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

Philippe Jehiel* Peter Katuščák† Fabio Michelucci‡

April 13, 2015

Abstract

Consider a long-term auctioneer who repeatedly sells identical or similar items and who might disclose selective information about past bidding. We present a theory that yields different predictions about bidding behavior depending on the information bidders are provided with, and then test it using a lab experiment. We focus on the disclosure of all bids and of winning bids only. Our theory is based on a selection bias: some of the bidders who are presented with historical winning bids mistakenly best-respond to that distribution, failing to realize that winning bids are not representative of all bids. In the steady state, this bias results in higher bids and auction revenue in comparison to the case when all bids are presented. Our experimental test confirms the qualitative predictions of the theory. On the theory side, our findings challenge the predictive power of Bayesian Nash Equilibrium based on rational bidders. On the market design side, they underline the role of historical market information as a key design choice.

Keywords: auctions, bidding, feedback, selection bias, mechanism design

JEL Classification: C91, C92, D44

An early draft. Please do not cite and do not distribute.

*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 Arkansas and University of Economics in Prague, Sam M. Walton College of Business, Business Building 301, Fayetteville, AR 72701, USA, Peter.Katuscak@gmail.com (corresponding author).

‡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, Fabio.Michelucci@cerge-ei.cz.

1 Introduction

Consider a long-term auctioneer who repeatedly sells identical or similar items. Typically, such auctioneer has gathered a lot of relevant information about 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 data on previous auctions of identical or similar items. Our work is motivated by the observation that a long term auctioneer might disclose selective information if such choice affects bidding behavior, and thus long-run expected revenue. We present a theory that yields different predictions depending on the information bidders are provided with, and then test it using a controlled lab experiment. In particular, we focus on first price auctions of two bidders and on 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, HI on only the winning bids is provided. Notice that if all bidders were fully rational, the type of disclosure policy would make no difference. This is because fully rational bidders can recover the distribution of all bids from observing the distribution of winning bids. Thus, under *standard* theory, the steady state of bidding coincides with the Bayesian Nash Equilibrium under either type of HI.

In comparison, our theory is based on a selection bias: some bidders mistakenly best respond to the historical distribution of winning bids because they fail to realize that winning bids are not representative of all bids. In the steady state, this bias results in higher bids and auction revenue in comparison to the case of HI on all bids. Our experiment confirms the qualitative predictions coming from the theory and implies that a long term auctioneer should choose to disclose HI on winning bids only over HI on all bids in order to generate higher revenues.

Our empirical findings have potentially profound implications. On the theory side, they challenge the predictive power of Bayesian Nash Equilibrium based on rational bidders. On the market design side, they open new ground for HI to have a key design role. In particular, they suggest that manipulation of HI is a virtually costless instrument to raise expected revenues.

We stress that our work is not directly applicable to situations in which strategic repeated bidding considerations are present, such as settings where the same pool of bidders interacts repeatedly over time. Furthermore, there is an implicit assumption that the disclosure policy is a choice variable of the auctioneer. In many real-world instances, this is indeed the case. However, in other instances, for example within public procurement, the exact disclosure policy might depend on a specific law or transparency regulations.[1] Additionally, in the

[1]For example, Bergemann and Hörner (2010) point out that the U.S. federal government is required to fully disclose identity of the bidders and the terms of each bid within its procurement auctions. In contrast to that, in auctions of mineral rights to U.S. government-owned land, only the winner's identity is revealed.

Dutch auction implementation of the first-price auction, winning bids are the only available HI option.

Examples of auction markets where our work should be relevant are internet auction markets matching buyers and sellers; instantaneous bidding by advertisers to place an advertisement onto a web page loaded by a consumer; those recurrent procurement auctions for which the set of bidders varies sufficiently over time.[2] In general, we have in mind auction markets in which each time the set of bidders is sufficiently diverse from the past, and thus strategically we can think of these auctions as one-shot auctions. This is an important observation to point out since the only robust empirical evidence about the impact of selective HI on bidding from the existing literature is for repeated bidding in which bidders receive feedback about the own past auctions. Thus, our work extends the role of HI beyond the applications for which it has been analyzed.

Having stressed the limits of applicability of our work, we also want to remark that we do not see anything special about first price auctions that would not generalize to other settings in which market participants need to form correct beliefs about the behavior of their opponents, and where the basis for forming such beliefs is the available HI coming from identical/similar market realizations in the past. In fact, we see it as a fruitful avenue of research to explore the impact of historical market information in other Bayesian games such as bargaining, search, and so forth.

Our work is by no means the first to address the role of disclosure policies about outcomes of past markets. However, it does so from a rather different angle compared to the existing literature. We point out how this is the case while commenting on the related literature in Section 2.

Finally, note that we focus on auctions with two bidders. We think that the two-bidder case is the most simple one to verify our theoretical claim that the Bayesian Nash Equilibrium is sensitive to the type of historical market information that bidders have access to. That said, from a market design perspective, we believe that the result that we stress should generalize to auctions with more than two bidders. In fact there are reasons to believe that for more than two bidders the difference between choosing winning bids versus all bids disclosure might be more pronounced. However, additional biases might be at play in this case given the higher computational complexity of the environment. We leave the discussion of this interesting theme to the concluding remarks at the end of the paper as a suggestion for future research. We think that more research in this area could unveil a broader picture about the biases and mistakes bidders are subject to in market environments.

The rest of the paper is structured as follows. Section 2 surveys related theoretical and experimental work and positions our paper relative to this literature. Section 3 presents the

[2]Our setting also requires bidders to not have dynamic purchase capacity (or procured good provision) constraints, so they do not condition their bidding on past auction outcomes or future bidding possibilities.

theoretical model. 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 Relation to Related Literature

In this section, we review related theoretical and empirical literature on the effect of HI in auctions, and we clarify how it differs from what we do.

A related idea to the one we present here can be found in the application of analogy-based expectation equilibrium (Jehiel 2005) to auctions (Jehiel 2011).[3] The common theme with what we do here is that in both cases bidders fictitiously best respond to the information they are given. The difference is that in Jehiel (2011) bidders receive aggregated HI on play in different types of auctions (that is, not only on the type they actually play), while in our setting bidders receive HI purely on the auction type they participate in.

Another closely related approach can be found in the literature that considers repeated interaction of a fixed set of bidders in an identical auction format with values being drawn anew for each auction repetition. In this context, Esponda (2008) applies the concept of self-confirming equilibrium (SCE) (Battigalli 1987, Fudenberg and Levine 1993, Dekel, Fudenberg and Levine 2004) to first-price auctions. In a SCE, beliefs are not required to be correct, they are only restricted to be consistent with feedback that players obtain about equilibrium outcomes. As put by Esponda (2008) himself, "suppose that a bidder always observes the [HI on] winning bid... SCE then requires her to have correct beliefs about the equilibrium distribution of the winning bid but not necessarily correct beliefs about the equilibrium distribution, say, of the second highest bid." For the private value environment,the author shows that if bidders receive HI on winning bids, the sets of SCE and BNE coincide.[4] Importantly for our purpose, given that players are assumed to best-respond to their beliefs in this approach, the result applies also to a setting where bidders form their beliefs based on observation of behavior of anonymous *other* bidders in previous auctions. Hence, unlike our theory, this theory predicts that bidders bid according to the Bayesian Nash equilibrium if receiving HI on winning bids.

