The Sample Complexity of Revenue Maximization*

Richard Cole†
Courant Institute
New York University

Tim Roughgarden‡
Stanford University

November 25, 2015

Abstract

In the design and analysis of revenue-maximizing auctions, auction performance is typically measured with respect to a prior distribution over inputs. The most obvious source for such a distribution is past data. The goal of this paper is to understand how much data is necessary and sufficient to guarantee near-optimal expected revenue.

Our basic model is a single-item auction in which bidders' valuations are drawn independently from unknown and non-identical distributions. The seller is given m samples from each of these distributions "for free" and chooses an auction to run on a fresh sample. How large does m need to be, as a function of the number k of bidders and $\epsilon > 0$, so that a $(1 - \epsilon)$-approximation of the optimal revenue is achievable?

We prove that, under standard tail conditions on the underlying distributions, $m = \text{poly}(k, \frac{1}{\epsilon})$ samples are necessary and sufficient. Our lower bound stands in contrast to many recent results on simple and prior-independent auctions and fundamentally involves the interplay between bidder competition, non-identical distributions, and a very close (but still constant) approximation of the optimal revenue. It effectively shows that the only way to achieve a sufficiently good constant approximation of the optimal revenue is through a detailed understanding of bidders' valuation distributions. Our upper bound is constructive and applies in particular to a variant of the empirical Myerson auction, the natural auction that runs the revenue-maximizing auction with respect to the empirical distributions of the samples.

Our sample complexity lower bound depends on the set of allowable distributions, and to capture this we introduce α-strongly regular distributions, which interpolate between the well-studied classes of regular ($\alpha = 0$) and MHR ($\alpha = 1$) distributions. We give evidence that this definition is of independent interest.

1 Introduction

Comparing the revenue of two different auctions requires an analysis framework for trading off performance on different inputs. For instance, in a single-item auction, a second-price auction with a reserve price $r > 0$ will earn more revenue than a second-price auction with no reserve price on some inputs, and less on others. Which auction is better?

The conventional approach in auction theory is Bayesian, or average-case, analysis. That is, bidders' valuations are assumed to be drawn from a distribution, and one auction is defined to be better than another if it has higher expected revenue with respect to this distribution.

*A preliminary version of this paper appeared in the 2014 ACM Symposium on the Theory of Computing.
†This work was supported in part by NSF Awards CCF-1217989 and CCF-1527568.
‡This research was done in part while visiting New York University, and was supported in part by NSF Awards CCF-1016885 and CCF-1215965 and an ONR PECASE Award.

The optimal auction is then the one with the highest expected revenue. The optimal auction depends on the assumed distribution, in some cases in a detailed way.

While there is now a significant body of work on worst-case revenue maximization (see [25]), a majority of modern computer science research on revenue-maximizing auctions uses Bayesian analysis to measure auction performance (see [24]). Since the comparison between auctions depends fundamentally on the assumed distribution, an obvious question is: *where does this prior distribution come from, anyway?*

In most applications, and especially in computer science contexts, the answer is equally obvious: *from past data.* For example, in Yahoo!'s keyword auctions, Bayesian analysis was used to provide guidance on how to set per-click reserve prices, and the valuation distributions used in this analysis are derived straightforwardly from bid data from the recent past [35]. This is a natural approach, but how well does it work?

1.1 The Model

The goal of this paper is to understand how much data is necessary and sufficient to guarantee near-optimal expected revenue. Our model is the following. There are k bidders in a single-item auction. The valuation (i.e., willingness-to-pay) of bidder i is a sample from a distribution F_i. The F_i's are independent but not necessarily identical.

The distribution $\mathbf{F} = F_1 \times \cdots \times F_k$ is unknown to the seller. The "data" comes in the form of m independent and identically distributed (i.i.d.) samples $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(m)}$ from $\mathbf{F}$ — equivalently, m i.i.d. samples from each of the k individual distributions $F_1, \ldots, F_k$. The seller observes the samples and then commits to a truthful auction $\mathcal{A}$.[1] We call this function from samples to auctions an *m-sample auction strategy.* The seller then earns the revenue of its chosen auction $\mathcal{A}$ on the "real" input, a fresh independent sample $\mathbf{v}^{(m+1)}$ from $\mathbf{F}$. See also Figure 1. We can state our main question as follows.

(*) *How many samples m are necessary and sufficient for the existence of an m-sample auction strategy that, for every distribution $\mathbf{F}$ in some class $\mathcal{D}$, has expected revenue at least $(1-\epsilon)$ times that of the optimal auction for $\mathbf{F}$?*

The expected revenue of an auction strategy is with respect to both the samples $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(m)}$ and the input $\mathbf{v}^{(m+1)}$ — i.e., over $m+1$ i.i.d. samples from $\mathbf{F}$. The expected revenue of an optimal auction is with respect to a single sample (the input) from $\mathbf{F}$. Our formalism is inspired by computational learning theory [38].

