How to Boost Revenues in FPAs? The Magic of Disclosing only Winning bids from Past Auctions

Philippe Jehiel* Peter Katuščák† Fabio Michelucci‡

January 25, 2014

Abstract

We show that an auctioneer (such as an auction house, or a procurement agency) should disclose historical information about winning bids from past auctions, as opposed to disclosing all bids, because this induces bidders to bid more, and raises revenues. We provide a theoretical explanation for our experimental findings based on a bias: at least some subjects may miss that winning bids are not representative of all bids when processing the available information about past bids.

Keywords: auctions, bidding, feedback, mechanism design

JEL Classification: C91, C92, D44

*Paris School of Economics and Department of Economics and ELSE, University College London, jehiel@enpc.fr

†CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Academy of Sciences of the Czech Republic, Politických vězňů 7, 111 21 Prague, Czech Republic, Peter.Katuscak@cerge-ei.cz.

‡CERGE-EI, Fabio.Michelucci@cerge-ei.cz

1 Introduction

What determines people's beliefs about the private information of their potential opponents? Game theoretic models standardly assume that players share common knowledge about such beliefs. In particular, auction theorists in a private value environment assume that bidders have common knowledge about the distribution function(s) from which their valuations are drawn. What do we think might happen in "real" auctions? Where do beliefs come from? In this paper we take the stand that bidders beliefs must come from some information about similar environments that comes from earlier auctions. This observation leads to the very important (and novel) market design issue of whether or not a market designer who has control about the information released about past auctions (such a procurement agency, or a auction house) should adopt a specific disclosure policy. We show that this is the case: in particular, providing information about winning bids dominates revenue-wise providing information about all bids. To sum up, one contribution we have is about the more fundamental problem of beliefs formation. The second one, is to the more applied market design question of whether the provision of specific feedback from past auctions affect bidding behavior in a systematic manner.

Our experimental evidence comes from bidding behavior in First Price Auctions (FPAs) in which bidders receive information about bids from previous, but otherwise identical, FPAs. A brief description of the experimental design is due to clarify where our contribution stands.[1] In each experimental session we had 24 subjects divided into 2 groups. Each subject was competing in several FPAs, each time against one opponent randomly selected from the subjects in his/her own group. The auctions were organized in 11 blocks of 6 auctions.

Block 1 was identical for the two groups: subjects had to bid only upon the knowledge that they were facing one opponent whose value was generated by the very same random number generator device as theirs. No information about the distribution nor the support of possible values was provided. Subjects played 6 auctions, that is they received 6 different values (sequentially) and they were asked to place 6 corresponding bids. Notice that Block 1 bidding differs from the standard implementation used in auction experiments where the exact distribution (and the support) from which values are drawn is communicated to subjects. Our alternative implementation was dictated by the fact that we wanted beliefs about the private information held by the opponents to arise endogenously (rather than being provided exogenously to the subjects) in order to check how bidders (learn to) best reply to selective information from past auctions. We also find it rather natural because we embrace the view of the (Bayes) Nash equilibrium of the game as the steady state of a learning process.

In Block 2 to 11 subjects went through the very same 6 auctions as in Block 1, so that the environment presented is stationary except in one dimension: in each of those blocks subjects received historical information about the bidding in the previous block, prior to their

[1] For the comprehensive description of the design, see Section 4.

bidding decision. Thus, our design allows to evaluate how initial beliefs evolve once historical information becomes available. In one treatment the historical information is about all the bids placed in the previous block's auctions (by subjects in the same group), while in the other the information is about the winning bids.

More importantly, the across treatment comparison allows us to evaluate if the provision of a specific type of information affects behavior in any systematic way.[2] The striking conclusion we reach is that the provision of feedback about winning bids increases average bid/value ratios by roughly 9%, which translates into a similar effect on revenues. This increase is the more remarkable once one observes that it comes at practically no cost. Our empirical analysis shows that there is a lot of initial heterogeneity in bidding behavior in Block 1 and that providing feedback on winning bids (relatively to providing feedback on all bids) induces those subjects who were shading more in Block 1 to bid substantially more aggressively. As for subjects that started with higher bid/value ratios in Block 1, we find that feedback on winning bids still has the effect of increasing bids (though to a lesser extent), while feedback on all bids if anything has the opposite effect of (mildly) reducing bid value/ratios.

The rest of the paper is structured as follows. Section 2 sums up related theoretical and experimental work about the provision of feedback in FPAs. Section3 presents an alternative theoretical explanation for the role of feedback in market design. Section 4 presents details of our experimental design. Section 5 presents our results, and Section 6 concludes. The Appendix contains the set of instructions and the demographic questionnaire we used at the end of the experiment.

2 Related Theories and Experimental results

In this section we briefly review the available experimental evidence about the effect of feedback. We do not enter into the details of the specific papers, but rather we contrast their approach with ours, since we look at the question of the effect of feedback from a rather different and novel point of view. We also review the available alternative theoretical explanations raised to explain such effect. The following section proposes an alternative explanation.

The existing experimental literature has looked at the effect of alternative feedback types on bidding using a variety of design protocols that differ regarding the type of feedback investigated, and the use of human or computerized opponents (and their numbers). The common feature among those papers (see, for instance, (Isaac and Walker 1985, Ockenfels and Selten 2005, Neugebauer and Selten 2006, Neugebauer and Perote 2008, Engelbrecht-Wiggans and Katok 2008)) and ours is the use of an environment where subjects are called to

[2]For such comparison we are interested in bidding behavior between the two treatment from the last blocks, as it is the one that is representative of the stationary equilibrium of the game.

bid in a sequence of repeated auctions upon observing some feedback.[3] The dividing feature is that all such papers provide feedback about own auction only, while we provide feedback about a set of past auctions. This difference in design has two crucial consequences: one more practical, and one more theoretical. The practical one is that our study predicts an effect of feedback even for one shot environments. This is so because under the interpretation of the (Bayes) Nash equilibrium as the steady state outcome of a learning process, we can take the bidding behavior in the last blocks of our experiment as a good approximation for what we should see in a one shot game if bidders have had the opportunity to fine tune their beliefs about the environment from some previous historical information. On the contrary, the predictions from the previous evidence only apply to auction markets where the participants bid in many subsequent auctions, and therefore have only more limited applicability from a market design perspective. The novel theoretical implications of our design come from the fact that we ensured that subjects could in no way verify until the very end of the experiment the payoff relevant outcomes. This implies that if they changed their decision over time, it was because of the informational content of the feedback they got, but not because of potential regret over past decisions. Thus, we provide additional input to the theoretical discussion regarding why the type feedback provision matters. Of the two explanations we are aware of, one is based on regret (Engelbrecht-Wiggans (1989), Filiz-Ozbay and Ozbay (2007)), and the other on learning (the so called impulse balance equilibrium, see Ockenfels and Selten (2005), Neugebauer and Selten (2006)). The basic idea behind regret theories is that giving feedback about the winning bid makes regret salient and induces bidders to bid more aggressively to avoid experiencing such emotional response. Under this explanation, if the regret is anticipated, one could have a feedback effect even in a one shot game (Katuščák et al. (2013) show experimentally that this is not the case). The learning explanation, as formulated in the papers cited above, works only if bidders play repeatedly and receive feedback on the outcome of the own auction. In the event a bidder loses (wins) an auction round he/she experiences an upward (downward) impulse, and in the long run it is assumed that a bidder bids in order to balance (in expectation) such impulses. As one can gather from this brief summary, the two approaches, while providing different interpretations, have in common that the key element is the experience of feedback on the outcome of own auction. Our design shows that feedback on winning bids raises bid/value ratios even if the necessary element advocated by the existing theories is missing. Therefore, a different explanation is at play in our environment. This does not exclude that previous experimental findings could be due to regret, or reasons in line with

