A/B Testing of Auctions

SHUCHI CHAWLA, University of Wisconsin - Madison
JASON HARTLINE, Northwestern University
DENIS NEKIPELOV, University of Virginia

X

For many application areas A/B testing, which partitions users of a system into an A (control) and B (treatment) group to experiment between several application designs, enables Internet companies to optimize their services to the behavioral patterns of their users. Unfortunately, the A/B testing framework can not be applied in a straightforward manner to applications like auctions where the users (a.k.a., bidders) submit bids before the partitioning into the A and B groups is made. This paper combines auction theoretic modeling with the A/B testing framework to develop methodology for A/B testing auctions. The accuracy of our methodis directly comparable to ideal A/B testing where there is no interference between A and B. Our results are based on an extension and improved analysis of the inference method of Chawla et al. [2014].

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

Additional Key Words and Phrases: econometrics, first price auction, all pay auction, revenue maximization

1. INTRODUCTION

A common method in the practice of large scale auction design, e.g., in auctions placing advertisements on online media and Internet search engines, is A/B testing. In A/B testing, the auction house is running an incumbent mechanism A, and would like to determine if a novel mechanism B obtains higher revenue. This is done by splitting the traffic so that most of it goes A and some of it, e.g., five to ten percent, goes to B.[1] An issue with this approach is that if the bidders are unaware of which mechanism their bid will be considered in, the bid equilibrium is neither for A nor B but for a mechanism C that is a convex combination of A and B.

A miscalculation sometimes performed in practice is to consider and compare average revenue from A (resp. B) from the times when A (resp. B) is run. This miscalculation is equivalent to simulating A on the bids in C and can often give the opposite conclusion; e.g., if A and B are one- and two-unit highest-bids-win pay-your-bid auctions, respectively, then B will always appear in this miscalculation to have higher revenue. For a fixed set of bids, a pay-your-bid mechanism's revenue is monotone in its allocation probabilities. Of course, in equilibrium, increased allocation probabilities can cause reduced revenue as bidders may lower their bids.

In this paper we develop a simple, direct method for inferring the revenues of mechanisms A and B from the equilibrium bids for mechanism C. To estimate the revenue of A (analogously for B):

(1) for any number N, the evaluation of a formula based on the definition of A and C gives N weights, and
(2) from a sorted list of N i.i.d. samples from the bid distribution, the estimated revenue of A is the weighted average of these bids.

This method of estimation compares favorably to *ideal A/B testing*, where the bids in A are in equilibrium for A, and respectively for B.

Our method applies generally to the position auction model popularized by the Varian [2007] and Edelman et al. [2007] analyses of auctions for sponsored search and now a fundamental model for the study of auction theory, see Hartline [2013]. To revisit the example auction scenario above, notice that the mechanism C, namely the 50–50 convex combination of one- and two-unit auctions, is mathematically equivalent to a position auction with weights 1 and 0.5 for the top two slots respectively, and zero for all remaining slots. Mechanisms A and B are also position auctions. More generally, a position auction is defined by a decreasing sequence of weights, bidders are assigned to positions in decreasing order of bids, and payments are charged. Typical payment rules are "generalized first price" and "generalized second price", the former requires bidders to pay their weighted bid, the latter requires bidders to pay the weighted bid of the next highest bidder. Position auctions with exogenously given weights are the standard model of auction for selling advertisements on Internet search engines. In this application, the weights correspond to the probability that the ad in each position on the search

[1] For very thin markets, the ratio of the split can even be 50-50.

results page will be clicked on (and payments are scaled by weights because bidders only pay if they are clicked on).

Our model of A/B testing is especially applicable to the following generalization of the classical problem with fixed weights. Consider the problem Google and Bing face in selecting the layout of the sponsored results on the results page of a search query. The main variable in the layout is the number of ads to display in the mainline, i.e., above the organic search results, versus the sidebar, i.e., to the right of the organic search results. Typical layouts include between zero and three advertisements in the mainline, and up to six on the sidebar. The weights, which correspond to click probabilities, depend on which of these four layout choices is selected. The selection of the number of mainline ads to show generalizes the opening example of selecting whether to sell one or two units of an item.

This paper extends the method and refines the analysis of Chawla et al. [2014]. A position auction can be viewed as a convex combination of multi-unit auctions. Chawla et al. [2014] give a method for estimating the revenue of the k-unit auction from the bids in any position auction. Our estimator can be viewed as the appropriate convex combination of the estimators for the multi-unit auction revenues. While a naive analysis bounds our estimation error by the worst estimation error of any of the multi-unit revenues; our main result is an improved analysis for the A/B testing environment where mechanism A (respectively, B) has a significant weight in the A/B test mechanism C. We also give an improved analysis of the estimation error for the multi-unit auction revenues.

In terms of the number N of samples from the bid distribution our estimation methods achieve the optimal rate of $\sqrt{N}$, i.e., the mean squared error is $O(1/\sqrt{N})$. When the A/B test is employed not to estimate revenues but to determine whether to replace incumbent auction A with novel auction B, the misclassification error is exponentially small in N. For our application to advertising on Internet search engines, deciding the number of ads to show on the mainline (between zero and four), again the misclassification probability is exponentially small in N. The precice theorems given in the paper give the dependence on the number of bidders in the auction and the fraction of traffic sent to each the A and B mechanisms.

Summary of Results. Our main results apply to first-price and all-pay position auctions and are:

(1) We improve the analysis of Chawla et al. [2014] for the error in estimating the revenue of a multi-unit auction from the bids in a position auction. (Sections 3 and 5.)
(2) We give a direct analysis of the estimation of the revenue of B from the bids of C which depends on the measure ϵ on B in the support of C. The improvement of (1) and this direct analysis enables the estimation of the revenue of B from C $= (1-\epsilon)$A $+ \epsilon$B with logarithmic dependence on $1/\epsilon$ which allows good error bounds even when ϵ is exponentially small. (Section 4.1.)
(3) If we do not wish to estimate revenues but only determine which of A or B will achieve higher revenue, we give an estimator that is accurate with probability approaching 1 at an exponential rate in N the number of observed bids from the A/B test mechanism C. In comparison, estimation of the auction revenues has quadratic rate. (Section 4.2.)
(4) We complement our theoretical analysis with simulation bounds that show that our methods are indeed very practical. (Section 6.)

Revenue versus Click Estimation. The primary focus of this paper and our discussion thus far is on A/B testing of auctions when the details of the auction are known. For our motivating application of position auctions when we vary the number of ads shown on the mainline (a.k.a., the configuration), we assumed that the click-through rates (a.k.a., position weights) of the positions for each configuration are known. These weights are generally not known, and can also be estimated. Estimation of these weights is much easier that estimating revenue and ideal A/B testing gives good estimates.

The distinction between the revenue estimation and click-through rate estimation problems is the following. Revenue is determined from bids that are submitted in advance of the partitioning choice of whether to run mechanism A or B; consequently, the bids are not in equilibrium for either A or B but for their convex combination C. Click-through rates arise from the behavior of the viewer of the search results page and this viewer's behavior is affected by the layout directly. Thus, while revenue must be estimated via inference methods like ours, the click-through rates can be estimated by ideal A/B testing.

In the *separable model*, where the click-through rates are exogenous to the layout and, importantly, do not depend on which ads are shown,[2] they can be estimated at the same time as we A/B test the auction formats for revenue. Our estimators are linear in the click-through rates; therefore, unbiased estimators for click-through rates can be used in place of exact click-through rates and the resulting revenue estimates will be unbiased. Furthermore, as the separable model posits that click rates are independent from bids, estimation errors in click-through rates can be easily factored in to the error of the revenue estimates.

The Generalized Second Price Auction. Our analysis does not extend to the generalized second price auction, i.e., to position auctions where each bidder pays the weighted bid of the next highest bidder. The first two issues that arise for the generalized second price auction are: (1) efficient symmetric equilibria may not exist [Gomes and Sweeney 2009]; (2) even when a symmetric equilibrium does exist, it may not be unique [Chawla and Hartline 2013]. For these reasons there is not a one-to-one correspondence between position weights and the allocation rule of the mechanisms.

Even restricted to environments where a unique efficient symmetric Bayes-Nash equilibrium exists, Bayes-Nash inference is challenging because bidder payments are not a simple function of their bids and inverting the bid function is problematic. Specifically, for the generalized first price auction, knowing the bid distribution exactly allows one to obtain the value distribution via a formula that for any quantile maps the bid and bid density at that quantile to the value at that quantile (see Equation (3)). On the other hand, for the generalized second price auction, estimating the value at some quantile requires knowing the bid and the bid density not only at that same quantile, but at all quantiles. Importantly, the function mapping bids to values is not linear in the bids. This makes inference challenging, and can potentially lead to large errors. Extending our methodology to this setting is an important open problem.

Related Work. Prior to our previous work [Chawla et al. 2014], to our knowledge, the problem of inferring parameters of one mechanism using observations drawn from another has not been considered from a theoretical perspective. A different line of work considers learning parameters of the environment from the outcome of a mechanism (e.g., Devanur and Kakade [2009]; Babaioff et al. [2009]; Gatti et al. [2012]), however, they either focus on truthful mechanisms or non-strategic settings, and do not require inference.

Within econometrics literature, several recent works have developed techniques for estimating the distribution of values from the distribution of bids of a given auction (see, e.g., Guerre et al. [2000], Paarsch and Hong [2006], and Bajari et al. [2013]). However, estimating parameters of one mechanism from the bids of another has not been considered. While these inference approaches form the basis over which we develop our methodology, our work shows that for certain quantities of interest such as the multi-unit auction revenues or the ranking of revenues of different mechanisms, we can bypass learning the underlying value distribution altogether, leading to better convergence bounds.

