Mechanism Design for Data Science

Shuchi Chawla[*] Jason Hartline[†] Denis Nekipelov[‡]

June 10, 2014

Abstract

Good economic mechanisms depend on the preferences of participants in the mechanism. For example, the revenue-optimal auction for selling an item is parameterized by a reserve price, and the appropriate reserve price depends on how much the bidders are willing to pay. A mechanism designer can potentially learn about the participants' preferences by observing historical data from the mechanism; the designer could then update the mechanism in response to learned preferences to improve its performance. The challenge of such an approach is that the data corresponds to the actions of the participants and not their preferences. Preferences can potentially be inferred from actions but the degree of inference possible depends on the mechanism. In the optimal auction example, it is impossible to learn anything about preferences of bidders who are not willing to pay the reserve price. These bidders will not cast bids in the auction and, from historical bid data, the auctioneer could never learn that lowering the reserve price would give a higher revenue (even if it would). To address this impossibility, the auctioneer could sacrifice revenue optimality in the initial auction to obtain better inference properties so that the auction's parameters can be adapted to changing preferences in the future. This paper develops the theory for optimal mechanism design subject to good inferability.

1 Introduction

The promise of data science is that if data from a system can be recorded and understood then this understanding can potentially be utilized to improve the system. Behavioral and economic data, however, is different from scientific data in that it is subjective to the system. Behavior changes with system changes, and to predict behavior for any given system change or to optimize over system changes, the model that generates the behavior must be inferred from the behavior. The ease with which this inference can be performed generally also depends on the system. Trivially, a system that ignores behavior does not admit any inference of a behavior generating model that can be used to predict behavior in a system that is responsive to behavior. To realize the promise of data science in economic systems, a theory for the design of such systems must also incorporate the desired inference properties.

Consider as an example the revenue-maximizing auctioneer. If the auctioneer has knowledge of the distribution of bidder values then she can run the first-price auction with a reserve price that is tuned to the distribution. Under some mild distributional assumptions, with the appropriate

[*]University of Wisconsin - Madison
[†]Northwestern University
[‡]University of California - Berkeley and University of Virginia

reserve price the first-price auction is revenue optimal (Myerson, 1981). Notice that the historical bid data for the first-price auction with a reserve price will never have bids for bidders whose values are below the reserve. Therefore, there is no data analysis that the auctioneer can perform that will enable properties of the distribution of bidder values below the reserve price to be inferred. It could be, nonetheless, that over time the population of potential bidders evolves and the optimal reserve price lowers. This change could go completely unnoticed in the auctioneer's data. The two main tools for optimizing revenue in an auction are reserve prices (as above) and ironing.[1] Both of these tools cause pooling behavior (i.e., bidders with distinct values take the same action) and economic inference cannot thereafter differentiate these pooled bidders. In order to maintain the distributional knowledge necessary to be able to run a good auction in the long term, the auctioneer must sacrifice the short-term revenue by running a non-revenue-optimal auction.

Economic inference in auctions is based on a very straightforward premise. We, the analyst, would like to infer the values of bidders from their bids. This is possible because a bidder's bid is in best response to the bid distribution and the bid distribution is observed in the data. Given any value for a bidder, and the empirical distribution of bids, we can solve for the bidder's utility maximizing bid. The resulting bid function can be easily obtained, and bids can be mapped to ranges of values via the inverse of this function. If the true equilibrium bid distribution is continuous with strictly positive density then both the noise in estimation and the width of the estimated value intervals vanish as the size of the observed data set increases. Notice that our earlier observation that nothing can be learned about the value distribution below the reserve price (or in an ironed interval) corresponds to the probability that a bidder is served being constant for bids within a given interval; The bid distribution would have zero density on this interval.

Our examples above that demonstrate a failure to perform economic inference are extreme in that the failure was due to total unresponsiveness of the mechanism. Intuitively, the degree to which good inference is possible should depend on the degree of responsiveness. The ideal econometric model for inference, however, does not separate mechanisms by the degree of responsiveness. As an example, consider auctioning a single item to one of 100 agents by the highest-bid-wins auction with either first-price or all-pay semantics. With first-price semantics only the winner pays his bid; with all-pay semantics all bidders pay their bids. Consider a bidder with value close to the median of the distribution. With the first-price auction the bidder bids by shading his value to the expected highest-other-bid given that his bid is the highest. For most continuous distributions, this is slightly below his value. With the all-pay auction, relative to the first-price bid, the bidder also reduces his bid by a multiplicative factor proportional to his likelihood of being the highest bidder. For a median value bidder with 99 other bidders, this probability is 2^{-99}. In contrast to the first-price auction where determining a reasonable bid is relatively easy, it is unlikely that a bidder would respond in an all-pay auction to a degree of accuracy that allows inference. As described above, the classical model for inference suggests that, in both cases, the value distribution can be inferred to a precision that vanishes with the number of observations. Our theory will enable the distinction of these two mechanisms as one with good inference and one with bad inference.

In summary, the goal of this work is in a theory for the design of mechanisms that perform

[1]Optimal auctions are constructed by mapping values to virtual values and then optimizing the virtual surplus (i.e., the sum of the the virtual values of agents who are served). Reserve prices exclude bidders with negative virtual values; bidders with value below the reserve price may as well not show up at the auction. Ironing comes from the process of constructing virtual values where a non-monotone virtual value function is "ironed" to be flat. When two or more bidders have the same virtual value they receive the same service probability (and payment); bidders in an ironed interval may as well make the same bid.

well in terms of revenue and inference. The motivating story that the theory should resolve is that of a revenue-maximizing auctioneer continuously adapting an auction in a slowly changing world. As the distribution of agent preferences evolves, the equilibrium behavior evolves, and the designer observes behavior and adapts the mechanism.

Model and Justification To explore this question of optimization of revenue of auctions with good inference we consider the following auction and bidding model. The auction design space is given by a position auction environment and either first-price or all-pay semantics. There are n agents and n positions and each position has a corresponding service probability. The auctioneer may modify the position weights by moving service probability from high positions to low positions (or discarding) and then runs a rank-by-bid position auction on the modified position weights. The bidders are single-dimensional with private values drawn independently and identically from a common prior distribution, have linear utility given by their value for service received minus any payment made, and bid in Bayes-Nash equilibrium (i.e., the model of classical auction theory). Discussion of these modeling choices follows.

We choose position auction environments because good revenue can only be obtained by optimization over the whole range of the distribution of values; therefore, good inference is especially important. For example, reserve pricing alone cannot achieve better than a logarithmic (in the number of agents) approximation to the optimal auction (Hartline and Yan, 2011). This choice, then places the problem in an environment where even obtaining a constant approximation to the optimal revenue is non-trivial. Aiming for a constant approximation will allow us to identify simple mechanisms that provide economic intuition; whereas, optimal mechanisms may be analytically intractable and economically opaque.

We choose to restrict the auctioneer to *rank-based* auctions where the position auction can be modified by shifting service probability downward (or discarding); importantly we disallow reserve prices or ironing (by value). Three reasons follow. First, as observed above, reserve prices and ironing do not permit good inference. Second, rank-based first-price and all-pay position auctions have a unique Bayes-Nash equilibrium that is symmetric and in which bids are always in the same order as values; therefore, the equilibrium allocation rule is predetermined by the position weights alone (Chawla and Hartline, 2013). Third, as a result that we will prove, both reserve prices and ironing can be simulated by shifting probability downward (or discarding) with at most a constant factor loss in revenue. Furthermore, the optimal rank-based auction (which obtains at least the revenue of this simulation) can be obtained by ironing by rank and discarding low ranks only. Ironing by rank considers a set of consecutive position and averages their service probabilities. Thus, this restriction is without loss up to a constant approximation.

We choose first-price and all-pay payment semantics because they are fundamental auction types with non-truthtelling equilibrium in the classical model for auction theory. This choice contrasts with the standard choice in the mechanism design literature where attention is often restricted to mechanisms with the truthtelling equilibrium. Of course, for truthful mechanisms in the classical model, bids are equal to values and inference is trivial. In practice, however, these nice-in-theory auctions are rarely employed. Ausubel and Milgrom (2006) give some explanation for this non-translation from theory to practice in an essay entitled "The Lovely but Lonely Vickrey Auction." Moreover, even if an auction possessing a truthtelling equilibrium in theory is employed in practice, truthtelling is unlikely. This loss of truthtelling could come from externalities, exposure, outside options, or privacy concerns. In contrast, our restriction to first-price and all-pay auctions permits the consideration of inference in the classical, and most fundamental, model of auction theory.

Results and organization Our results are as follows.

- In Section 2 we show how to estimate the bid distribution from a finite sample of observed bids. Then we provide a general principle for inferring values from the estimated distribution of bids and its derivative for a non-truthful mechanism. We illustrate this principle in application to a first price and an all-pay mechanism.