The main difference between these two approaches and our approach is that in the former case (some) bidders have incorrect beliefs due to insufficiently rich HI, while in the latter case

[3]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.

[4]Intuitively, even though a bidder does not observe the highest bid of his opponents when winning the auction, he can extrapolate this information from knowing the winning bids when he himself bids low. This is because, under independence, such bids are identical to the would-be highest competing bids he does not get to observe in case of winning.

the incorrect beliefs are driven by bounded rationality of bidders in the presence of sufficient information for forming correct beliefs.

Another strand of literature focuses on the role of HI on *own* past auctions and how bidders adjust their bidding strategies in response. Conceptually, the main difference with our approach is that this line of research treats HI more as a way of behaviorally perfecting one's bidding strategy rather than as a source of information for forming beliefs about the relevant distribution needed to form a best reply. 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 relation to one's value) in future auctions. On the other hand, having won but learning that the second highest bid (if such information is available) was much 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 type of impulse can be stimulated by providing feedback on winning bids, whereas the latter type can be stimulated by providing feedback on the second highest bids. One way of interpreting the origin of these bidding impulses is that they are triggered by an *ex post* experience of loser and winner regret, respectively. A set of related papers (Engelbrecht-Wiggans 1989, Filiz-Ozbay and Ozbay 2007) proposes that this idea extends in the form of *ex ante* regret anticipation even to a one-shot setting. According to this theory, knowing that one will learn about the winning (second highest) bid makes one to anticipate loser (winner) regret and hence 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.[5] Experimental evidence from one shot bidding is far less robust, however. Regarding the feedback on the winning bid as opposed to only learning whether one won or not, Filiz-Ozbay and Ozbay (2007) find bids and revenues to be higher. However, Katuščák, Michelucci and Zajíček (2013), in an experiment with a much larger sample size and employing various experimental protocols, including the one used by Filiz-Ozbay and Ozbay (2007), do not find any effect. Ratan (2013) obtains a similar conclusion.[6]

By assuming that bidders fictitiously best-respond to HI irrespective of outcomes of their own past auctions, we ignore the impulse balance and anticipated regret considerations in our theoretical development. This is not to say that such considerations are necessarily absent in field applications. Rather, our objective is to isolate the effect of anonymous HI on bidding through beliefs. This is because our setting is meant to represent situations in which bidders

[5]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.

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

participate in auctions only once or infrequently, and hence the formation of bidders' beliefs has to rely on information about how *other* bidders behaved in previous auctions. In parallel with this shift of focus, our experimental design carefully departs from experiments on repeated bidding by withdrawing any feedback on the outcome of one's own past auctions (even whether one won or lost) until the very end of the experiment and adopts further measures to avoid any regret considerations (see Section 4 for details).[7]

Notice that our work implicitly raises the possibility that the similar treatment effect found by experiments with repeated bidding and feedback on own auction might be explained by subjects using the limited information on own auction to form beliefs about the distribution of the highest bid of the opponents, rather than by the proposed regret or learning considerations. It might then be an interesting research question to try to disentangle the effect of regret or learning from the one coming from beliefs in those settings.[8]

Another experimental paper that looks at the effect of HI is Sonsino and Ivanova-Stenzel (2006). The only available type of HI in their study is vectors of bids and the associated values. The focus of the paper is on bidders' endogenous choice of how much HI to sample and how that correlates with bidder performance. That is why they have no variation in HI type unlike most of the other studies mentioned above (and ours). Furthermore, the availability of information on values alongside the information on bids makes this study more difficult to be directly compared with the other studies including ours.

Finally, there is also literature that is more loosely connected to our approach but that nevertheless suggests some important considerations. One such consideration is dynamic collusion-building. Dufwenberg and Gneezy (2002) consider a repeated common value auction in which each bidding group is randomly drawn from a fixed population. The common value is constant and publicly known. If all bids from previous auctions are disclosed, bidders might attempt to bid low early on in order to induce other bidders to also bid low in future auctions. This option is absent if there is HI on the winning bids only.[9] If there are finitely many auction rounds, though, this argument is purely empirical. Theoretically, any possible low-bid cooperation unravels irrespective of the type of HI. Experimentally, the authors find suggestive evidence that bids are on average lower under non-anonymous (i.e. containing experimental identity) HI on all bids as opposed to non-anonymous HI on winning bids and interpret it as an attempt to collude.[10] With regard to our setting, we believe that such collusion-building

[7] 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.

[8] Note that the explanation sketched above is consistent with the fact that the evidence from one-shot auction seems inconclusive.

[9] The auction in Dufwenberg and Gneezy (2002) is framed as a procurement auction. We reframed it here as a purchase auction.

[10] Even though the authors use the term "signaling" instead of "collusion" and they contrast their interpretation with collusion in a setting with a fixed set of bidders that interact repeatedly, the reasoning is quite similar, and we unify it under the name "collusion."

is unlikely to operate for two reasons. The most important reason is that in a private value environment, a low bid might very well be competitive if the bidder's (unobserved) value is low. A secondary reason is that if the HI is anonymous, it is impossible to identify whether a particular group of bidders is attempting to jump-start collusion by consistently, i.e. for all kind of values, bidding low. This hampers formation of a critical mass of bidders to sustain a collusive outcome.

Another repeated bidding consideration is signaling of values through bids in case values of a bidder are positively serially correlated. 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. Being aware of this revelation, bidders change their bidding strategically in order to inhibit others' learning about their value and to speed up their learning about other bidders' values. For example, with values of a bidder being identical over time, Cason et al. (2011) show theoretically that when all bids are revealed after an auction, high-value bidders tend to decrease their bids early on by pooling with low-value bidders to avoid revealing their type (deception effect). When just the winning bid is revealed, high-value bidders tend to decrease their bids early on as well, but this time they do so to reveal more information about the type of their opponents (extraction effect).[11] With regard to our setting, this consideration is absent since bidding groups are drawn from a large pool of bidders and HI is anonymous. However, since values of a bidder over time are likely to be serially correlated in field settings, this line of research underlines the difference between making HI anonymous as opposed to non-anonymous. As discussed above, a non-anonymous HI tends to reduce bids and hence revenue. As a result, if the auctioneer is interested in maximizing revenue, it is better to provide anonymous HI.

3 Theory

In this section, we model a steady state of a learning process in which bidders form their belief about the DHCB based on the provided HI.

Time is discrete and there is a continuum of auctions for the same or similar object happening in each time period. All the auctions have two participating bidders. Bidders have access to HI from the previous period. We consider two types of HI: on all bids and on winning bids. Let $G(\cdot)$ be the steady-state distribution of all bids. Then the steady-state DHCB is also given by $G(\cdot)$, whereas the steady-state distribution of the winning bids is given by $G^2(\cdot)$. Under HI on all bids, given a value of v and a utility function of $u(\cdot)$ (normalized so that $u(0) = 0$), all bidders correctly perceive the DHCB and choose their bid b to maximize $u(v - b)G(b)$.