The answer to the question (*) could be a function of up to three different parameters: the error tolerance ϵ, the number k of bidders, and the set $\mathcal{D}$ of allowable distributions[2]. It is clear that some restriction on $\mathcal{D}$ is necessary for the question (*) to be interesting: without any restriction, no finite number of samples is sufficient to guarantee near-optimal revenue, even when there is only one bidder.[3]

[1] An auction is *truthful* if truthful bidding is a dominant strategy for every bidder. That is: for every bidder i, and all possible bids by the other bidders, i maximizes its expected utility (value minus price paid) by bidding its true value. For single-item auctions, the optimal expected revenue of any (possibly non-truthful) auction, measured at a Bayes-Nash equilibrium with respect to the prior distribution, is no larger than that of the optimal truthful auction. Also, the restriction to dominant strategies is natural given our assumption of an unknown distribution.

[2] As the distribution $\mathbf{F}$ is unknown, we seek uniform sample complexity bounds, meaning bounds that depend only on $\mathcal{D}$ and not on $\mathbf{F}$.

[3] To see this, consider all distributions that take on a value M^2 with probability $\frac{1}{M}$ and 0 with probability $1 - \frac{1}{M}$. The optimal auction for such a distribution earns expected revenue M. It is not difficult to prove that, for every m, there is no m-sample auction strategy with near-optimal revenue for every such distribution — for sufficiently large M, all m samples are 0 with high probability and the auction strategy has to resort to an uneducated guess for M.

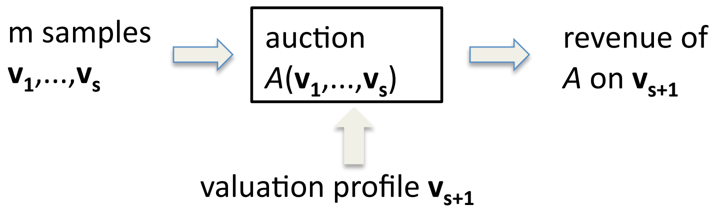

Figure 1: A single-item auction A is chosen as a function of s i.i.d. samples $\mathbf{v}_1, \ldots, \mathbf{v}_s$ from an unknown distribution $F_1 \times \cdots \times F_n$, and applied to a fresh sample $\mathbf{v}_{s+1}$ from the same distribution. The benchmark is the expected revenue of the Myerson-optimal auction for $F_1, \ldots, F_n$.

Our model can be viewed as an interpolation between worst-case and average-case analysis. It inherits much of the robustness of the worst-case model, since we demand guarantees for every underlying F, while allowing very good approximation guarantees with respect to a strong benchmark.

1.2 Distributional Assumptions

Two distributional assumptions that have been extensively used (see e.g. [24]) are the *regularity* and *monotone hazard rate (MHR)* conditions. The former asserts that the "virtual valuation" function $v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ is nondecreasing, where f_i is the density of F_i, while the second imposes the strictly stronger condition that $\frac{f_i(v_i)}{1-F_i(v_i)}$ is nondecreasing. The "most tail-heavy" regular distribution has the distribution function $F_i(v_i) = 1 - \frac{1}{v_i+1}$, while the most tail-heavy MHR distributions are the exponential distributions.

Our lower bound on the sample complexity of revenue maximization depends on the set of allowable distributions, and to capture this we introduce a parameterized condition that interpolates between the regularity and MHR conditions; this condition is also useful in other contexts (see Section 4).

Definition 1.1. (α-Strongly Regular Distribution) *Let F be a distribution with positive density function f on its support $[a, b]$, where $0 \le a < \infty$ and $a \le b \le \infty$. Let $\varphi(v) = v - \frac{1-F(v)}{f(v)}$ denote the corresponding virtual valuation function. F is* α-strongly regular *if*

$$\varphi(y) - \varphi(x) \ge \alpha(y - x) \tag{1}$$

whenever $y > x \ge 0$.

For distributions with a differentiable virtual valuation function φ, condition (1) is equivalent to $\frac{d\varphi}{dv} \ge \alpha$. Regular and MHR distributions are precisely the 0- and 1-strongly regular distributions, respectively. A product distribution $\mathbf{F} = F_1 \times \cdots \times F_k$ is called α-strongly regular if each F_i is α-strongly regular. For the lower bound, we take the set $\mathcal{D}$ of allowable distributions in (*) to be the α-strongly regular distributions for a parameter $\alpha \in [0, 1]$.