- We prove in Section 3 that rank based auctions are close to optimal for revenue: for every value distribution and every symmetric auction facing a position feasibility constraint[2], there exists a rank based auction satisfying the same constraint that achieves a 4-approximation (and in most cases, a close-to-1 approximation). Further, we prove that optimizing for revenue over the class of rank based auctions requires knowing n parameters of the value distribution that we call the multi-unit revenues. These multi-unit revenues are a discrete analog to the revenue curve, the derivative of which defines virtual values for revenue-optimal auctions.

- We show how to estimate the multi-unit revenues from samples drawn from the bid distribution of a mechanism in Section 4. We give bounds on the error in estimation in terms of the allocation rule of the mechanism, and the number of bid samples. We show that the error decreases as the inverse square root of the number of bid samples.

- Also in Section 4, we show that rank based auctions achieve a good revenue versus inference tradeoff: for every $\epsilon > 0$, there exists a rank based auction that obtains a $(1 - \epsilon)$ fraction of the optimal rank based revenue, and achieves error in inferring the multi-unit revenues that scales as $O(1/\epsilon)$.[3]

- We show how to estimate the revenue curve for the underlying value distribution from samples drawn from the bid distribution of a mechanism in Section 5. We give bounds on the error in estimation in terms of the allocation rule of the mechanism, the number of bid samples, and the measurement error in bids. These bounds require the allocation function to have a minimum and a maximum slope at all quantiles. An implication is that the revenue-optimal auction (with a mostly "flat" allocation rule) is generally poor at inference and requires substantially larger samples to achieve the inference qualtify of the iron by rank auctions.

  We show that the dependence of the estimation error on the aforementioned parameters is polynomially worse in this setting as compared to those for the problem of learning just the multi-unit revenues.

Related Work The inference approaches that we use in this paper are related to the recent work in econometrics literature on the so-called plug-in estimation of strategic response models. For instance, Guerre et al. (2000) analyzes the estimation of the distribution of values from the distribution of bids in the first-price auction. An overview of related empirical models applied to other auction environments can be found in Paarsch and Hong (2006). The principle for inference in

[2]For irregular value distributions, we need the assumption that the auction does not iron over the values corresponding to the highest $1/n$ quantiles.

[3]Note that this scaling in error can be offset by obtaining a large number of bid samples, and by reducing measurement error in bids.

these models is based on the two-step approach discussed in detail in Bajari et al. (2013). The idea of the two step method is based on estimation of the empirical distribution of equilibrium outcome directly from the data in the first step. In the second step this empirical distribution is plugged into the first-order condition for each player to obtain this player's payoff (or value in case of auctions). This literature, however, does not consider the question of comparing different mechanisms in terms of the possibility to infer the parameters of another mechanism. The mechanism that generates the data is usually taken as given. Such an analysis has, however, been considered in non-strategic environments in Chernozhukov et al. (2013). The non-strategic structure of the environment makes it significantly easier to analyze, provided that one does not need to recover the best response correspondence from the empirical observations, which is one of the key components of our analysis.

In the mechanism design literature, the problem of designing mechanisms to enable learning the parameters of a market has not been considered from a theoretical perspective previously. Several works have considered the problem of learning optimal pricing schemes in an online setting (e.g., Babaioff et al. (2012)). However, these works assume non-strategic behavior on part of the agents, which makes the inference much simpler. Other works consider the problem of learning click-through-rates in the context of a sponsored search auction (a generalization of the position environment we study) while simultaneously obtaining good revenue (e.g., Devanur and Kakade (2009); Babaioff et al. (2009); Gatti et al. (2012)), however, they restrict attention to truthful mechanisms, and again do not require inference.

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.

Finally, the theory that we develop for optimizing revenue over the class of iron by rank auctions is isomorphic to the theory of envy-free optimal pricing developed by Hartline and Yan (2011).

2 Preliminaries

2.1 Auction Theory

A standard auction design problem is defined by a set $[n] = \{1, \dots, n\}$ 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. If x_i indicates 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 $\boldsymbol{b} = (b_1, \dots, b_n)$ from the agents and maps the vector $\boldsymbol{b}$ of bids to an allocation $\tilde{\boldsymbol{x}}(\boldsymbol{b}) = (\tilde{x}_1(\boldsymbol{b}), \dots, \tilde{x}_n(\boldsymbol{b}))$, specifying the probability with which each agent is served, and prices $\tilde{\boldsymbol{p}}(\boldsymbol{b}) = (\tilde{x}_1(\boldsymbol{b}), \dots, \tilde{x}_n(\boldsymbol{b}))$, specifying the expected amount that each agent is required to pay. An auction is denoted by $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{p}})$.

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, $\tilde{p}_i(\boldsymbol{b}) = b_i$.

Feasibility in position auction environments Auction designers face a feasibility constraint that restricts the allocations $\boldsymbol{x}$ that the mechanism may produce. In this paper we focus on position auction environments. In such environments, the feasibility constraint is given by *position weights* $1 \geq w_1 \geq w_2 \geq \cdots \geq w_n \geq 0$. An allocation function assigns agents (randomly) to positions 1 through n, and an agent assigned to position i gets allocated w_i. In other words, an allocation $\boldsymbol{x} = (x_1, \ldots, x_n)$, sorted so that $x_1 \geq x_2 \geq \cdots \geq x_n$, is feasible if and only if it can be obtained by shifting weight downward (or discarding), i.e., for all $k \in [n]$, $\sum_{j \leq k} x_j \leq \sum_{j \leq k} w_j$.

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:

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

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 $n\,\mathbf{Rev}[x]$.

Rank-based auctions In a rank-based auction, the allocation to an agent depends solely on the rank of his bid among the other agents' bids, and not on the actual bid. For example, a k-highest-bids-win auction is a rank-based auction, however, a k-highest-bids-win-subject-to-reserve-bid-r auction is not a rank-based auction.

2.2 Inference

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

$$u(q,z) = (v(q) - z)\, x(b^{-1}(z)). \tag{3}$$

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 $\frac{d}{dz}u(z,q)$ is 0 at $z = b(q)$. Rarranging, we obtain:

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

Inference for all-pay auctions We repeat the calculation above for rank-based all-pay auctions; the starting equation (3) is replaced with the analogous equation for all-pay auctions:

$$u(q, z) = v(q)\, x(b^{-1}(z)) - z. \tag{5}$$

Differentiating with respect to z we obtain:

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

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

$$v(q) = \tfrac{b'(q)}{x'(q)}. \tag{6}$$

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.

2.3 Statistical Model and Methods

In this section we describe the errors in the estimated bid distribution and standard analyses for rates of convergence. The main error in estimation of the bid distribution is the *sampling error* due to drawing only a finite number of samples from the bid distribution.

The analyst obtains N samples from the bid distribution. Each sample is the corresponding agent's best response to the *true* bid distribution. We assume that the number of samples N is roughly polynomial in n, the number of agents in a single auction.

We can estimate the equilibrium bid distribution $b(q)$ as follows. Let $\hat{b}_1, \cdots, \hat{b}_N$ denote the N samples drawn from the bid distribution. Sort the bids so that $\hat{b}_1 \leq \hat{b}_2 \leq \cdots \leq \hat{b}_N$ and define the *estimated bid distribution* $\hat{b}(\cdot)$ as

$$\hat{b}(q) = \hat{b}_i \qquad\qquad \forall i \in N, q \in [i-1, i)/N \tag{7}$$

Definition 1. *For function $b(\cdot)$ and estimator $\hat{b}(\cdot)$, the* mean squared error *as a function of the number of samples N is*

$$\mathrm{MSE}_b(N) = \mathbf{E}\Big[\sup_q \big|b(q) - \hat{b}(q)\big|^2\Big]^{1/2}.$$

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., n) constant, that is, $r(N)\,\mathrm{MSE}(N) = \Theta(1)$.*

We will be interested in comparing the error bounds achieved by different algorithms for inference. Accordingly, we will state these bounds in terms of the rate of convergence. We will also state bounds on the mean squared error of the quantities we estimate, in terms of the mechanism that generates the bids, in order to optimize for the mechanism.

Lemma 2.1. *The estimator* $\hat{b}(\cdot)$ *defined directly from the bids by equation* (7) *has a rate of convergence of* $r(N) = \sqrt{N}$ *and mean squared error* $\mathrm{MSE}_b(N) = O(\sup_q b'(q)/\sqrt{2N})$. *Recalling that* $v(q) \leq 1$ *for all quantiles* q, *we obtain the following expressions for mean squared error in terms of the allocation function.*

(i) If the samples are generated from the first-price auction, then the mean squared error can be evaluated as $\mathrm{MSE}_b(N) = O(\sup_q \left\{\frac{x'(q)}{x(q)}\right\} / \sqrt{2N})$.

(ii) If the samples are generated from the all-pay auction, then the mean squared error can be evaluated as $\mathrm{MSE}(N) = O(\sup_q x'(q)/\sqrt{2N})$.