[11] The auction in Cason et al. (2011) is framed as a procurement auction. We reframed it as a purchase auction.

Under HI on winning bids, *sophisticated* bidders (fraction $1 - \lambda$ of the population) correctly perceive that the HI is not representative of all bids and extact the correct DHCB from it. As a result, they maximize $u(v - b)G(b)$. On the other hand, *naive* bidders (fraction $\lambda > 0$ of the population) fail to perceive the bias and believe that the DHCB is given by $G^2(\cdot)$. As a result, they maximize $u(v - b)G^2(b)$.

3.1 Historical Information on All Bids

We begin the analysis by considering the easier case of HI on all bids. Let $G_A(\cdot)$ be the steady-state distribution of bids under this HI type and let $b_A(\cdot)$ be the associated bidding function. As outlined above, given a value of v, all bidders choose their bid b to maximize $u(v - b)G_A(b)$, or, equivalently, $\ln u(v - b) + \ln G_A(b)$. The optimal bidding function $b_A(v)$ is characterized by the first-order condition

$$\frac{u'[v - b_A(v)]}{u[v - b_A(v)]} = \frac{G'_A[b_A(v)]}{G_A[b_A(v)]}. \tag{1}$$

Defining $w(\cdot) \equiv u'(\cdot)/u(\cdot)$ and noting that $w(\cdot)$ is strictly decreasing for both risk-neutral and risk-averse bidders, we obtain that

$$v = b_A(v) + w^{-1}\left(\frac{G'_A[b_A(v)]}{G_A[b_A(v)]}\right). \tag{2}$$

Assuming that $b_A(v)$ is strictly increasing, this characterization is equivalent to

$$b_A^{-1}(x) = x + w^{-1}\left(\frac{G'_A(x)}{G_A(x)}\right) \tag{3}$$

for any x in the range of $b_A(v)$. With $F(\cdot)$ being the distribution of valuations, it follows that

$$\begin{aligned} G_A(x) &= \Pr\{b_A(v) \leq x\} \\ &= \Pr\left\{v \leq b_A^{-1}(x)\right\} \\ &= F\left\{x + w^{-1}\left(\frac{G'_A(x)}{G_A(x)}\right)\right\}. \end{aligned} \tag{4}$$

In order to achieve tractability in case of HI on winning bids, we now specialize the setup by assuming the valuations are distributed uniformly on $[0, 1]$ and the utility function has the constant relative risk-aversion form $u(c) = c^\alpha$ with $\alpha > 0$. Here, $\alpha = 1$ corresponds to risk-neutrality, $\alpha \in (0, 1)$ to risk-aversion and $\alpha > 1$ to risk-loving. Under such conditions, we obtain that:

Proposition 1. *If* $u(c) = c^\alpha$ *and bidders' values are distributed uniformly on* $[0, 1]$*, then under*

HI on all bids, the steady-state distribution of bids is given by

$$G_A(x) = (1+\alpha)\, x \tag{5}$$

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

$$b_A(v) = \frac{1}{1+\alpha} v. \tag{6}$$

Both the steady-state distribution of bids and the bidding function coincide with analogous objects under the Bayesian Nash equilibrium.

3.2 Historical Information on Winning Bids

Now consider the case of HI on winning bids. Let $G_W(\cdot)$ be the steady-state distribution of bids under this HI type. In this case, there are two types of bidders: sophisticated and naive. Let $b_S(\cdot)$ and $b_N(\cdot)$ be the steady-state bidding functions of the sophisticated and naive bidders, respectively. As outlined above, given a value of v, sophisticated bidders choose their bid b to maximize $u(v-b)G_W(b)$, or, equivalently, $\ln u(v-b) + \ln G_W(b)$. On the other hand, naive bidders choose their bid b to maximize $u(v-b)G_W^2(b)$, or, equivalently, $\ln u(v-b) + 2\ln G_W(b)$. Then following the same steps as in the previous subsection, the optimal bidding function $b_N(v)$ of a naive bidder is characterized by

$$b_N^{-1}(x) = x + w^{-1}\left(2\frac{G'_W(x)}{G_W(x)}\right) \tag{7}$$

and the optimal bidding function $b_S(v)$ of a sophisticated bidder is characterized by

$$b_S^{-1}(x) = x + w^{-1}\left(\frac{G'_W(x)}{G_W(x)}\right) \tag{8}$$

for any x in the range of $b_N(\cdot)$ and $b_S(\cdot)$, respectively. Note that, since $w(\cdot)$ is strictly decreasing and of both $b_N(\cdot)$ and $b_S(\cdot)$ are assumed to be strictly increasing, (7) and (8) imply that the naive bidders bid more given their value that the sophisticated bidders do. Using these two equations, it follows that $G_W(\cdot)$ is given by

$$\begin{aligned} G_W(x) &= \lambda \Pr\{b_N(v) \le x\} + (1-\lambda)\Pr\{b_S(v) \le x\} \\ &= \lambda \Pr\left\{v \le b_N^{-1}(x)\right\} + (1-\lambda)\Pr\left\{v \le b_S^{-1}(x)\right\} \\ &= \lambda F\left\{x + w^{-1}\left(2\frac{G'_W(x)}{G_W(x)}\right)\right\} + (1-\lambda)F\left\{x + w^{-1}\left(\frac{G'_W(x)}{G_W(x)}\right)\right\}. \end{aligned} \tag{9}$$

As in the previous subsection, we now specialize the setup by assuming the valuations are distributed uniformly on $[0, 1]$ and the utility function has the constant relative risk aversion form $u(c) = c^\alpha$ with $\alpha > 0$. Under such conditions we obtain that:

Proposition 2. *If $u(c) = c^\alpha$ and bidders' values are distributed uniformly on $[0, 1]$, then under HI on winning bids, for $x \in [0, 1/(1+\alpha)]$, the steady-state distribution of bids is given by*

$$G_W(x) = \left(1 + \alpha - \frac{\alpha\lambda}{2}\right) x. \tag{10}$$

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

$$G'_W(x) = \frac{\lambda\alpha}{2} \frac{G_W(x)}{G_W(x) - \lambda x - (1-\lambda)} \tag{11}$$

with the initial condition

$$G_W\left(\frac{1}{1+\alpha}\right) = 1 - \frac{\lambda\alpha}{2(1+\alpha)}. \tag{12}$$

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

$$b_S(v) = \frac{1}{1+\alpha} v. \tag{13}$$

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

$$b_N(v) = \frac{2}{2+\alpha} v \tag{14}$$

For $v > (2+\alpha)/(2+2\alpha)$, its inverse is given by

$$b_N^{-1}(x) = \frac{G_W(x) - (1-\lambda)}{\lambda}. \tag{15}$$

Proof. The proof follows from specializing (7)-(9) to the case of CRRA preferences and uniform distribution of valuations. It is available from the authors upon request. □

Note that the sophisticated bidders bid the same way under HI on winning bids as bidders bid under HI on all bids. This is a consequence of the functional form assumptions we have applied. Compared to the sophisticated bidders, the naive bidders bid more. We are able to obtain a partially closed solution for the latter's bidding strategy. For illustration, Figure 1 plots the steady-state bidding functions for both types of bidders.