1.3 Our Results

Our main result is that $m = \text{poly}(k, \frac{1}{\epsilon})$ samples are necessary and sufficient for the existence of an m-sample auction strategy that, for every strongly regular distribution $\mathbf{F}$, has expected revenue at least $(1 - \epsilon)$ times that of an optimal auction.

Both our upper and lower bounds on the sample complexity of revenue maximization are significant. For the lower bound, it is far from obvious that the number of samples per bidder needs to depend on k at all, let alone polynomially. Indeed, for many relaxations of the problem we study, the sample complexity is a function of ϵ only.

- If there is an unlimited supply of items (digital goods), then the problem reduces to separate single-bidder problems, for which poly$(\frac{1}{\epsilon})$ samples suffice for a $(1-\epsilon)$-approximation for all regular distributions [18, Lemma 4.1].
- If bidder valuations are independent and *identical* draws from an unknown regular distribution, then poly$(\frac{1}{\epsilon})$ samples suffice for a $(1-\epsilon)$-approximation [18, Theorem 4.3].
- If only a $\frac{1}{4}$-approximation of the optimal expected revenue is required, then only a *single* sample is required. This follows from a generalization of the Bulow-Klemperer theorem [7] to non-i.i.d. bidders [26, Theorem 4.4].

Thus, the necessary dependence on k fundamentally involves the interplay between bidder competition, non-identical distributions, and a very close (but still constant) approximation of the optimal revenue.

On a conceptual level, our lower bound shows that designing c-approximate auctions for constants c sufficiently close to 1 is a qualitatively different problem than for more modest constants like $\frac{1}{4}$. For example, previous work has demonstrated that auctions with reasonably good approximation factors are possible with minimal dependence on the valuation distributions (e.g. [9, 26, 2]) or even, when there is no bidder with a unique valuation distribution, with *no* dependence on the valuation distributions [17, 18, 36]. Another interpretation of some previous results, such as [11, 9], is the existence of constant-factor approximate auctions that derive no benefit from bidder competition. Our lower bound identifies, for the first time, a constant approximation threshold beyond which "robustness" and "prior-independence" results of these types cannot extend. Our argument formalizes the idea that, with two or more non-identical bidders, the *only way* to achieve a sufficiently good constant approximation of the optimal revenue is through a detailed understanding of bidders' valuation distributions and an essentially optimal resolution of bidder competition.

We provide an upper bound on the number of samples needed for near-optimal approximation by analyzing a natural auction strategy. Recall that for a distribution $\mathbf{F}$ that is known a priori, Myerson's optimal auction gives the item to the bidder with the highest virtual valuation $\varphi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$, or to no one if all virtual valuations are negative [33]. The *empirical Myerson auction* is the obvious analog when one has data rather than distributional knowledge: define $\bar{F}_i$ as the empirical distribution of the samples from F_i, and run the optimal auction for $\bar{\mathbf{F}}$.[4] We prove that a variant on the empirical Myerson auction has expected revenue at least $(1-\epsilon)$ times optimal provided it is given a sufficiently large polynomial number of samples.[5] A key aspect in our analysis is identifying the (non-pointwise) sense in which empirical virtual valuation functions approximate the actual virtual valuation functions; this is non-trivial even for the special case of MHR distributions.

[4]Since the empirical distributions are generally not regular even when the underlying distributions $\mathbf{F}$ are, a standard extra "ironing" step is required; see Section 2 for details.

[5]Left unmodified, the empirical Myerson mechanism can be led astray by poor approximations at the upper end of the valuation distributions caused by a small sample effect. We prove that excluding the very highest samples from the empirical distributions addresses the problem.

1.4 Technical Approach

The proofs of our upper and lower bounds are fairly technical, so we provide here an overview of the main ideas. We begin with the upper bound, which roughly consists of the following steps.

1. (Lemma 6.2) For some fixed bidder with distribution F, consider the corresponding m samples $v_1 \geq \cdots \geq v_m$. Define the "empirical quantile" $\bar{q}_j$ of v_j as $\frac{2j-1}{2m}$, the expected quantile of the jth order statistic. Taking a "net" of quantiles and applying standard large deviation bounds shows that all but the bottom $\hat{\xi}$ fraction of the empirical quantiles are good multiplicative approximations of their expectations with high probability; here $\hat{\xi} > 0$ is a key parameter that will depend on k and ϵ.

