982) we parameterize the degree of risk aversion using CRRA functions, letting $u(z) = z^{\alpha}$ for some $\alpha \leq 1$. Reflecting our assumption of anonymous matching in the experiment, we assume that the parameter α is distributed symmetrically across the two bidders according to a density on $(0, 1)$ with cumulative F. This distribution will be estimated from our data.

The Bayesian game is thus defined by the rules of the FPA together with the distributions f and g according to which the (payoff) types (valuation v and risk parameter α) of the two bidders are drawn.

Our main interest lies in understanding the effect of providing different feedback to bidders. In one case, referred to as the A environment, bidders are given access to a complete data set of past bids (from identical auctions played in the past). In another case, referred to as W, bidders are given access only to a selection of the database, namely only the winning bids.

In our experiment, subjects were not told the distribution of valuations and they were not told either anything about the risk attitude of other subjects (experimentalists did not know anything about the latter either). That is, bidders had to make their choice of bids based only on the feedback. Accordingly, we propose that bidders bidding in period t fictitiously best reply to the Historical Information provided to them concerning the bids in the period $t - 1$ auctions (all of them, not only those auctions in which they participated).[11]

Specifically, let us define as $A_{t-1}(b)$ the cumulative distribution of all bids that a bidder observes in the A environment at a generic time $t - 1$, and let $W_{t-1}(b)$ be the cumulative distribution of winning bids in the W environment at a generic time $t - 1$ (assume an arbitrarily large dataset of bids so that both $A_{t-1}(.)$ and $W_{t-1}(.)$ can be taken to be continuously differentiable).

In environment A, bidders with valuation v and utility $u(\cdot)$ best respond to the distribution A_{t-1} by picking the bid

$$\underset{b}{\operatorname{argmax}}\, u(v - b)A_{t-1}(b).$$

In environment W, deriving what is the best strategy from the distribution of winning bids

[11]Note that since we are mainly interested in steady state comparisons, we do not model in details the learning process. In particular, we are not specific about how bidders should bid at an initial time for which no previous HI is available. For instance, one could assume that a bidders bid is a random draw from $U[0, v_i]$.

is less straightforward. We assume that some bidders are naive in the sense that they consider the distribution W_{t-1} to be representative of all bids, thereby choosing

$$\underset{b}{\operatorname{argmax}}\, u(v-b)W_{t-1}(b)$$

as their bid. Such bidders are subject to selection bias.

The other bidders are sophisticated in the sense of being able to infer the distribution of all bids from the distribution of winning bids and best-responding to it. That is, sophisticated bidders with valuation v and utility $u(\cdot)$ select

$$\underset{b}{\operatorname{argmax}}\, u(v-b)\sqrt{W_{t-1}(b)}$$

given that the cumulative distribution of the maximum of two realizations drawn from independent distributions with cumulative $\sqrt{W_{t-1}(\cdot)}$ is precisely $W_{t-1}(\cdot)$.

We allow the share of naive bidders to be related to the degree of risk aversion, and we let $\lambda(\alpha)$ denote the share of naive bidders having utility $u(v-b)=(v-b)^{\alpha}$. While the profile of $\lambda(\cdot)$ affects the analysis in environment W, observe that it has no impact in environment A. Moreover, if all bidders were sophisticated in W, there would be no difference between environments A and W (ignoring the possible effect of small sample size). We will thus attribute the observed difference between A and W as resulting from the fact that $\lambda(\cdot)$ is bounded away from 0 and we will suggest an empirical method to estimate $\lambda(\cdot)$.

A simple observation is that, in environment W, a naive bidder with payoff type (v_i, α_i) bids more than a sophisticated bidder with the same payoff type irrespective of the feedback W_{t-1}. To see this, observe that the problem of a *sophisticated* bidder with payoff type (v_i, α_i) is to maximize $\alpha_i \ln(v_i-b)+1/2\ln W_{t-1}(b)$ with respect to b, which is the same as maximizing $2\alpha_i \ln(v_i-b)+\ln W_{t-1}(b)$, the problem of a *naive* bidder with payoff type $(v_i, 2\alpha_i)$. That is a naive bidder bids as a sophisticated bidder with half his value of α_i, which implies that fixing v_i and α_i the naive bids more than the sophisticated irrespective of W_{t-1}.[12]

A byproduct of the above observation is that in general when observing a bid b in environment W it is impossible to disentangle whether the bidder who placed the bid is naive or sophisticated. Since, subjects did not participate in both A and W environments, there is no way to identify whether a given subject is naive or not from our data. What we do in the estimation section is to estimate the actual $F(\alpha)$ in our population of bidders in environment A and to propose an identification strategy (that works at the aggregate level) to provide an estimated value for $\lambda(\alpha)$ based on the assumption that the distribution of α is the same in environments A and W (which is legitimate given the homogeneity of the population of subjects

[12]Formally, this is because $\frac{\partial^2}{\partial\alpha\partial b}\alpha\ln(v-b)=\frac{-1}{v-b}<0$ for all $b<v$ so that a smaller α bidder always prefers a larger b for a given v irrespective of W. (In other words, a single crossing condition holds in our setup.)

in treatments A and W).

From a policy perspective, our main interest lies in understanding the long run behaviors of such learning processes which we will describe with the steady state, letting the distributions of bids be the same in all periods.[13]

The steady states in the A environment correspond to the Bayes Nash equilibria of the above Bayesian game and these have been analyzed by Cox et al. (1982). While no close form solution is available, we will refer to $BNE(F)$ as the corresponding Bayes Nash Equilibrium (assumed to be unique) where recall that $F(\alpha)$ is the distribution of risk attitude in the population.[14]

The proposition below describes how the steady state in environment W is related to the one in environment A.

Proposition 1. *Let $BNE(F)$ be the steady state in environment A where $F(\alpha)$ is the cumulative of risk parameters. Then the steady state equilibrium in environment W with the same distribution of risk parameter is given by $BNE(F^*)$, where $F^*(\alpha) = F(\alpha) + \int_{\alpha}^{2\alpha} \lambda(z) f(z) dz$.*

Proof. This is an immediate consequence of the observation that in environment W a naive bidder with risk type α behaves like a sophisticated bidder with risk type $\alpha/2$, thereby implying that the equilibrium in the W environment coincides with an equilibrium in the A environment in which the mass of bidders with risk parameter less than α is given by $F^*(\alpha)$. **Q. E. D.**

Ideally, one would like to compare the revenue in the steady state of W and A. Intuitively, given that naive bidders bid more aggressively than sophisticated bidders in W, one would expect that bids and thus revenues are higher in W than in A. While we expect this to be true for a wide range of distributions of payoff types, it need not hold for all distributions given that an upward shift of bid distributions (in the first-order stochastic dominance sense) may sometimes lead to a downward shift of best-response.[15] In the next example, we illustrate how the equilibrium construction works in a simple situation in which revenue is higher in W than in A.

Example 1. *Assume that v_i is uniformly distributed on $[0, 100]$, and all bidders have the same risk parameter α, and let $\lambda(\alpha) = \lambda \in (0, 1)$. Then:*

In environment A,

$$A(b) = (1 + \alpha)\, b \tag{1}$$

[13]More work is required to analyze under what circumstances, the proposed learning dynamics converges, but clearly if it converges it must be to a staedy state. In the data, we will check whether the distribution of bids varies toward the end of the repetition as to check whether the steady state assumption is legitimate.

[14]The distribution of valuations G is not varied in our analysis, thereby explaining why we do not index the BNE by G.