To estimate the value distribution, as is evident from equations (4) and (6), an estimator for the derivative of the bid function, or equivalently, for the density of the bid distribution, is needed. Estimation of densities is standard; however, they require assumptions on the distribution, e.g., continuity, and the convergence rates are strictly slower. Our main results do not need to explicitly estimate the value distribution and therefore, we defer these standard methods to Section 5 where our technique is compared with the straightforward approach of estimating the value distribution explicitly.

3 Rank-based auctions

One of the main contributions of this paper is to introduce a restricted class of rank-based auctions for position environments that simultaneously have good performance and good econometric properties. Recall that in a rank-based auction the allocation to an agent depends solely on the rank of his bid among other agents' bids, and not on the actual bid. For a position environment, a rank-based auction assigns agents (potentially randomly) to positions based on their ranks.

Consider a position environment given by non-increasing weights $\boldsymbol{w} = (w_1, \ldots, w_n)$. For notational convenience, define $w_{n+1} = 0$. Define the cumulative position weights $\boldsymbol{W} = (W_1, \ldots, W_n)$ as $W_k = \sum_{j=1}^{k} w_j$, and $W_0 = 0$. We can view the cumulative weights as defining a piece-wise linear, monotone, concave function given by connecting the point set $(0, W_0), \ldots, (n, W_n)$.

A random assignment of agents to positions based on their ranks induces an expected weight to which agents of each rank are assigned, e.g., $\bar{w}_k$ for the kth ranked agent. These expected weights can be interpreted as a position auction environment themselves with weights $\bar{\boldsymbol{w}}$. As for the original weights, we can define the cumulative position weights $\bar{\boldsymbol{W}}$ as $\bar{W}_k = \sum_{j=1}^{k} w_j$. An important issue for optimization of rank-based auctions is to characterize the inducible class of position weights.

Lemma 3.1 (e.g., Devanur et al., 2013)**.** *There is a rank-based auction with induced position weights* $\bar{\boldsymbol{w}}$ *for position environment with weights* $\boldsymbol{w}$ *if and only if their cumulative weights satisfy* $\bar{W}_k \leq W_k$ *for all* k, *denoted* $\bar{\boldsymbol{W}} \leq \boldsymbol{W}$.

Any feasible weights $\bar{\boldsymbol{w}}$ can be constructed from a sequence of the following two operations.

rank reserve For a given rank k, all agents with ranks between $k+1$ and n are rejected. The resulting weights $\bar{\boldsymbol{w}}$ are equal to $\boldsymbol{w}$ except $\bar{w}_{k'} = 0$ for $k' > k$.

iron by rank Given ranks $k' < k''$, the ironing-by-rank operation corresponds to, when agents are ranked, assigning the agents ranked in an interval $\{k', \ldots, k''\}$ uniformly at random to these same positions. The ironed position weights $\bar{\boldsymbol{w}}$ are equal to $\boldsymbol{w}$ except the weights on the ironed interval of positions are averaged. The cumulative ironed position weights $\bar{\boldsymbol{W}}$ are equal to $\boldsymbol{W}$ (viewed as a concave function) except that a straight line connects $(k'-1, \bar{W}_{k'-1})$ to $(k'', \bar{W}_{k''})$. Notice that concavity of $\boldsymbol{W}$ (as a function) and this perspective of the ironing procedure as replacing an interval with a line segment connecting the endpoints of the interval implies that $\boldsymbol{W} \geq \bar{\boldsymbol{W}}$ coordinate-wise, i.e., $W_k \geq \bar{W}_k$ for all k.

Multi-unit highest-bids-win auctions form a basis for position auctions. Consider the marginal position weights $\boldsymbol{w}' = (w'_1, \ldots, w'_n)$ defined by $w'_k = w_k - w_{k+1}$. The allocation rule induced by the position auction with weights $\boldsymbol{w}$ is identical to the allocation rule induced by the convex combination of multi-unit auctions where the k-unit auction is run with probability w'_k.

In this section we develop a theory for optimizing revenue over the class of all rank-based auctions that resembles Myerson's theory for optimal auction design. Where Myerson's theory employs ironing by value and value reserves, our approach analogously employs ironing by rank and rank reserves. We then show that the revenue of rank-based auctions is close to optimal for position environments. Our econometric study of rank-based auctions is deferred to Section 4.

3.1 Optimal rank-based auctions

In this section we describe how to optimize for expected revenue over the class of iron by rank auctions. Recall that iron by rank auctions are linear combinations over k-unit auctions. The characterization of Bayes-Nash equilibrium, cf. equation (2), shows that revenue is a linear function of the allocation rule. Therefore, the revenue of a position auction can be calculated as the convex combination of the revenues from the k-unit highest-bids-win auctions.

The revenue from a k-unit n-agent highest-bids-win auction with agent values drawn i.i.d. from distribution F can be calculated in terms of the the agents revenue curve $R(q)$ and the allocation rule $x^{(k:n)}(q)$ of the k-highest-bids-win auction (for n agents). 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.

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

Notice that $x^{(0:n)}(q) = 0$ and $x^{(n:n)}(q) = 1$. The *per-agent* revenue obtained is $P_k = \mathbf{E}_q\big[R'(q)\, x^{(k:n)}(q)\big]$. Notice that $P_0 = P_n = 0$.

Given the multi-unit revenues, $\boldsymbol{P} = (P_0, \ldots, P_n)$, the problem of designing the optimal rank-based auction is well defined: given a position environment with weights $\boldsymbol{w}$, find the weights $\bar{\boldsymbol{w}}$ for an rank-based auction with cummulative weights $\bar{\boldsymbol{W}} \leq \boldsymbol{W}$ maximizing the sum $\sum_k (\bar{w}_k - \bar{w}_{k+1}) P_k$. This optimization problem is isomorphic to the theory of envy-free optimal pricing developed by Hartline and Yan (2011). We summarize this theory below; a complete derivation can be found in Appendix A.

Define the *multi-unit revenue curve* as the piece-wise linear function connecting the points $(0, P_0), \ldots, (n, P_n)$. This function may or may not be concave. Define the *ironed multi-unit revenues* as $\bar{\boldsymbol{P}} = (\bar{P}_0, \ldots, \bar{P}_n)$ according to the smallest concave function that upper bounds the multi-unit revenue curve. Define the multi-unit marginal revenues, $\boldsymbol{P}' = P'_1, \ldots, P'_n$ and $\bar{\boldsymbol{P}}' = \bar{P}'_1, \ldots, \bar{P}'_n$, as the left slope of the multi-unit and ironed multi-unit revenue curves, respectively. I.e., $P'_k = P_k - P_{k-1}$ and $\bar{P}'_k = \bar{P}_k - \bar{P}_{k-1}$.

Theorem 3.2. *Given a position environment with weights $\boldsymbol{w}$, the revenue-optimal iron-by-rank auction is defined by position weights $\bar{\boldsymbol{w}}$ that are equal to $\boldsymbol{w}$, except ironed on the same intervals as $\boldsymbol{P}$ is ironed to obtain $\bar{\boldsymbol{P}}$, and positions k for which $\bar{P}'_k$ is negative are discarded (by setting $\bar{w}_k = 0$).*

As is evident from this description of the optimal iron-by-rank auction, the only quantities that need to be ascertained to run this auction is the multi-unit revenue curve defined by $\boldsymbol{P}$. Therefore, an econometric analysis for optimizing iron-by-rank auctions need not estimate the entire value distribution; estimation of the multi-unit revenues is sufficient.

3.2 Optimal rank-based auctions with strict monotonicity

Position auctions, by definition, have non-increasing position weights $\boldsymbol{w}$. The ironing in the iron-by-rank optimization of the preceding section was to convert the problem of optimizing multi-unit marginal revenue subject to non-increasing position weight, to a simpler problem of optimizing multi-unit marginal revenue without any constraints. In this section, we describe the optimization of rank-based auctions (i.e., ones for which position weights can be shifted only downwards or discarded) subject to *strictly decreasing* position weights. This strictness is needed for insuring good inference properties, the details of which are formalized in Section 4.

As described by Lemma 3.1, position weights $\bar{\boldsymbol{w}}$ are feasible as a rank-based auction in position environment $\boldsymbol{w}$ if the cumulative position weights satisfy $W_k \geq \bar{W}_k$ for all k. Suppose we would like to optimize $\bar{\boldsymbol{w}}$ subject to a strict monotonicity constraint $\bar{w}'_k = \bar{w}_k - \bar{w}_{k+1} \geq \epsilon$. As non-trivial ironing by rank always results in consecutive positions with the same weight, i.e., $\bar{w}'_k = 0$ for some k, the optimal rank-based mechanism will require overlapping ironed intervals.