2. (Lemma 6.4) Recall that the expected revenue of an auction equals its expected virtual welfare [33]. Myerson's optimal auction maximizes virtual welfare pointwise, whereas our auction maximizes (ironed) empirical virtual welfare pointwise. In a perfect world, we would be able to argue that the empirical virtual valuation functions are good pointwise approximations of the true virtual valuation functions, and hence the expected virtual welfare of our auction is close to that of Myerson's auction. Unfortunately, good relative approximation of quantiles does not necessarily translate to good relative approximation of virtual valuations. The reason is that a virtual valuation function $v - \frac{1-F(v)}{f(v)}$ can change arbitrarily rapidly in a region where the density changes rapidly (even for MHR distributions).

   We instead prove a different sense in which empirical virtual values approximate actual virtual values, working in the (quantile) domain as well as in the range of the virtual valuation functions. Recall that the quantile $q(v)$ is defined as $1-F(v)$. We show that for suitable $\Delta_1, \Delta_2 > 0$, for all but the top $\hat{\xi}$ fraction of quantiles in $[0,1]$, with high probability the empirical virtual value $\bar{\varphi}(q)$ is sandwiched between $\varphi(q(1+\Delta_1))$ and $\varphi(q/(1+\Delta_2))$, modulo small additive factors. (By $\varphi(q)$ we mean $\varphi(F^{-1}(1-q))$.) The additive factors are functions of $1/q$, as well as of k and ϵ, which complicates the analysis.

3. (Lemmas 6.8 and 6.9) Consider a fixed bidder i. By the previous step, up to additive factors, we can lower bound the virtual welfare contributed by a bidder i with a quantile $q_i = 1 - F_i(v_i)$ outside the top $\hat{\xi}$ fraction in the empirical Myerson mechanism by the virtual value contributed by i in the optimal auction when it has a quantile of $q_i(1+\Delta_1)$. Or not quite: an additional issue is that the empirical virtual valuation of a different bidder j with quantile q_j might be larger than its true virtual value, leading the empirical Myerson auction to allocate to j over the rightful winner i, and resulting in a reduction in the total virtual welfare being accumulated as compared to the actual Myerson auction. The difference in virtual values is bounded by the additive factors described in (2); as the outcomes are determined in the empirical auction, we will need two sets of factors: for bidder i, the factors from the lower sandwiching bound, and for bidder j, the factors from the upper sandwiching bound. We will show that there is only a small probability of large additive factors, which suffices to bound the expected reduction in revenue when the additive factors are large. When the additive factors are small, their contribution to the reduction in revenue is also small. All the reductions end up being polynomial functions of k and ϵ.

4. (Lemma 6.6 and 6.7) There is one more issue. Because of the shift in quantile space — we compare the virtual value in the Myerson auction at quantile $q_i(1+\Delta_1)$ to the virtual value in the empirical auction at quantile q_i — and also because the reserve prices in the two auctions may differ, we also have to analyze the revenue loss at the lower end of

the distributions, or more precisely, around the reserve prices. This too is a polynomial function of k and ϵ.

We now discuss the lower bound proof. This involves arguing that, if the number of samples is too small, then for every auction strategy, there exists a distribution for which the auction strategy's expected revenue is not near-optimal. We prove this by exhibiting a "distribution of distributions" and proving that every auction strategy has expected revenue — where the expectation is now with respect to both the initial random choice of the valuation distributions, and then with respect to both the m samples and the input — bounded away from the expected revenue of an optimal auction (where the expectation is over both the choice of distributions and the input). We are unaware of any other lower bounds in auction theory that have this form.

Our construction involves taking a base set of "worst-case" α-strongly regular distributions and truncating them at random points. A key observation is that, when such a distribution is truncated at a point H_i, the corresponding virtual valuation function is linear with coefficient α except at the truncation point, where the virtual valuation jumps to H_i. The high-level intuition is that, when confronted with valuations that are higher than those seen in any of the samples, no auction can know whether a high valuation v corresponds to a truncation point (with virtual value v) or not (with virtual value only $\alpha(v-1)$). Properly implemented, this idea can be used to prove that every auction strategy errs with constant probability on precisely the set of inputs that contribute the lion's share of the optimal revenue. The lower bound follows.

1.5 Prior and Concurrent Related Work

We provide detailed comparisons only with the papers most closely related to the present work. For previously studied models about revenue-maximization with an unknown distribution, which differ in various respects from our model, see [4, 29]. For other uses of samples in auction design that differ from ours, see Fu et al. [21], who use samples to extend the Crémer-McLean theorem [15] to partially known valuation distributions, and Chawla et al. [10], who design non-truthful auctions that both have equilibrium revenue within a constant of optimal and enable accurate inference about the valuation distribution from samples. For other ways to parameterize partial knowledge about valuations, see e.g. Azar et al. [3] and Chiesa et al. [12]. For asymptotic optimality results in various symmetric settings, which identify conditions under which the expected revenue of some auction of interest (e.g., second-price) approaches the optimal with an increasing number of (i.i.d.) bidders, see Neeman [34], Segal [37], Baliga and Vohra [6], and Goldberg et al. [22]. For applications of learning theory concepts to prior-free auction design in unlimited-supply settings, see Balcan et al. [5]. Finally, the technical issue of ironing from samples comes up also in Ha and Hartline [23], in the context of prior-free mechanism design, and the aforementioned Chawla et al. [10]. Our goal of obtaining a $(1-\epsilon)$-approximation of the maximum revenue achieved by any auction is impossible in the more demanding settings of [23, 10].