[15]An upward shift of best-response is related to a downward shift of the inverse hazard rate of bids $A(b)/a(b)$ where $a(b) = A'(b)$.

for $b \in [0, 1/(1+\alpha)]$ *and the steady-state bidding function is given by*

$$b_A(v_i) = \frac{1}{1+\alpha} v_i \tag{2}$$

In environment W*, for* $b \in [0, 1/(1+\alpha)]$

$$W(b) = \left(1 + \alpha - \frac{\alpha\lambda}{2}\right) b. \tag{3}$$

For $b > 1/(1+\alpha)$*, it is characterized by the differential equation*

$$W'(b) = \frac{\lambda\alpha}{2} \frac{W(b)}{W(b) - \lambda b - (1-\lambda)} \tag{4}$$

with the initial condition

$$W\left(\frac{1}{1+\alpha}\right) = 1 - \frac{\lambda\alpha}{2(1+\alpha)}. \tag{5}$$

The steady-state bidding function for the sophisticated bidders is given by

$$b_S(v_i) = \frac{1}{1+\alpha} v_i. \tag{6}$$

For the naive bidders, for $v_i \leq (2+\alpha)/(2+2\alpha)$*, it is given by*

$$b_N(v_i) = \frac{2}{2+\alpha} v_i \tag{7}$$

For $v_i > (2+\alpha)/(2+2\alpha)$*, it is (implicitly) given by:*

$$b_N(v_i) = W^{-1}(\lambda v + 1 - \lambda). \tag{8}$$

Revenue is higher in W *than in* A *because the bids of the sophisticated in* W *are the same as in* A *and naive bidders bid more agressively than sophisticated bidders in* W*.*

4 Experimental Design

Our experiment involves repeated two-bidder First Price Auctions (FPA) in two treatments: HI on all bids (treatment A) and HI on winning bids (treatment W). We describe first a representative session from the perspective of a subject receiving information about the experiment. After this description, we discuss specific design choices, and noise reduction measures.

After entering the lab and taking their seats, subjects in both treatments were given the very same printed instructions. They also received a sheet of paper and a pencil in case they wanted

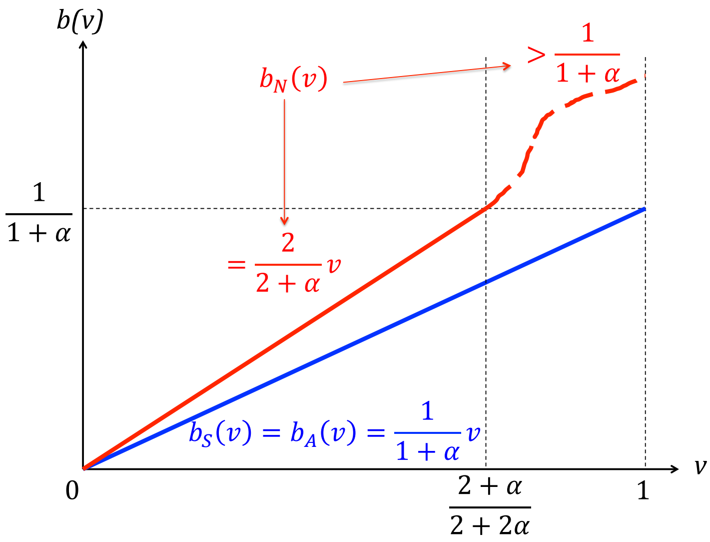

Figure 1: Example 1. $b_N(\cdot)$ and $b_S(\cdot)$: steady-state bidding functions of naive and sophisticate bidders in W. $b_A(\cdot)$: steady-state bidding function of all bidders in A.

to take notes. The instructions explained (in non technical terms) that the 24 participants in the session were divided into 2 groups of 12 and that members of each group would interact only among themselves. This interaction consisted of rounds of first-price auctions, each time against one randomly determined opponent from the same group. The subjects were explained the rules of the auction, the fact that they would play 11 blocks of 6 auction rounds each, and that everyone would get paid on the basis of the same 5 randomly chosen rounds at the exchange rate of 1 point=10 CZK (Czech crowns). On top of that, each subject knew he/she would be paid a 100 CZK show-up fee.[16]

The instructions informed subjects that values in block 1 were generated independently across subjects using the same random number generator. They also stated that value generation for later blocks would be explained at the end of block 1, that there would be no feedback on the outcome of the individual auctions until the very end of the experiment, and that the bidding would be followed by a demographic questionnaire. Finally, a shot of a bidding screen used in block 1 was provided.[17] We then provided subjects with a quiz consisting of eight multiple-choice questions testing subjects' understanding of the instructions.[18] In block 1, each subject was presented with six consecutive bidding screens. Each screen listed the subject's value realization for that round and the subject was asked to submit a bid.

[16]At the time of the experiment, the exchange rate was approximately 27.5 CZK/EUR and 20.1 CZK/USD.

[17]The screen contained a link to a built-in calculator that the subjects could use while bidding.

[18]The quiz is presented in the Appendix. After subjects answered all the questions, their answers were checked by an experimenter. Incorrect answers were infrequent. In case of an incorrect answer, an explanation was provided and a subject was given additional time to correct the answer. The experiment proceeded into block 1 only after all the subjects answered all the questions correctly.

Block 1 was identical in treatments A and W. After bidding in block 1 was completed, a second set of printed instructions was distributed. This time, the instructions were treatment-specific. In both treatments, subjects were told that their value realizations in all subsequent blocks would be the same, including the order, as their value realizations in block 1. At the same time, they were reminded that the matching into bidding pairs was random and independent in each bidding round within and across blocks. In treatment A, subjects were told that, from block 2 onwards, they would receive historical information about 72 bids from the auctions played in their group in the previous block. In treatment W, subjects were told that, from block 2 onwards, they would receive historical information about the 36 winning bids from the auctions played in their group in the previous block. Then a detailed explanation of the instruments used to display historical information from the previous block was given (see later for details) together with a shot of a bidding screen used in blocks 2 through 11. It should be stressed that, in treatment W, we considered two different pairing of subjects: One for the determination of winning bids shown in the HI -this is referred to as the Information Pairing and only the winning bids from the pairs in the Information pairing were sown In HI; and one for the determination of the actual payoffs -this is referred to as the Payoff Pairing. We will later develop the reasons for this design choice. The instructions were explicit about the difference between the *Earnings Pairing* and the *Information Pairing*.

After subjects had an opportunity to go over the instructions at their own pace and privately ask questions, we provided subjects with treatment-specific quizzes consisting of ten multiple-choice questions.[19] The bidding then continued for blocks 2 through 11.[20] When subjects completed bidding in all 66 rounds, we provided subjects with a demographic questionnaire. Then, subjects were presented with a feedback screen showing their and their opponents' bids and the winner in each of the 66 auctions they played. The computer then determined the 5 payoff-relevant auction rounds and displayed the results of those, including the subject's earnings in points and in CZK, on a new screen.

We select some design aspects to be discussed in greater detail.

Information vs payoff pairing:

The main reason why we used two different pairing in W is to avoid that subjects could infer from the HI whether their bids in the previous block were winning bid. As a result, any difference of bidding between A and W can be attributed to the difference of beliefs about the distribution of opponent's bids (regret theory has no bite in explaining the difference between A and W given that subjects cannot even observe whether their bid was winning or not). An alternative to exclude effects related to feedback from own past auctions would be to exclude the own auction outcomes from HI provided to a subject, but this has the undesirable effect of

[19]The quizzes from both treatments are presented in the Appendix. Again, when checking on the quiz answers, we observed a very good understanding of the instructions.

[20]Because of the necessity to collect all the bidding information for the previous block, bidding moved to the next block only when all the subjects in a given group had completed bidding in the previous block.

having subjects being exposed to different HI (see the discussion below for elaboration).

HI generation and its richness:

Two reasons led us to form HI at the level of a block and to use 6 bidding rounds in a block. First, we want each subject to bid for several different values given a certain realization of HI. This allows to estimate a bidding strategy for each subject given that realization of HI. Second, we need to form a rich enough history of past bids. With 6 rounds of bidding in a block, a history of all bids consists of 72 bids and a history of winning bids consists of 36 bids. This represents a good compromise between achieving those objectives and not overloading subjects with too many auctions.[21]

We chose to include one own's bid in the HI. The main reason for doing this (rather than elimination from the HI shown to a subject his/her own bids) is that in this way all subjects face the same HI (which facilitates the comparison of bid shading across subjects). Another reason is that in field applications, auctioneers would most likely provide aggregate HI, without filtering out previous bids of a given bidder.