Several works have considered the problem of empirically optimizing the reserve price of an auction in an online repeated auction setting (e.g., Reiley [2006]; Brown and Morgan [2009]; Ostrovsky and Schwarz [2011]). The most notable of these is the work of Ostrovsky and Schwarz [2011]. Ostrovsky and Schwarz [2011] adapt their mechanism over time to respond to empirical data by determining the optimal reserve price for the empirically observed distribution, and then setting a reserve price that is slightly smaller. This allows for inference around the optimal reserve price and ensures that the mechanism quickly adapts to changes in the distribution.

2. PRELIMINARIES

A standard auction design problem is defined by a set $[n] = \{1, \ldots, n\}$ of $n \geq 2$ agents, each with a private value v_i for receiving a service. The values are bounded: $v_i \in [0, 1]$; They are independently and identically distributed according to a continuous distribution F. For x_i indicating the probability of service and p_i the expected payment required, agent i has linear utility $u_i = v_i x_i - p_i$. An auction elicits bids $\boldsymbol{b} = (b_1, \ldots, b_n)$ from the agents and maps the vector $\boldsymbol{b}$ of bids to an allocation $\tilde{\boldsymbol{x}}(\boldsymbol{b}) = (\tilde{x}_1(\boldsymbol{b}), \ldots, \tilde{x}_n(\boldsymbol{b}))$, specifying the probability with which each agent is served, and prices $\tilde{\boldsymbol{p}}(\boldsymbol{b}) = (\tilde{x}_1(\boldsymbol{b}), \ldots, \tilde{x}_n(\boldsymbol{b}))$, specifying the expected amount that each agent is required to pay. An auction is denoted by $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$.

[2]Though inaccurate, this is one of the most standard models under which the search advertising auction is studied.

Standard payment formats. In this paper we study two standard payment formats. In a *first-price* format, each agent pays his bid upon winning, that is, $\tilde{p}_i(\boldsymbol{b}) = b_i\,\tilde{x}_i(\boldsymbol{b})$. In an *all-pay* format, each agent pays his bid regardless of whether or not he wins, that is, $\tilde{p}_i(\boldsymbol{b}) = b_i$.

Bayes-Nash equilibrium. The values are independently and identically distributed according to a continuous distribution F. This distribution is common knowledge to the agents. A strategy s_i for agent i is a function that maps the value of the agent to a bid. The distribution of values F and a profile of strategies $\boldsymbol{s} = (s_1, \cdots, s_n)$ induces interim allocation and payment rules (as a function of bids) as follows for agent i with bid b_i.

$$\tilde{x}_i(b_i) = \mathbf{E}_{\boldsymbol{v}_{-i}\sim F}[\tilde{x}_i(b_i, \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}))] \text{ and}$$
$$\tilde{p}_i(b_i) = \mathbf{E}_{\boldsymbol{v}_{-i}\sim F}[\tilde{p}_i(b_i, \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}))]\,.$$

Agents have linear utility which can be expressed in the interm as:

$$\tilde{u}_i(v_i, b_i) = v_i\tilde{x}_i(b_i) - \tilde{p}_i(b_i).$$

The strategy profile forms a *Bayes-Nash equilibrium* (BNE) if for all agents i, values v_i, and alternative bids b_i, bidding $s_i(v_i)$ according to the strategy profile is at least as good as bidding b_i. I.e.,

$$\tilde{u}_i(v_i, s_i(v_i)) \geq \tilde{u}_i(v_i, b_i). \tag{1}$$

A symmetric equilibrium is one where all agents bid by the same strategy, i.e., $\boldsymbol{s}$ statisfies $s_i = s$ for some s. For a symmetric equilibrium, the interim allocation and payment rules are also symmetric, i.e., $\tilde{x}_i = \tilde{x}$ and $s_i = s$ for all i. For implicit distribution F and symmetric equilibrium given by stratey s, a mechanism can be described by the pair $(\tilde{x}, \tilde{p})$. Chawla and Hartline [2013] show that the equilibrium of every auction in the class we consider is unique and symmetric.

The strategy profile allows the mechanism's outcome rules to be expressed in terms of the agents' values instead of their bids; the distribution of values allows them to be expressed in terms of the agents' values relative to the distribution. This later representation exposes the geometry of the mechanism. Define the *quantile* q of an agent with value v to be the probability that v is larger than a random draw from the distribution F, i.e., $q = F(v)$. Denote the agent's value as a function of quantile as $v(q) = F^{-1}(q)$, and his bid as a function of quantile as $b(q) = s(v(q))$. The outcome rule of the mechanism in quantile space is the pair $(x(q), p(q)) = (\tilde{x}(b(q)), \tilde{p}(b(q)))$.

Revenue curves and auction revenue. Myerson [1981] characterized Bayes-Nash equilibria and this characteriation enables writing the revenue of a mechanism as a weighted sum of revenues of single-agent posted pricings. Formally, the *revenue curve* $R(q)$ for a given value distribution specifies the revenue of the single-agent mechanism that serves an agent with value drawn from that distribution if and only if the agent's quantile exceeds q: $R(q) = v(q)\,(1-q)$. $R(0)$ and $R(1)$ are defined as 0. Myerson's characterization of BNE then implies that the expected revenue of a mechanism at BNE from an agent facing an allocation rule $x(q)$ can be written as follows:

$$P_x = \mathbf{E}_q[R(q)x'(q)] = -\mathbf{E}_q[R'(q)x(q)] \tag{2}$$

where x' and R' denote the derivative of x and R with respect to q, respectively.

The expected revenue of an auction is the sum over the agents of its per-agent expected revenue; for auctions with symmetric equilibrium allocaton rule x this revenue is $n\,P_x$.

Position environments. A *position environment* expresses the feasibility constraint of the auciton designer in terms of *position weights* $\boldsymbol{w}$ satisfying $1 \geq w_1 \geq w_2 \geq \cdots \geq w_n \geq 0$. A *position auction* assigns agents (potentially randomly) to positions 1 through n, and an agent assigned to position i gets allocated with probability w_i. The *rank-by-bid position auction* orders the agents by their bids, with ties broken randomly, and assigns agent i, with the ith largest bid, to position i, with allocation probability w_i. *Multi-unit environments* are a special case and are defined for k units as $w_j = 1$ for $j \in \{1, \ldots, k\}$ and $w_j = 0$ for $j \in \{k+1, \ldots, n\}$. The *highest-k-bids-win* multi-unit auction is the special case of the rank-by-bid position auction for the k-unit environment.

Rank-by-bid position auctions can be viewed as convex combinations of highest-bids-win multi-unit auctions. The *marginal weights* of a position environment are $\boldsymbol{w}' = (w'_1, \ldots, w'_n)$ with $w'_k = w_k - w_{k+1}$. Define $w'_0 = 1 - w_1$ and note that the marginal weights $\boldsymbol{w}'$ can be interpreted as a probability distribution over $\{0, \ldots, n\}$. As rank-by-bid position auctions are efficient, the rank-by-bid position

auction with weights $\boldsymbol{w}$ has the exact same allocation rule as the mechanism that draws a number of units k from the distribution given by $\boldsymbol{w}'$ and runs the highest-k-bids-win auction.

In our model with agent values drawn i.i.d. from a continuous distribution, rank-by-bid position auctions with either all-pay or first-price payment semantics have a unique Bayes-Nash equilibrium and this equilibrium is symmetric and efficient, i.e., in equilibrium, the agents' bids and values are in the same order [Chawla and Hartline 2013]. Denote the highest-k-bids-win allocation rule as $x^{(k:n)}$ and its revenue as $P_k = P_{x^{(k:n)}} = \mathbf{E}_q\big[-R'(q)\, x^{(k:n)}(q)\big]$. This allocation rule is precisely the probability an agent with quantile q has one of the highest k quantiles of n agents, or at most $k-1$ of the $n-1$ remaining agents have quantiles greater than q. Formulaically,

$$x^{(k:n)}(q) = \sum_{i=0}^{k-1} \tbinom{n-1}{i} q^{n-1-i}(1-q)^i.$$

Importantly, the allocation rule of a rank-by-bid position auction does not depend on the distribution at all. The allocation rule x of the rank-by-bid position auction with weights $\boldsymbol{w}$ is:

$$x(q) = \sum\nolimits_k w'_k\, x^{(k:n)}(q).$$

By revenue equivalence [Myerson 1981], the per-agent revenue of the rank-by-bid position auction with weights $\boldsymbol{w}$ is:

$$P_x = \sum\nolimits_k w'_k\, P_k.$$

Of course, $P_0 = P_n = 0$ as always serving or never serving the agents gives zero revenue.

Inference. The distribution of values, which is unobserved, can be inferred from the distribution of bids, which is observed. This derivation is summarized by Lemma 2.1, below. Interested readers can find a detailed derivation in Chawla et al. [2014] or Appendix A.

LEMMA 2.1. *For an auction with first-price payment format, the value function $v(\cdot)$ satisfies,*

$$v(q) = b(q) + \tfrac{x(q)\, b'(q)}{x'(q)}. \tag{3}$$

For an auction with all-pay payment format, it satisfies,

$$v(q) = \tfrac{b'(q)}{x'(q)}. \tag{4}$$

Recall that the functions $x(q)$ and $x'(q)$ are known precisely for position auctions, while functions $b(q)$ and $b'(q)$ are observed in the data and can be estimated.