Elkind [20] studies a learning problem closely related to ours, in the restricted setting of discrete distributions with known finite supports but with unknown probabilities. In the model in [20], learning is done using an oracle that compares the expected revenue of pairs of auctions, and $O(n^2K^2)$ oracles calls suffice to determine the optimal auction (where n is the number of bidders and K is the support size of the distributions). Elkind [20] notes that such oracle calls can be implemented approximately by sampling (with high probability), but no specific sample complexity bounds are stated.

Dhangwatnotai et al. [18], motivated by "prior-independent" auctions that are simultaneously approximately optimal for a wide range of valuation distributions, implicitly studied the single-bidder version of the learning problem we that study. With one bidder, the goal is to learn approximately the monopoly price of an unknown distribution from samples. Their results imply sample complexity upper bounds for this problem of $O(\epsilon^{-2})$ and $O(\epsilon^{-3})$ for MHR and regular distributions, respectively.[6] As our results show, the learning problem is quite different and more delicate with multiple non-i.i.d. bidders.

The papers of Cesa-Bianchi et al. [8] and Medina and Mohri [30] give algorithms for learning the optimal reserve-price-based single-item auction. Our problem of learning the best single-item auction — whether reserve-based or otherwise — is harder. With non-i.i.d. bidders, there need not be a reserve price for which the second-price auction has expected revenue more than 50% times the optimal [26].

Concurrently with our work, Dughmi et al. [19] proved negative results (exponential sample complexity) for learning near-optimal mechanisms in multi-parameter settings that are much more complex than the single-item auctions studied here. The paper also contains positive results for restricted classes of mechanisms.

1.6 Subsequent Related Work

The preliminary version of this paper [13] motivated several follow-up works. Huang et al. [28] study the single-bidder version of our problem, studied implicitly in [18], and give optimal sample complexity bounds under several different distributional assumptions. Both the upper and lower bounds in [28] improve, in terms of the dependence on ϵ^{-1}, over those implied by the present work. This is also the only paper other than the present work that proves any sample complexity lower bounds.

Morgenstern and Roughgarden [32] adapt tools from statistical learning theory [1] to give general sample complexity upper bounds that cover all single-parameter settings for bounded or MHR valuation distributions. The results in [32] apply to many more environments than single-item auctions, and improve over the sample complexity upper bound of the present work (even for single-item auctions), but unlike the present work, their results do not apply to regular distributions and do not result in computationally efficient learning algorithms.

Very recently, Devanur et al. [16] devised a different learning algorithm for our problem which provides a strict improvement on our upper bound. The sample complexity bound in [16] is roughly the same as in [32], but the learning algorithm is computationally efficient and also accommodates regular distributions.

Finally, Cole and Shravas investigated further applications of α-strongly regular distributions [14].

2 Preliminaries

This section reviews Myerson's optimal single-item auction [33] for the case of known distributions. There are k bidders, and for each bidder i there is a distribution F_i from which its valuation is drawn.[7]

For each buyer i, the auctioneer computes a virtual valuation $\varphi_i(v) = v - [1 - F_i(v)]/f_i(v)$, where f_i is the density function corresponding to F_i. $\varphi_i(v)$ is required to be a non-decreasing

[6]Similarly, our sample complexity upper bound naturally leads to a prior-independent single-item auction. This auction achieves a $(1-\epsilon)$-approximation of the optimal auction when bidders' valuations are drawn from different regular distributions $F_1, \dots, F_k$ and there are sufficiently many bidders of each type.

[7]Our results extend to the case of k groups of an arbitrary number of bidders, where all the bidders from group i have i.i.d. valuations drawn from F_i. Then the sample bounds are a function of the number n of bidders, rather than the number k of groups, but the bound is on the number of samples needed from each group of bidders.

function of v. This holds by definition for regular distributions; in general, if this does not hold φ_i can be modified (or *ironed*) so that it does hold, as implicitly explained in the next paragraph. Next, the auctioneer runs an analog of a second-price auction on the virtual values of the bids (virtual bids for short): the bidder, if any, with the highest non-negative virtual bid wins the auction (ties are broken arbitrarily) and is charged the minimum bid needed to win (or at least to tie for winning). More precisely, let i be the winning bidder and let b_2 be the second highest virtual bid. Then the price is ${\varphi_i}^{-1}(\min\{0, b_2\})$. We note that ${\varphi_i}^{-1}(0)$ can be viewed as a bidder-specific reserve price for i; it is also called the *monopoly price* for i.