To our knowledge, performance optimization subject to a strict monotonicity constraint has not previously been considered in the literature. At a high level our approach is the following. We start with $\boldsymbol{w}$ which induces the cumulative position weights $\boldsymbol{W}$ which constrain the resulting position weights $\bar{\boldsymbol{w}}$ (via its cumulative $\bar{\boldsymbol{W}}$) of any feasible rank-based auction. We view $\bar{\boldsymbol{w}}$ as the combination of two position auctions. The first has weakly monotone weights $\bar{\boldsymbol{y}} = (\bar{y}_1, \ldots, \bar{y}_n)$; the second has strictly monotone weights $((n-1)\epsilon, (n-2)\epsilon, \ldots, \epsilon, 0)$; and the combination has weights $\bar{w}_k = \bar{y}_k + (n-k)\epsilon$ for all k. The revenue of the combined position auction is the sum of the revenues of the two component position auctions. Since the second auction has fixed position weights, its revenue is fixed. Since the first position auction is weakly monotone and the second is strictly, the combined position auction is strictly monotone and satisfies the constraint that $\bar{w}'_k \geq \epsilon$ for all k.

This construction focuses attention on optimization of $\bar{\boldsymbol{y}}$ subject to the induced constraint imposed by $\boldsymbol{w}$ and after the removal of the ϵ strictly-monotone allocation rule. I.e., $\bar{\boldsymbol{w}}$ must be feasible for $\boldsymbol{w}$. The suggested feasibility constraint for optimization of $\bar{\boldsymbol{y}}$ is given by position weights $\boldsymbol{y}$ defined as $y_k = w_k - (n-k)\epsilon$. Notice that, in this definition of $\boldsymbol{y}$, a lesser amount is subtracted from successive positions. Consequently, monotonicity of $\boldsymbol{w}$ does not imply monotonicity of $\boldsymbol{y}$.

To obtain $\bar{\boldsymbol{y}}$ from $\boldsymbol{y}$ we may need to iron for two reasons, (a) to make $\bar{\boldsymbol{y}}$ monotone and (b) to make the multi-unit revenue curve monotone. In fact, both of these ironings are good for revenue. The ironing construction for monotonizing $\boldsymbol{y}$ constructs the concave hull of the cumulative position weights $\boldsymbol{Y}$. This concave hull is strictly higher than the curve given by $\boldsymbol{Y}$ (i.e., connecting $(0, Y_0), \ldots, (n, Y_n)$). Similarly the ironed multi-unit revenue curve given by $\bar{\boldsymbol{P}}$ is the concave hull of the multi-unit revenue curve given by $\boldsymbol{P}$. The correct order in which to apply these ironing procedures is to first (a) iron the position weights $\boldsymbol{y}$ to make it monotone, and second (b) iron the multi-unit revenue curve $\boldsymbol{P}$ to make it concave. This order is important as the revenue of the position auction with weights $\bar{\boldsymbol{y}}$ is only given by the ironed revenue curve $\bar{\boldsymbol{P}}$ when the $\bar{\boldsymbol{y}}' = 0$ on the ironed intervals of $\bar{\boldsymbol{P}}$.

Theorem 3.3. *The optimal ϵ strictly-monotone rank-based auction for position weights $\boldsymbol{w}$ has position weights $\bar{\boldsymbol{w}}$ constructed by*

1. *defining $\boldsymbol{y}$ by $y_k = w_k - (n-k)\epsilon$ for all k.*
2. *averaging position weights of $\boldsymbol{y}$ on intervals where $\boldsymbol{y}$ should be ironed to be monotone.*
3. *averaging the resulting position weights on intervals where $\boldsymbol{P}$ should be ironed to be concave to get $\bar{\boldsymbol{y}}$*
4. *setting $\bar{\boldsymbol{w}}$ as $\bar{w}_k = \bar{y}_k + (n-1)\epsilon$.*

Proof. The proof of this theorem follows directly by the construction and its correctness. □

The rank-based auction given by $\bar{\boldsymbol{w}}$ in position environment given by $\boldsymbol{w}$ can be implemented by a sequence of iron-by-rank and rank-reserve operations. Such a sequence of operations can be found, e.g., via an approach of Alaei et al. (2012) or Hardy et al. (1929).

3.3 Approximation via rank-based auctions

In this section we show that the revenue of optimal rank-based auction approximates the optimal revenue (over all auctions) for position environments. Instead of making this performance directly we will instead identify a simple non-optimal rank-based auction that approximates the optimal auction. Of course the optimal rank-based auction of Theorem 3.2 has revenue at least that of this simple rank-based auction, thus its revenue also satisfies the same approximation bound.

Our approach is as follows. Just as arbitrary rank-based mechanisms can be written as convex combinations over k-unit highest-bids-win auctions, the optimal auction can be written as a convex combination over optimal k-unit auctions. We begin by showing that the revenue of optimal k-unit auctions can be approximated by multi-unit highest-bids-win auctions when the agents' values are distributed according to a regular distribution (Lemma 3.4, below). In the irregular case, on the other hand, rank-based auctions cannot compete against arbitrary optimal auctions. For example, if the agents' value distribution contains a very high value with probability $o(1/n)$, then an optimal auction may exploit that high value by setting a reserve price equal to that value; On the other hand, a rank-based mechanism cannot distinguish very well between values correspond to quantiles above $1 - 1/n$. We show that rank-based mechanisms can approximate the revenue of any mechanism that does not iron the quantile interval $[1-1/n, 1]$ (but may arbitrarily optimize over the remaining quantiles). Theorem 3.6 presents the precise statement.

Lemma 3.4. *For regular k-unit n-agent environments, there exists a $k' \leq k$ such that the highest-bid-wins auction that restricts supply to k' units (i.e., a rank reserve) obtains at least half the revenue of the optimal auction.*

Proof. This lemma follows easily from a result of Bulow and Klemperer (1996) that states that for agents with values drawn i.i.d. from a regular distribution the revenue of the k'-unit n-agent highest-bid-wins auction is at least the revenue of the k'-unit $(n-k')$-agent optimal auction. To apply this theorem to our setting, let us use $\mathbf{OPT}(k, n)$ to denote the revenue of an optimal k-unit n-agent auction, and recall that nP_k is the revenue of a k-unit n-agent highest-bids-win auction.

When $k \leq n/2$, we pick $k' = k$. Then,

$$nP_k \geq \mathbf{OPT}(k, n-k) \geq \frac{(n-k)}{n}\mathbf{OPT}(k, n) \geq \frac{1}{2}\mathbf{OPT}(k, n),$$

and we obtain the lemma. Here the first inequality follows from Bulow and Klemperer's theorem and the third from the assumption that $k \leq n/2$. The second inequality follows via by lower bounding $\mathbf{OPT}(k, n-k)$ by the following auction which has revenue exactly $\frac{(n-k)}{n}\mathbf{OPT}(k, n)$: simulate the optimal k-unit n-agent on the $n-k$ real agents and k fake agents with values drawn independently from the distribution. Winners of the simulation that are real agents contribute to revenue and the probability that an agent is real is $(n-k)/n$.

When $k > n/2$, we pick $k' = n/2$. As before we have:

$$nP_{n/2} \geq \mathbf{OPT}(n/2, n/2) = \frac{1}{2}\mathbf{OPT}(n, n) \geq \frac{1}{2}\mathbf{OPT}(k, n).$$ □

Lemma 3.5. *For any number of agents n, distribution with revenue curve $R(\cdot)$, and quantile $q \leq 1 - 1/n$, there exists an integer $k \leq (1-q)n$ such that the k-unit n-agent highest-bid-wins auction is at least a quarter of $nR(q)$, the revenue from posting price $v(q)$.*

Proof. First we get a lower bound on P_k for any k. For any value v', the total expected revenue of the k-unit n-agent highest-bid-wins auction is at least $v'\,k$ times the probability that $k+1$ agents have value at least v'. The median of a binomial random variable corresponding to n Bernoulli trials with with success probability $(k+1)/n$ is $k+1$. Thus, the probability that this binomial is at least $k+1$ is at least $1/2$. Combining these observations by choosing $v' = v(1-(k+1)/n)$ we have,

$$n\,P_k \geq v(1-(k+1)/n)\,k/2.$$

Choosing $k = \lfloor(1-q)n\rfloor - 1$, for which $v(1-(k+1)/n) \geq v(q)$, the bound simplifies to,

$$n\,P_k \geq v(q)\,k/2.$$

The ratio of P_k and $R(q) = (1-q)\,v(q)$ is therefore at least $k/2(1-q)n > k/2(k+2)$, which for $q \leq 1 - 3/n$ (or, $k \geq 2$) is at least $1/4$.