[3] The only exception to this environment are Filiz-Ozbay and Ozbay (2007), Katuščák, Michelucci and Zajíček (2013), and Ratan (2013), who look at a one shot environment. Filiz-Ozbay and Ozbay (2007) is the first paper to look at such an environment (with 4 human bidders) and to find that the type of feedback matters. However, Katuščák et al. (2013) have shown that their finding is not robust using a larger sample, and a variety of different treatments that include both settings with human opponents and computerizes opponents. Ratan (2013) also finds no effect of feedback in a one shot game, but his paper looks at computerized opponents only.

the the impulse balance equilibrium. But it also opens to the possibility that the main force at play in those environments could be the same behind our findings here.

In the next section, we elaborate that in the winning bid treatment subjects might erroneously best reply to the distribution of winning bids as if it was the distribution of all bids. Such type of mistake drives up bid/value ratios. We conjecture that something similar could be happening in other experimental settings as well. As a matter of fact, in a sense what our environment gives is just a richer feedback (a full distribution) rather than a single observation each time.

3 An Alternative Explanation

The present section proposes that the specific types of feedback administered in our experimental setting might yield different bidding functions if bidders are not fully rationale. In particular, we call a bidder naive if, under the winning bids feedback, he/she mistakenly best replies to the empirical distribution of winning bids as if it were the distribution of all bids. It is not surprising that naive bidding results in higher bids: by assumption, bidders are best responding to more aggressive bids as compared to the all bids feedback (for which no mistake or bias is assumed). In fact, under risk neutrality and uniform distribution, they behave as if they were in the all bids feedback with 3 rather than 2 bidders. Below we illustrate that the same qualitative change goes through if bidders are not fully naive and/or not risk neutral.

3.1 All Bids Treatment

Suppose that G is the steady-state cdf of bids, with the associated pdf being g. Suppose that bidders best-respond to this distribution. Then, given a value v, a bidder solves

$$\begin{aligned} b(v) &= \arg\max_b (v-b)^\alpha G(b) \\ &= \arg\max_b \alpha \ln(v-b) + \ln G(b), \end{aligned}$$

where $\alpha > 0$ indexes the bidder's risk attitude. $\alpha = 1$ corresponds to risk-neutrality, $\alpha \in (0,1)$ to risk-aversion and $\alpha > 1$ to risk-loving. The necessary FOC is

$$-\frac{\alpha}{v-b(v)} + \frac{g[b(v)]}{G[b(v)]} = 0,$$

which gives

$$v = \alpha \frac{G[b(v)]}{g[b(v)]} + b(v).$$

Assuming that $b(\cdot)$ is strictly increasing, we have that

$$\begin{aligned} G(x) &= \Pr\{b(v) \leq x\} \\ &= \Pr\{v \leq b^{-1}(x)\} \\ &= \Pr\left\{v \leq \alpha\frac{G(x)}{g(x)} + x\right\} \\ &= F\left(\alpha\frac{G(x)}{g(x)} + x\right), \end{aligned}$$

where F is the cdf of the value distribution. This differential eqution characterizes the steady-state bid distribution as a function of the value distribution.

If F is uniform on $[0, 1]$, then we have that

$$G(x) = \alpha\frac{G(x)}{g(x)} + x$$

with the boundary condition $G(0) = 0$. Guessing the solution to be $G(x) = \gamma x$, we have

$$\gamma x = \alpha x + x,$$

giving

$$G(x) = (1 + \alpha)x$$

with the support $x \in [0, 1/(1 + \alpha)]$. That is, bidders bid uniformly on $[0, 1/(1 + \alpha)]$. In fact,

$$\begin{aligned} b^{-1}(x) &= \alpha\frac{G(x)}{g(x)} + x \\ &= (1 + \alpha)x, \end{aligned}$$

which gives the equilibrium bidding function

$$b(v) = \frac{1}{1 + \alpha}v.$$

The standard risk-neutral Nash is a special case with $\alpha = 1$.

For completeness, note that, under this solution, the bidder's objective is strictly concave in b, so the FOC is sufficient as well.

3.2 Winning Bids Treatment

Suppose that G is the steady-state cdf of bids, with the associated pdf being g. But instead of being given G, bidders are given the distribution of winning bids, i.e., G^2. Suppose that bidders best-respond to the perceived distribution of bids being G^β, with $\beta \in [1, 2]$. $\beta = 1$

corresponds to full rationality (i.e., correclty backward-engineering the observed distribution to obtain G), $\beta = 2$ corresponds to full naivite (treating G^2 as if it were the distribution of bids), with $\beta \in (1, 2)$ corresponding to partial naivite. Then, given a value v, a bidder solves

$$\begin{aligned} b(v) &= \arg\max_b (v-b)^\alpha G^\beta(b) \\ &= \arg\max_b \alpha \ln(v-b) + \beta \ln G(b), \end{aligned}$$

where $\alpha > 0$ indexes the bidder's risk attitude. $\alpha = 1$ corresponds to risk-neutrality, $\alpha \in (0, 1)$ to risk-aversion and $\alpha > 1$ to risk-loving. The necessary FOC is

$$-\frac{\alpha}{v - b(v)} + \beta \frac{g[b(v)]}{G[b(v)]} = 0.$$

The rest of the solution follows along the same lines as above with α replaced by α/β. That is, (partial) naivite has the same effect as additional risk aversion. The steady-state bid distribution is then given by

$$G(x) = \frac{\alpha + \beta}{\beta} x$$

with the support $x \in [0, \beta/(\alpha + \beta)]$. That is, bidders bid uniformly on $[0, \beta/(\alpha + \beta)]$. In fact, the equilibrium bidding function is

$$b(v) = \frac{\beta}{\alpha + \beta} v.$$

Hence, naive bidding increases bids.

4 Experimental Design

The focus of our experimental design is the comparison of bidding behavior in two treatments, which from now on we refer to as treatment A and treatment W. In both treatments, subjects had to play repeatedly in FPAs markets consisting of 2 bidders (that is one single opponent). The only difference between the treatments is that in treatment A subjects received feedback about bidding in past auction about all bids, while in treatment W they received feedback about winning bids only. Subject could proceed through the FPAs rounds at their own pace. The more detailed description of the experimental design is the following. We ran 8 sessions of 24 subjects each, for a total of 192 subjects, of which 96 in treatment A, 96 in W. Rather than running 4 sessions of treatment A and 4 of W, in order to avoid the possible issue of endogenous selection due to the day/time of the day, we opted to run 8 sessions, each with 12 subjects in treatment A and 12 subjects in W. Thus, each session was identical and here we can focus on the description of a representative one.