Eliminating repeated bidding effects:

In our experiment subjects bid repeatedly rather than just once. We use two measures to avoid the occurrence of repeated bidding effects. First, each subject's opponent is chosen randomly and anonymously from among the other 11 group members. Second, subjects do not receive any feedback on the outcome of the auctions they have played until the very end of the experiment. An alternative way of precluding potential repeated bidding effects would be to use cohorts of one-shot bidders, each having access to HI on bidding of the previous cohort. However, such design would be very demanding in terms of the number of subjects and therefore exceedingly costly.[22]

Display of HI to subjects:

Here the main tradeoff is between presenting something easy to understand and preserving detailed information about the distribution of (winning) bids. We achieve both objectives by presenting the HI in two formats. In order to provide an overall assessment of what the (winning) bid distribution in the previous block was, the subjects are provided with a histogram with 10 bars, each representing the percentage of all (winning) bids placed within the bins $[0, 10]$, $(10, 20]$,...,$(80, 90]$ and $[90, \infty)$. We use a histogram since it should be familiar to most subjects from textbooks, magazines, newspapers or introductory statistical classes. In order to allow subjects to recover arbitrarily fine details of the distribution a second format of HI is used. By inputting numbers into two boxes, subjects can recover an exact percentage of (win-

[21]We use 11 blocks because our extensive pilot experiment testing has shown that the distribution of slopes stabilizes around blocks 9 and 10 and there might sometimes be a last-block effect in terms of subjects bidding slightly less.

[22]The reason is that to form a rich-enough HI, one would need many subjects in each cohort. Also, it takes at least several cohorts for bidding behavior to stabilize. On top of that, each sequence of cohorts observing HI on the previous cohort would form one statistically independent cluster of data. To gain statistical power, one would need at least several such clusters, hence further increasing the required number of subjects.

ning) bids between the two numbers. Moreover, entering only one number in the left (right) box, they can obtain an exact percentage of (winning) bids at and above (below) that number.

Lack of information about distribution of opponents' values:

Recall that while each subject's value is drawn from the uniform distribution on $[0, 100]$, subjects are not given such information. This departure from the more standard practice of telling subjects how values are drawn is motivated by two considerations. First, we want to stay close to the theoretical model introduced in the previous section in which bidders can only use HI on bids to form their beliefs. Providing information about the distribution of values interferes with this objective by giving subjects an additional source of information for forming their beliefs. Second, in the majority of field applications, bidders do not have information about the distribution of opponents' valuations. If desired, such distribution would typically have to be inferred from other sources, predominantly from HI on bids in past auctions.

Subjects'payments:

Subject payments are determined by 5 randomly chosen rounds out of the 66 bidding rounds. This is a reasonable compromise between minimizing hedging effects on the one hand and providing sufficiently strong incentives and reducing luck-driven variance of subject earnings on the other hand.

4.1 Noise Reduction Measures

First, we would like to reduce noise in the estimation of the slopes of the bidding strategies. Such noise could result, for instance, from all 6 value realizations being clustered around the same low or medium values. To avoid this, when drawing values in a given block, instead of drawing the six values for a subject independently from the uniform distribution on $[0, 100]$, they are drawn from uniform distributions on $[0, 100/6]$, $(100/6, 200/6]$, ..., $(500/6, 100]$, and then randomly scrambled. This way the realized values provide an even coverage of the support of the value distribution for each subject in each block while preserving the form of the compound distribution (uniform on $[0, 100]$). Of course, the values of a single subject are not independent across the six rounds. However, independence of generation across subjects and random scrambling of values imply that knowing one's value in a given round does not help to predict the values of the opponent in that or any other round.[23] Furthermore, for each subject, it is indeed true that both his and his opponent's value in a given round are independently drawn from the uniform distribution on $[0, 100]$. Second, we would like the variation in HI *from one block to another* to be attributable purely to changes in bidding behavior rather than changes in value realizations or, in treatment W, changes in the formation of bidding pairs. We take two steps to achieve this. For each subject, we repeat the value realizations and their ordering in each block. Moreover, in each block, we also repeat the pairing of subjects for the purpose

[23] Subjects are explicitly informed about this fact in the instructions.

of constructing the HI on winning bids in treatment W (Information Pairing).[24] Third, we would like any differences in bidding behavior *between the two treatments* to be attributable purely to differences in bidding behavior rather than changes in value realizations. For this reason, we use the very same value realizations, subject-by-subject and round-by-round, in all bidding groups in both treatments. That is, in any session, a subject sitting at a terminal notionally labelled $i \in \{1, .., 12\}$ in treatment A has the same value realizations in each round of a block as a subject sitting at a terminal notionally labelled i in treatment W.Fourth, in order to limit noise originating from different timing of sessions for the two treatments, each session consists of 24 subjects split into two independent groups of 12 subjects, one in each treatment. Of the 24 subjects in each session, the allocation to the two treatments is random, subject only to gender balancing (see below). Fifth, we also want to avoid potential noise from different gender composition of subjects in the two treatments. This is because previous literature (Chen, Katuščák and Ozdenoren 2013, Pearson and Schipper 2013) has documented that there are gender differences in bidding in first-price auctions with independent private values, with women typically bidding more than men. To achieve this, for each session, we recruited the same number of men and women by means of separate recruitment campaigns. Due to random no-show-up, this resulted in approximately gender-even pool of subjects coming to the lab. In all but two sessions, we were able to utilize 12 men and 12 women.[25] Of these, 6 men and 6 women were assigned to each treatment. Moreover, men (women) were directed to occupy the same terminals across the two treatments. This means that men (women) had the same value realizations in the two treatments.[26]

4.2 Logistics and Subjects

Altogether, we ran 10 laboratory sessions of 24 subjects each, utilizing 240 subjects in total. In each session, half of the subjects participated in treatment A and the other half in treatment W. That is, we have data on 10 independent bidding groups of 12 subjects (120 subjects in total) in each treatment. In terms of gender composition, out of the 20 bidding groups, 17 (8 in treatment A and 9 in treatment W) consisted of 6 men and 6 women, 2 (one in each treatment) of 5 men and 7 women and 1 (in treatment A) of 7 men and 5 women. Altogether, we utilized 119 male and 121 female subjects. All the sessions were conducted between December

[24]The former measure gives another layer of importance to the even coverage of the support of the value distribution. Since value realizations repeat block after block, the even value generation avoids a scenario in which a subject could face unluckily low value realizations round after round, block after bloc. Such subject could easily end up frustrated due to having little chances of winning (with a big surplus), which could introduce noise into the data.

[25]One session utilized 10 men and 14 women and one session utilized 13 men and 11 women.

[26]In the session with 10 men and 14 women, 5 men and 7 women were allocated to each treatment. In particular, in each of the two treatments, one woman took a terminal usually allocated to men. In the session with 13 men and 11 women, 7 men and 5 women we allocated to treatment A and 6 men and 6 women were allocated to treatment W. In the former, one man took a terminal usually allocated to women.