We can also describe the auction in terms of a revenue function. This also allows for situations where $\varphi_i(v)$ is not a nondecreasing function of v. The revenue function is computed in quantile space: $q_i(v) = 1 - F_i(v)$ is the probability that i will have a valuation of at least v. Now we view v as a function of q_i. We introduce the expected revenue function, $R_i(q_i)$. It is a function of the quantile q_i: $R_i(q_i) = v(q_i) \cdot q_i$ is the expected revenue if i is the sole bidder and $v(q_i)$ is the price being charged. The auctioneer computes the smallest concave upper bound $\mathrm{CR}_i(q)$ of $R_i(q)$. Now $\varphi_i(v(q_i))$ is defined to be the slope of $\mathrm{CR}_i(q_i)$ (this yields an increasing function φ_i, which coincides with the previous definition in the case of regular distributions). At points where there is no unique slope we choose $\varphi_i(v(q)) = \lim_{(q'>q)\to q} \varphi_i(v(q'))$. The auction then proceeds as before. Henceforth, overloading notation, we write $\varphi_i(q_i)$ rather than $\varphi_i(v(q_i))$.

Myerson [33] proved that for every auction, the expected virtual welfare equals the expected revenue. This result is important because, in many situations, it is much easier to reason about expected virtual welfare than directly about expected revenue.

Theorem 2.1 (Myerson)**.** *The expected revenue of any single-item auction is given by*

$$\sum_{i=1}^{k} E_{q_i}[\varphi_i(q_i) \cdot x_i(q_i)],$$

where $x_i(q_i)$ is the probability (over others' valuations and any coin flips by the auction) that i wins the item with a bid at quantile q_i in F_i.

Let $\mathbf{q} = (q_1, q_2, \ldots, q_k)$ be a vector of quantiles drawn from $F_1 \times F_2 \times \ldots \times F_k$. We can rewrite the expected revenue as

$$\sum_{i=1}^{k} E_{q_i}[\varphi_i(q_i) \cdot x_i(q_i)], = \sum_i \int_{\mathbf{q}} \varphi_i(q_i) I_i(\mathbf{q})\ d\mathbf{q}, \tag{2}$$

where $I_i(\mathbf{q})$ is the indicator function showing whether i wins when the bids are at quantiles $\mathbf{q}$ (or more generally, the probability that it wins). This immediately implies that allocating to a bidder with the highest virtual value, i.e. Myerson's auction, is optimal.

3 Statement of Main Results

We formally state our upper and lower bound results in turn.

Theorem 3.1. *In a single-item auction with k bidders with independent regular valuation distributions, if $m = \Omega(\frac{k^{10}}{\epsilon^7} \ln^3 \frac{k}{\epsilon})$, then there is an m-sample auction strategy with expected revenue at least $1 - \epsilon$ times that of an optimal auction.*

The auction strategy in Theorem 3.1 is a variant on the "empirical Myerson auction," described in detail in Section 6.1.

Our lower bound result has an analogous form, although the polynomial in k and ϵ is considerably smaller. Our lower bound grows larger as $\alpha \geq 0$ grows smaller.

Theorem 3.2. *For every auction strategy Σ, for every $k \geq 2$, for every sufficiently small $\epsilon > 0$, for every $\alpha \geq 0$ and m satisfying:*

i. $\alpha = 1$ *and* $m \leq \left(\frac{1-\ln 2}{96e^3 \min\{1,\frac{k}{e}\} \ln \max\{e,k\}}\right)^{1/2} \frac{k}{\sqrt{\epsilon}}$;

ii. $0 < \alpha < 1$, $\alpha^{1/(1-\alpha)} \geq \frac{1}{k}$, *and* $m \leq \left(\frac{1-\alpha 2^{1-\alpha}}{96e^3}\right)^{1/(1+\alpha)} \frac{k}{\epsilon^{1/(1+\alpha)}}$;

iii. $0 < \alpha < 1$, $\frac{1}{2m} < \alpha^{1/(1-\alpha)} < \frac{1}{k}$, *and* $m \leq \left(\frac{1-\alpha 2^{1-\alpha}}{96e^3}\right)^{1/(1+\alpha)} \left(\frac{1}{k\alpha^{1/(1-\alpha)}}\right)^{\alpha/(1+\alpha)} \frac{k}{\epsilon^{1/(1+\alpha)}}$;

iv. $0 < \alpha < 1$, $\alpha^{1/(1-\alpha)} \leq \frac{1}{2m}$, *and* $m \leq \frac{(1-\alpha 2^{1-\alpha})2^\alpha}{96e^3} \frac{k}{\epsilon}$;

v. $\alpha = 0$ *and* $m \leq \frac{1}{96e^3} \frac{k}{\epsilon}$,

there exists a set $F_1, \ldots, F_k$ of α-strongly regular valuation distributions such that the expected revenue of Σ (over the m samples and the input) is less than $1 - \epsilon$ times that of an optimal auction for $F_1, \ldots, F_k$.