For $q \in (1-3/n, 1-1/n]$, we pick $k = 1$. Then, P_1 is at least $1/n$ times $v(q)$ times the probability that at least two agents have a value greater than or equal to $v(q)$. We can verify for $n \geq 2$ that

$$P_1 \geq \frac{v(q)}{n}\left(1 - q^n - n(1-q)q^{n-1}\right) \geq \frac{1}{4}(1-q)\,v(q).$$

□

Theorem 3.6. *For regular agents and position environments, the optimal rank-based auction obtains at least half the revenue of the optimal auction. For (possibly irregular agents) and position environments, optimal rank-based auction obtains at least a quarter of the revenue of the optimal auction that does not iron or set a reserve price on the quantile interval $[1-1/n, 1]$.*

Proof. In the regular setting, the theorem follows from Lemma 3.4 by noting that the optimal auction (that irons by value and uses a value reserve) in a position environment is a convex combination of optimal k-unit auctions: since the revenue of each of the latter can be approximated by that of a k'-unit highest-bids-win auction with $k' \leq k$, the revenue of the convex combination can be approximated by that of the same convex combination over k'-unit highest-bids-win auctions; the resulting convex combination over k'-unit auctions satisfies the same position constraint as the optimal auction.

In the irregular setting, once again, any auction in a position environment is a convex combination of optimal k-unit auctions. The expected revenue of any k-unit auction is bounded from above by the expected revenue of the optimal auction that sells at most k items in expectation. The per-agent revenue of such an auction is bounded by $\overline{R}(1-k/n)$, the revenue of the optimal allocation rule with ex ante probability of sale k/n. Here $\overline{R}(\cdot)$ is the ironed revenue curve (that does not iron on quantiles in $[1-1/n, 1]$). $\overline{R}(1-k/n)$ is the convex combination of at most two points on the revenue curve $R(a)$ and $R(b)$, $a \leq 1-k/n \leq b < 1-1/n$. Now, we can use Lemma 3.5 to obtain an integer $k_a < n(1-a)$ such that P_{k_a} is at least a quarter of $R(a)$, likewise k_b for b. Taking the appropriate convex combination of these multi-unit auctions gives us a 4-approximation to the optimal auction k-unit auction (that does not iron over the quantile interval $[1-1/n, 1]$). Finally, the convex combination of the multi-unit auctions with k_a and k_b corresponds to a position auction with that is feasible for a k unit auction (with respect to serving the top k positions with probability one, service probability is only shifted to lower positions). □

4 Inference in rank-based auctions

Recall that the performance of any rank-based auction is governed by the multi-unit revenues $P_1, \ldots, P_n$ with P_k equal to the per-agent revenue of the highest-k-agents-win auction. In order to optimize over the class of rank based auctions, then, we need to estimate the n quantities P_k. We now describe how to estimate these quantities from the observed bids, and how the error in the estimation of the bid distribution translates into errors in the estimated multi-unit revenues.

Let x denote the allocation rule of the auction that we run, and let b denote the bid distribution in BNE of this auction. Recall that $x^{(k:n)}(\cdot)$ denotes the allocation rule of the highest-k-agents-win auction. In the following, we use x_k as a short-form for $x^{(k:n)}$. Then, the per-agent revenue of this auction is given by:

$$P_k = \mathbf{E}_q\left[x_k'(q)R(q)\right] = \mathbf{E}_q\left[x_k'(q)v(q)(1-q)\right]$$

We will now perform our analysis for the all-pay and first-price auction formats separately, using the respective bid-to-value conversion equations from Section 2.

4.1 Inference for an all-pay 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_k above we get

$$P_k = \mathbf{E}_q\left[x_k'(q)(1-q)\frac{b'(q)}{x'(q)}\right] = \mathbf{E}_q\left[Z_k(q)b'(q)\right]$$

where $Z_k(q) = (1-q)\frac{x_k'(q)}{x'(q)}$.

Writing the expectation as an integral and integrating by parts we obtain the following lemma. Here we note that $b(0) = 0$ and $Z_k(1) = 0$.

Lemma 4.1. *The per-agent revenue of the highest-k-agents-win auction can be written as a linear combination of the bids in an all-pay auction:*

$$P_k = \mathbf{E}_q\big[-Z_k'(q)b(q)\big]$$

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

This formulation allows us to express the error in estimation of P_k in terms of the error in estimating the bid distribution. In particular, let $\hat{P}_k$ denote the estimate of P_k obtained by plugging the bid estimator $\widehat{b}(\cdot)$ in the formula given by Lemma 4.1. Then we can write the error in P_k as:

$$|\hat{P}_k - P_k| \quad = \quad \mathbf{E}_q\Big[\Big|-Z_k'(q)(\widehat{b}(q) - b(q))\Big|\Big] \leq \mathbf{E}_q\big[|Z_k'(q)|\big] \sup_q |\widehat{b}(q) - b(q)|$$

Lemma 2.1 then gives the following bound on the mean squared error for P_k:

$$\mathrm{MSE}_{P_k}(N) \quad \leq \quad \frac{\sup_q b'(q)}{\sqrt{2N}} \mathbf{E}_q\big[|Z_k'(q)|\big]$$

We now proceed to bound $\mathbf{E}_q[|Z_k'(q)|]$. To this end we first note that if x is a convex combination over the allocation rules of the multi-unit highest-bids-win auctions, then Z_k has a single local maximum (see the appendix for a proof).

Lemma 4.2. *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 $q \in [0,1]$.*

Let $Z_k^* = \sup_q Z_k(q)$. Then, we can bound $\mathbf{E}_q[|Z_k'(q)|]$ by $2Z_k^* - Z_k(1) - Z_k(0) \leq 2Z_k^*$. We get the following theorem:

Theorem 4.3. *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 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:*

$$\mathrm{MSE}_{P_k}(N) \leq \sqrt{\frac{2}{N}} \sup_q \{x'(q)\} \sup_q \left\{\frac{(1-q)x_k'(q)}{x'(q)}\right\}$$

.

4.2 Inference from a first-price auction

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_k we get:

$$P_k = \mathbf{E}_q\left[(1-q)x_k'(q)b(q) + \frac{(1-q)x_k'(q)x(q)b'(q)}{x'(q)}\right] = \mathbf{E}_q\big[(1-q)x_k'(q)b(q) + Z_k(q)x(q)b'(q)\big]$$

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

Integrating the second expression by parts, we get

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

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

Lemma 4.4. *The per-agent revenue of the highest-k-agents-win auction can be written as a linear combination of the bids in a first-pay auction:*

$$P_k = \mathbf{E}_q\big[-x(q)Z'_k(q)b(q)\big]$$

where $Z_k(q) = (1-q)\frac{x'_k(q)}{x'(q)}$ *and* $x(q)$ *are known precisely.*

To bound the error in estimating P_k, once again we need to bound the integral $\int_0^1 x(q)|Z'_k(q)|\,dq$. Recall that Z_k has a single local maximum for $q \in [0,1]$ when x is a convex combination over the multi-unit auctions. This implies the following lemma (see the appendix for a proof).

Lemma 4.5. $\int_0^1 x(q)|Z'_k(q)|\,dq \leq 2x(q^*_k)Z_k(q^*_k) + 1$ *where* $q^*_k = \operatorname{argmax}_q Z_k(q)$.

Using this lemma and applying Lemma 2.1 from Section 2 we get the following theorem.

Theorem 4.6. *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 for all k, the mean squared error in estimating* P_k *from N samples from the bid distribution for a first-price auction with allocation rule* x *is:*

$$\mathrm{MSE}_{P_k}(N) \leq \sqrt{\frac{2}{N}} \sup_q \left\{\frac{x'(q)}{x(q)}\right\} \sup_q \left\{\frac{(1-q)x(q)x'_k(q)}{x'(q)}\right\}$$

.

4.3 Revenue versus inference tradeoff for rank-based auctions

We now consider optimizing for expected revenue over the class of rank based auctions subject to good inferability of the parameters P_k. Recall that the revenue of a rank based auction with position weights $\boldsymbol{w}$ and marginal weights $\boldsymbol{w}' = w_k - w_{k+1}$ is given by $\sum_k w'_k P_k$. On the one hand, estimating the P_k's well is important to be able to optimize $\boldsymbol{w}$ – we should place the most marginal weight on positions with high P_k's. On the other hand, the weights $\boldsymbol{w}$ determine the allocation rule x as a weighted sum of the k-unit allocation rules x_k, which in turn via Theorems 4.3 and 4.6 determine the error in the P_k's – we should ensure that all positions get some minimum marginal weight. This is the problem of finding the optimal ϵ strictly monotone rank-based auction that we discussed in Section 3.2 (see Theorem 3.3).

We now claim that the auction returned by Theorem 3.3 obtains revenue close to the optimal rank-based auction. In particular, one way of obtaining an ϵ strictly monotone auction given the estimates $\widehat{P_k}$ is to run the optimal auction with probability $1-\epsilon$ and with probability ϵ run the auction that assigns equal marginal weight to every position. In particular, for every k, $w'_k \geq \epsilon/n$. For this auction, recall that $x = \sum_k w'_k x_k$, and $x' = \sum_k w'_k x'_k$. Therefore, for any quantile q,

$$x'_k(q)/x'(q) \leq 1/w'_k \leq n/\epsilon.$$

We obtain the following theorem.