Statistical Methods. The errors in the estimated bid distribution follow from standard analyses. For N draws from the bid distribution, define the *estimated bid distribution* $\hat{b}(\cdot)$ as

$$\hat{b}(q) = \hat{b}_i \qquad\qquad \forall i \in N, q \in [i-1, i)/N \tag{5}$$

where $\hat{b}_i$ is the ith smallest bid in the sample. For function $b(\cdot)$ and estimator $\hat{b}(\cdot)$, the *absolute error* as a function of the number of samples N is

$$\mathbf{E}\Big[\sup_q \big|b(q) - \hat{b}(q)\big|\Big].$$

The *rate of convergence* of an unbiased estimator, $r(N)$, captures the dependence of the mean squared error in terms of the number of samples N, keeping all other quantities (including, e.g., the number of agents n) constant. The following lemma gives the standard bound on the mean squared error for the estimated bid distirbution. Here, G and $\hat{G}$ denote the real and empirical cumulative distribution of bids, respectively. Interested readers can find a detailed derivation in Chawla et al. [2014].

LEMMA 2.2. *The absolute error in the bid distribution is bounded as: $\mathbf{E}\Big[\sup_q \Big|\hat{G}(b(q)) - G(b(q))\Big|\Big] \leq \frac{1}{2\sqrt{N}}$. Consequently, the estimator $\hat{b}(\cdot)$ defined directly from the bids by equation* (5) *has a rate of convergence of $r(N) = \sqrt{N}$ and absolute error $\mathbf{E}\Big[\sup_q \big|b(q) - \hat{b}(q)\big|\Big] < \sup_q b'(q)/2\sqrt{N}$. Recalling that $v(q) \leq 1$*

for all quantiles q, we obtain the following expressions for absolute error in terms of the allocation function.

(i) If the samples are generated from the first-price auction, then the absolute error in bids can be evaluated as $\mathbf{E}\left[\sup_q \left|b(q) - \hat{b}(q)\right|\right] < \sup_q \left\{\frac{x'(q)}{x(q)}\right\} /2\sqrt{N}$.

(ii) If the samples are generated from the all-pay auction, then the absolute error in bids can be evaluated as $\mathbf{E}\left[\sup_q \left|b(q) - \hat{b}(q)\right|\right] < \sup_q x'(q)/2\sqrt{N}$.

Lemma 2.1 enables the value function, or equivalently, the value distribution, to be estimated from the estimated bid function and an estimator for the derivative of the bid function, or equivalently, the density of the bid distribution.

3. INFERENCE METHODOLOGY AND ERROR BOUNDS FOR ALL-PAY AUCTIONS

We will now develop a methodology and error bounds for estimating the revenue of one rank-based auction using bids from another rank-based auction. There are two reasons behind our assumption that the auction that we run (that generates the observed bids) is a rank-based auction. First, the allocation rule (in quantile space) of a rank based auction is independent of the bid and value distribion; therefore, it is known and does not need to be estimated. Second, the allocation rules that result from rank-based auctions are well behaved, in particular their slopes are bounded, and our error analysis makes use of this property.

Recall from Section 2 that the revenue of any rank-based auction can be expressed as a linear combination of the multi-unit revenues $P_1, \ldots, P_n$ with P_k equal to the per-agent revenue of the highest-k-bids-win auction. Therefore, in order to estimate the revenue of a rank-based auction, it suffices to estimate the P_ks accurately.

In the following section we describe a function mapping the observed bids to the revenue estimate. In Sections 3.2 and 3.3 we develop error bounds on our estimate for the revenue of a multi-unit auction and a general position auction respectively.

3.1. Inference equation

Consider estimating the revenue of an auction with allocation rule y from the bids of an all-pay position auction. The per-agent revenue of the allocation rule y is given by:

$$P_y = \mathbf{E}_q[y'(q)R(q)] = \mathbf{E}_q[y(q)v(q)(1-q)]$$

Let x denote the allocation rule of the auction that we run, and b denote the bid distribution in BNE of this auction. Recall that for an all-pay auction format, we can convert the bid distribution into the value distribution as follows: $v(q) = b'(q)/x'(q)$. Substituting this into the expression for P_y above we get

$$P_y = \mathbf{E}_q\left[y(q)(1-q)\frac{b'(q)}{x'(q)}\right] = \mathbf{E}_q[Z_y(q)b'(q)]$$

where $Z_y(q) = (1-q)\frac{y(q)}{x'(q)}$.

This expression allows us to derive the revenue using the empirical bid distribution. The direct approach via an estimation of the derivative of the bid distirbution $b'(\cdot)$ leads to a large error because estimating derivatives cannot be done with a good convergence rate. In Chawla et al. [2014] we observe that integration by parts yields an estimator that is a simple weighted average of the empirical bid distribution.

LEMMA 3.1 (CHAWLA ET AL. [2014]). *The per-agent revenue of a rank-based auction with allocation rule y can be written as a linear combination of the bids in an all-pay auction:*[3]

$$P_y = \mathbf{E}_q\left[-Z'_y(q)b(q)\right]$$

where $Z_y(q) = (1-q)\frac{y(q)}{x'(q)}$ depends on the allocation rule x of the mechanism and is known precisely.

[3]Note that $b(0) = 0$ and $Z_k(1) = 0$.

This formulation allows us to estimate P_y directly as a weighted average of the observed bids. Specifically, recalling that $\hat{b}(\cdot)$ denotes the estimated bid distribution, we can write the estimate of P_y as:

$$\hat{P}_y = \mathbf{E}_q\Big[-Z'_y(q)\hat{b}(q)\Big] \tag{6}$$

and, using that $\hat{b}$ is a step function,

$$= \sum_{i=1}^{N} \hat{Z}_{y,i}\hat{b}_i \tag{7}$$

where $\hat{b}_i$ is the ith smallest of the N bids obtained from auction x, and,

$$\hat{Z}_{y,i} = \left(1 - \frac{i}{N}\right)\frac{y'(\frac{i}{N})}{x'(\frac{i}{N})} - \left(1 - \frac{i+1}{N}\right)\frac{y'(\frac{i+1}{N})}{x'(\frac{i+1}{N})}.$$

A consequence of (7) is that our estimator for P_y is linear in y', which, if y is a position auction, in turn depends linearly on the position weights. Therefore a relative error in estimating the position weights in y translates into the same relative error in estimating P_y, in addition to the error bounds given below.

3.2. Error bound for the multi-unit revenues

We will now develop an error bound for the estimator for Lemma 3.1 for the case of the multi-unit revenues. In the following, denote the the allocation rule of the highest-k-bids-win auction as x_k (for an implicit number n of agents). We can therefore express the error in the estimation of P_k in terms of the error in estimating the bid distribution.

$$|\hat{P}_k - P_k| \;=\; \left|\mathbf{E}_q\Big[-Z'_k(q)(\widehat{b}(q) - b(q))\Big]\right| \tag{8}$$

In Chawla et al. [2014] we gave a simple but weak analysis of the error, obtaining the following bound:

$$\mathbf{E}_{\hat{b}}\Big[|\hat{P}_k - P_k|\Big] \le \sqrt{\frac{2}{N}}\;\sup_q\{x'(q)\}\;\sup_q\left\{\frac{x'_k(q)}{x'(q)}\right\}$$

To understand this error bound, we note that the maximum slope of the multi-unit allocation rules x_k, and therefore also that of any rank-based auction, is always bounded by n, the number of agents in the auction (summarized as Fact 3.2, below). This bounds the $\sup_q x'(q)$ term.

FACT 3.2. *The maximum slope of the allocation rule x of any n-agent position auction is bounded by n: $\sup_q x'(q) \le n$. The maximum slope of the allocation rule x_k for the n-agent highest-k-bids-win auction is bounded by*

$$\sup_q x'_k(q) \in \left[\frac{1}{\sqrt{2\pi}}, \frac{1}{\sqrt{\pi}}\right]\frac{n-1}{\sqrt{\min\{k-1, n-k\}}} = \Theta\left(\frac{n}{\sqrt{\min\{k, n-k\}}}\right).$$

The quantity $\sup_q\left\{\frac{x'_k(q)}{x'(q)}\right\}$ on the other hand may be rather large, even unbounded, if $x'(\cdot)$ is near zero at some q. The approach taken in Chawla et al. [2014] is to explicitly design x with position weight $w'_k > \epsilon$ (i.e., to mix x_k into x) to ensure this latter term is bounded. We will now develop a stronger bound on the error in P_k with better dependence on the quantity $\sup_q\left\{\frac{x'_k(q)}{x'(q)}\right\}$.

We start with Equation (8) and partition the expectation on the right hand side of the equation as follows for some $\alpha > 0$:

$$\begin{aligned}|\hat{P}_k - P_k| \;&=\; \left|\mathbf{E}_q\Big[-Z'_k(q)(\widehat{b}(q) - b(q))\Big]\right| \\ &\le\; \mathbf{E}_q\left[\frac{(\log(1+Z_k(q)))^\alpha}{Z_k(q)}|Z'_k(q)|\right]\sup_q\left|\frac{Z_k(q)}{(\log(1+Z_k(q)))^\alpha}(\hat{b}(q) - b(q))\right|\end{aligned}$$

An appropriate choice of α gives us the following theorem. We defer the proof to the appendix.