At the two extreme points of MHR distributions ($\alpha = 1$) and regular distributions ($\alpha = 0$), the lower bound in Theorem 3.2 is $\Omega(\frac{k}{\sqrt{\epsilon \ln k}})$ and $\Omega(\frac{k}{\epsilon})$, respectively. In all cases, the dependence on the number of bidders is linear or near-linear.

We prove these two theorems in Sections 6 and 5 respectively. Before that, in the next section, we briefly indicate two applications of α-strong regularity.

4 Applications of Strong Regularity

We believe our definition of α-strongly regular distributions is of independent interest. Almost all previous expected revenue approximation guarantees for auctions in Bayesian settings apply to one of three sets of valuation distributions: all distributions, all regular distributions, or all MHR distributions (see e.g. [24]). Strongly regular distributions interpolate between regular and MHR distributions, and should broaden the reach of many existing approximation bounds that are stated only for MHR distributions. To prove this point, we mention a couple of examples of such extensions; we are confident that many others are possible, as has been shown subsequently by Cole and Shravas [14].

The following property of MHR distributions is well known [27, Lemma 4.1].

Lemma 4.1 ([27])**.** *Let F be an MHR distribution with monopoly price r. If $q(r)$ is the quantile of valuation r in the distribution F, then $q(r) \geq \frac{1}{e}$.*

We next show how to generalize this result to α-strongly regular distributions.

Lemma 4.2. *Let F be an α-strongly regular distribution with $\alpha \in (0, 1)$ and monopoly price r. If $q(r)$ is the quantile of valuation r in the distribution F, then $q(r) \geq \alpha^{1/(1-\alpha)}$.*

Proof. Set $\lambda = 1 - \alpha$. Let $h(\cdot)$ denote the hazard rate of F, and choose c so that $h(r) = \frac{1}{\lambda r + c}$. Recall that $\varphi(v) = v - \frac{1}{h(v)}$. Since $\varphi(r) = 0$ and hence $h(r) = 1/r$, we have $c = r(1 - \lambda)$.

The α-strong-regularity condition, $\frac{d\varphi}{dv} \geq \alpha$, implies that $1 + \frac{1}{h^2}\frac{dh}{dv} \geq \alpha$, or $\frac{d}{dv}\left(\frac{1}{h}\right) \leq \lambda$. It follows that, for all $v \leq r$, $h(v) \leq \frac{1}{\lambda v + c}$ and hence

$$h(v) \leq \frac{1}{\lambda(v - r) + r}.$$

Now write $H(x) = \int_0^x h(v)dv$; it is well known and easy to verify that $q(v) = e^{-H(v)}$. We complete the proof by deriving

$$
\begin{aligned}
q(r) &= e^{-H(r)} = e^{-\int_0^r h(v)dv} \\
&\geq e^{-\left[\frac{1}{\lambda}\log(r+\lambda(v-r))\,|_0^r\right]} \\
&= e^{-\frac{1}{\lambda}\log\frac{r}{r(1-\lambda)}} = e^{\log(1-\lambda)^{1/\lambda}} = (1-\lambda)^{1/\lambda} = \alpha^{1/(1-\alpha)}.
\end{aligned}
$$

□

Hartline et al. [27, Theorem 4.2] study a revenue maximization problem in social networks, and give a mechanism with approximation guarantee

$$\frac{1}{4-\frac{2}{e}} \approx .306 \tag{3}$$

when players' private valuations are drawn from MHR distributions.[8] The MHR assumption is used only in applying Lemma 4.1. Relaxing the distributional assumption to α-strong regularity and reoptimizing the proof in [27, Theorem 4.2] using Lemma 4.2 extends this approximation guarantee accordingly, with the term $1/e$ in (3) replaced by $\alpha^{1/(1-\alpha)}$.