For each session, we recruited males and females separately to ensure a balanced gender composition and to be able to control for any eventual gender effect.[4] All subjects were given the very same written instructions at the start of the experiment. They also received a white piece of paper and a pencil in case they wanted to take notes. The instructions explained (in non technical terms) that the participants in the experiment were divided into 2 groups of 12 and that each of them would compete in a series of FPAs each time against one opponent, which was randomly selected via an independent draw from the other members of the same group. Subjects were also explained the rules of the FPA, the fact that they would play 11 blocks of 6 auctions each, and that they would get paid on the basis of 5 randomly chosen auctions at the exchange rate of 1 point=10 CZK (on top of this each subject knew he/she would be paid 100 CZK show-up fee). There are two main reasons why a block consists of multiple auctions in our design. First, we wanted each subject to experience bidding for different values so that we could estimate the individual bidding functions (see value generation below for details). Second, we needed to form a rich enough history of past bids. Six auctions was a good compromise between achieving those objectives and not overloading subjects with too many auctions. The reason why we chose 11 rather than perhaps the more natural choice of 10 blocks is because we had some indication of a last block effect from some pilots we ran to fine tune the design. Subjects were also told that they would not get any feedback regarding whether they won or not any given auction until the end of the experiment. Recall that this is one of the main novelty of our design as it allows to disentangle the informational impact of a specific feedback type from potential regret effects. Further, they knew they would receive new instructions at the end of Block 1 to explain the rules for Block 2 to 11. In terms of information we provided subjects in Block 1, this consisted only of telling them that for every auction they would receive a valuation determined by a random number generator, and that for each auction a new independent value generation would be applied. Subjects were not told neither the distribution, nor the support that was used for the random draw of their valuations (however, to ensure that the environment was fair and symmetric we told them that the number random generator used was the same for all of them).

In practice, we applied for each subject in a group 6 draws where each draw i, $i = 0, .., 5$, was from a uniform distribution with support $[(100/6)i, 100/6)(i + 1)]$. The 6 draws were independent across subjects. We prefer this way of drawing subjects value over drawing the 6 values independently from a uniform distribution with support $[0, 100]$ because it guarantees noise reduction in the estimation of the bidding function.

The very same value realizations randomly drawn for a treatment were replicated for the other treatment. Moreover, we ensured that not only for any auction in a treatment there was an identical auction in the other treatment, but further that the subjects playing a specific set

[4]Out of 8 sessions, 6 had 12 males/12 females, 1 had 13males/11females, 1 had 14 females/10 males.

of values across the two treatments were of the same gender.[5] All the above measures were taken to have the closest possible environment between the two treatments up to the feedback variation. Before bidding in Block 1, all subjects were administered a quiz (see Appendix) to verify their understanding of the instructions. One of the experimenters checked the answers and ensured the subject had a good understanding. Incorrect answers were infrequent. Notice that there are no treatment differences in Block 1, thanks to the fact that we opted to use 2 sets of instructions sequentially. The reason behind this design choice was to have subjects starting from a perfectly symmetric environment so that bidding in Block 1 could not be influenced by the anticipation of a different feedback later on.

After bidding in Block 1 was completed a second set of instructions was administered, this time differing according to whether the subject belonged to the group playing treatment A or W. In this second set of instructions, subjects were told that in all subsequent blocks they would play exactly the same auctions (that is with the very same value realizations) they played in Block 1 except that some information about bidding in the previous block would be available. In particular, in treatment A subjects were told that from Block 2 onwards they would receive information about all bids from all the auctions played in the previous block by their group.

The representation of the information from past block was a key element to ensure the success of our design. We paid special attention both to the graphical layout of the onscreen information, as well as to the written explanation of how to make use of the information. We decided to present the information about bids from the previous block's 36 auctions (6 auction times 6 rounds, per group) using two instruments. One histogram with 10 bars, each representing the percentage of all bids (for treatment A) placed within the interval specified by the bar. This was to have a simple representation of the empirical distribution of past bids via an instrument subjects should be familiar with (histograms are commonly used in newspapers and magazines). A possible drawback of using such instrument is that, regardless of the choice of how many bars to adopt, it does not allow to check more detailed information about the empirical distribution. To circumvent this problem, we also gave subjects the possibility to enter 2 numbers into 2 boxes to get the exact information about the percentage of all bids that lied between those two numbers (entering only one number in the left (right) box would give the percentage of all bids above (below) such number). Finally, and on built calculator was available to subjects by pressing on an icon. The printed instructions provided subjects with a screenshot of the screen they would see so that they could familiarize with the interface containing the information about past bids. Before restarting the experiment to collect bidding for Blocks 2 to 11, a quiz (see Appendix) was distributed to both groups to check the understanding of the second part of the instructions. Again, when verifying the quizzes we could

[5]There is only one exception to this pattern, because in one session we had an odd number of males and females (13 and 11, respectively). Thus, the statement is true for 95 out of 96 cases.

observe a very good understanding from subjects. Notice that the since all subjects played with the same values across blocks, in case the information they saw changed it must have been because of changes in bidding behavior.

All that we said for treatment A holds true for the W up to two differences. The first one is the natural one that whenever above we referred to "all bids", for the W treatment it should be read as "winning bids". The second one is a design choice we have opted for in order to preclude that a subject could find out whether his/her bids from the previous block were winning bids or not (by using accordingly the 2 boxes provided). This is important to avoid that any possible consideration related to regret might influence bidding behavior. The way this was done was by explaining subjects that two different pairings were used to form an auction pair. For what really mattered to them, that is earnings, we used what we referred to as Earning Pairing. To compute the information they received after each block, we used what we referred to as Information Pairing. The Earning pairing is identical to the pairing used in the A treatment. The Information Pairing uses a random matching of subjects into auction pairs that is independent from the Earning pairings. It is employed to determine the information about Block 1 that is provided in Block 2, and the very same pairing is repeated for all subsequent blocks. This to ensure that, exactly as in treatment A, in case the information provided to subjects changed, it would be solely due to a change in bidding behavior. The other alternative to exclude regret effects would have been to exclude the own auction from information provided to a subject. We did not opt for this because it would have implied providing a different feedback to each bidder. This concludes the description of a representative session. All other sessions were identical. In fact, for all 8 sessions we used exactly the same set of values. This is also a design choice dictated by the fact that we are interested in computing and comparing individual bidding functions and to reduce noise we want to compute them exploiting the same set of values across sessions.

4.1 Other Stages

When a subject had completed bidding in the 66 auctions (11 blocks of 6 auction rounds), we administered a demographic questionnaire in which we collected information about age, country of origin, number of siblings, academic major, the highest achieved academic degree, self-reported risk-tolerance (on a scale of 1 to 7) and previous experience with online and offline auctions (note that, by the protocol of the sampling procedure, we already knew each subject's gender).

Finally, the subjects were presented with feedback about each of the 66 auctions they played. The feedback consisted of a recapitulation for each auction round of whether they won or not the round, their bid, and their opponents bid. After this information screen, the computer determined the 5 payoff relevant auction rounds. Then, a new screen would display a similar information to the one described above, plus the conversion of points into CZK, and the final monetary payoff of the subject inclusive of the show-up fee.