THEOREM 3.3. *Let x_k denote the allocation function of the k-highest-bids-win auction and x be the allocation function of any rank-based auction. Then for all k, the mean squared error in estimating P_k from N samples from the bid distribution for an all-pay auction with allocation rule x is bounded by:*

$$\mathbf{E}_{\hat{b}}\left[|\hat{P}_k - P_k|\right] \leq \frac{20}{\sqrt{N}} \sup_q\{x'_k(q)\} \log\max\left\{\sup_{q:x'_k(q)\geq 1} \frac{x'(q)}{x'_k(q)}, \sup_q \frac{x'_k(q)}{x'(q)}\right\}.$$

Invoking Fact 3.2, we note that the error in P_k given by Theorem 3.3 can be bounded by

$$\mathbf{E}_{\hat{b}}\left[|\hat{P}_k - P_k|\right] = O\left(\frac{n}{\sqrt{\min\{k, n-k\}}}\right)\frac{1}{\sqrt{N}} \log\max\left\{\sup_{q:x'_k(q)\geq 1} \frac{x'(q)}{x'_k(q)}, \sup_q \frac{x'_k(q)}{x'(q)}\right\}.$$

To understand the above error bound better, we make the following observations:

— When $x = x_k$, we get an error bound of $O(\sup_q x'_k(q)/\sqrt{N})$, which is the same (within constant factors) as the statistical error in bids.
— We also get a good error bound when x and x_k are close enough without being identical: when $\epsilon x'_k \leq x' \leq x'_k/\epsilon$, we get a bound of $O(\log(1/\epsilon)\sup_q x'_k(q)/\sqrt{N})$.
— Finally, as long as $x' \geq \epsilon x'_k$, that is, the highest-k-bids-win auction is mixed in with ϵ probability into x, we observe via Fact 3.2 that $\sup_{q:x'_k(q)\geq 1} x'(q)/x'_k(q) \leq \sup_q x'(q) \leq n$, and obtain an error bound of $O(\log(n/\epsilon)\sup_q x'_k(q)/\sqrt{N})$.

3.3. Error bound for arbitrary rank-based revenues

We now develop an error bound for our estimator for the revenue, P_y, of an arbitrary position auction with allocation rule y from the bids of another position auction x. Let us write y as a position auction with weights $\boldsymbol{w}$:

$$y = \sum\nolimits_k w'_k\, x_k$$
$$P_y = \sum\nolimits_k w'_k\, P_k$$

Accordingly, the error in P_y is bounded by a weighted sum of the error in P_k:

$$\mathbf{E}_{\hat{b}}\left[|\hat{P}_y - P_y|\right] \leq \sum\nolimits_k w'_k\, \mathbf{E}_{\hat{b}}\left[|\hat{P}_k - P_k|\right]$$

applying Theorem 3.3,

$$\begin{aligned}
&\leq \frac{20}{\sqrt{N}} \sum\nolimits_k w'_k \sup_q\{x'_k(q)\}\left(\log n + \log\frac{1}{w'_k} + \log\sup_q \frac{y'(q)}{x'(q)}\right)\\
&= O\left(\frac{n\log n}{\sqrt{N}} \log\sup_q \frac{y'(q)}{x'(q)}\right) \qquad (9)
\end{aligned}$$

Unfortunately, the above bound can be quite loose, as the following simple example demonstrates. Suppose that $x = y$ and $w'_k = 1/n$ for all k. Then the above approach (via a tighter bound on the sum over k) leads to an error bound of $O(\sqrt{n}\log n/\sqrt{N})$, whereas, the true error bound should be $O(1/\sqrt{N})$, arising due to the statistical error in bids. Furthermore, it is desirable to obtain an error bound that depends directly on $\sup_q y'(q)$, rather than on the constituent $\sup_q x'_k(q)$; The latter can be much larger than the former. Below, we analyze the error in P_y directly, leading to a slightly tighter bound.

THEOREM 3.4. *The expected absolute error in estimating the revenue of a position auction with allocation rule y using N samples from the bid distribution for an all-pay position auction with allocation*

rule x is bounded as below; Here n is the number of positions in the two position auctions.

$$\mathbf{E}_{\hat{b}}\Big[|\hat{P}_y - P_y|\Big] \leq \frac{20}{\sqrt{N}}\sqrt{n\log n}\ \sup_q\{y'(q)\}\ \log\max\left\{\sup_{q:y'(q)\geq 1}\frac{x'(q)}{y'(q)}, \sup_q \frac{y'(q)}{x'(q)}\right\} + \frac{O(1)}{N}\sup_q\{x'(q)\}\ \sup_q\left\{\frac{y'(q)}{x'(q)}\right\}$$

Note that the first term in the error bound in Theorem 3.4 dominates, and this term is identical to the bound in Theorem 3.3, except for an extra $\sqrt{n\log n}$ term. Moreover, when $\sup_q y'(q) < \sqrt{n}$, Theorem 3.4 gives us a tighter error bound than Equation (9). We will now prove the theorem. Here we provide an outline for the proof; The complete argument can be found in the appendix.

PROOF SKETCH: As for the multi-unit revenues,

$$|\hat{P}_y - P_y| = \Big|\mathbf{E}_q\Big[Z_y(q)(\widehat{b'}(q) - b'(q))\Big]\Big| = \Big|\mathbf{E}_q\Big[Z'_y(q)(\widehat{b}(q) - b(q))\Big]\Big|$$

In order to simplify our analysis of the error in P_y, we will break up the error into two components: the bias in the estimator $\hat{P}_y$ and the deviation of $\hat{P}_y$ from its mean.

$$\begin{aligned}|\hat{P}_y - P_y| &\leq \Big|\hat{P}_y - \mathbf{E}\Big[\hat{P}_y\Big]\Big| + \Big|\mathbf{E}\Big[\hat{P}_y\Big] - P_y\Big| \\ &= \Big|\mathbf{E}_q\Big[Z'_y(q)(\widehat{b}(q) - \tilde{b}(q))\Big]\Big| + \Big|\mathbf{E}_q\Big[Z'_y(q)(\tilde{b}(q) - b(q))\Big]\Big| \qquad (10)\end{aligned}$$

Here, $\tilde{b}$ is a step function that equals the expectation of the empirical bid function $\widehat{b}$:

$$\tilde{b}(q) = \mathbf{E}\Big[\widehat{b}(q)\Big]$$

The bias of the estimator $\hat{P}$ (i.e. the second term in (10)) is easy to bound. We prove in the appendix that $\tilde{b}(q) - b(q)$ is at most $O(1)/N$ times $\sup_q\{x'(q)\}$. This implies the following lemma.

LEMMA 3.5. *With $\tilde{b}$ defined as above,*

$$\Big|\mathbf{E}\Big[\hat{P}_y\Big] - P_y\Big| = \Big|\mathbf{E}_q\Big[Z'_y(q)(\tilde{b}(q) - b(q))\Big]\Big| = \frac{O(1)}{N}\sup_q\{x'(q)\}\ \sup_q\left\{\frac{y'(q)}{x'(q)}\right\}.$$

We now focus on the first term in (10), namely the integral over the quantile axis of $Z'_y(q)(\widehat{b}(q) - \tilde{b}(q))$. The approach of Section 3.2 does not provide a good upper bound on this quantity, because a counterpart of Lemma B.1 (see Appendix B) fails to hold for Z_y. Instead, we will express the integral as a sum over several independent terms, and show that it is small in expectation.

To this end, we first identify the set of quantiles at which the function $\widehat{b}$ "crosses" the function $\tilde{b}$ from below. This set is defined inductively. Define $i_0 = 0$. Then, inductively, let i_ℓ be the smallest integer strictly greater than $i_{\ell-1}$ such that

$$\widehat{b}\left(\frac{i_\ell - 1}{N}\right) \leq \tilde{b}\left(\frac{i_\ell - 1}{N}\right) \text{ and } \widehat{b}\left(\frac{i_\ell}{N}\right) > \tilde{b}\left(\frac{i_\ell}{N}\right).$$

Let i_{k-1} be the last integer so defined, and let $i_k = N$. Let I denote the set of indices $\{i_0, \cdots, i_k\}$. Let $\mathbf{T}_{i,j}$ denote the following integral:

$$\mathbf{T}_{i,j} = \int_{q=i/N}^{q=j/N} Z'_y(q)(\widehat{b}(q) - \tilde{b}(q))\,\mathrm{d}q$$

Then, our goal is to bound the quantity $\mathbf{E}_{\hat{b}}[|\mathbf{T}_{0,N}|]$ where $\mathbf{T}_{0,N}$ can be written as the sum:

$$\mathbf{T}_{0,N} = \sum_{\ell=0}^{\ell=k-1} \mathbf{T}_{i_\ell, i_{\ell+1}}.$$

We now claim that conditioned on I and the maximum bid error, this is a sum over independent random variables. In the following, let G denote the bid distribution, and $\hat{G}$ the empirical bid distribution.

LEMMA 3.6. *Conditioned on the set of indices I and $\Delta = \sup_q |\hat{G}(b(q)) - G(b(q))|$, over the randomness in the bid sample, the random variables $\mathbf{T}_{i_\ell, i_{\ell+1}}$ are mutually independent.*

Then we apply Chernoff-Hoeffding bounds, coupled with the approach from Section 3.2 to bound each individual $\mathbf{T}_{i_\ell, i_{\ell+1}}$, to obtain a bound on the proability that $\mathbf{E}_{\hat{b}}[|\mathbf{T}_{0,N}| \,|\, I, \Delta]$ exceeds some value $a > 0$.

LEMMA 3.7. *With $I = \{i_0, \cdots, i_k\}$ and $T_{i,j}$ defined as above, for any $a > 0$,*

$$\Pr\left[|\mathbf{T}_{0,N}| \geq a \,|\, I, \Delta\right] \leq \textit{exp}\left(-\frac{a^2}{n(40\Delta C)^2}\right), \quad \textit{where } C = \sup_q\{y'(q)\} \log\max\left\{\sup_{q: y'(q)\geq 1} \frac{x'(q)}{y'(q)}, \sup_q \frac{y'(q)}{x'(q)}\right\}.$$

Combining this lemma with an absolute bound on $|\mathbf{T}_{0,N}|$, and removing the conditioning on I and Δ, we obtain the Theorem 3.4. □

4. APPLICATIONS TO A/B TESTING

We now discuss applications of the inference approach we developed in Section 3.