2013 and February 2014 at the Laboratory of Experimental Economics at the University of Economics in Prague.[27] The experiment was conducted using a computerized interface programmed in z-Tree (Fischbacher 2007). Subjects were recruited using the Online Recruitment System for Economic Experiments (Greiner 2004) from our subject database. Most of our subjects were students from the University of Economics in Prague. A minority were students of other universities in Prague. According to the demographic questionnaire, at the time of the experiment, 45 percent of the subjects did not hold any degree, 46 percent held a bachelor's degree and 9 percent held a master's degree. Regarding the field of study, 6 percent had a mathematics or statistics major, 7 percent had a science, engineering or medicine major, 67 percent had an economics or business major, 8 percent had a social science major other than economics or business, and 12 percent had a humanities or some other major. Almost 99 percent of our subjects were between 18 and 28 years old, with three subjects being older (up to 39). Also, 47 percent of the subjects claimed to have had some experience with online auctions, 3 percent with offline auctions and 10 percent claimed experience with both types of auctions. Subjects were paid in cash in Czech crowns (CZK) at the end of their session. Each session lasted approximately 2 hours with an average earning of 421 CZK.[28]

5 Results

This section presents our empirical findings. Subsection 5.1 discusses "steady-state" treatment effects on bidding behavior. Subsection 5.2 analyzes "steady-state" treatment effects on average auction revenue and efficiency. Before we continue, let us mention some common features of our analysis. We often pay special attention to blocks 1, 10 and 11. Block 1 is treatment-free, so focusing on this block lets us observe whether the two subject groups differ due to non-treatment reasons. We take blocks 10 and 11 to approximate the steady-state bidding behavior under the two feedback types. The reason for why we focus on block 10 alongside with block 11 is that behavior in block 11 may be affected by last-period effects, whereas behavior in block 10 should be less so.[29] Looking at both of these blocks lets us see whether our results are sensitive to such potential effects. When considering statistical significance, unless noted otherwise, we employ two-sided tests at 95% significance level. In block 1 comparisons, standard errors are adjusted for clustering at subject level. For all other comparisons, standard errors are adjusted for clustering at bidding group level.

[27] See http://www.vse-lee.cz/eng/about-lee/about-us.

[28] For a comparison, an hourly wage that students could earn at the time of the experiment in research assistant or manual jobs typically ranged from 75 to 100 CZK.

[29] We observed some indication of a last block effect in some pilot experiments we ran to fine tune our design.

Figure 2: Average Bidding Function Slopes by Treatment and Block

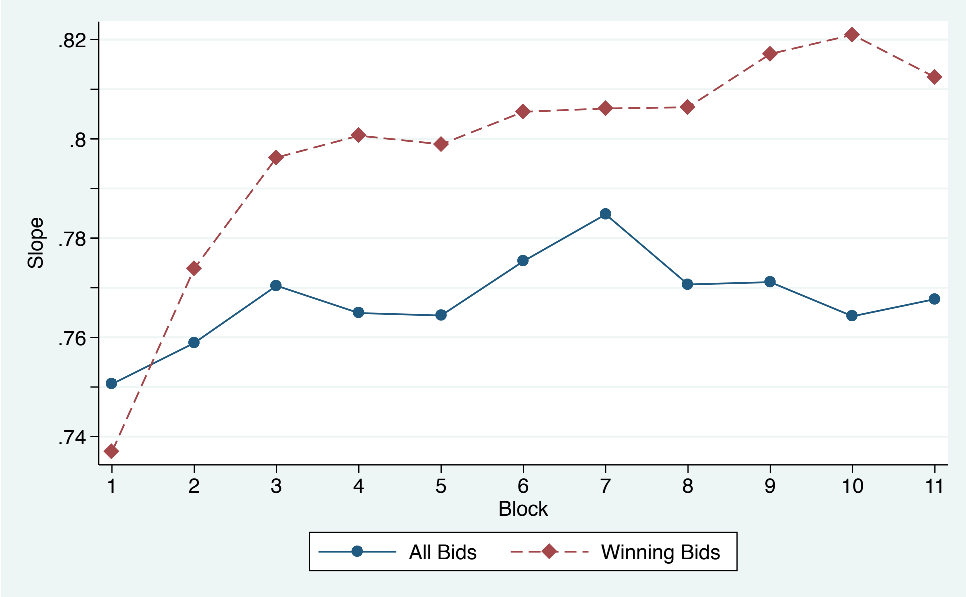

5.1 Bidding

As the initial step of the analysis, we estimate the slope of the bidding function for each individual subject in each block using OLS. This measure is a summary statistic of the behavior of a subject in a particular block. We assume that the bidding function is linear and has a zero-intercept. With v_{ijt} denoting the value and b_{ijt} denoting the corresponding bid of subject i in round j of block t, the estimate of slope for subject i in block t is given by

$$\widehat{slope}_{it} = \frac{\sum_{j=1}^{6} v_{ijt} b_{ijt}}{\sum_{j=1}^{6} v_{ijt}^2} = \sum_{j=1}^{6} \left(\frac{v_{ijt}^2}{\sum_{k=1}^{6} v_{ikt}^2} \right) \frac{b_{ijt}}{v_{ijt}}. \tag{9}$$

That is, the estimated slope is a square-value-weighted average of the six individual bid/value ratios in a given block. We exclude 9 subjects from the subsequent analysis due to slopes of their bidding strategies systematically exceeding 1.

Figure 2 plots the average of the bidding function slopes by treatment and block. This figure captures the overall findings. In block 1, which is treatment-free, behavior in the two treatments should differ only by noise. Indeed, the average slopes are very close in the two treatments at 0.751 (with the standard error of 0.012) in A and 0.737 (0.017) in W. Notice that such values are higher than in previous experiments (for instance, Katuščák et al. (2015), who use the same estimation procedure, find an average bid value ratio of 0.69). This could be due to the fact that in the current experiment we do not provide any information on the distribution of values, which implies a more uncertain (and ambiguous) environment for the subjects in the initial block. In the following blocks, a discrepancy between the two treatments arises.

Figure 3: Effect of Treatment (W-A)on the Average Bidding Function Slopes by Block

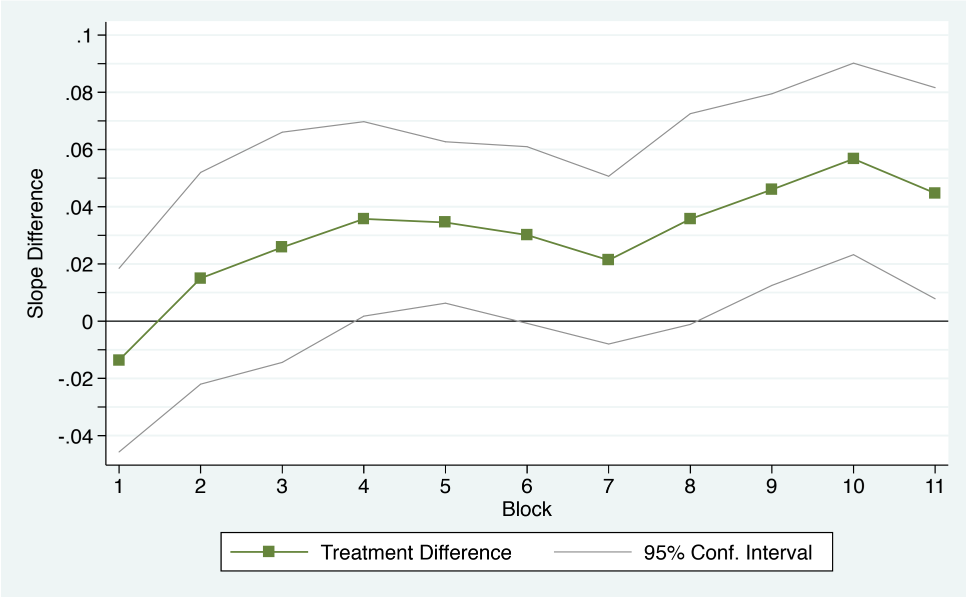

Subjects in W start bidding more than their counterparts in A. In the final blocks 10 and 11, the average slope in A is 0.764 (0.010) and 0.768 (0.011), respectively, changed little from block 1 (t-test p-values of 0.116 and 0.055, respectively). On the other hand, the corresponding average slopes in W are 0.821 (0.016) and 0.812 (0.018), respectively. In case of W, this constitutes a statistically significant increase over the average slope in block 1 (t-test p-values of 0.000 in both cases). In proportional terms, bids in W in blocks 10 and 11 are 7.5 percent and 5.7 percent, respectively, higher than bids in in the same blocks of A.