Theorem 4.7. *For every $\epsilon > 0$, there exists a rank based auction on n agents that obtains a $1-\epsilon$ approximation to the optimal rank based revenue. Furthermore, from N samples of the bid distribution, we can estimate parameters P_k for $k \in [n]$ with error bounds as below:*

$$\text{For the first price format: } |\hat{P}_k - P_k| \leq \sqrt{\frac{2}{N}}\,\frac{n}{\epsilon}\,\sup_q \left\{\frac{x'(q)}{x(q)}\right\}$$

$$\text{For the all pay format: } |\hat{P}_k - P_k| \leq \sqrt{\frac{2}{N}}\,\frac{n}{\epsilon}\,\sup_q \left\{x'(q)\right\}$$

We remark that while the theorem above gives the same upper bound on the error in estimation for P_ks for both the first price and all pay auction formats, comparing the bounds in Theorems 4.3 and 4.6 shows that the first price format is better at inference than the all pay format. Note that the error bound can be made arbirarily small by picking a large enough sample size N.

5 Inferring the revenue curve

In the previous section we considered the problem of inferring the parameters of the position auction from samples obtained from a first price and an all-pay auction. We now consider the problem of infering the entire revenue curve $R(\cdot)$ from bid samples by first infering the value distribution. This is relevant, for instance, if we want to estimate the revenue of an arbitrary mechanism, and if we want to optimize for revenue over the class of all mechanisms. We find that the inference problem becomes harder and a tight bound on revenue requires polynomially more samples as compared to the previous setting.

5.1 Propagation of errors in inference of value distribution

In order to infer the value distribution from the bid distribution, as given in Section 2 by equations (4) and (6), we need to estimate the derivative of the bid function $b(\cdot)$. Let $G(\cdot)$ denote the c.d.f. for bids, that is, $G(z) = b^{-1}(z)$ is the probability that a random bid is no more than z. Let $g(z) = \frac{d}{dz}b^{-1}(z)$ denote the corresponding density function. Note that $g(b(q)) = 1/b'(q)$. We will therefore focus on estimating $g(\cdot)$.

The density $g(\cdot)$ cannot be estimated directly from the empirical bid distribution of equation (7) because the derivative of that distribution is undefined. Nonetheless, a number of standard estimators are available to estimate $g(\cdot)$. Using such an estimator, $\hat{g}(\cdot)$, we obtain an estimator for the derivative as follows:

$$\hat{b}'(q) = 1/\hat{g}(\hat{b}(q)). \tag{8}$$

In Appendix B we formally state the requirements for an estimator of $b'(\cdot)$. We assume that we know the rate of convergence for this estimator, i.e. the sequence $r(N)$ with $r(N) \to \infty$ as we obtain more samples, such that:

$$E\left[\sup_b |\hat{g}(b) - g(b))|^2\right]^{1/2} = O(1/r(N)).$$

For instance, if one uses the histogram-based estimator for the density of bids, then

$$\frac{1}{\hat{b}'(q)} = \frac{1}{Nh}\sum_{i=1}^{N} \mathbf{1}\left\{|\hat{b}(q) - b_i| \leq h\right\},$$

where h is the bandwidth, which is selected such that $h\,N/\log(N) \to \infty$ as we get more samples. In this case given that the class of indicators $\mathbf{1}\{|b-t| \leq h\}$ when $t \in [b-\epsilon, b+\epsilon]$ (which depends on N) has a metric entropy of order $O(\epsilon)$, then the estimator for the derivative of the quantile function of the bid distribution converges at rate $r(N) = \sqrt{N\,h}$.

Functional objects, such as distribution densities, can be estimated using many estimators. If we restrict ourselves to feasible estimators (that are completely data-driven) and avoid oracle estimators, we can talk about an feasible estimator that achieves the fastest convergence rate. Such a convergence rate is called the *optimal convergence rate*. Stone (1980) established the optimal conevergence rate for estimation of one-dimensional density, which we formulate here without proof.

Lemma 5.1. *Suppose that the density of bids $g(\cdot)$ has k derivatives. Then the optimal convergence rate for the estimator for the density $\hat{g}(\cdot)$ is $r(N) = N^{k/(1+2k)}$.*

This theorem implies that there is a lower bound on the convergence rate equal to $N^{1/3}$ for estimation of distribution densities that have one derivative. At the same time, for functions that are very smooth, the optimal convergence rate can approach the maximum rate of $N^{1/2}$.

We now establish the mean-squared error of the estimator for the derivative of the bid distribution.

Theorem 5.2. *The mean-squared error for the estimator (8) for the derivative of the bid function at quantile q can be represented as*

$$\mathrm{MSE}_{b'(q)}(N) = E\left[(\hat{b}'(q) - b'(q))^2\right]^{1/2} = O\left(\frac{b'(q)^2}{r(N)} + \frac{|b''(q)|}{\sqrt{2N}}\right)$$

The value function can now be estimated using Equation (4) or (6), as applicable.

5.2 Inference from an all pay auction

Recall that for rank based mechanisms, we know the allocation function $x(q)$ and its derivative $x'(q)$ precisely. For all pay auctions Equation (6) allows us to relate the value distribution to the bid distribution: $v(q) = b'(q)/x'(q)$.

Let $\mathrm{MSE}_{b'(q)}(N) = E\left[(\hat{b}'(q) - b'(q))^2\right]^{1/2}$ and $\mathrm{MSE}_{v(q)}(N) = E\left[|\hat{v}(q) - v(q)|^2\right]^{1/2}$ denote the mean squared error in b' and v respectively.

We can express the mean squared error in R, $\mathrm{MSE}_{R(q)}(N) = E\left[(\hat{R}(q) - R(q))^2\right]^{1/2}$, in terms of the error in b' as follows. We write the revenue at quantile q as $R(q) = (1-q)v(q)$ meaning that we can estimate the revenue as $\hat{R}(q) = (1-q)\hat{v}(q)$ by replacing the true value with its estimated counterpart. then

$$\frac{\mathrm{MSE}_{R(q)}(N)}{R(q)} = \frac{(1-q)\,\mathrm{MSE}_{v(q)}(N)}{(1-q)v(q)} = \frac{\mathrm{MSE}_{b'(q)}(N)/x'(q)}{b'(q)/x'(q)} = \frac{\mathrm{MSE}_{b'(q)}(N)}{b'(q)}.$$

Our goal is to bound this relative error by a quantile-dependent quantity, $\epsilon(q)$. Formally, we require that

$$\frac{\mathrm{MSE}_{R(q)}(N)}{R(q)} \le \epsilon(q),$$

assuming that $R(q) > 0$. Our discussion above demonstrates that this error in turn can be expressed in the relative error of estimation of the derivative of the bid function $\mathrm{MSE}_{b'(q)}(N)/b'(q)$. We can now use Theorem 5.2 to obtain an error bound for the revenue curve[4] and derive conditions on the allocation rule x that bound the relative error in $b'(q)$ by $\epsilon(q)$ at all quantiles q:

Theorem 5.3. *Suppose that the allocation rule $x(\cdot)$ of an all-pay auction along with its second and first derivatives satisfies for all quantiles q:*

$$\frac{\Omega(1)|x''(q)|}{\epsilon(q)\sqrt{N}} \le x'(q) \le O(1)\frac{\epsilon(q)r(N)}{v(q)}, \text{ and, } N \ge \frac{1}{2}\left(\frac{v'(q)}{v(q)\epsilon(q)}\right)^2.$$

where $r(N)$ is the convergence rate for the estimator of the bid density $g(\cdot)$. Then, the relative error in estimating the revenue curve from N samples of the bid distribution is bounded by function $\epsilon(q)$.

Let us consider the bounds on x' closely. The lower bound of this expression is determined by the curvature of the allocation function and the number of samples. It requires the allocation rule to "separate" bids within the range $|x''(q)|/\sqrt{2N}$. The upper bound in this expression is driven by the sampling noise in the inference of the density of bids: if the allocation rule "jumps" at a certain quantile, the density of bids at that quantile is low and the relative error in bid density due to sampling becomes large. Note also that the lower bound on the number of samples required is determined by the slope of the value function, with more samples required for more concentrated distributions.

From the upper bound on $x'(q)$ we note that the error in estimating the revenue curve is at best $v(q)x'(q)/r(N)$, that is, we estimate the revenue at a rate of at most $r(N)$.[5] Recall from Lemma 5.1 that $r(N)$ can be as small as $N^{1/3}$ if x is not sufficiently smooth. This rate is much slower than the parametric convergence rate $\sqrt{N}$ derived in Theorem 4.7 and means that inference for the multi-unit revenues can be performed with much fewer samples than inference for general mechanisms.

5.3 First price auctions

The analysis for the first-price auctions follows closely our analysis for all-pay auctions. Recall that the value function can be obtained from the bid function and its derivative as

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

[4]Proofs for this section can be found in the appendix.
[5]This simplification ignores whether an allocation rule achieving this rate exists.