4.1. Estimating revenues of novel mechanisms

Let us consider the setup described in the introduction where an auction house running auction A would like to determine the revenue of a novel mechanism B. The typical approach for doing so is to run the auction B with some probability $\epsilon > 0$ and A with the remaining probability. Ideally, if in doing so, the auction house obtains ϵN bids in response to the auction B out of a total of N bids, the revenue of B can be estimated within an error bound of

$$\Theta\left(\frac{1}{\sqrt{\epsilon}}\right) \frac{\sup_q\{x'_B(q)\}}{\sqrt{N}} \tag{11}$$

where x_B denotes the allocation rule corresponding to B.

In practice, however, instead of obtaining bids in response to B alone, the seller obtains bids in response to the aggregate mechanism $C = (1-\epsilon)A + \epsilon B$. We can then use (7) to estimate the revenue of B. As a consequence of Theorem 3.4, and noting that for positions auctions with n positions, $x'_C(q) \leq n$ and $x'_B(q)/x'_C(q) \leq 1/\epsilon$ for all quantiles q, we obtain the following error bound

COROLLARY 4.1. *The revenue of a rank based mechanism B can be estimated from N bids of a rank-based mechanism $C = (1-\epsilon)A + \epsilon B$ with absolute error bounded by*

$$O(1)\sqrt{n \log n} \log(n/\epsilon)\ \frac{\sup_q\{x'_B(q)\}}{\sqrt{N}}. \tag{12}$$

Relative to the ideal situation described above, our error bound has a better dependence on ϵ and a worse dependence on n. Note that when ϵ is very small, our error bound (12) may be smaller than the ideal bound in (11). This is not surprising: the ideal bound ignores information that we can learn about the revenue of B from the $(1-\epsilon)N$ bids obtained when B is not run.

When B is a multi-unit auction, we obtain a better error bound which is closer to the ideal bound in (11).

COROLLARY 4.2. *The revenue of the highest-k-bids-win mechanism B can be estimated from N bids of a rank-based mechanism $C = (1-\epsilon)A + \epsilon B$ with absolute error bounded by*

$$O(1)\log(n/\epsilon)\ \frac{\sup_q\{x'_B(q)\}}{\sqrt{N}}. \tag{13}$$

Finally, we note that we can simultaneously infer the revenues of *all* positions auctions by learning the revenues of the multi-unit auctions.

COROLLARY 4.3. *N bids from a mechanism C with $x_C = (1-\epsilon)x_A + \sum_{k=1}^{n} \frac{\epsilon}{n} x_k$ can be used to simultaneously estimate* all *of the multi-unit revenues, and consequently all position auction revenues with absolute error bounded by*

$$O(1)\frac{n\log(n/\epsilon)}{\sqrt{N}}.$$

4.2. Comparing revenues

We have considered the case where the empirical task was to recover the revenues for one mechanism (y) using the sample of bids responding to another mechanism (x). In many practical situations the empirical task is simply the verification of whether the revenue from a given mechanism is higher than the revenue from another mechanism. Or, equivalently, the task could be to verify whether one mechanism provides revenue which is a certain percentage above that of another mechanism. We now demonstrate that this is a much easier empirical task in terms of accuracy than the task of inferring the revenue.

Suppose that we want to compare the revenues of mechanisms B_1 and B_2 by mixing them in to an incumbent mechanism A, and running the composite mechanism $M = \epsilon B_1 + \epsilon B_2 + (1-2\epsilon)A$. Specifically, we would like to determine whether $P_{B_1} > \alpha P_{B_2}$ for some $\alpha > 0$. Consider a binary classifier $\widehat{\gamma}$ which is equal to 1 when $P_{B_1} > \alpha P_{B_2}$ and 0 otherwise. Let $\gamma = \mathbf{1}\{P_{B_1} - \alpha\, P_{B_2} > 0\}$ be the corresponding "ideal" classifier for the case where the distribution of bids from mechanism M is known precisely. To evaluate the accuracy of the classifier, we need to evaluate the probability $\Pr(\widehat{\gamma} = 1|\gamma = 0)$, and likewise, $\Pr(\widehat{\gamma} = 0|\gamma = 1)$. The classifier will give the wrong output if the sampling noise in estimating $\hat{P}_{B_1} - \alpha\, \hat{P}_{B_2}$ is greater than $|P_{B_1} - \alpha\, P_{B_2}|$.

Our main result of this section says that keeping α, the difference $|P_{B_1} - \alpha\, P_{B_2}|$, and the number of positions n constant, the probability of incorrect output decreases exponentially with the number of bids N.

THEOREM 4.4. *Suppose that N bids from a mechanism $M = \epsilon B_1 + \epsilon B_2 + (1-2\epsilon)A$ for arbitrary rank-based mechanisms B_1, B_2, and A, are used to estimate the classifier $\gamma = \mathbf{1}\{P_{B_1} - \alpha\, P_{B_2} > 0\}$ that establishes whether the revenue of mechanism B_1 exceeds α times the revenue of mechanism B_2. Then the error rate of the binary classifier is bounded from above by*

$$\exp\left(-O\left(\frac{Na^2}{\alpha^2 n^3 \log(n/\epsilon)}\right)\right),$$

where $a = |P_{B_1} - \alpha\, P_{B_2}|$. In other words, once the number of samples is polynomially large in n, the error rate decreases exponentially with the number of samples.

A complete proof of the theorem appears in the appendix. Here we provide a proof sketch.

PROOF SKETCH: We need to bound the probability that the error in estimating $\hat{P}_{B_1} - \alpha\hat{P}_{B_2}$ is greater than $|P_{B_1} - \alpha\, P_{B_2}|$. This error can in turn be decomposed into the error in estimating P_{B_1} and that in estimating P_{B_2}. Denote $a = |P_{B_1} - \alpha\, P_{B_2}| > 0$. Then,

$$\Pr\left(|(\widehat{P}_{B_1} - \alpha\, \widehat{P}_{B_2}) - (P_{B_1} - \alpha\, P_{B_2})| > a\right) \leq \Pr\left(|\widehat{P}_{B_1} - P_{B_1}| > a/2\right) + \Pr\left(|\widehat{P}_{B_2} - P_{B_2}| > a/2\alpha\right).$$

Let x denote the allocation rule of the mechanism M that we are running, and let b be the corresponding bid function. Now, recall that for

$$\Delta = \sup_q |\widehat{G}(b(q)) - G(b(q))|, \text{ and } C = \sup_q\{x'_{B_1}(q)\} \log \max\left\{\sup_{q:x'_{B_1}(q)\geq 1} \frac{x'(q)}{x'_{B_1}(q)}, \sup_q \frac{x'_{B_1}(q)}{x'(q)}\right\},$$

Equation (10), Lemma 3.5, and Lemma 3.7 together imply that (conditional on Δ)

$$\Pr\left(|\widehat{P}_{B_1} - P_{B_1}| > a/2\right) \leq 2\exp\left(-\frac{1}{n(40\,\Delta\, C)^2}\left(\frac{a}{2} - O\left(\frac{n}{\epsilon N}\right)\right)^2\right).$$

Now recall that $C < n\log(n/\epsilon)$, and $\Delta <$ constant$/\sqrt{N}$ with high probability (Lemma 2.2). As a result, we establish that for $a = \omega(n/\epsilon N)$,

$$\Pr\left(|\widehat{P}_{B_1} - P_{B_1}| > a/2\right) \leq \exp\left(-O\left(\frac{Na^2}{n^3\log(n/\epsilon)}\right)\right).$$

Likewise,

$$\Pr\left(|\widehat{P}_{B_2} - P_{B_2}| > a/2\alpha\right) \leq \exp\left(-O\left(\frac{Na^2}{\alpha^2 n^3\log(n/\epsilon)}\right)\right).$$

□

We obtain a similar error bound when our goal is to estimate which of r different novel mechanisms obtains the most revenue, for any $r > 1$:

COROLLARY 4.5. *Suppose that our goal is to determine which of r position auctions, $B_1, B_2, \cdots, B_r$, obtains the most revenue while running incumbent mechanism A, by running each of the novel mechanisms with probability ϵ/r. Then the error probability of the corresponding classifier constructed using N bids from composite mechanism $M = \sum_{i=1}^{r} \epsilon/r B_i + (1-\epsilon)A$ is bounded from above by*

$$r \exp\left(-O\left(\frac{Na^2}{n^3 \log(rn/\epsilon)}\right)\right),$$

where a is the absolute difference between the revenue obtained by the best two of the r mechanisms.

5. INFERENCE METHODOLOGY AND ERROR BOUNDS FOR FIRST-PRICE AUCTIONS

While most of our analysis so far has focused on the case of all-pay auctions, our methodology and results extend as-is to first-price auctions as well. Here we sketch the differences between the two cases.

Recall that in a first-price auction, we can obtain the value distribution from the bid distribution as follows: $v(q) = b(q) + x(q)b'(q)/x'(q)$. Substituting this into the expression for P_y we get:

$$P_y = \mathbf{E}_q\left[(1-q)y'(q)b(q) + \frac{(1-q)y'(q)x(q)b'(q)}{x'(q)}\right] = \mathbf{E}_q[(1-q)y'(q)b(q) + Z_y(q)x(q)b'(q)]$$

where, as before, $Z_y(q) = \frac{(1-q)y'(q)}{x'(q)}$.