Figure 3 plots the estimate of the treatment effect (W minus A) on the average slope by block, together with its 95% confidence interval. The estimated treatment effect is positive in blocks 2 through 11 and it is statistically significant in blocks 4, 5 and 9 through 11. By the final blocks 10 and 11, the treatment effect reaches 0.057 (0.016) and 0.045 (0.018), respectively. Moreover, block-by-block, we compute the average bidding function slope in each bidding group (10 in each treatment) and we perform the Mann-Whitney ranksum test for the equality of distributions of the average slopes.[30] Starting from block 4 and with the exception of blocks 6 and 7, we reject the null hypothesis in favor of the distribution under W first-order stochastically dominating the one under A.

Figure 2 also shows that subjects in treatment W bid slightly less on average that subjects in treatment A do in block 1. This suggests that the former might bid less than the latter do if they were subjected to treatment A. In order to see what the treatment effect looks like after removing this non-treatment-related bidding difference, Figure 4 plots the difference in differences in average bidding strategy slopes between block i and block 1. By construction,

[30]Note that subjects do not interact across the individual bidding groups, so each bidding group presents a statistically independent observation.

Figure 4: Effect of Treatment (W-A) on the Average Bidding Function Slopes by Block after Subtracting the Treatment Effect Estimate in Block 1

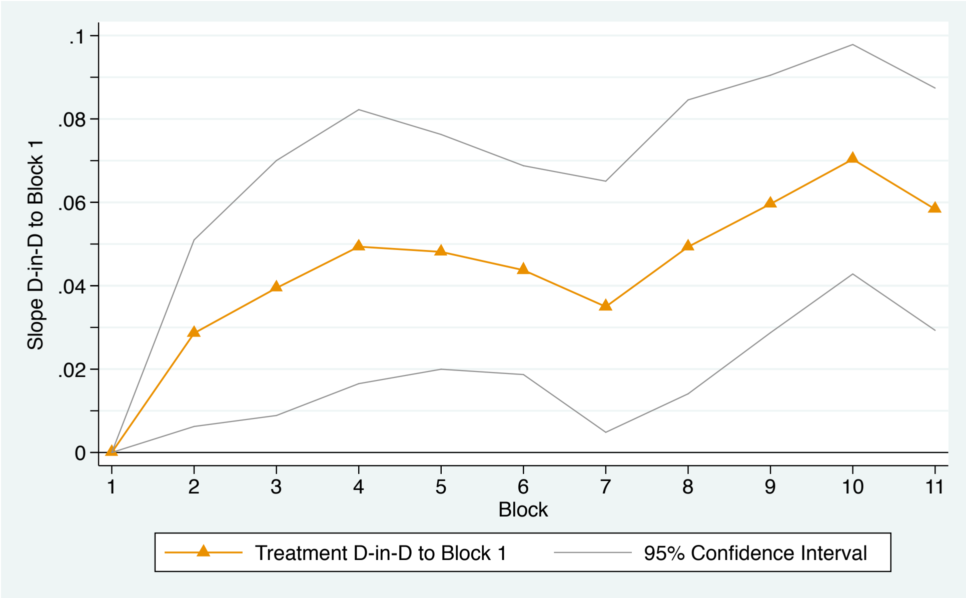

the difference-in-differences is zero in block 1. In all the remaining blocks, it is positive and statistically significant. This suggests that the lack of statistical significance for the treatment effect in some blocks as visible in Figure 2 might indeed be due to the heterogeneity of subjects across the two treatments.

To picture the overall impact of the treatment on bidding behavior, Figure 5 presents kernel estimates of the probability density function (pdf) of the bidding function slopes as well as its empirical cumulative distribution function (cdf) by the two treatments for blocks 1, 10 and 11. There is very little observable difference between the distributions under A and W in block 1 (Kolmogorov-Smirnov test p-value of 0.645). In contrast to that, in blocks 10 and 11, with an exception of a few slope realizations below 0.5, the distribution under W first-order stochastically dominates the distribution under A. Indeed, the Kolmogorov-Smirnov test rejects the null hypothesis of no difference between the two distributions in all blocks from block 3 onward (with the p-value being 0.001 or less from block 6 onward).

Put together, these findings show that in the "steady state", as approximated by blocks 10 and 11, W robustly and significantly shifts the distribution of bidding strategy slopes to the right relative to A. That is, given the very same values, subjects tend to bid more under W than they do under A. This finding supports the corresponding theoretical prediction that we laid out in Section 3.

As a side note, a corollary implication that arises from our results is that the "overbidding puzzle" observed in previous experiments on first-price auctions may be aggravated in those experimental designs in which subjects are given HI on past winning bids.

Figure 5: Distribution of Bidding Function Slopes by Treatment in Blocks 1, 10 and 11