The most important prediction from the the comparison of the two types of HI is that the distribution of bids under HI on winning bids is shifted to the right in comparison to HI on all bids. As a consequence, the expected revenue under the former HI type is higher. Our experiment aims to test these theoretical predictions.

Figure 1: Steady-state Bidding Functions

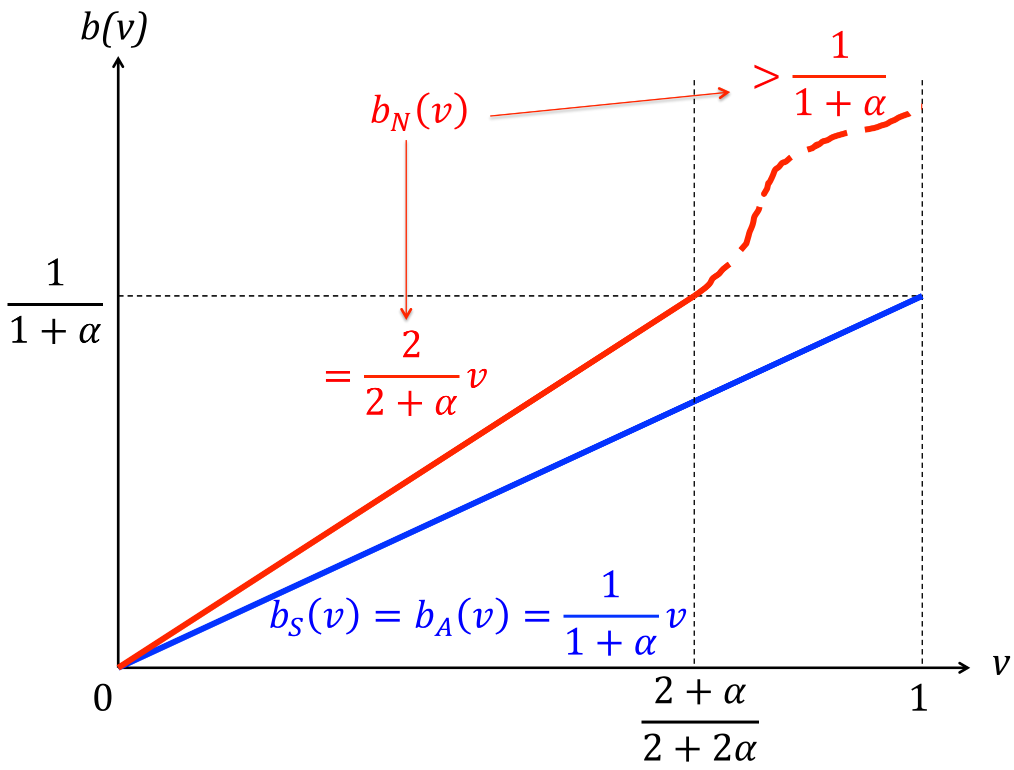

Note: $b_N(\cdot)$ and $b_S(\cdot)$ are the steady-state bidding functions of naive and sophisticate bidders, respectively, under HI on winning bids. $b_A(\cdot)$ is the steady-state bidding function of all bidders under HI on all bids.

4 Experimental Design

This section describes our experimental design. In Subsection 4.1, we discuss the main features of our design. Subsection 4.2 focuses on those details of the design aimed at reducing noise in the data. Subsection 4.3 describes the instructions and the execution of the experiment. Finally, Subsection 4.4, discusses the logistics of the experiment and the demographic characteristics of our subjects.

4.1 Main Features of the Design

Our experiment involves repeated two-bidder auctions in two treatments: HI on all bids (denoted "A") and HI on winning bids (denoted "W"). In each treatment, bidders have access to the respective type of HI for a series of previous two-bidder auctions under identical value distribution. Our objective is to track the bidding behavior over time until it stabilizes, and then to compare it across the two treatments.

The design of the experiment requires making choices regarding several aspects. We select four main ones to be discussed in detail below: repeated vs. one-shot bidding, how to form the HI and how rich it should be, how to present it, and how to filter out potential effects of HI not connected to the formation of beliefs about the DHCB. Each of these choices involves a certain trade-off. Other aspects of the design such as how subjects' values are generated, what subjects know about the generation process and how subjects are paid are discussed at the end of this subsection.

We begin by presenting an overview of the design. Bidders interact in groups of 12 subjects. They bid repeatedly. In each bidding round, the opponent is randomly and anonymously chosen from among the other 11 bidders in the group. Each bidder's value is drawn from the uniform distribution on $[0, 100]$, without subjects having knowledge of this distribution or its support. HI is incremented in "blocks" rather than "rounds" (a round of bidding for a bidder is one auction). Each block consists of 6 bidding rounds, and there are 11 blocks in total. In block $t \in \{2, ..., 11\}$, subjects have access to HI on bidding in their group in block $t - 1$. Note that there is no HI in block 1 and therefore no HI base for different bidding across the two treatments. In fact, our objective is to seed the learning process with a treatment-independent distribution of bids. We therefore make the two treatments identical until the end of block 1. This involves handing out two sets of instructions. The first set, handed out at the beginning of the experiment, is identical across the two treatments. These instructions do not mention any specifics regarding the type of HI to be received from block 2 onward.[12]

We now move to the four main design choices outlined above. First, in our experiment subjects bid repeatedly rather than just once. In principle, the fact that subjects bid repeatedly might lead to repeated bidding effects with subjects trying to influence each other's behavior and/or coordinating on some particular behavior. We use two measures to avoid the occurrence of those 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. 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. Furthermore, it is not clear that employing one-shot bidders would be ideal as subjects might require a few blocks to get familiar with the experimental interface and how to make use of the HI.

[12]The instructions are described in more detail in Subsection 4.3.

Second, we comment on the way we decided to generate the HI and its richness. There are two reasons that lead 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 way we can estimate a linear, zero intercept, bidding strategy for each subject given that realization of HI. This is because we want to track the distribution of the slopes of the bidding strategies over time until it stabilizes, and then compare the distributions across the two treatments. 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. Using 6 rounds in a block represents a good compromise between achieving those objectives and not overloading subjects with too many auctions. Also, 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.[13]

An additional remark regarding our construction of HI is to note that what subjects see about the previous block is influenced both by the others' and by one's own bidding behavior. We prefer providing HI based on everyone's bids over the alternative of omitting own bids for two reasons. First, in order to be close to the theoretical model, we want to preserve the symmetry of the HI. Moreover, an asymmetry may also slow down convergence of bidding behavior. Second, in terms of speaking toward auction design, an auctioneer in field applications would most likely provide aggregate HI, without filtering out previous bids of a given bidder.

Third, we comment on the way HI is communicated to subjects. Here the main tradeoff is between presenting something easy to understand (and therefore use) and preserving detailed information about the distribution of (winning) bids in the previous block. We aim to 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 such way of representing a distribution should be familiar to most subjects from textbooks, magazines, newspapers or introductory statistical classes. Even though a histogram provides a very good quick overview of the underlying distribution, its drawback is that, regardless of how many bins it is based on, it is at least somewhat coarse about fine details of the distribution. To remedy this drawback, subjects are also provided with a second format of HI. By inputting numbers into two boxes, they can recover an exact percentage of (winning) 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. This way, subjects can

[13]We ran five different pilot experiments to learn about how various features of the design impact subject behavior.

recover arbitrary details of the underlying distribution.