4.2 Logistics and Subject Pool

The total subject pool consisted of 192 subjects, over 8 sessions of 24 subjects each, of which 96 in treatment A, 96 in W. The subjects were recruited using the Online Recruitment System for Economic (Greiner 2004) among students from the University of Economics in Prague.

Of all subjects, 43 percent do not hold any degree, 48 percent hold a bachelor's degree and 8 percent hold a master's degree. Regarding the field of study, 6 percent have a mathematics or statistics major, 9 percent have a science, engineering or medicine major, 67 percent have an economics or business major, 7 percent have a social science major other than economics or business, and 11 percent have a humanities or some other major. Over 98 percent of our subjects are between 18 and 28 years old, with the remainder being older (up to 39). Also, 46 percent of subjects claim to have a experience with online auctions, 4 percent with offline and 9 percent claim experience with both types. The subjects were paid in cash in Czech crowns (CZK) at the end of their session. Each session lasted approximately 120 minutes with an average earning of 438 CZK, of which 100 CZK was the show-up fee. For a comparison, an hourly wage that students can earn in research assistant or manual jobs typically ranges from 50 to 100 CZK.

5 Results

This section presents our empirical findings. Subsection 5.1 discusses treatment effects on bidding behavior. Subsection 5.2 analyzes treatment effects on average auction revenue and efficiency. Before we continue, let us mention some common features of our analysis. We will 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 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.[6] Looking at both of these blocks lets us see whether our results are sensitive to such potential effects. When considering statistical significance, we universally 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, they are adjusted for clustering at bidding group level.

[6]We observed some indication of a last block effect from some pilots we ran to fine tune our design.

Figure 1: Average Bidding Function Slopes by Treatment and Block

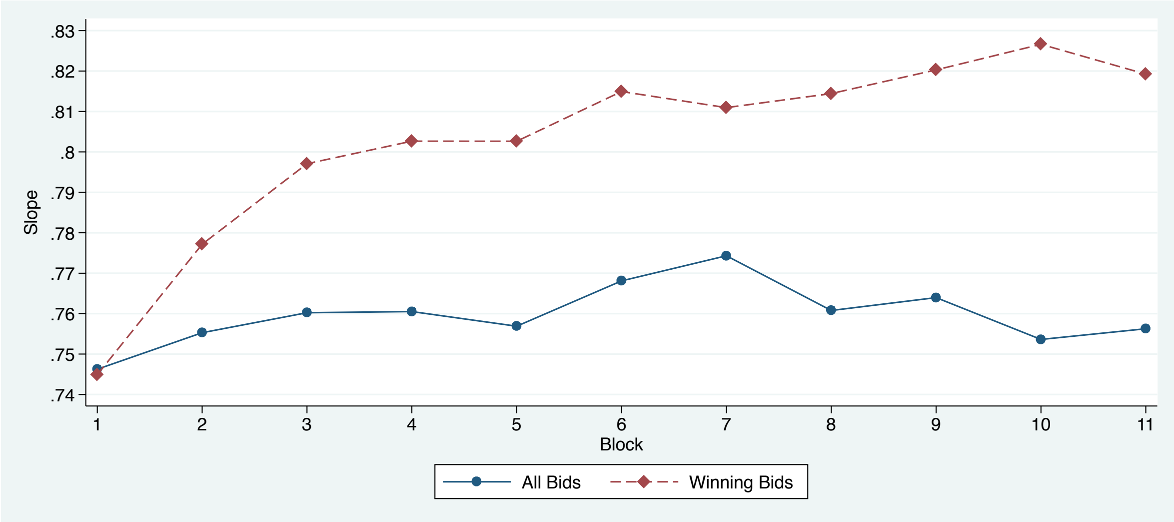

5.1 Bidding

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 bidder in a particular block. We assume that this function is linear and has a zero-intercept. With v_{ijt} denoting the values and b_{ijt} denoting the corresponding bids of subject i in round j of block t, the estimate of slope for subject i in block t is given by

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

That is, the estimated slope is a square-value-weighted average of the six individual bid/value ratios in a given block.

Figure 1 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 almost identical in the two treatments at 0.746 (with the standard error of 0.015) in A and 0.745 (0.013) in W. Notice that such values are higher than in previous experiments (for instance Katuščák et al. (2013), 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 initially we do not provide any information on the distribution of bidders private information (nor the support), which implies a more uncertain (and ambiguous) environment for the subjects at the start. If anything, this should go in the direction of reinforcing the findings below. This is because it is all the more surprising to get a strong feedback difference when initial bidding is already so aggressive. In the following blocks, a discrepancy between the two treatments arises. Subjects in W start bidding more

Figure 2: Effect of Treatment on the Average Bidding Function Slopes by Block

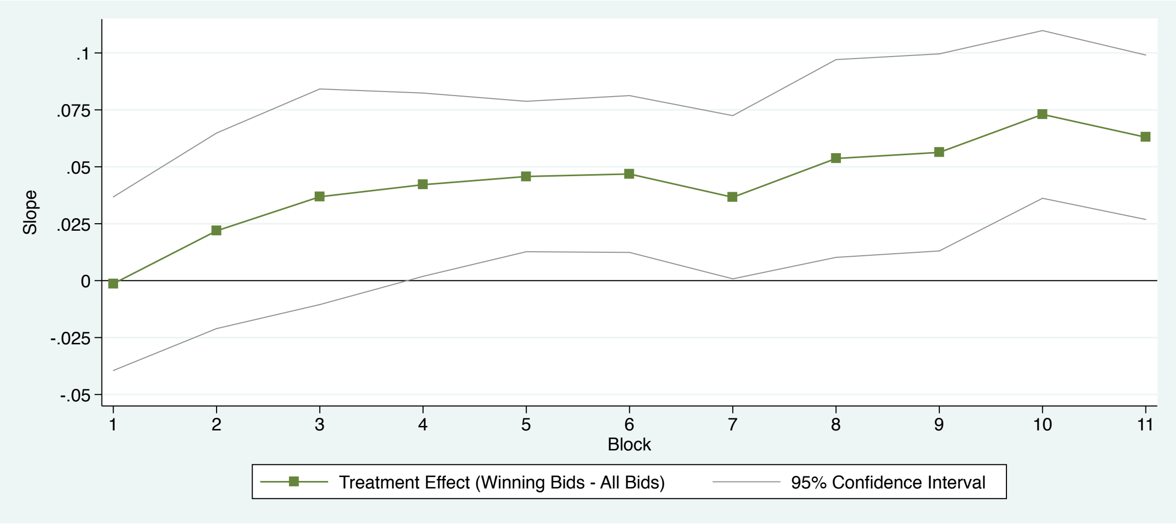

than their counterparts in A. In the final blocks 10 and 11, the average slope in A is 0.754 (0.010) and 0.756 (0.010), respectively, almost unchanged from block 1. On the other hand, the corresponding average slopes in W are 0.827 (0.015) and 0.819 (0.015), respectively. In case of W, this constitutes a statistically significant increase over the average slope in block 1 (t-test p-values of 6.02 and 5.17, respectively).