The value function is estimated by replacing the bid function and its derivative with their estimated counterparts. We further notice that the relative impact on the expected revenue can be bounded in the same fashion as for the all-pay auctions, meaning that

$$\frac{\mathrm{MSE}_{R(q)}(N)}{R(q)} = \frac{\mathrm{MSE}_{v(q)}(N)}{v(q)}.$$

This allows us to write an analog of Theorem 5.3 for the first-price auctions.

Theorem 5.4. *Suppose that the allocation rule $x(\cdot)$ of a first-price auction along with its second and first derivatives satisfies for all quantiles q:*

$$\frac{\Omega(1)|x''(q)|}{\sqrt{N}\epsilon(q)} \leq x'(q) \leq O(1)x(q)r(N)\epsilon(q), \quad N \geq \frac{1}{2}\left(\frac{v'(q)}{\epsilon(q)}\right)^2.$$

where $r(N)$ is the convergence rate for the estimator of the bid density $g(\cdot)$. Then the relative error in estimating the revenue curve from N samples of the bid distribution is bounded by function $\epsilon(q)$.

Once again from the upper bound on $x'(q)$, we get that the error $\epsilon(q)$ is at least $x'(q)/(r(N)x(q))$, that is, the revenue curve can be estimated at best at a rate of $r(N)$, which by Lemma 5.1 can be as small as $N^{1/3}$.

6 Discussion and Conclusions

We conclude with some observations and discussion.

- Good inference requires careful design of the mechanism. Perfect inference and perfect optimality cannot be achieved together.
- We cannot achieve good accuracy in infering the revenue of an arbitrary mechanism, or in infering the entire revenue curve. In contrast, the multi-unit revenues P_k are special functions that depend linearly on the bid distribution (and not, for example, on bid density). This property enables them to be learned accurately.
- Rank based mechanisms achieve a good tradeoff between revenue optimality and quality of inference in position environments: (1) They are close to optimal regardless of the value distribution; (2) Optimizing over this class for revenue requires estimating only n parameters P_k that, by our observation above, are "easy" to estimate accurately; (3) Rank based mechanisms satisfy the necessary conditions on the slope of the allocation function that enable good inference.

References

Alaei, S., Fu, H., Haghpanah, N., Hartline, J., and Malekian, A. 2012. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*.

Ausubel, L. M. and Milgrom, P. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg, Eds. MIT Press, Cambridge, MA, Chapter 1, 17–40.

Babaioff, M., Dughmi, S., Kleinberg, R., and Slivkins, A. 2012. Dynamic pricing with limited supply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. ACM, New York, NY, USA, 74–91.

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.

Bulow, J. and Klemperer, P. 1996. Auctions vs negotiations. *American Economic Review 86,* 1, 180–194.

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.

Chernozhukov, V., Fernández-Val, I., and Melly, B. 2013. Inference on counterfactual distributions. *Econometrica 81,* 6, 2205–2268.

Devanur, N. R., Ha, B. Q., and Hartline, J. D. 2013. Prior-free auctions for budgeted agents. In *ACM Conference on Electronic Commerce*. 287–304.

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.

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.

Guerre, E., Perrigne, I., and Vuong, Q. 2000. Optimal nonparametric estimation of first-price auctions. *Econometrica 68,* 3, 525–574.

Hardy, G., Littlewood, J., and Pólya, G. 1929. Some simple inequalities satisfied by convex functions. *Messenger of Math 58*, 145–152.

Hartline, J. and Yan, Q. 2011. Envy, truth, and profit. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 243–252.

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.

Stone, C. J. 1980. Optimal rates of convergence for nonparametric estimators. *The Annals of Statistics 8,* 6, 1348–1360.

A Finding the optimal iron by rank auction

Recall that iron by rank auctions are weighted sums of multi-unit auctions. Therefore, their revenue can be expressed as a weighted sum over the revenues P_k of k-unit auctions. We consider a position environment given by non-increasing weights $\boldsymbol{w} = (w_1, \ldots, w_n)$, with $w_0 = 0$, $w_1 = 1$, and $w_{n+1} = 0$. Define the cumulative position weights $\boldsymbol{W} = (W_1, \ldots, W_n)$ as $W_k = \sum_{j \leq i} w_j$.

Define the *multi-unit revenue curve* as the piece-wise constant function connecting the points $(0, P_0, \ldots, (n, P_n)$. This function may or may not be concave. Define the *ironed multi-unit revenue curve* as $\bar{\boldsymbol{P}} = (\bar{P}_1, \ldots, \bar{P}_n)$ the smallest concave function that upper bounds the multi-unit revenue curve. Define the multi-unit marginal revenues as $\boldsymbol{P'} = P'_1, \ldots, P'_n$ and $\bar{\boldsymbol{P}}' = \bar{P}'_1, \ldots, \bar{P}'_n$ as the left slope of the multi-unit and ironed multi-unit revenue curves, respectively. I.e., $P'_k = P_k - P_{k-1}$ and $\bar{P}'_k = \bar{P}_k - \bar{P}_{k-1}$.

We now see how the revenue of any position auction can be expressed in terms of the multi-unit revenue curves and marginal revenues.

$$\begin{aligned}
\mathbf{E}[\text{revenue}] &= \sum_{k=0}^{n} P_k\, w'_k = \sum_{k=0}^{n} P'_k\, w_k \\
&\leq \sum_{k=0}^{n} \bar{P}_k\, w'_k = \sum_{k=0}^{n} \bar{P}'_k\, w_k.
\end{aligned}$$

The first equality follows from viewing the position auction with weights $\boldsymbol{w}$ as a convex combination of multi-unit auctions (where its revenue is the convex combination of the multi-unit auction revenues). The second and final inequality follow from rearranging the sum (an equivalent manipulation to integration by parts). The inequality follows from the fact that $\bar{\boldsymbol{P}}$ is defined as the smallest concave function that upper bounds $\boldsymbol{P}$ and, therefore, satisfies $\bar{P}_k \geq P_k$ for all k. Of course the inequality is an equality if and only if $w'_k = 0$ for every k such that $\bar{P}'_k > P'_k$.

We now characterize the optimal ironing-by-rank position auction. Given a position auction weights $\boldsymbol{w}$ we would like the ironing-by-rank which produces $\bar{\boldsymbol{w}}$ (with cumulative weights satisfying $\boldsymbol{W} \geq \bar{\boldsymbol{W}}$) with optimal revenue. By the above discussion, revenue is accounted for by marginal revenues, and upper bounded by ironed marginal revenues. If we optimize for ironed marginal revenues and the condition for equality holds then this is the optimal revenue. Notice that ironed revenues are concave in k, so ironed marginal revenues are monotone (weakly) decreasing in k. The position weights are also monotone (weakly) decreasing. The assignment between ranks and positions that optimizes ironed marginal revenue is greedy with positions corresponding to ranks with negative ironed marginal revenue discarded. Tentatively assign the kth rank agent to slot k (discarding agents that correspond to discarded positions). This assignment indeed maximizes ironed marginal revenue for the given position weights but may not satisfy the condition for equality of revenue with ironed marginal revenue. To meet this condition with equality we can randomly permute (a.k.a., iron by rank) the positions that corresponds to intervals where the revenue curve is ironed. This does not change the surplus of ironed marginal revenue as the ironed marginal revenues on this interval are the same, and the resulting position weights $\bar{\boldsymbol{w}}$ satisfy the condition for equality of revenue and ironed marginal revenue.

B Requirements for the estimator of the derivative of the bid function

In general, we can express

$$p'(b) = \Phi_n(b; G, g), \text{ and } x'(b) = \Psi_n(b; G, g),$$

where $G(\cdot)$ is the cdf of the distribution of bids and $g(\cdot)$ is the pdf. For instance, when the mechanism $\mathcal{M}$ is the first-price auction, then $\Phi_n(b; G, g) = n\, G^{n-1}(b)\, g(b)$ and $\Psi_n(b; G, g) = G^n(b) + n\, b\, G^{n-1}(b) g(b)$. For the all-pay auction $\Phi_n(b; G, g) = n\, G^{n-1}(b)\, g(b)$ and $\Psi_n(b; G, g) = 1$.

Assumption 1. *Suppose that*

(i) Suppose that the density of bids g is bounded by a universal constant $\bar{g}$. There exists an estimator $\widehat{g}$ for which $r\,(\widehat{g} - g)(\cdot)$ converges to a tight stochastic process with convergence rate r such that $r \to \infty$ $r/\sqrt{nT} \to 0$

(ii) $\Phi_n(b; G, g)$ and $\Psi_n(b; G, g)$ are smooth functionals of G and g for each b such that for any two pairs (G_1, g_1) and (G_2, g_2) with $\|G_1 - G_2\| \leq \varepsilon_1$ and $\|g_1 - g_2\| \leq \varepsilon_2$:

$$\sup_{b \in [0, \bar{v}]} |\Phi_n(b; G_1, g_1) - \Phi_n(b; G_2, g_2)| \leq J_n^{\Phi,1} \varepsilon_1 + J_n^{\Phi,2} \varepsilon_2$$

and

$$\sup_{b \in [0, \bar{v}]} |\Psi_n(b; G_1, g_1) - \Psi_n(b; G_2, g_2)| \leq J_n^{\Psi,1} \varepsilon_1 + J_n^{\Psi,2} \varepsilon_2$$

We imposed this high-level assumption to facilitate a wide range of estimators that can be used to estimate the distribution and the density of the distribution of bids.