Fourth, a distinctive feature of our design is that we want to channel the effect of HI on previous auctions to work exclusively through beliefs about the DHCB. Thus, we would like to filter out other potential effects of the way we present HI. For this purpose, we augment the design by using two different pairings of subjects in treatment W. One is used to construct the HI and it remains the same block after block (see the next subsection). We call this pairing "Information Pairing." Another one is used to determine the actual auction winners and subject earnings. Under this pairing, the subjects are paired independently across all 66 rounds of bidding. We call this pairing "Earnings Pairing." It is identical to the pairing used in treatment A. The two pairings are determined independently. To see why we added this feature, note that if the two pairings coincide, the following issue might arise in treatment W. Because we provide subjects with sufficient tools to access arbitrarily detailed information about the distribution of winning bids from the previous block, a subject could in principle determine which of his/her bids from the previous block were winning and which were not.[14] This possibility is undesirable for the purpose of our experiment because observation of outcomes of own past auctions can affect future bidding behavior in ways that are very different from simply using the HI to form beliefs about the DHCB. Several authors (Ockenfels and Selten 2005, Neugebauer and Selten 2006) have considered how feedback on whether one won or not can affect future bidding via balancing impulses to bid more and to bid less. Such impulses could be interpreted as resulting from loser and winner regret, respectively (Engelbrecht-Wiggans 1989, Filiz-Ozbay and Ozbay 2007, Engelbrecht-Wiggans and Katok 2008). Even though these effects are hypothesised to be the strongest if the highest competing bid is known, which is not the case in our experiment, we consider such feedback on own past auctions undesirable.[15]

In what follows, we discuss design details concerning bidders' valuations and determination of bidders' payoffs from the experiment. As we mentioned earlier, each subject's value is drawn from the uniform distribution on $[0, 100]$. However, the subjects are informed neither about the shape of this distribution, nor about its support. They only know that, in a given bidding round, the values are generated independently across the two bidding opponents using the same random number generator (see the next subsection for more details). 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

[14]This could be done by using the two-number-box HI interface by typing in one's bid from the previous block into both boxes. From debriefing subjects participating in our pilot experiments, we learned that a few of them did this.

[15]Another alternative to exclude impulse balance/regret effects related to feedback from own past auctions would be to exclude the own auction outcomes from HI provided to a subject. We did not follow this option because it would have implied providing a different realization of HI to each bidder, which we wanted to avoid (see the discussion above).

source of information for forming their beliefs. Second, we believe that in 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.

Subject payments are determined by 5 randomly chosen rounds out of the 66 bidding rounds. We believe that 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.

4.2 Noise Reduction Measures

We use several techniques to reduce noise in our data. 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 bidder 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 bidder 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.[16] Furthermore, for each bidder, 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 of constructing the HI on winning bids in treatment W (Information Pairing).[17]

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-

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

[17]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.

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.[18] 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.[19]

4.3 Instructions and Execution of the Experiment

After entering the lab and assuming 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 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 during the experiment. 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.[20]

The instructions also informed the subjects that each bidder's values in block 1 were generated independently across subjects using the same random number generator. However, subjects were told neither the distribution from which the values were drawn, nor its support.

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

[19] 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.

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

The instructions 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, the instructions presented a shot of a bidding screen used in block 1, which contained a link to a built-in calculator that the subjects could use while bidding.[21]

Subjects had the opportunity to go over the instructions at their own pace and privately ask questions. We then administered a quiz consisting of eight multiple-choice questions aimed at testing subjects' understanding of the instructions.[22] 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. 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.

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 (W), subjects were told that, from block 2 onwards, they would receive historical information about 72 bids (36 winning bids) from the auctions played in their group in the previous block. The instructions then provided a detailed explanation of the histogram and the two-box distribution tracker we described in Subsection 4.1 together with a shot of a bidding screen used in blocks 2 through 11. This was the same screen as the one used in block 1 with the addition of the HI presented in the two formats mentioned above. Additionally, in treatment W, the instructions explained the difference between the *Earnings Pairing* and the *Information Pairing* as described in Subsection 4.1. The subjects were explicitly told that the HI they would receive could not be used to determine whether a particular bid they placed in the previous block had or had not been a winning bid. They were also told that the HI generated based on the *Information Pairing* was as informative about the bidding behavior in the group as the one based on the *Earnings Pairing* would have been. Finally, in treatment A, the instructions also explicitly mentioned that since all subjects were assigned the same values in each block, if the HI had changed from one block to another, it must have been because the group members changed how they bid and not for any other reason (e.g., different value realizations in different blocks). In treatment W, this statement was expanded

[21]The instructions are presented in the Appendix.

[22]The quiz is presented in the Appendix.

to include the reason that the same pairing was used for the purpose of generating the HI.[23]

After subjects had an opportunity to go over the instructions at their own pace and privately ask questions, we administered treatment-specific quizzes consisting of ten multiple-choice questions.[24] Again, when checking on the quiz answers, we observed a very good understanding of the instructions. The bidding then continued for blocks 2 through 11. 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. When subjects completed bidding in all 66 rounds, we administered a demographic questionnaire in which we collected information about gender, age, country of origin, number of siblings, academic major, the highest achieved academic degree, self-reported monetary risk-tolerance (on a scale of 1 to 7), previous experience with online and offline auctions and the average amount of monthly spending. After that, 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.

4.4 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 2013 and February 2014 at the Laboratory of Experimental Economics at the University of Economics in Prague.[25] 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

[23]The instructions are presented in the Appendix.

[24]The quizzes from both treatments are presented in the Appendix.

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

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.[26]

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.[27] 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.

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 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

$$\widehat{slope}_{it} = \frac{\sum_{j=1}^{6} v_{ijt} b_{ijt}}{\sum_{j=1}^{6} v_{ijt}^2} = \sum_{j=1}^{6} \left(\frac{v_{ijt}^2}{\sum_{k=1}^{6} v_{ikt}^2} \right) \frac{b_{ijt}}{v_{ijt}}. \tag{16}$$

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

[26] 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.

[27] 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

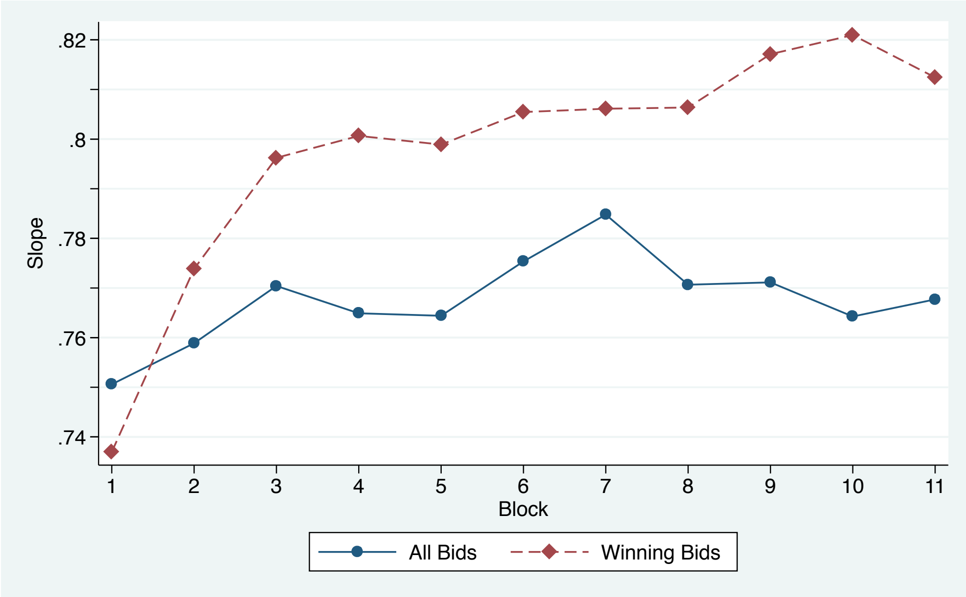

their bidding strategies systematically exceeding 1.