Figure 2 plots the estimate of the treatment effect (W minus A) on the average slope by block, together with its 95% confidence interval. The treatment effect is positive in all blocks 2 through 11 and it is statistically significant starting from block 4. By the final blocks 10 and 11, the treatment effect reaches 0.073 (0.017) and 0.063 (0.017), respectively. Moreover, block-by-block, we compute the average bidding function slope in each bidding group (8 in each treatment) and we perform a Mann-Whitney ranksum test for the equality of distributions of the average slopes.[7] Starting from block 5 and with the exception of block 7, we reject the null hypothesis in favor of distribution under W dominating the one under A.

To picture the overall impact of the treatment on bidding behavior, Figure 3 presents kernel estimates of the probability density function (pdf) of the bidding function slope 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. 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.

In order to obtain a more detailed treatment comparison of evolutions of the two distributions block-by-block, Figure 4 displays 10th, 25th, 50th, 75th and 90th percentiles of the two distributions in each block. Analogously to Figure 1, we observe that, with a possible excep-

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

Figure 3: Distribution of Bidding Function Slopes by Treatment in Blocks 1, 10 and 11

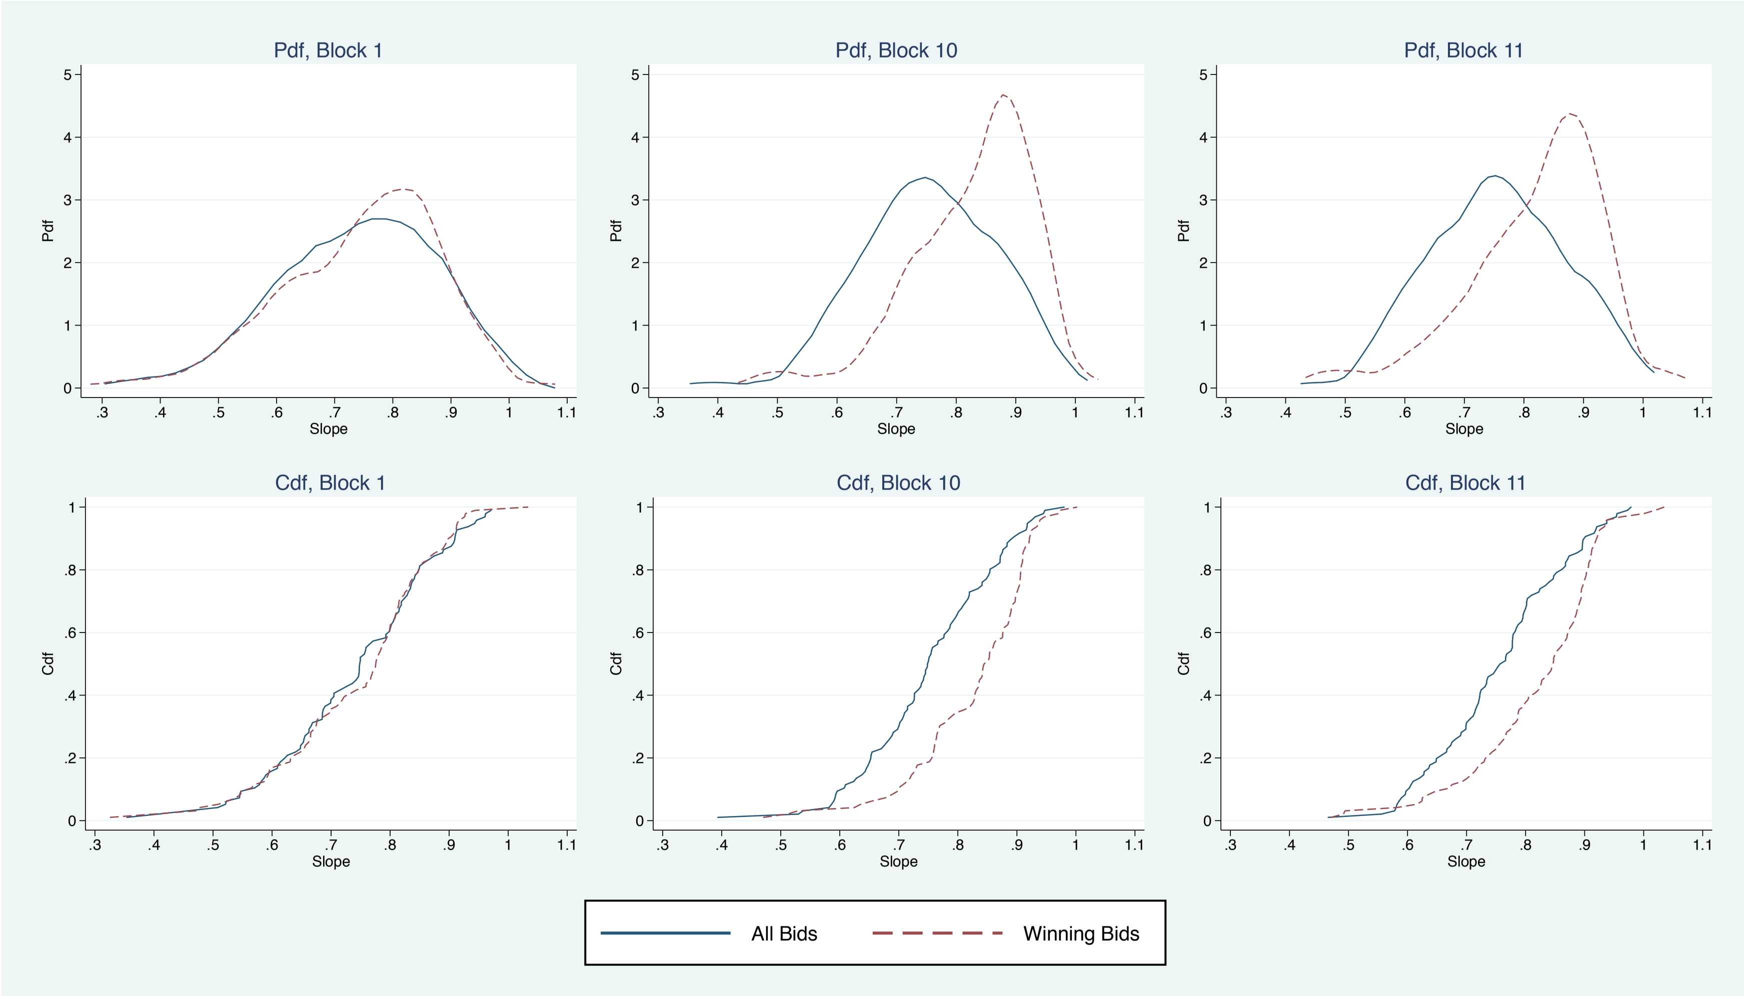

Figure 4: Percentiles of the Distribution of Bidding Function Slopes by Treatment from Blocks 1to 11