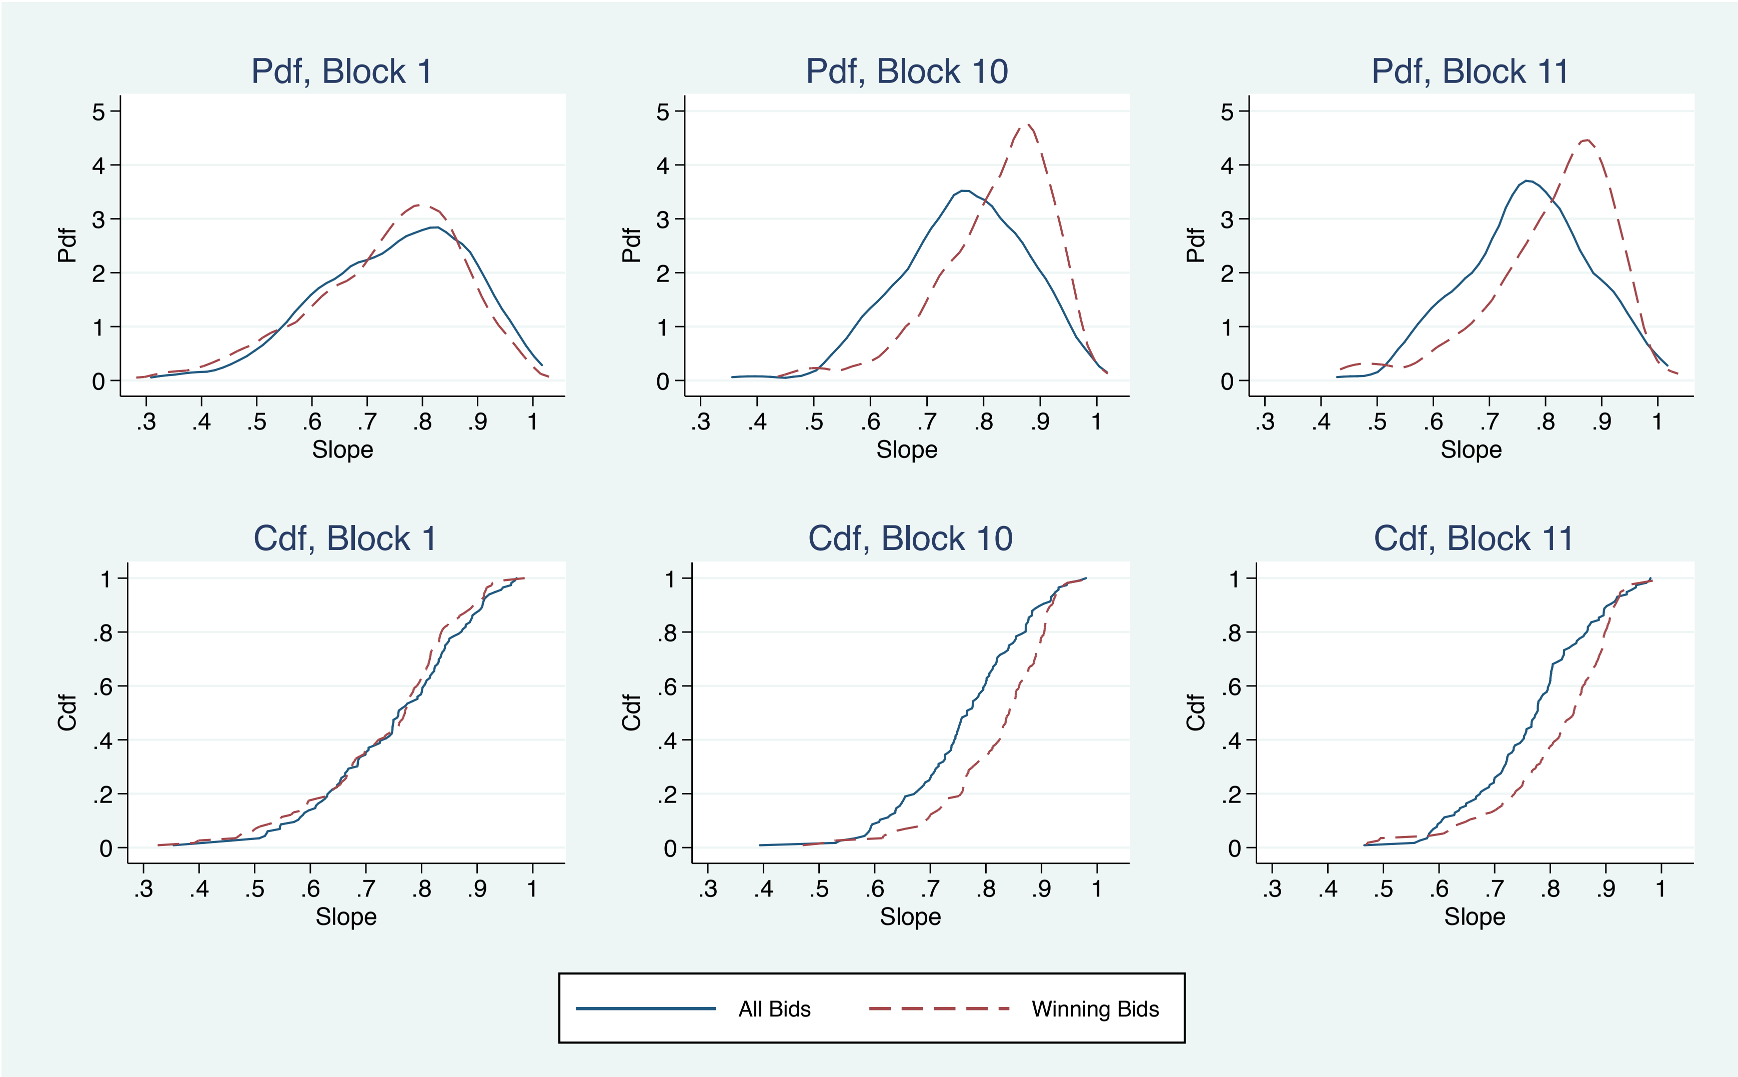

5.2 Average Revenue and Efficiency

Having discussed individual bidding behavior, we now switch our attention to group-level outcomes, namely average auction revenue and efficiency. Given a set of two-bidder auctions, we define average revenue as the average of the winning bids in these auctions. We define average efficiency as the ratio of the actually realized aggregate value and the maximum realizable aggregate value in these auctions. The latter is given by the sum of maximum values across the individual auctions.

In each block, there are many possible ways of matching subjects into pairs within a bidding group. Just in any single round, there are $11 \times 9 \times 7 \times 5 \times 3 \times 1 = 10,395$ unique ways of matching subjects into pairs. Moreover, since the order of value presentation is randomized across different rounds (within a block) and subjects see bidding feedback only after each block (rather than round), one should also consider matches across different rounds of a given block. To estimate the treatment effect on average revenue and efficiency in a given block, we generate 1,000 bootstrap draws from the data on values and bids in this block. Each draw is generated as follows. At the level of a bidding group, we first randomly (with the uniform distribution) draw one of the 6 rounds, separately and independently for each subject. Next, we randomly pair the 12 members of the bidding group into pairs. Any pattern of pairing is equally likely. We then use the very same pattern of selected rounds and pairs in each bidding group. Finally, we use the values and bids from the chosen rounds and the pairing pattern to determine the auction winners and average revenue and efficiency separately in each treatment.[31] Each bootstrap draw hence generates a matched pair of average revenues and a matched pair of average efficiencies, one for each treatment.

The resulting bootstrap data for blocks 1, 10 and 11 is plotted in Figure 6. Regarding the average revenue, the scatterplot for block 1 is symmetric around the diagonal. Precisely, the fraction of bootstrap draws for which the average revenue in W exceeds the one in A plus one half of the fraction of cases in which the two average revenues are equal, a measure we call "fraction revenue $W > A$", is equal to 0.394. In the absence of any difference between the two treatments in block 1, we would expect this measure to be close to 0.5. The fact that the measure is below 0.5 is due to the fact that subjects in W appear to bid a bit less in block 1 than subjects in A do (see Figure 2).[32] In contrast to that, the vast majority of points in the scatterplots for blocks 10 and 11 lie above the diagonal. Precisely, fraction revenue $W > A$ is equal to 0.894 in block 10 and 0.855 in block 11. Hence, average revenue appears to be systematically higher in W than in A. To capture the magnitude of the difference, the average of the ratio of revenue under W to that under A is 1.068 and 1.058 in blocks 10 and 11, respectively. Also, regressing the revenue in W on that in A and imposing a zero

[31] In case a bootstrap draw involves one of the 9 subjects who has been excluded due to overbidding, their matched data is excluded from computing the averages.

[32] See below for a formal statistical test of the null hypothesis that the fraction is equal to 0.5.

Figure 6: Average Revenue and Average Efficiency across the Treatments in Blocks 1, 10 and 11