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. (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 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. 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

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

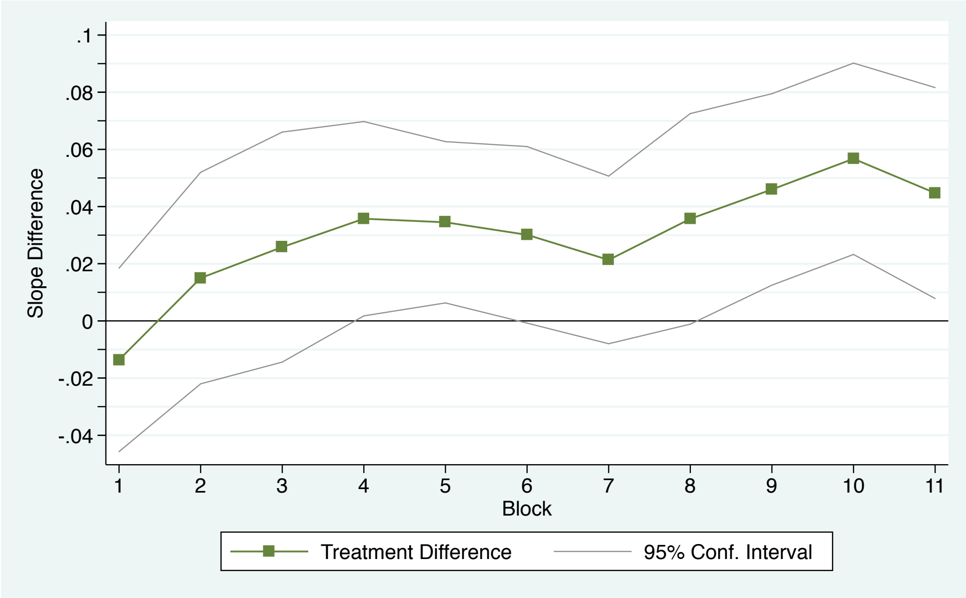

bidding group (10 in each treatment) and we perform the Mann-Whitney ranksum test for the equality of distributions of the average slopes.[28] Starting from block 4 and with the exception of blocks 6 and 7, we reject the null hypothesis in favor of the distribution under W first-order stochastically 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, 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 bidders 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

[28]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

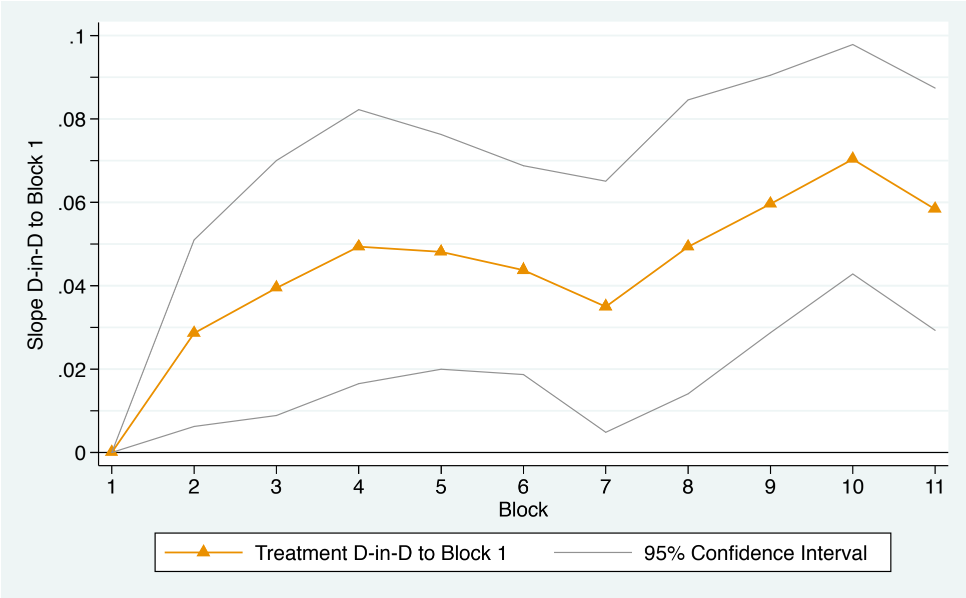

block 3 onward (with the p-value being 0.001 or less from block 6 onward).

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, bidders 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 bidders are given HI on past winning bids.