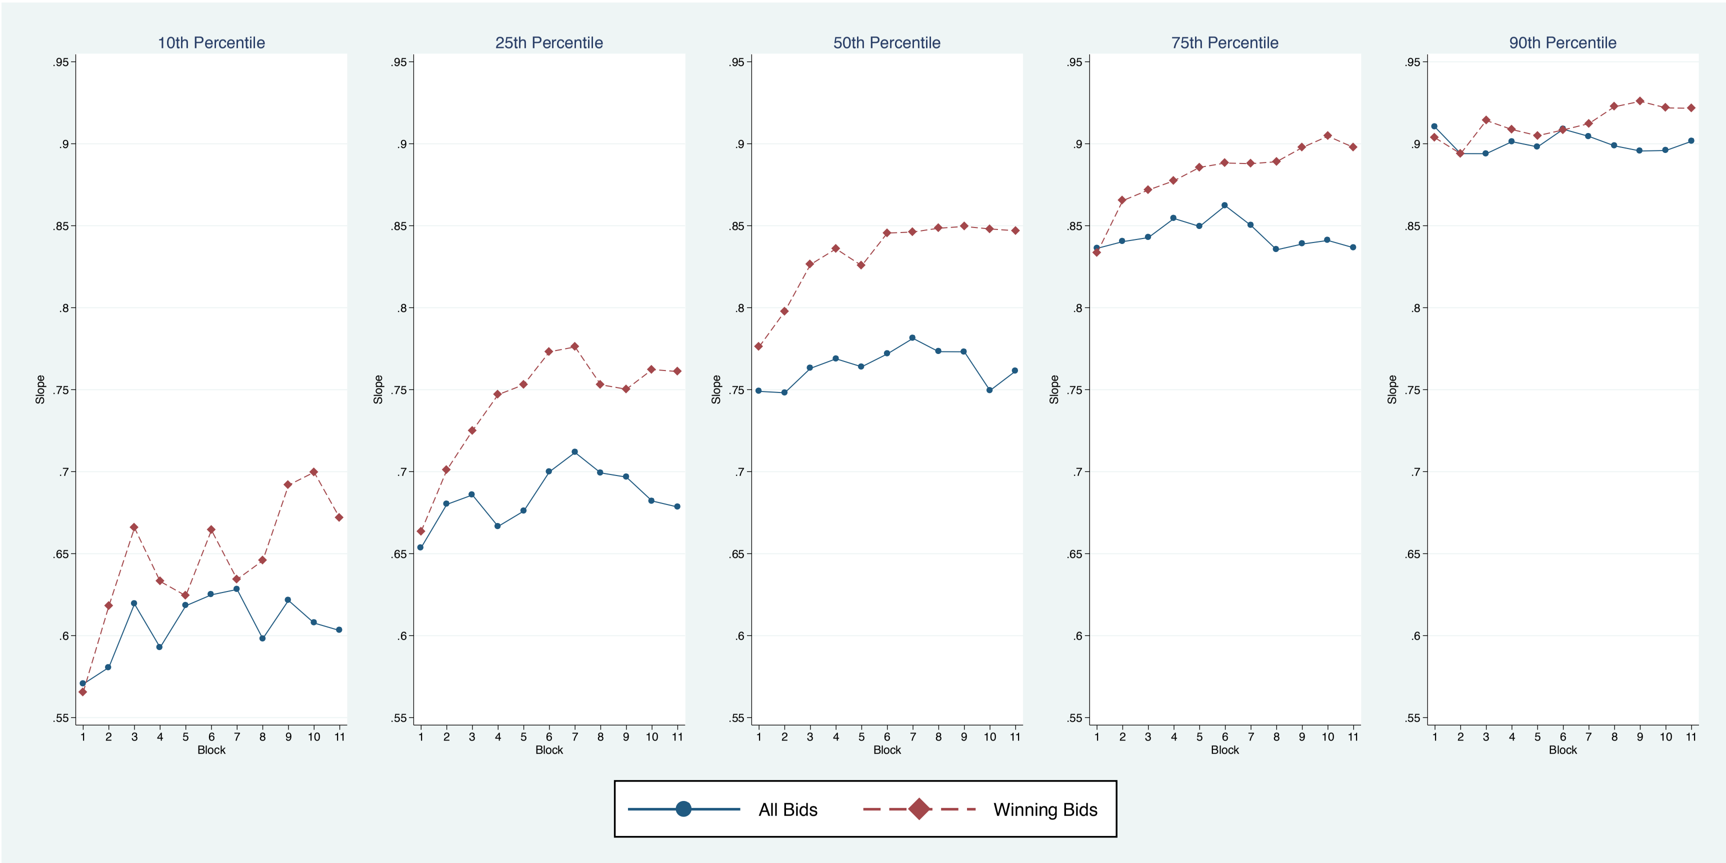

tion of the median, there is little difference in these percentiles between A and W in block 1. The estimated treatment effect is positive in all blocks 2 through 11 for all five percentiles, with the exception of the 90th percentile being essentially the same across the two treatments in blocks 2 and 6. Moreover, the treatment effect is statistically significant starting from block 10 for the 10th percentile, from block 4 (with the exception of blocks 8 and 9) for the 25th percentile, from block 3 for the median, from block 5 for the 75th percentile and in blocks 8 and 10 for the 90th percentile. Moreover, to allow for a possibility that a part of the difference between the medians of the two treatments in blocks 10 and 11 is driven by initial subject heterogeneity (as suggested by the gap in the median of the two groups in block 1), we also run a difference-in-differences analysis between blocks 1 and 10 and 1 and 11, respectively. That is, we subtract the median difference in block 1 from the median difference in block 10 or 11 and test whether the resulting difference is significantly different from 0. For both pairs for blocks, we indeed find that the difference is statistically significant.[8] This analysis documents that the treatment effect on the average slope captured by Figure 1 is shared across various percentiles of the slope distribution except for the 90th one, most robustly so in the middle of the distribution.

Put together, these finding show that in the steady state, as approximated by blocks 10 and 11, W robustly and significantly shifts most of the distribution of bidding strategy slopes to the right relative to A. That is, given the very same values, bidders tend to bid more under W than they do under A. This finding supports the theoretical hypothesis that we laid out in Section 3. To put a quantitative measure on this shift, in block 10 and 11, relative to the average bidding function slope of about 0.755 under A, bidders under W bid on average 8.3 to 9.7 percent higher.

In order to shed further light on the sources of this shift, Figure 5 presents a within-subject scatterplot of the slope in block 1 and block 10 and another one for blocks 1 and 11. In both graphs, we also plot a Epanechnikov kernel-weighted local polynomial-smoothed prediction line from the regression of slope in block 10 (11) on the slope in block 1. In both cases, we eliminated several outliers to make the plot more compact and to avoid the predictions being driven by these outliers.[9] The prediction lines in both plots show that subjects who bid relatively low in block 1 tend to increase their bids in the late blocks. However, the "low" range ends at about 0.73 in A, after which subjects bid similarly in the late blocks as they do in block 1. In contrast to that, for the same slope in the "low" range in block 1, subjects in W increase their bids by about 0.1 more than subjects in A. Moreover, the "low" range

[8]This is true also for the 75th percentile, whereas the difference-in-differences is not significant in either pair for the 90th percentile. For the 10th and the 25th percentile, one of the differences is robustly and the other one marginally statistically different from zero.