Integrating the second expression by parts, we get

$$\begin{aligned}\int_0^1 Z_y(q)x(q)b'(q)\,dq &= Z_y(q)x(q)b(q)|_0^1 - \int_0^1 (Z_y'(q)x(q) + Z_y(q)x'(q))b(q)\,dq \\ &= -\int_0^1 Z_y'(q)x(q)b(q)\,dq - \int_0^1 (1-q)y'(q)b(q)\,dq\end{aligned}$$

When we put this back in the expression for P_y two of the terms cancel, and we get the following lemma.

LEMMA 5.1. *The per-agent revenue of an auction with allocation rule y can be written as a linear combination of the bids in a first-pay auction:*

$$P_y = \mathbf{E}_q\big[-x(q)Z_y'(q)b(q)\big]$$

where $Z_y(q) = (1-q)\frac{y'(q)}{x'(q)}$ and $x(q)$ are known precisely.

As in the case of the all-pay auction format, we can write the error in P_y as:

$$\begin{aligned}|\hat{P}_y - P_y| &= \mathbf{E}_q\Big[\Big|-x(q)Z_y'(q)(\widehat{b}(q) - b(q))\Big|\Big] \\ &\le E\left[\frac{(\log(1+Z_y(q)))^\alpha}{Z_y(q)}|Z_y'(q)|\right] \sup_q \left|x(q)\frac{Z_y(q)}{(\log(1+Z_y(q)))^\alpha}(\hat{b}(q) - b(q))\right|\end{aligned}$$

With an appropriate choice of α we obtain the following theorem.

THEOREM 5.2. *The expected absolute error in estimating the revenue of a position auction with allocation rule y using N samples from the bid distribution for a first-pay position auction with allocation rule x is bounded as below; Here n is the number of positions in the two position auctions.*

$$\begin{aligned}\mathbf{E}_{\hat{b}}\Big[|\hat{P}_y - P_y|\Big] \le \frac{20}{\sqrt{N}}\sqrt{n\log n}\ \sup_q\{y'(q)\}\ \log\max\left\{\sup_{q:y'(q)\ge 1}\frac{x'(q)}{y'(q)}, \sup_q\frac{y'(q)}{x'(q)}\right\} \\ + \frac{O(1)}{N}\sup_q\{x'(q)\}\ \sup_q\left\{\frac{y'(q)}{x'(q)}\right\}.\end{aligned}$$

When y is the highest-k-bids-win allocation rule, the error improves to:

$$\mathbf{E}_{\hat{b}}\left[|\hat{P}_k - P_k|\right] \leq \frac{20}{\sqrt{N}} \sup_q\{x'_k(q)\} \log\max\left\{\sup_{q:x'_k(q)\geq 1}\frac{x'(q)}{x'_k(q)}, \sup_q \frac{x'_k(q)}{x'(q)}\right\}.$$

Because the error bounds in Theorem 5.2 are identical to those in Theorems 3.3 and 3.4, Corollaries 4.1, 4.2, 4.3, and Theorem 4.4 continue to hold when bids are drawn from a first-price auction.

6. SIMULATION EVIDENCE

We now present some simulation evidence to support our theoretical results. Our focus in this exercise is the inference for P_k for a particular k. Specifically, for different values of n and $\epsilon = 0.01$, we consider the problem of inferring the revenue P_k from the bids of a mechanism with allocation rule $x = \epsilon x^{(k:n)} + \sum_{l\neq k} w_l x^{(l:n)}$ with the weights w_l picked uniformly at random so as to sum to $1-\epsilon$. Values are generated from the beta distribution with parameters $\alpha = \beta = 2$. This distribution of values is supported on $[0,1]$, it is unimodal with the mode and the mean at $1/2$ and it is symmetric about the mean. We transform this distribution by shrinking its support to ensure that the optimal revenue from this value distribution is equal to 1.

Specifics of the simulation. The simulation is designed in the following way. We consider a grid over quantiles with the cell size Δ such that each grid point $p \in \{0, \cdots, 1/\Delta\}$ corresponds to the quantile $q_{(p)} = \Delta\, p$. We will first describe how we generate the allocation functions x and $x^{(k:n)}$, as well as the multipliers $Z_{k,i}$, in order to infer P_k via Equation (7).

For a given number of players n we compute the allocation functions for each k-unit auction with $k = 1, \ldots, n$ on the grid yielding $x^{(k:n)}_{(p)} = x^{(k:n)}(q_{(p)})$ for each grid point. Likewise, we compute

$$x_{(p)} = \epsilon x^{(k:n)}_{(p)} + (1-\epsilon)\sum_{l\neq k} w_l x^{(l:n)}_{(p)}.$$

The derivatives of the allocation rules $x(\cdot)$ and $x^{(k:n)}(\cdot)$ are computed from the adjacent grid points as

$$x^{(k:n)\prime}_p = \frac{x^{(k:n)}_{(p+1)} - x^{(k:n)}_{(p-1)}}{2\Delta} \text{ and } x'_p = \frac{x_{(p+1)} - x_{(p-1)}}{2\Delta}.$$

Function $Z_k(\cdot)$ and its first derivative are approximated on the same grid as

$$Z_{k(p)} = (1-\Delta\, p)\frac{x^{(k:n)\prime}_{(p)}}{x'_{(p)}} \text{ and } Z'_{k(p)} = \frac{Z_{k(p+1)} - Z_{k(p-1)}}{2\Delta}.$$

Next we describe how we generate bids from the symmetric equilibrium for x. We first compute the quantile function of the value distribution on the grid as $v_{(p)} = v(\Delta\, p)$. Then for the all-pay auction with the allocation rule $x(\cdot)$ we compute the approximated quantile function of the bid distribution $b(\cdot)$ as

$$b_{(p)} = \sum_{l\leq p} v_{(l)}(x_{(l)} - x_{(l-1)}) \approx \int_0^{\Delta\, p} v(q)x'(q)\, dq.$$

Using the same grid summation technique we compute the "true" revenue from the k-unit auction of interest as

$$P_k = \sum (1-\Delta\, p) v_{(p)}(x^{(k:n)}_{(p)} - x^{(k:n)}_{(p-1)}).$$

In each simulation we generate a sample of $N = nT$ bids, where T denotes the number of auctions run (each with n bidders). To perform simulations, at each simulation round, we generate a sample of N bids by drawing uniformly a vector of $N = T \times n$ integers s between 1 and $[1/\Delta]$ and for each such integer s, the corresponding simulated draw of the bid is $b_{(s)}$. In each N-sample, $\{b_i\}_{i=1}^N$, we order the bids and then compute the estimator for P_k as

$$\widehat{P}_k = -\frac{1}{N}\sum_{i=1}^N Z'_{k(i)} b^{(i:N)}.$$

Table I. This table presents the mean absolute deviation for $\widehat{P}_{\lfloor n/2 \rfloor}$ across Monte-Carlo simulations (overall 1000 simulations for each choice of n and T) from $P_{\lfloor n/2 \rfloor}$ computed from the true distribution of values via grid approximation, multiplied by $\sqrt{n \times T}$.

	$T =$					
$n =$	10	20	30	40	50	60
10	0.2955	0.2869	0.2862	0.3108	0.2798	0.2880
15	0.3389	0.3422	0.3423	0.3378	0.3375	0.3377
20	0.4259	0.3967	0.4087	0.4084	0.4211	0.4139
25	0.3547	0.3467	0.3532	0.3471	0.3577	0.3584
30	0.4366	0.4692	0.4438	0.4271	0.4540	0.4225
35	0.4323	0.4317	0.417	0.4377	0.4386	0.4200

Table II. This table presents the mean absolute deviation for $\widehat{P}_{\lfloor n/2 \rfloor}$ across Monte-Carlo simulations (overall 1000 simulations for each choice of n and T) from $P_{\lfloor n/2 \rfloor}$ computed from the true distribution of values via grid approximation.

	$T =$					
$n =$	10	20	30	40	50	60
10	0.0296	0.0205	0.0164	0.0147	0.0139	0.013
15	0.0249	0.0177	0.0141	0.0127	0.0108	0.0101
20	0.0246	0.0178	0.0149	0.0124	0.0108	0.0101
25	0.0231	0.016	0.0124	0.0111	0.0101	0.0096
30	0.0213	0.0148	0.0127	0.0108	0.0093	0.0086
35	0.0197	0.0134	0.0115	0.0103	0.009	0.0081

For each choice of the number of players n and the number of auctions T we re-sample the bids from the bid distribution 1000 times and then compute the average mean absolute deviation of the estimated $\widehat{P}_k$ from the computed P_k across simulations. The simulation results for various choices of the number of players and the number of auctions are presented in Tables I and II.

Discussion of results. We choose $k = \lfloor n/2 \rfloor$. Recall that with these choices, Theorem 3.3 tells us that the absolute error in estimating P_k should be bounded by $O(\sqrt{n}\log(n/\epsilon)/\sqrt{N})$.

Table I reports the averaged absolute error in P_k normalized by the factor $\sqrt{n \times T} = \sqrt{N}$. Per our prediction, the numbers reported in this table should depend on n alone, and should increase with n roughly as $\sqrt{n}\log n$. We note that for a fixed n as T is varied, the errors in Table I roughly stay constant; This validates our prediction that $\sqrt{N}\,\mathbf{E}\left[|\hat{P}_k - P_k|\right]$ is a function of n alone. In other words, this supports our finding of the $\sqrt{N}$ convergence rate for the estimator $\widehat{P}_k$. Next, note that as n increases, the errors in Table I increase but at a very slow rate. Rather than increasing in proportion to $\sqrt{n}$, the errors appear to increase in proportion to $\log n$. While further simulations are needed to validate this observation, this indicates that the true estimation error is even better than our bound indicates.