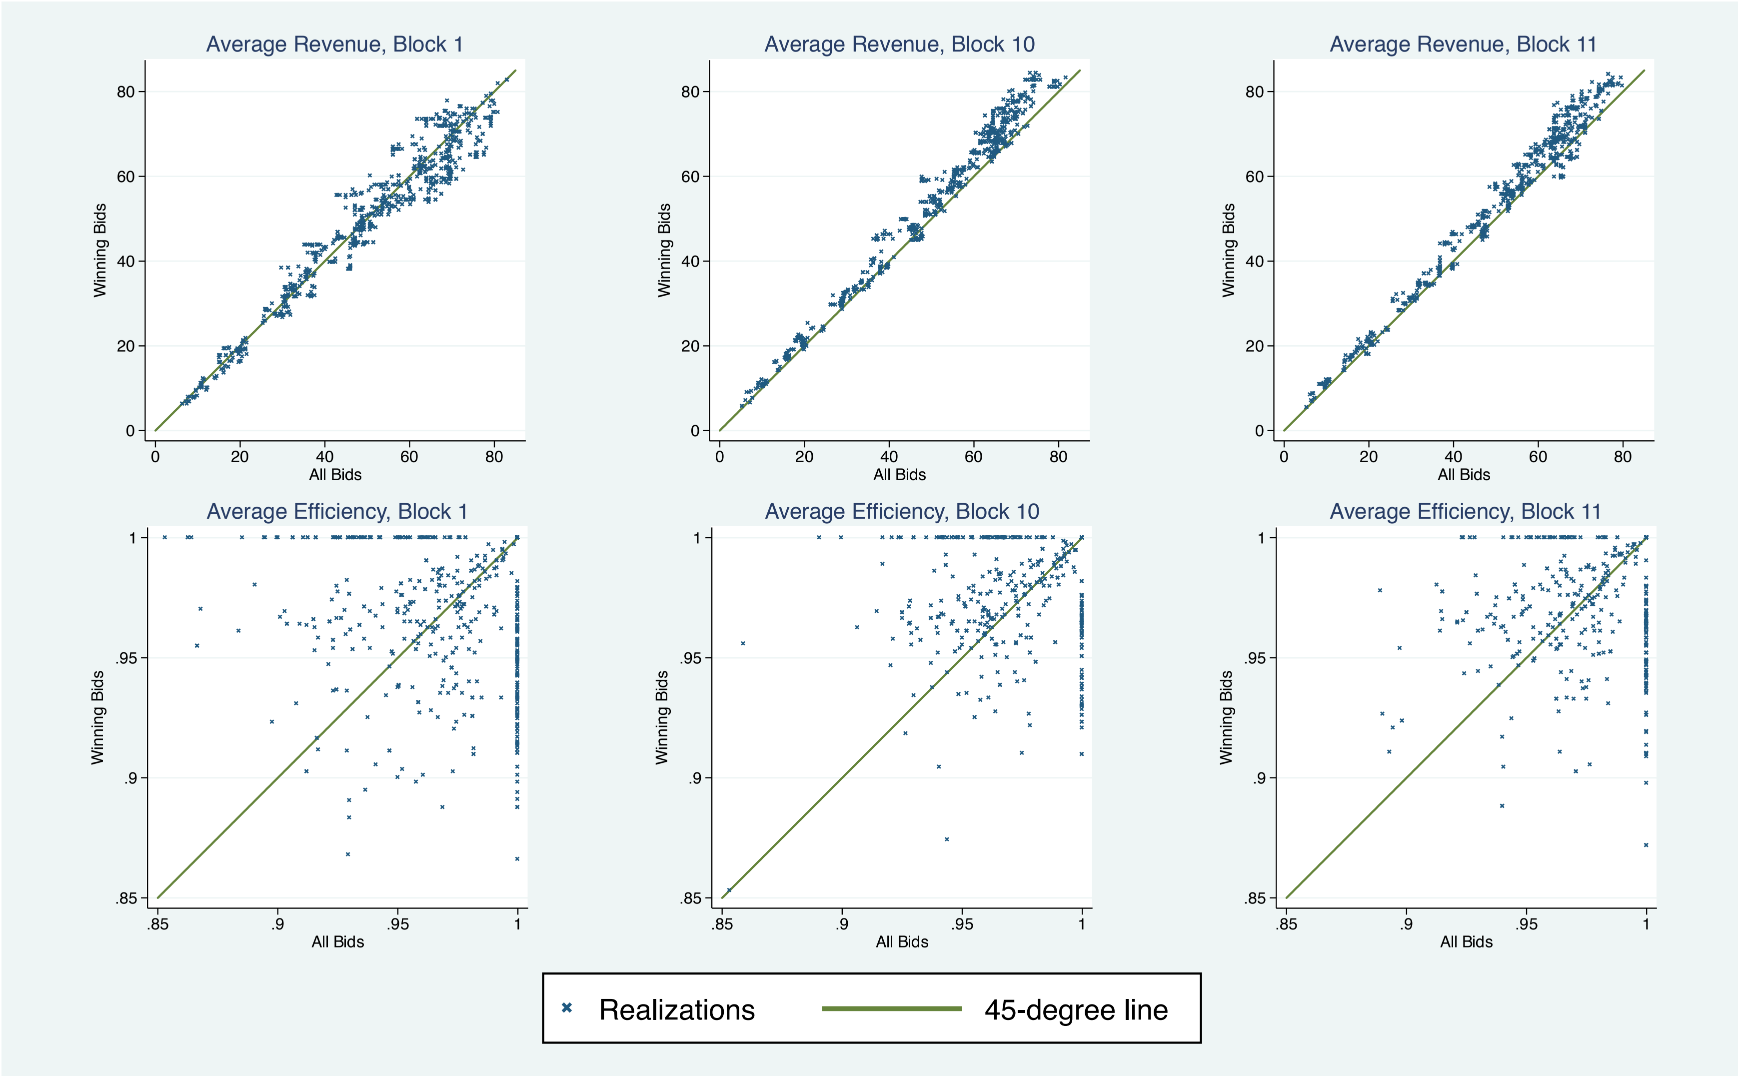

Table 1: Results of the “Superbootstrap” Tests

	Block 1	Block 10	Block 11
Fraction revenue	0.394	0.894	0.855
	(0.260)	(0.000)	(0.000)
Fraction efficiency	0.486	0.539	0.495
	(0.678)	(0.186)	(0.822)
Average revenue ratio	0.985	1.068	1.058
	(0.284)	(0.002)	(0.010)
Revenue regression slope	0.974	1.067	1.056
	(0.286)	(0.004)	(0.016)

Note: P-values for two-sided tests of the null hypothesis in parentheses.

intercept gives a coefficient of 1.067 and 1.056 in blocks 10 and 11, respectively.[33] Hence, on average, the steady-state average revenue is about 6 percent higher in W than it is in A. For a comparison, the analogous measures for block 1 are 0.985 (average revenue ratio) and 0.974 (revenue regression slope).

In contrast to the apparent treatment effect on average revenue, the impact on average efficiency is much less profound, if any. In all three blocks displayed in Figure 6, significant parts of the scatterplots are located on both sides of the diagonal. Precisely, defining “fraction efficiency $W > A$” analogously to fraction revenue $W > A$, the fractions are 0.486, 0.539 and 0.495 in blocks 1, 10 and 11, respectively.

We also conduct formal tests to examine these informal conclusions. The key is to obtain distributions of the various measures reported in the previous two paragraphs under the null hypothesis of no treatment effect.[34] To do that, we employ a “superbootstrap” procedure that works as follows. In each superbootstrap draw, we first randomly allocate the data from the 20 bidding groups into two groups of 10 and pretend that the first group corresponds to A and the second group to W. Then, based on this artificial treatment assignment, we obtain 1,000 bootstrap draws and compute the average revenue and average efficiency for each bootstrap draw using the same procedure as described above for the original bootstrap draw. Next, we compute all four statistics of interest (fraction revenue $W > A$, fraction efficiency $W > A$, average revenue ratio, revenue regression slope). We repeat the whole procedure 1,000 times, thus obtaining 1,000 superbootstrap realizations of the four statistics under the null hypothesis. We then examine where in these distributions the original statistic realizations are located and compute P-values for two-sided tests of the respective null hypotheses. The null hypotheses

[33]Note that the regression coefficient is simply a square A-revenue-weighted average of the 1,000 revenue ratios. The reasoning is analogous to the one following equation 9.

[34]Note that using t-tests based on conventional standard errors for averages and regression coefficients is inappropriate since the 1,000 bootstrap draws are not statistically independent.

are that that fraction revenue $W > A$ is equal to 0.5, fraction efficiency $W > A$ is equal to 0.5, the average revenue ratio is equal to 1 and the revenue regression slope is equal to 1. The results are reported in Table 1. In block 1, we do not reject the null hypothesis for any of the four measures. Also, we do not reject the null hypothesis for fraction efficiency $W > A$ in either of the blocks 10 and 11. To the contrary, we do reject the null hypothesis for fraction revenue $W > A$, the average revenue ratio and the revenue regression slope in blocks 10 and 11. Hence the formal test results confirm the informal conclusions drawn above.

Overall, these findings show that the in the "steady state", as approximated by blocks 10 and 11, W significantly increases the average revenue realizations in comparison to A. That is, given the very same values, the auctioneer realizes a higher auction revenue on average. This finding supports the corresponding theoretical prediction that we laid out in Section 3.