[9]In particular, we exclude all subjects for whom the slope in either of the two blocks is below 0.45 or above 1. This eliminates 10 subjects (5 from each treatment) in the scatterplot of blocks 1 and 10 and 11 subjects (7 from A and 4 from W) in the scatterplot of blocks 1 and 11.

Figure 5: Bidding Dynamics on Individual Subjects by Treatment

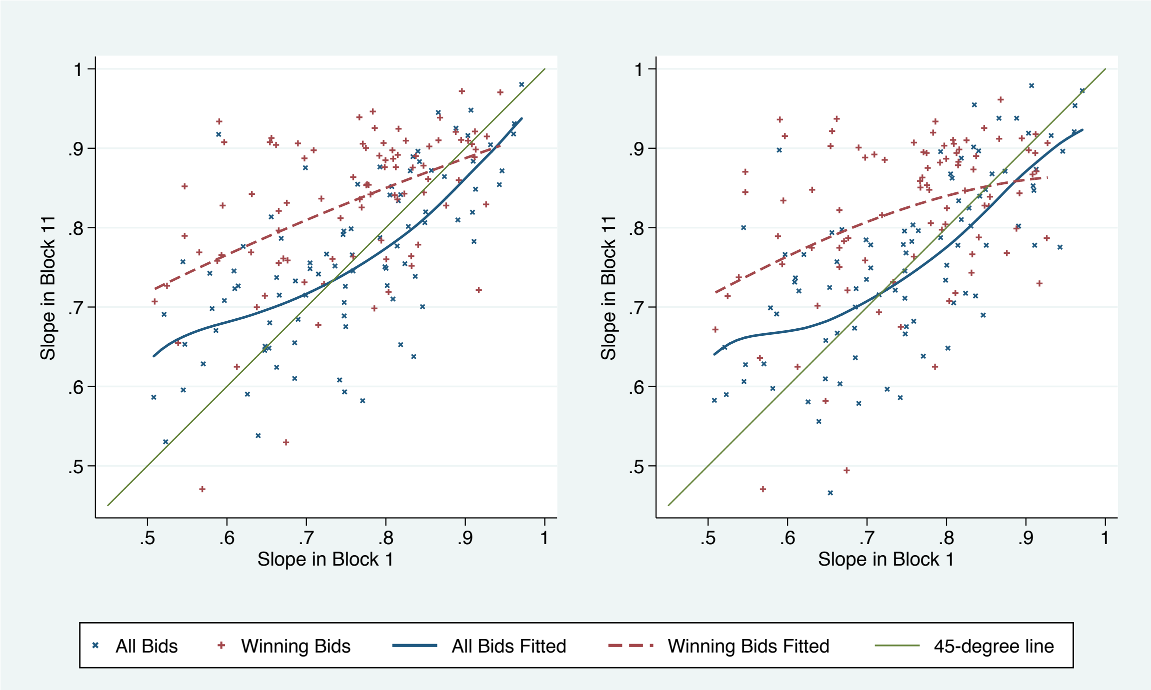

ends at about 0.86, after which subjects bid less in late blocks than they do in block 1. The figure therefore reveals that the main driver of the treatment effect is that lower-slope bidders in block 1 are switching to slopes in the late blocks that are systematically higher under W than they are under A. For high-slope bidders in block 1, any treatment difference is much smaller or absent.

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 bidders 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 blcok. Each draw