Table II shows the unnormalized averaged absolute deviation of the estimated $\widehat{P}_{\lfloor n/2 \rfloor}$ from $P_{\lfloor n/2 \rfloor}$. The overall absolute error is below 3% for all choices of numbers of auctions and the numbers of players. In cases where the number of players is large, the error falls to below 1% of the optimal revenue with a sample of over 40 auctions.

REFERENCES

BABAIOFF, M., SHARMA, Y., AND SLIVKINS, A. 2009. Characterizing truthful multi-armed bandit mechanisms: Extended abstract. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. EC '09. ACM, New York, NY, USA, 79–88.

BAJARI, P., HONG, H., AND NEKIPELOV, D. 2013. Game theory and econometrics: A survey of some recent research. *Advances in Economics and Econometrics*.

BROWN, J. AND MORGAN, J. 2009. How much is a dollar worth? tipping versus equilibrium coexistence on competing online auction sites. *Journal of Political Economy 117,* 4, 668–700.

CHAWLA, S., HARTLINE, J., AND NEKIPELOV, D. 2014. Mechanism design for data science. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 711–712.

CHAWLA, S. AND HARTLINE, J. D. 2013. Auctions with unique equilibria. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. ACM, New York, NY, USA, 181–196.
DEVANUR, N. R. AND KAKADE, S. M. 2009. The price of truthfulness for pay-per-click auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. EC '09. ACM, New York, NY, USA, 99–106.
EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review 97,* 1, 242–259.
GATTI, N., LAZARIC, A., AND TROVÒ, F. 2012. A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. ACM, New York, NY, USA, 605–622.
GOMES, R. AND SWEENEY, K. 2009. Bayes-nash equilibria of the generalized second price auction. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 107–108.
GUERRE, E., PERRIGNE, I., AND VUONG, Q. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica 68,* 3, 525–574.
HARTLINE, J. D. 2013. Bayesian mechanism design. *Foundations and Trends in Theoretical Computer Science 8,* 3, 143–263.
MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research 6*, 58–73.
OSTROVSKY, M. AND SCHWARZ, M. 2011. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM Conference on Electronic Commerce*. EC '11. ACM, New York, NY, USA, 59–60.
PAARSCH, H. J. AND HONG, H. 2006. *An introduction to the structural econometrics of auction data*. Vol. 1. The MIT Press.
REILEY, D. H. 2006. Field experiments on the effects of reserve prices in auctions: more magic on the internet. *The RAND Journal of Economics 37,* 1, 195–211.
VARIAN, H. 2007. Position auctions. *International Journal of Industrial Organization 25,* 6, 1163–1178.

A. INFERENCE FOR AUCTIONS

The distribution of values, which is unobserved, can be inferred from the distribution of bids, which is observed. Once the value distribution is inferred, other properties of the value distribution such as its corresponding revenue curve, which is fundamental for optimizing revenue, can be obtained. In this section we describe the basic premise of the inference assuming that the distribution of bids known exactly.

The key idea behind the inference of the value distribution from the bid distribution is that the strategy which maps values to bids is a best response, by equation (1), to the distribution of bids. As the distribution of bids is observed, and given suitable continuity assumptions, this best response function can be inverted.

The value distribution can be equivalently specified by distribution function $F(\cdot)$ or value function $v(\cdot)$; the bid distribution can similarly be specified by the bid function $b(\cdot)$. For rank-based auctions (as considered by this paper) the allocation rule $x(\cdot)$ in quantile space is known precisely (i.e. it does not depend on the value function $v(\cdot)$). Assume these functions are monotone, continuously differentiable, and invertible.

Inference for first-price auctions. Consider a first-price rank-based auction with a symmetric bid function $b(q)$ and allocation rule $x(q)$ in BNE. To invert the bid function we solve for the bid that the agent with any quantile would make. Continuity of this bid function implies that its inverse is well defined. Applying this inverse to the bid distribution gives the value distribution.

The utility of an agent with quantile q as a function of his bid z is

$$u(q, z) = (v(q) - z)\, x(b^{-1}(z)). \tag{14}$$

Differentiating with respect to z we get:

$$\tfrac{d}{dz} u(q, z) = -x(b^{-1}(z)) + \big(v(q) - z\big)\, x'(b^{-1}(z))\, \tfrac{d}{dz} b^{-1}(z).$$

Here x' is the derivative of x with respect to the quantile q. Because $b(\cdot)$ is in BNE, the derivative $\frac{d}{dz}u(q,z)$ is 0 at $z = b(q)$. Rarranging, we obtain:

$$v(q) = b(q) + \frac{x(q)\, b'(q)}{x'(q)}$$

Inference for all-pay auctions. We repeat the calculation above for rank-based all-pay auctions; the starting equation (14) is replaced with the analogous equation for all-pay auctions:

$$u(q,z) = v(q)\, x(b^{-1}(z)) - z. \tag{15}$$

Differentiating with respect to z we obtain:

$$\frac{d}{dz}u(q,z) = v(q)\, x'(b^{-1}(z))\, \frac{d}{dz} b^{-1}(z) - 1,$$

Again the first-order condition of BNE implies that this expression is zero at $z = b(q)$; therefore,

$$v(q) = \frac{b'(q)}{x'(q)}.$$

Known and observed quantities. Recall that the functions $x(q)$ and $x'(q)$ are known precisely: these are determined by the rank-based auction definition. The functions $b(q)$ and $b'(q)$ are observed. The calculations above hold in the limit as the number of samples from the bid distribution goes to infinity, at which point these obserations are precise.

B. MISSING PROOFS

PROOF OF THEOREM 3.3. We will make use of the following lemma from Chawla et al. [2014].

LEMMA B.1 (CHAWLA ET AL. [2014]). *Let x_k denote the allocation function of the k-highest-bids-win auction and x be any convex combination over the allocation functions of the multi-unit auctions. Then the function $Z_k(q) = (1-q)\frac{x_k'(q)}{x'(q)}$ achieves a single local maximum for $q \in [0,1]$.*

Recall that for $\alpha > 0$ we can write[4]

$$|\hat{P}_k - P_k| \le E\left[\frac{(\log(1+Z_k(q)))^\alpha}{Z_k(q)}|Z_k'(q)|\right] \sup_q \left|\frac{Z_k(q)}{(\log(1+Z_k(q)))^\alpha}(\hat{b}(q) - b(q))\right|$$

We start by considering the first term. Lemma B.1 shows that $Z_k'(\cdot)$ changes sign only once. Consider the region where the sign of $Z_k'(\cdot)$ is constant and make the change of variable $t = Z_k(q)$. Recall that $Z_k^* = \sup_q Z_k(q)$ and we note that $\inf_q Z_k(q) \ge 0$. Then we can evaluate the first term as

$$E\left[\frac{(\log(1+Z_k(q)))^\alpha}{Z_k(q)}|Z_k'(q)|\right] \le 2\int_0^{Z_k^*} \frac{(\log(1+t))^\alpha}{t}\, dt$$

Note that for any $t > 0$, $\log(1+t) \le t$. Thus,

$$\int_0^\delta \frac{(\log(1+t))^\alpha}{t}\, dt < \frac{\delta^\alpha}{\alpha}$$

Now split the integral into two pieces as

$$\int_0^{Z_k^*} \frac{(\log(1+t))^\alpha}{t}\, dt = \int_0^1 \frac{(\log(1+t))^\alpha}{t}\, dt + \int_1^{Z_k^*} \frac{(\log(1+t))^\alpha}{t}\, dt$$

We just proved that the first piece is at most $1/\alpha$. Now we upper bound the second piece and consider the integrand at $t \ge 1$. First, note that

$$(\log(1+t))^\alpha = \left(\log t + \log(1+\frac{1}{t})\right)^\alpha \le \left(\log t + \frac{1}{t}\right)^\alpha \le (\log t + 1)^\alpha.$$

[4]In this entire proof the logarithms are natural logarithms.

Thus, the integral behaves as

$$\int_1^{Z_k^*} \frac{(\log(1+t))^\alpha}{t}\,dt \le \int_1^{Z_k^*} \frac{(\log(t)+1)^\alpha}{t}\,dt = \frac{1}{1+\alpha}(\log\, Z_k^*+1)^{1+\alpha}.$$

Thus, we just showed that

$$E\left[\frac{(\log(1+Z_k(q)))^\alpha}{Z_k(q)}|Z_k'(q)|\right] \le \frac{2}{\alpha} + \frac{2}{1+\alpha}(\log\, Z_k^*+1)^{1+\alpha},$$

which is at most $2(1+e)/\alpha$ for $\alpha < 1/\log\, Z_k^*$.