C Proofs

Proof of Lemma 2.1. Consider estimation of the bid function using the sorted bids $b^{(1)} \geq b^{(2)} \geq \ldots \geq b^{(N)}$. Then the bid function is estimated as

$$\widehat{b}(q) = b^{([qN])},$$

where $[\cdot]$ is the floor integer. We can equivalently express this function as a solution of the following equation

$$\frac{1}{N} \sum_{i=1}^{N} \mathbf{1}\{b_i \leq \widehat{b}(q)\} = q + o_p(1/\sqrt{N}),$$

where $o_p(1/\sqrt{N})$ corresponds to the error that arizes because the empirical cdf is a step function and identifies the true cdf in the steps of size $1/N$. Let $G(\cdot)$ be the cdf of bids and $\widehat{G}(\cdot)$ be the empirical cdf. Then the equation above can be rewritten as $\widehat{G}(\widehat{b}(q)) - G(b(q)) = o_p(1/\sqrt{N})$. Now we decompose this expression as

$$\widehat{G}(\widehat{b}(q)) - G(b(q)) = \widehat{G}(\widehat{b}(q)) - G(\widehat{b}(q)) + G(\widehat{b}(q)) - G(b(q)).$$

By the Donsker theorem $\sqrt{N}(\widehat{G}(t) - G(t))$ converges to a tight mean zero stochastic process $\mathbb{G}(t)$ over t with covariance function such that $H(t,t) = G(t)(1-G(t))$. Note that

$$\sup_t |H(t,t)| \le \frac{1}{4}.$$

This means that

$$E\left[\left(\sqrt{N}\sup_t(\widehat{G}(t) - G(t))\right)^2\right] \le \frac{1}{4}.$$

Next, consider the following expansion:

$$G(\widehat{b}(q)) - G(b(q)) = g(b(q))(\widehat{b}(q) - b(q)) + o(|\widehat{b}(q) - b(q)|^2).$$

Combining this result together with the decomposition above and recalling that $b'(q) = 1/g(b(q))$, we write

$$\sqrt{N}(\widehat{b}(q) - b(q)) = -b'(q)\sqrt{N}(\widehat{G}(\widehat{b}(q)) - G(\widehat{b}(q))) + o_p(1).$$

Then we write

$$E\left[\left(\sqrt{N}\sup_q(\widehat{b}(q) - b(q))\right)^2\right] \le \sup_q b'(q) E\left[\left(\sqrt{N}\sup_t(\widehat{G}(t) - G(t))\right)^2\right]^{1/2}.$$

This means that

$$\mathrm{MSE}_b(N) \le \frac{\sup_q b'(q)}{2}\frac{1}{\sqrt{N}}.$$

Then, recalling that $v(q) \le 1$ and $v(q) - b(q) \le 1$, we can replace the upper bound on $b'(q)$ by $x'(q)$ for an all-pay auction and by $x'(q)/x(q)$ for the first-price auction. □

Proof of Lemma 4.2. Consider the function $A(q) = 1/Z_k(q) = x'(q)/(1-q)x'_k(q)$. $x'(q)$ is a weighted sum over $x'_j(q)$ for $j \in \{1, \cdots, n-1\}$. So, $A(q)$ is a weighted sum over terms $x'_j(q)/(1-q)x'_k(q)$. Let us look at these terms closely.

$$\frac{x'_j(q)}{(1-q)x'_k(q)} = \alpha_{k,j} q^{k-j}(1-q)^{j-k-1}$$

where $\alpha_{k,j}$ is independent of q. The functions $q^{k-j}(1-q)^{j-k-1}$ are convex. This implies that $A(q)$ which is a weighted sum of convex functions is also convex. Consequently, it has a unique minimum. Therefore, $Z_k(q) = 1/A(q)$ has a unique maximum. □

Proof of Lemma 4.5. Recall that $Z_i(q)$ has a single local maximum at quantile q_i^*. So we get

$$\int_0^1 x(q)|Z'_i(q)|\,dq = \int_0^{q_i^*} x(q)Z'_i(q)\,dq - \int_{q_i^*}^1 x(q)Z'_i(q)\,dq$$

Integrating by parts,

$$\int x(q)Z_i'(q)\,dq = x(q)Z_i(q) - \int x'(q)Z_i(q)\,dq = x(q)Z_i(q) - \int (1-q)x_i'(q)\,dq$$

Therefore,

$$\int_0^1 x(q)|Z_i'(q)|\,dq = 2x(q_i^*)Z_i(q_i^*) - \int_0^{q_i^*}(1-q)x_i'(q)\,dq + \int_{q_i^*}^1 (1-q)x_i'(q)\,dq < 2x(q_i^*)Z_i(q_i^*) + 1$$

□

Proof of Theorem 5.2. Consider the difference

$$\hat{b}'(q) - b'(q) = \frac{1}{\hat{g}(\hat{b}(q))} - \frac{1}{g(\hat{b}(q))} + \frac{1}{g(\hat{b}(q))} - \frac{1}{g(b(q))}.$$

Using the Taylor expansion, with probability approaching 1 we can bound

$$|\hat{b}'(q) - b'(q)| \leq \frac{1}{g(b(q))^2}|\hat{g}(\hat{b}(q)) - g(\hat{b}(q))| + |\frac{g'(b(q))}{g^2(b(q))}||\hat{b}(q) - b(q)|.$$

Note that $b'(q) = 1/g(b(q))$. Also note that if we differentiate both sides of this expression with respect to q, we obtain

$$b''(q) = g'(b(q))b'(q)/g^2(b(q)).$$

By our assumption, $E[(|\hat{g}(\hat{b}(q)) - g(\hat{b}(q)))^2]^{1/2} = O(1/r(N))$. Also, by Lemma 2.1, $|\hat{b}(q) - b(q)| = O(b'(q)/(2\sqrt{N}))$. Combining these results, we obtain that

$$E[(\hat{b}'(q) - b'(q))^2]^{1/2} = O\left(\frac{b'(q)^2}{r(N)} + \frac{|b''(q)|}{2\sqrt{N}}\right).$$

□

Proof of Theorem 5.3. The proof of this theorem reduces to substitution of appropriate expressions of $b'(\cdot)$ and $b''(\cdot)$ into Theorem 5.2. For all-pay auctions $b'(q) = x'(q)v(q)$ and $b''(q) = x''(q)v(q)x'(q)v'(q)$. Therefore, we can express

$$\frac{MSE_{R(q)}(N)}{R(q)} = O\left(\frac{x'(q)v(q)}{r(N)} + \frac{|x''(q)v(q) + x'(q)v'(q)|}{2\sqrt{N}x'(q)v(q)}\right).$$

We guarantee the bound $\epsilon(q)$ for this expression if each term is bounded by $\epsilon(q)$:

$$\frac{x'(q)v(q)}{r(N)},\ \frac{|x''(q)|}{2\sqrt{N}x'(q)},\ \frac{v'(q)}{2\sqrt{N}v(q)} \leq \epsilon(q).$$

□

Proof of Theorem 5.4. Following the analysis of Theorem 5.3, we substitute the appropriate expressions of $b'(\cdot)$ and $b''(\cdot)$ into Theorem 5.2. Note that for the first-price auction

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

and

$$b''(q) = v'(q)\frac{x'(q)}{x(q)} + (v(q) - b(q))\frac{x''(q)}{x(q)} - 2(v(q) - b(q))\left(\frac{x'(q)}{x(q)}\right)^2.$$

This means that we can evaluate

$$\begin{aligned}\frac{MSE_{R(q)}(N)}{R(q)} = O\Bigg(&\frac{(v(q) - b(q))}{r(N)}\frac{x'(q)}{x(q)} \\ &+ \frac{\left|v'(q)\frac{x'(q)}{x(q)} + (v(q) - b(q))\frac{x''(q)}{x(q)} - 2(v(q) - b(q))\left(\frac{x'(q)}{x(q)}\right)^2\right|}{2\sqrt{N}\frac{(v(q)-b(q))}{r(N)}\frac{x'(q)}{x(q)}}\Bigg).\end{aligned}$$

To guarantee the bound $\epsilon(q)$, we need to bound each of the terms by $\epsilon(q)$. Thus, we require that

$$\frac{(v(q) - b(q))}{r(N)}\frac{x'(q)}{x(q)}, \frac{v'(q)}{(v(q) - b(q))2\sqrt{N}}, \frac{|x''(q)|}{x'(q)2\sqrt{N}}, \frac{x'(q)}{\sqrt{N}x(q)} \leq \epsilon(q).$$

□