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

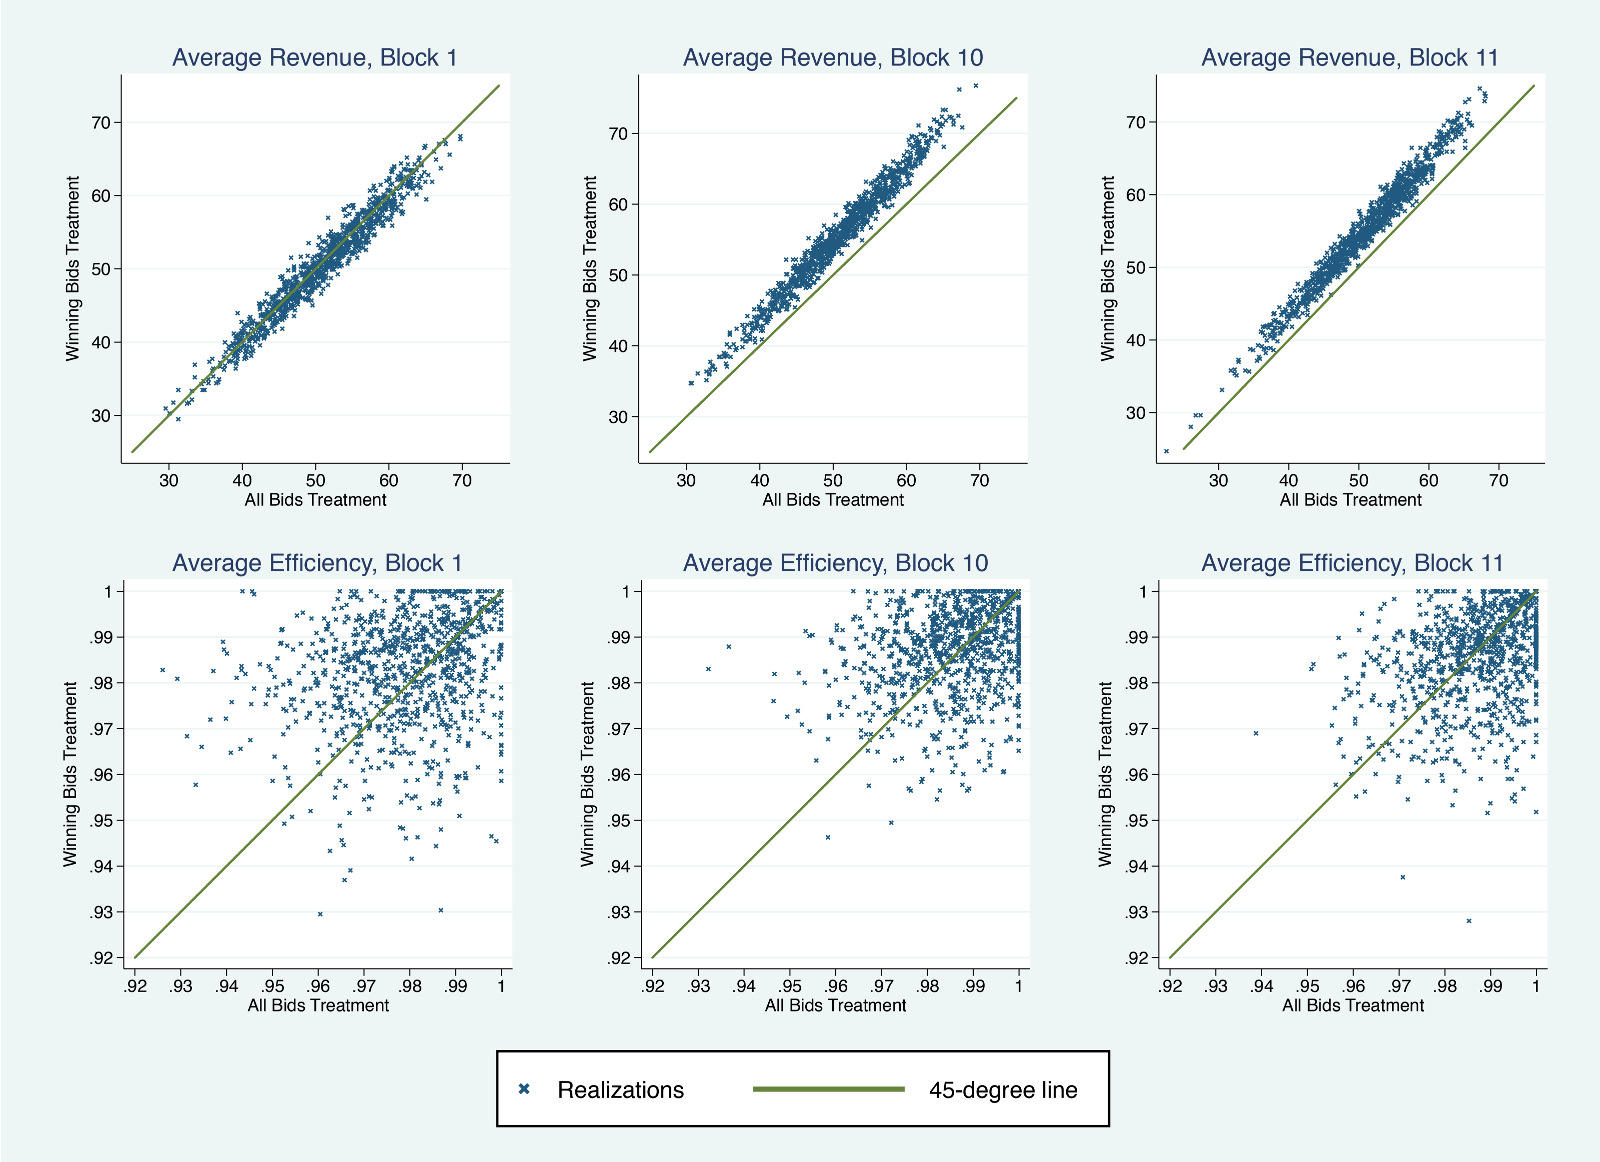

is generated as follows. At the level of a bidding group, we first randomly (with the uniform distribution) draw one of the 6 rounds, separately and independently for each bidder. Next, we randomly pair the 12 members of the bidding group into pairs. Any pattern of pairing is equally likely. We then 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. 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 average revenue, the scatterplot for block 1 is symmetric around the diagonal. Indeed, in the absence of any difference between the two treatment in block 1, this is what we would expect. In contrast to that, *all* points in the scatterplots for blocks 10 and 11 lie above the diagonal. Hence, without any statistical testing, we can reject the hypothesis of zero average revenue impact of the treatment. Rather, in line with the results presented in the previous subsection, average revenue is higher under W in comparison to A. In particular, the mean (median) of the ratio of revenue under W to that under A in block 10 is 1.092 (1.091). Also, regressing the revenue under W on that under A in block 10 and imposing a zero intercept gives a coefficient of 1.091. The analogous figures for block 11 are 1.079 (1.080) and 1.079. Hence we conclude that, in the steady-state, the average revenue is about 8 to 9 percent higher under W than it is under A.

In contrast to a clear treatment effect on average revenue, the impact on average efficiency is much less profound, if any. In all three blocks, significant parts of the scatterplots are located on both sides of the diagonal. We therefore conclude that there is no significant effect of treatment on average efficiency.

6 Conclusions

The main focus of this paper is the important market design issue of whether an auction house should disclose a specific type of historical information about past auctions or not. We compare the two natural alternatives of providing information about all bids versus providing information about winning bids only. We find that historical feedback affects bidding behavior in a systematic manner, with winning bids feedback yielding a very considerable increase in average bid/value ratios (compared to all bids feedback) in the amount of roughly 9%. A similar effect holds for revenues. A few studies had looked at the effect of feedback on bidding in FPAs markets, but from a rather different angle than ours. In fact, previous studies have looked at the effect of providing a bidder with feedback about the outcome of auctions he/she played in the past. Thus, unlike our, those studies can have some valid predictions only for environments where bidders play repeatedly. The theoretical explanations advanced by these previous studies are either based on experienced regret, or on a learning based theory (impulse balance equilibrium), both of which need feedback on own auction to operate. Our design, by disentangling the informational effect of feedback from potential regret considerations highlights that a different explanation must be at play in our environment. We propose one based on bidders naively best replying (in the winning bids feedback) to the information they receive as if it was about all bids. Our design cannot rule out that in the previous experiments regret consideration were important, but it suggests that it is unclear whether regret it is a key factor at play when discussing the role of feedback in bidding behavior in FPAs.

A corollary contribution we have is on the understanding of how feedback affects beliefs formation. In this paper we take the view that beliefs must come from some historical information bidders have been exposed to. For this reason at the beginning of the experiment, unlike in standard experiments, we do not provide any information about the distribution from which the private information of the opponents is drawn (nor the support of the distribution). Perhaps because of this, initial bid/value ratios are higher than what normally found. Also, not surprisingly, there is a lot of initial heterogeneity. We observe that behavior stabilizes after a few blocks, with the effect of feedback on winning bids raising average bid/value ratios lasting for a longer period. Such feedback has also a slightly stronger effect on homogenizing behavior due to the stronger impact on the behavior of those subjects that had started with lower bid/value ratios.

Finally, we think that our study, by illustrating a novel channel for the role of feedback, might stimulate further work on several directions. In fact, our study used what we thought was the most natural setting to start thinking about the role of historical information in affecting bidding behavior. That said, we see three main avenues of departure from our design for future research. First, we drew valuations from a Uniform distribution with support $[0, 100]$ mainly to be comparable with previous experimental studies. Since our subjects are not told such information, one might opt for different distributions and/or supports and see whether our findings are sensitive to any such variation. Second, one might consider whether an alternative disclosure of historical information could be superior to the "only winning bids" one, perhaps in the form of some specific statistic. Third, and perhaps more interestingly, note that FPAs were a natural target of previous studies because they provide a setting were there is almost always some ex-post regret. Given that our findings operate in a regret free environment, one might expect historical information to matter also in other market institutions. For instance, providing a specific feedback might matter even in private value SPAs if entry decisions are endogenous. Apart from the study of entry decisions, feedback might matter in other important contests such as search models. We leave these questions open for future research.

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 GACR, or CERGE-EI.

References

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.

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

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.

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

Katuščák, Peter, Fabio Michelucci, and Miroslav Zajíček, "Does Anticipated Regret Really Matter? Revisiting the Role of Feedback in Auction Bidding," 2013. CERGE-EI Working Paper no. 487, available from http://www.cerge-ei.cz/pdf/wp/Wp487.pdf.

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.

Ratan, Anmol, "Is Decision Regret Relevant for Auctions? A Reconsideration of Anticipated Regret in First Price Auctions," 2013. Mimeo, Monash University.