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

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

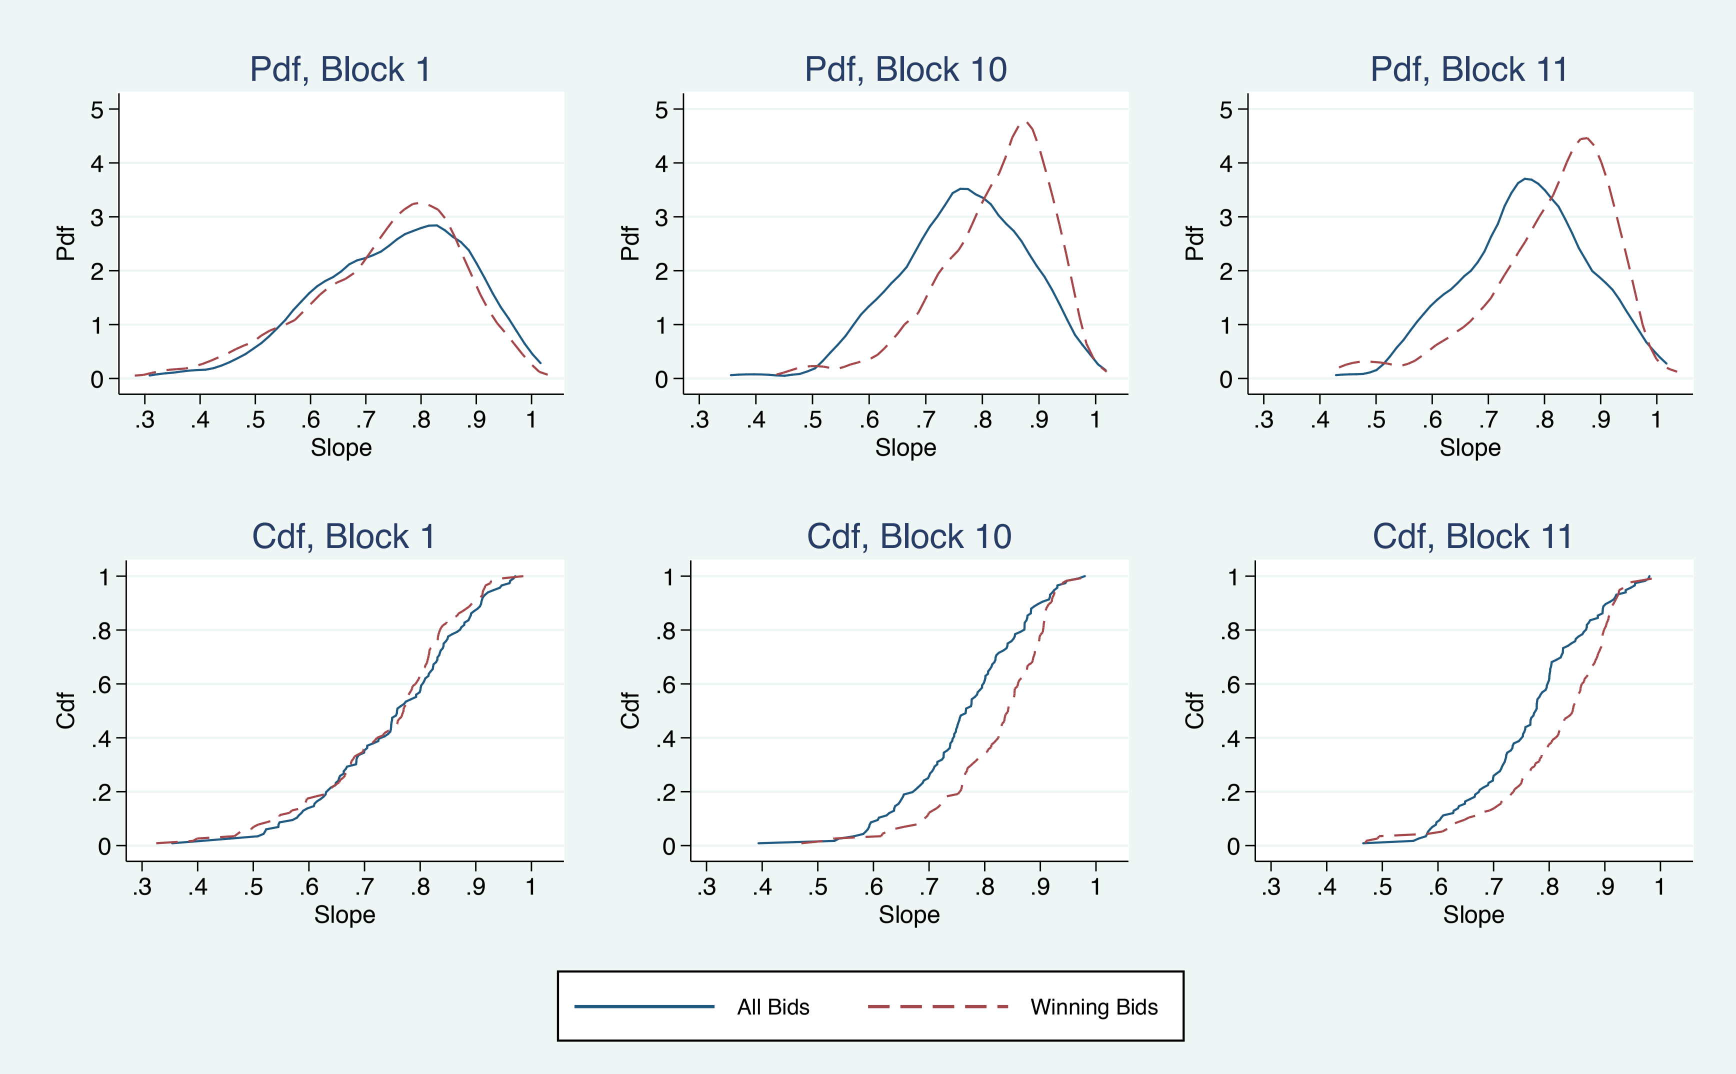

Figure 6: Average Revenue and Average Efficiency across the Treatments in Blocks 1, 10 and 11