Now consider the term

$$\sup_q \left|\frac{Z_k(q)}{(\log(1+Z_k(q)))^\alpha}(\hat{b}(q)-b(q))\right|$$

Note that $\log(1+t) \ge \min\{1,t\}/2$. So the first term can be bounded from above as

$$\frac{Z_k(q)}{(\log(1+Z_k(q)))^\alpha} \le 2^\alpha \max\left\{Z_k(q),\, (Z_k(q))^{1-\alpha}\right\}$$

The second term behaves as

$$(\hat{b}(q)-b(q)) = -x'(q)v(q)(\hat{G}(b(q))-G(b(q)))$$

where G and $\hat{G}$ are the real and empirical bid distributions, respectively. We recall that $\mathbf{E}\left[\sup_q \left|\hat{G}(b(q))-G(b(q))\right|\right] \le \frac{1}{2\sqrt{N}}$. Thus

$$\begin{aligned}\mathbf{E}\left[\sup_q \left|\frac{Z_k(q)}{(\log(1+Z_k(q)))^\alpha}(\hat{b}(q)-b(q))\right|\right] &\le 2^\alpha \sup_q \left(\max\left\{x_k'(q),\, (x_k'(q))^{1-\alpha}(x'(q))^\alpha\right\}\right)\frac{1}{2\sqrt{N}} \\ &\le 2^\alpha \sup_q (x_k'(q)) \left(\underbrace{\max\left(1,\, \sup_{q:x_k'(q)\ge 1}\frac{x'(q)}{x_k'(q)}\right)}_{=:A}\right)^\alpha \frac{1}{2\sqrt{N}}\end{aligned}$$

Now we combine the two evaluations together and pick $\alpha = \min\{1/\log A, 1/\log Z_k^*\}$, with A defined as above, to obtain

$$\begin{aligned}\mathbf{E}\left[|\hat{P}_k - P_k|\right] &\le \frac{2(1+e)}{\alpha}2^\alpha A^\alpha \frac{1}{2\sqrt{N}}\sup_q (x_k'(q)) \\ &\le \frac{20}{\sqrt{N}}\ \sup_q\{x_k'(q)\}\ \max\left\{\log A, \log\sup_q\left\{\frac{x_k'(q)}{x'(q)}\right\}\right\}\end{aligned}$$

□

PROOF OF LEMMA 3.5. We can write the function $\tilde{b}(i/N)$ as

$$\begin{aligned}\tilde{b}\left(\frac{i}{N}\right) &= \frac{N!}{(i-1)!(N-i)!}\int G(t)^{i-1}(1-G(t))^{N-i}g(t)t\,dt \\ &= \frac{N!}{(i-1)!(N-i)!}\int t^{i-1}(1-t)^{N-i}b(t)\,dt\end{aligned}$$

Note that

$$\frac{N!}{(i-1)!(N-i)!}t^{i-1}(1-t)^{N-i}$$

is the density of the beta-distribution with parameters $\alpha = i$ and $\beta = N-i+1$. Denote this density $f(t;\alpha,\beta)$. Then we can write

$$\tilde{b}\left(\frac{i}{N}\right) = \int_0^1 b(t)f(t;\alpha,\beta)\,dt.$$

Now let $q \in [i/N,\ (i+1)/N]$, and consider an expansion of $b(t)$ at q such that

$$b(t) = b(q) + b'(q)(t-q) + O((t-q)^2)$$

Now we substitute this expansion into the formula for $\tilde{b}(\cdot)$ above to get

$$\tilde{b}\left(\frac{i}{N}\right) = b(q) + b'(q)\int_0^1 (t-q) f(t;\alpha,\beta)\,dt + O(\int_0^1 (t-q)^2 f(t;\alpha,\beta)\,dt)$$

The mean of the beta distribution is $\alpha/(\alpha+\beta)$ and the variance is $\alpha\beta/((\alpha+\beta)^2(\alpha+\beta+1))$. This means that

$$\tilde{b}\left(\frac{i}{N}\right) - b(q) = b'(q)\left(\frac{i}{N+1} - q\right) + O\left(\frac{1}{N^2}\right).$$

Thus

$$\sup_{q\in[i/N,(i+1)/N]} \left|\tilde{b}\left(\frac{i}{N}\right) - b(q)\right| \le \sup_q b'(q)\frac{2}{N} + O\left(\frac{1}{N^2}\right).$$

Therefore, the expectation $\left|\hat{P}_y - \mathbf{E}\left[\hat{P}_y\right]\right|$ is at most $O(1)/N\ \sup_q\{x'(q)\}\ \sup_q Z_y(q)$. □

PROOF OF LEMMA 3.6. Fix I and ℓ, and note that the function $\tilde{b}$ is fixed (that is, it does not depend on the empirical bid sample). Then, the sum $\mathbf{T}_{i_\ell,i_{\ell+1}}$ depends only on the empirical bid values $\widehat{b}(q)$ for quantiles in the interval $[i_\ell/N, i_{\ell+1}/N)$. By the definition of I, we know that the smallest i_ℓ bids in the sample are all smaller than $\tilde{b}((i_\ell-1)/N) \le \tilde{b}(i_\ell/N)$, and the largest $N - i_{\ell+1}$ bids in the sample are all larger than $\tilde{b}(i_{\ell+1}/N) \ge \tilde{b}((i_{\ell+1}-1)/N)$. On the other hand, the empirical bids $\widehat{b}(q)$ for $q \in [i_\ell/N, i_{\ell+1}/N)$ lie within $[\tilde{b}(i_\ell/N), \tilde{b}((i_{\ell+1}-1)/N)]$. Therefore, conditioned on i_ℓ and $i_{\ell+1}$, the latter set of empirical bids is independent of the former set of empirical bids. □

PROOF OF LEMMA 3.7 AND THEOREM 3.4. We will use Chernoff-Hoeffding bounds to bound the expectation of $\mathbf{T}_{0,N}$ over the bid sample, conditioned on I and $\Delta = \sup_{q\in[0,1]} |\hat{G}(b(q)) - G(b(q))|$. We first note that $\mathbf{T}_{0,N}$ has mean zero because for any integer $i \in [0,N]$, $\mathbf{E}_{\text{samples}}\left[\widehat{b}(i/N)\right] = \tilde{b}(i/N)$.

Next we note that the $\mathbf{T}_{i,j}$'s are bounded random variables. Specifically, let Q be an interval of quantiles over which the difference $\widehat{b}(q) - \tilde{b}(q)$ does not change sign. Then, following the proof of Theorem 3.3, we can bound

$$\begin{aligned}|\mathbf{T}_Q| &= \left|\int_Q Z_y'(q)(\widehat{b}(q) - \tilde{b}(q))\,\mathrm{d}q\right| \\ &\le 20\Delta \underbrace{\sup_q\{y'(q)\}\ \log\max\left\{\sup_{q:y'(q)\ge 1}\frac{x'(q)}{y'(q)}, \sup_q \frac{y'(q)}{x'(q)}\right\}}_{=:\,C}.\end{aligned}$$

Likewise, over an interval Q where Z_y' does not change sign, we again get $|\mathbf{T}_Q| \le 20\Delta C$ with C defined as above. Moreover, for an interval Q over which Z_y' changes sign at most t times, we have

$$\int_Q |Z_y'(q)(\widehat{b}(q) - \tilde{b}(q))|\,\mathrm{d}q \le t \cdot 20\Delta C.$$

Finally, noting that Z_y is a weighted sum over the n functions Z_k defined for the k-unit auctions, and that by Lemma B.1 each Z_k has a unique maximum, we note that Z_y' changes sign at most $2n$ times.

We now apply Chernoff-Hoeffding bounds to bound the probability that the sum $\sum_{\ell=0}^{\ell=k-1} \mathbf{T}_{i_\ell,i_{\ell+1}}$ exceeds some constant a. With τ_ℓ denoting the upper bound on $|\mathbf{T}_{i_\ell,i_{\ell+1}}|$, this probability is at most

$$\exp\left(-\frac{a^2}{\sum_\ell \tau_\ell^2}\right).$$

By our observations above, for all ℓ, $\tau_\ell \le 40\Delta C$, and $\sum_\ell \tau_\ell \le \int_0^1 |Z'_y(q)(\widehat{b}(q) - \tilde{b}(q))|\,\mathrm{d}q \le 40n\Delta C$. Therefore, $\sum_\ell \tau_\ell^2 \le n(40\Delta C)^2$. We can now choose $a = \sqrt{n \log n}\, 40\Delta C$ to make the above probability at most $1/n$.

Putting everything together, we get that conditioned on I and Δ, the expected value of $|\mathbf{T}_{0,N}|$ over the bid sample is at most $a + 1/n \cdot 40n\Delta C = O(1)\sqrt{n \log n}\Delta C$. Since this bound is independent of I, the same bound holds when we remove the conditioning on I. The theorem now follows by plugging in the expected value of Δ from Lemma 2.2. □

PROOF OF THEOREM 4.4. Denote $a = -(P_y - \alpha\, P_x) > 0$. Then we note that

$$P\left(|(\widehat{P}_y - \alpha\,\widehat{P}_x) - (P_y - \alpha\, P_x)| > a)\right) \le P\left(|\widehat{P}_y - P_y| > a\right) + P\left(|\widehat{P}_x - P_x| > a/\alpha\right)$$

Now, recall that for $C_1 = \sup_q |\widehat{G}(b(q)) - G(b(q))|$ and $C_2 = \sup_q\{y'(q)\} \max\left\{\log \sup_q x'(q),\, \log \sup_q \left\{\frac{y'(q)}{x'(q)}\right\}\right\}$ we established the result that (conditional on C_1)

$$P\left(|\widehat{P}_y - P_y| > a\right) \le 2\exp\left(-\frac{a^2}{n(40\, C_1\, C_2)^2}\right).$$

Now recall that $C_1 = \epsilon/\sqrt{N}$, where random variable ϵ has a bounded variance. As a result, we establish that

$$P\left(|\widehat{P}_y - P_y| > a\right) \le \exp\left(-O\left(\frac{N}{n}\right)\right).$$

We note that $\widehat{P}_x = \frac{1}{N}\sum_{i=1}^N b_i$, which is the sample average of bids. Thus, to evaluate the second probability we can apply the Chernoff-Hoeffding bound directly to establish that

$$P\left(|\widehat{P}_x - P_x| > a/\alpha\right) \le 2\exp\left(-\frac{2Na^2}{\alpha^2}\right) = \exp\left(-O(N)\right)$$

□