5.3 Estimation of Fraction of Naive Bidders (In Progress, Available Soon)

6 Concluding Remarks

We have compared two disclosure of historical information (HI) in first-price auctions: all bids and winning bids only. Under *standard* theory of Bayesian Nash Equilibrium with fully rational bidders, the choice among the two disclosure policies should be irrelevant in the steady state. Instead, our experimental test shows that actual bidding behavior differs and that HI on winning bids should be preferred to HI on all bids in terms of revenue maximization. We have proposed that, under HI on winning bids, a fraction of bidders might mistakenly best respond to the distribution of winning bids rather than to the one of all bids. Our work, on the theory side, challenges the prediction that historical market information should not matter for equilibrium outcomes as long as it is rich enough for forming correct beliefs. On the market design side, it suggests that, in terms of revenue, a long term auctioneer should disclose past winning bids only. Even though our experimental test focuses on auctions with two-bidders, the market design implications should extend beyond that case and also to other market institutions in which the determination of beliefs about the behavior of the opponent(s) plays a crucial role.

A sufficient statistic for the formulation of a bidder's best response is the distribution of the highest competing bid (DHCB). In the case of two bidders only, the HI that bidders get under HI in all bids is, in the steady state, exactly the DHCB. That is, the bidders are given exactly the distribution they need to best-respond to in the Bayesian Nash Equilibrium. Hence the two-bidder case provides the natural benchmark because the bounded rationality that we envision should be limited only to the case of HI on winning bids.

The selection bias theory that we propose can be extended to settings with more than two bidders. For example, even though the naive bidders fail to understand that bids and winning bids are different objects, one could assume that they otherwise perform all the necessary steps, including computation of order statistic distributions, to determine the DHCB.[35] Under this assumption, naive bidders, when presented with HI on all bids, $A(\cdot)$, form correct beliefs about the DHCB, $A^{n-1}(\cdot)$. Instead, when presented with HI on winning bids they use $W(\cdot)$ in place of $W^{\frac{1}{n}}(\cdot)$ to form beliefs about the DHCB. After computing the order statistic distribution, they believe that the DHCB is given by $W^{n-1}(\cdot)$ rather than $W^{\frac{n-1}{n}}(\cdot)$. As a result, they end up bidding more than their best response under correct beliefs.

An alternative type of mistake that bidders might make is to confuse the HI they are provided with for the DHCB. This type of mistake would imply that, under HI on all bids, naive bidders best-respond to $A(\cdot)$, whereas under HI on winning bids, they best-respond to $W(\cdot)$. As a result, naive bidders end up bidding less than their best response under correct beliefs in the former case and more in the latter case.

The predictions based on these two alternative motivations are indistinguishable in auctions with two bidders. For more than two bidders they differ.[36] Importantly, for a market design perspective, they share the feature that the predicted effect of HI on bidding increases with the number of bidders.[37]

The discussion above hints at the variety of biases that might be at play in more complex environments than the one we looked at. It also implies that there is a vast scope for future research in order to pin down more precisely the biases bidders are subject to in various market environments in response to different types of historical market information.

7 Acknowledgements

Special thanks go to Tomáš Miklánek for his excellent job with programming and organizing the experiment. This study was supported by the Grant Agency of the Czech Republic (GACR 402/11/1726). All opinions expressed are those of the authors and have not been endorsed by GAČR, Paris School of Economics, University College London, University of Arkansas, University of Economic in Prague or CERGE-EI.

[35]Notice that for the two-bidder case, it does not matter if bidders cannot compute order statistic distributions, as they are given the DHCB directly.

[36]The first approach predicts that the revenue difference between the two HI types is entirely driven by incorrect DHCB beliefs of naive bidders under HI on winning bids. Moreover, the overbidding in the latter case grows larger with an increasing number of bidders. Conversely, under the second approach, the overbidding compared to the best reply under HI on winning bids diminishes with the number of bidders. On the other hand, underbidding compared to the best reply under HI on all bids increases with the number of bidders.

[37]This hinges on the fraction of naive bidders being constant as n grows larger, which admittedly could be unrealistic over some range.

References

Battigalli, Pierpaolo, "Comportamento Razionale ed Equilibrio nei Giochi e Nelle Situazioni Sociali," Dissertation, Bocconi University, Milano 1987.

Bergemann, Dirk and Johannes Hörner, "Should Auctions be Transparent?," 2010. COWLES FOUNDATION DISCUSSION PAPER NO. 1764.

Cason, Timothy, Karthik Kannan, and Ralph Siebert, "An Experimental Study of Information Revelation Policies in Sequential Auctions," *Management Science*, 2011, *57* (4), 667–688.

Chen, Yan, Peter Katuščák, and Emre Ozdenoren, "Why Can't a Woman Bid More Like a Man?," *Games and Economic Behavior*, January 2013, *77* (1), 181–213.

Cox, James C., Vernon Smith, and James M. Walker, "Auction Market Theory of Heterogenous Bidders," *Economics Letters*, 1982, *9*, 319–325.

__ , __ , and __ , "Theory and Misbehavior of First-Price Auctions: Comment," *The American Economic Review*, December 1992, *82* (5), 13921412.

Dekel, Eddie, Drew Fudenberg, and David K. Levine, "Learning to Play Bayesian Games," *Games and Economic Behavior*, February 2004, *46* (2), 282–303.

Dufwenberg, Martin and Uri Gneezy, "Information disclosure in auctions: an experiment," *Journal of Economic Behavior and Organization*, 2002, *48*, 431–444.

Engelbrecht-Wiggans, Richard, "The Effect of Regret on Optimal Bidding in Auctions," *Management Science*, 1989, *35* (6), 685–692.

__ and Elena Katok, "Regret and Feedback Information in First-Price Sealed-Bid Auctions," *Management Science*, April 2008, *54* (4), 808–819.

Esponda, Ignacio, "Information Feedback in First Price Auctions," *RAND Journal of Economics*, Summer 2008, *39*, 491–508.

Filiz-Ozbay, Emel and Erkut Y. Ozbay, "Auction with Anticipated Regret: Theory and Experiments," *American Economic Review*, April 2007, *54* (4), 1407–1418.

Fischbacher, Urs, "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Fudenberg, Drew and David K. Levine, "Self-Confirming Equilibrium," *Econometrica*, May 1993, *61* (3), 523–545.

Füllbrunn, Sascha and Tibor Neugebauer, "Varying the number of bidders in the first-price sealed-bid auction: experimental evidence for the one-shot game," *Theory and Decision*, 2013, *75* (3), 421–447.

Greiner, Ben, "An Online Recruitment System for Economic Experiments," *Kurt Kremer, Volker Macho (Eds.); Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63, Göttingen : Ges. für Wiss. Datenverarbeitung*, 2004, pp. 79–93.

Huck, Steffen, Philippe Jehiel, and Tom Rutter, "Feedback spillover and analogy-based expectations: A multi-game experiment," *Games and Economic Behavior*, 2011, *71* (2), 351–365.

Isaac, Mark R. and James M. Walker, "Information and Conspiracy in Sealed Bid Auctions," *Journal of Economic Behavior and Organization*, 1985, *6*, 139–159.

Jehiel, Philippe, "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 2005, *123*, 81–104.

__ , "Manipulative Auction Design," *Theoretical Economics*, 2011, *6* (2), 185–217.

__ **and Frédéric Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Katuščák, Peter, Fabio Michelucci, and Miroslav Zajíček, "Does Feedback Really Matter in One-Shot First-Price Auctions?," *Journal of Economic Behavior and Organization*, 2015, *Forthcoming.*

Neugebauer, Tibor and Javier Perote, "Bidding 'as if' risk neutral in experimental first price auctions without information feedback," *Experimental Economics*, 2008, *11*, 190–202.

__ **and Reinhard Selten**, "Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets," *Games And Economic Behavior*, 2006, *54*, 183–204.

Ockenfels, Axel and Reinhard Selten, "Impulse Balance Equilibrium and Feedback in First Price Auctions," *Games And Economic Behavior*, 2005, *51* (1), 155–170.

Pearson, Matthew and Burkhard C. Schipper, "Menstrual Cycle and Competitive Bidding," *Games and Economic Behavior*, March 2013, *78* (2), 1–20.

Sonsino, Doron and Radosveta Ivanova-Stenzel, "Experimental internet auctions with random information retrieval," *Experimental Economics*, 2006, *9* (4), 323–341.