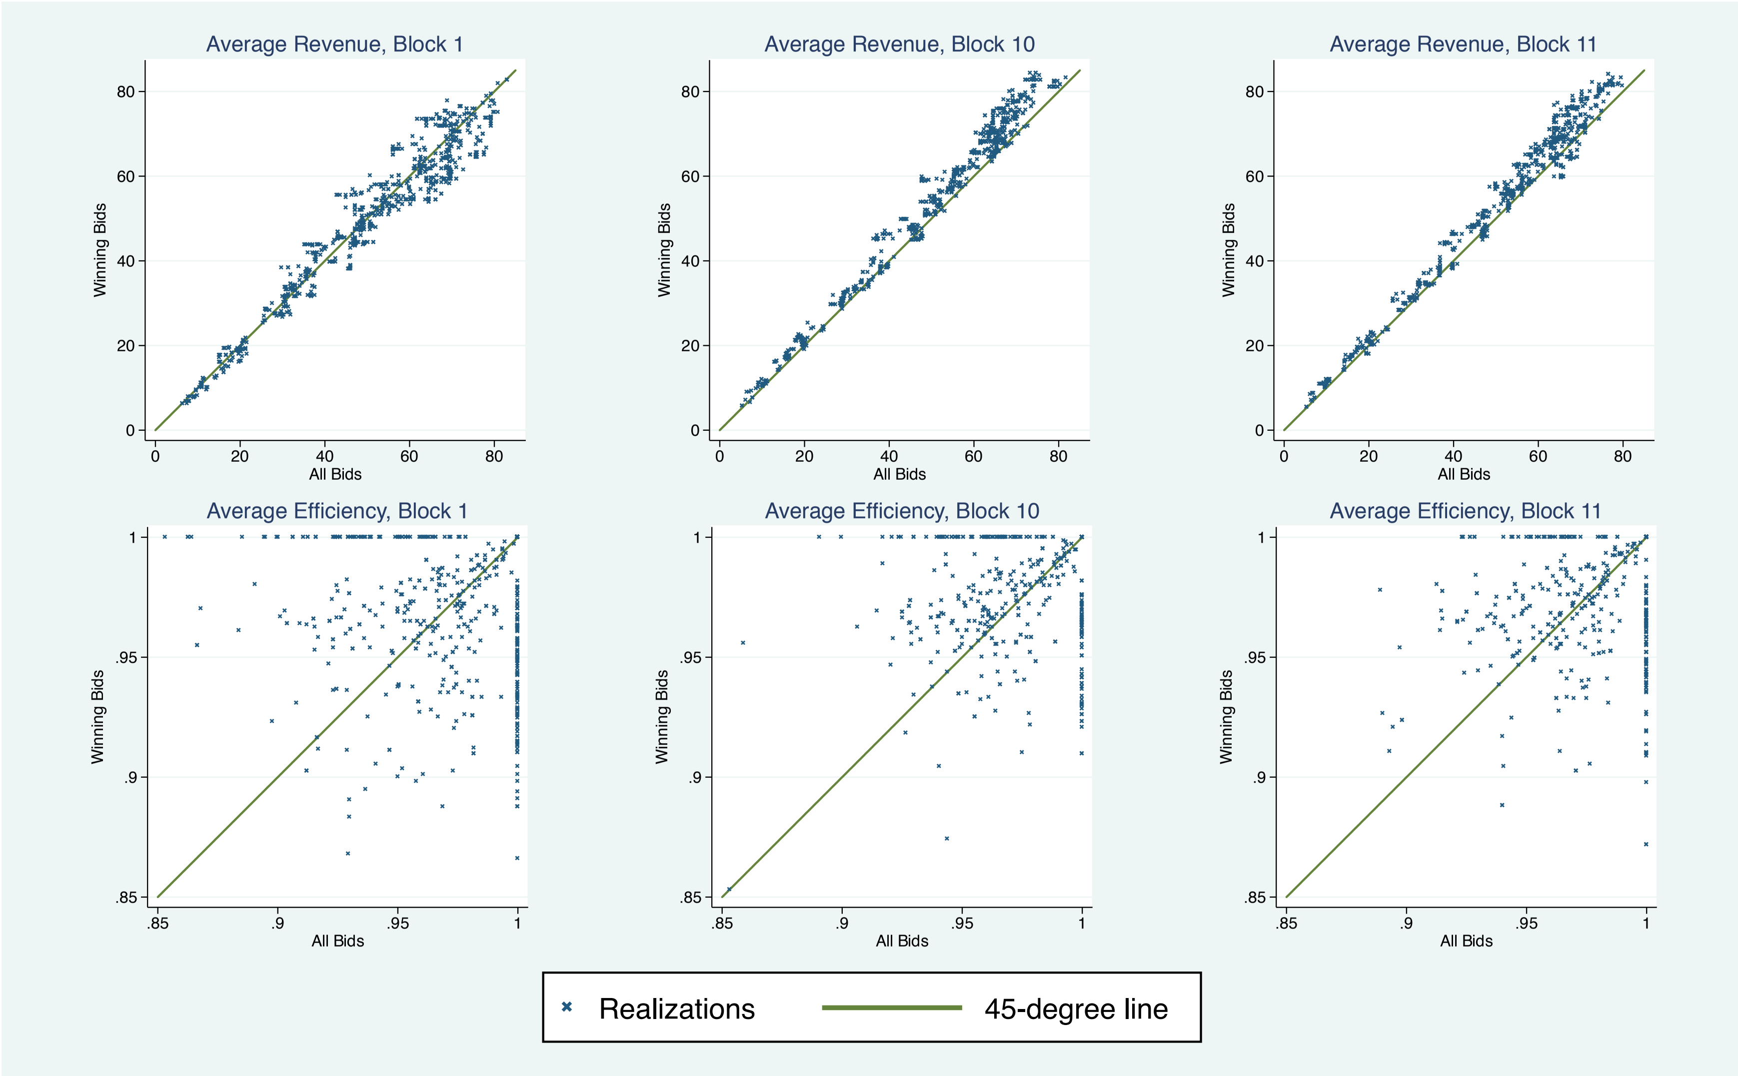

block (rather than round), one should also consider matches across different rounds of a given block. To estimate the treatment effect on average revenue and efficiency in a given block, we generate 1,000 bootstrap draws from the data on values and bids in this block. Each draw is generated as follows. 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 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.[29] 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).[30] 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 intercept gives a coefficient of 1.067 and 1.056 in blocks 10 and 11, respectively.[31] 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

[29] 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.

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

[31] 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 16.

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.

distributions of the various measures reported in the previous two paragraphs under the null hypothesis of no treatment effect.[32] 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 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.

[32] 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.

6 Concluding Remarks

We have compared two policies on the 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 for the auction outcome in the steady state. The reason is that bidders should be able to recover one distribution of bids from the other one. Instead, we have proposed that, in the context of two-bidder auctions, 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 experimental test confirms the prediction of the theory that HI on winning bids should be preferred to HI on all bids in terms of revenue maximization. Thus, 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 opt to disclose past winning bids only. Even though our experimental test focuses on two-bidder first price auctions, we believe that the market design implications extend beyond the two-bidder case and also extend to other market institutions in which the determination of beliefs about the behavior of the opponent(s) has a crucial role.

That said, we want to comment on why focusing on the two-bidder case is particularly important, and what further issues arise once moving away from it. 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 a useful benchmark in that the bounded rationality that we envision should be limited only to the case of HI on winning bids.

The motivation we proposed when describing the two-bidder case was based on a selection bias argument: some bidders (*naive* bidders) best respond to the distribution of winning bids failing to realize that it is not representative of the distribution of all bids. This reasoning can be extended to settings with more than two bidders with some additional assumptions. 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.[33] Under this assumption, naive bidders, when presented with HI on all bids, $G(\cdot)$, form correct beliefs about the DHCB, $G^{n-1}(\cdot)$. However, they overestimate the DHCB when presented with HI on winning bids. Because of the selection bias, they use $G^n(\cdot)$ in place of $G(\cdot)$ to form beliefs

[33] 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.

about the DHCB. After computing the order statistic distribution, they believe that the DHCB is given by $G^{n(n-1)}(\cdot)$ rather than $G^{n-1}(\cdot)$. As a result, they end up bidding more than their best response.[34]

An alternative type of mistake that bidders might be making 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 $G(\cdot)$, whereas under HI on winning bids, they best-respond to $G^n(\cdot)$. As a result, naive bidders end up bidding less than their best response (which is based on the DHCB given by $G^{n-1}(\cdot)$) in the former case and more than their best response in the latter case.

Notice that the predictions based on these two alternative motivations are indistinguishable in auctions with two bidders. They both predict correct DHCB beliefs under HI on all bids and, for the naive bidders, DHCB beliefs given by $G^2(\cdot)$ instead of $G(\cdot)$ under HI on winning bids. For auctions with more than two bidders, the additional assumption in 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.

Note that under both motivating stories, the predicted effect of HI on bidding increases with the number of bidders. In fact, though, it might be unrealistic to assume that the same fraction of naive bidders keeps confusing different distributions of bids as their difference magnifies. It might be more reasonable to expect that the incidence of such naivete diminishes as the size of the resulting mistake becomes larger.

The discussion above suggests 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

[34]This can easily be seen by noting that maximizing $u(v-b)G^k(b)$ is equivalent to maximizing $ln[u(v-b)G^k(b)]$. The derivative of this expression with respect to b is given by $-u'(v-b)/u(v-b)+kG'(b)/G(b)$. Assuming that the distribution of bids is continuous with a positive density, $G'(b)/G(b)$ is positive. As a result, the marginal benefit of increasing one's bid is strictly increasing in k. Hence, any bidder who perceives the DHCB to be $G^{n(n-1)}(\cdot)$ ($k=n(n-1)$) ends up bidding more than a bidder who perceives the DHCB to be $G^{n-1}(\cdot)$ ($k=n-1$).

by GAČR, Paris School of Economics, University College London, University of Arkansas, University of Economic in Prague or CERGE-EI.

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.

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 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.

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

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

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