For a second example, Hartline and Roughgarden [26, Theorem 3.2] consider downward-closed single-parameter environments[9] and prove that, when bidders' valuations are drawn from MHR distributions, the VCG mechanism with "eager" monopoly reserve prices[10] has expected revenue at least $\frac{1}{2}$ times that of an optimal mechanism. For an α-strongly regular distribution F with monopoly price r, and $v \geq r$, we have $\varphi(v) \geq \alpha(v-r)$ and hence

$$r + \frac{1}{\alpha}\varphi(v) \geq v;$$

this inequality generalizes Lemma 3.1 in [26]. Following the proof in [26, Theorem 3.2] shows that, for every downward-closed single-parameter environment with bidders valuations drawn from α-strongly regular distributions, the VCG mechanism with eager monopoly reserves has expected revenue at least $\frac{\alpha}{\alpha+1}$ times that of an optimal mechanism.

5 The Lower Bound: Proof of Theorem 3.2

Formal Statement Fix $\alpha \geq 0$ and $0 < \delta \leq 1$, where δ is sufficiently small. We show that for every auction strategy Σ, there exists a set $F_1, \ldots, F_k$ of α-strongly regular distributions such that the expected revenue of the auction strategy (over the samples and the input) is at most

[8] They also give a $\frac{1}{4}$-approximation algorithm for arbitrary valuation distributions.

[9] A (binary) single-parameter environment is specified by a set of bidders and the feasible subsets of bidders that can simultaneously win. For example, if each bidder i wants a known bundle S_i of items, then the feasible subsets are those in which the bundles of the chosen bidders are pairwise disjoint. Such an environment is downward closed if every subset of a feasible set is again feasible.

[10] In more detail, one reserve price r_i per bidder i is fixed in advance. A bid is collected from each bidder. Bidders who bid below their reserve prices are removed from further consideration. From the remaining bidders, the mechanism chooses winners to maximize the sum of their bids, subject to feasibility. The mechanism charges the unique prices for which losing bidders pay 0 and truthful bidding is a dominant strategy. See [26] for details.

the following fraction of the expected revenue of the optimal auction for $F_1, \ldots, F_k$:

$$
\begin{aligned}
1-\epsilon(1,\delta) &= 1-\frac{1-\ln 2}{96e^3\min\{1,\frac{k}{e}\}\ln\max\{e,k\}}\delta^2 && \text{for } \alpha=1\\
1-\epsilon(\alpha,\delta) &= 1-\frac{1-\alpha 2^{1-\alpha}}{96e^3}\delta^{1+\alpha} && \text{for } \alpha<1 \text{ and } \tfrac{1}{k}\le \alpha^{1/(1-\alpha)}\\
1-\epsilon(\alpha,\delta) &= 1-\frac{1-\alpha 2^{1-\alpha}}{96e^3}\delta^{1+\alpha}\frac{1}{(k\alpha^{1/(1-\alpha)})^{\alpha}} && \text{for } \alpha<1 \text{ and } \tfrac{\delta}{2k}<\alpha^{1/(1-\alpha)}<\tfrac{1}{k}\\
1-\epsilon(\alpha,\delta) &= 1-\frac{1-\alpha 2^{1-\alpha}}{96e^3}2^{\alpha}\delta && \text{for } \alpha<1 \text{ and } \alpha^{1/(1-\alpha)}\le\tfrac{\delta}{2k}\\
1-\epsilon(\alpha,\delta) &= 1-\frac{1}{96e^3}\delta && \text{for } \alpha=0
\end{aligned}
$$

We note that if $\alpha < 1$, then $\alpha 2^{1-\alpha} < 1$ also. In addition, for fixed k, $\lim_{\alpha\to 0}[k\alpha^{1/(1-\alpha)}]^{\alpha} = 1$. Substituting k/m for δ yields the bounds in Theorem 3.2. For $\frac{1}{k} \le \alpha < 1$ and sufficiently small constant $\epsilon > 0$, $\Omega(k/\epsilon^{1/(1+\alpha)})$ samples are necessary for a $(1-\epsilon)$-approximation. For the MHR ($\alpha = 1$) case, $\Omega(k/\sqrt{\epsilon \ln k})$ samples are necessary, and for the regular ($\alpha = 0$) case, $\Omega(k/\epsilon)$ samples are needed.

The Base Distributions We identify the worst-case distributions for a given $\alpha \ge 0$. Specifically, for $v \in [0,\infty)$, consider

$$
\begin{aligned}
F^{\alpha}(v) &= 1-\left(\frac{1}{1+(1-\alpha)v}\right)^{\frac{1}{1-\alpha}} && \text{for } 0\le\alpha<1;\\
F^{1}(v) &= 1-e^{-v} && \text{for } \alpha=1;\\
f^{\alpha}(v) &= \left(\frac{1}{1+(1-\alpha)v}\right)^{\frac{2-\alpha}{1-\alpha}} && \text{for } 0\le\alpha<1;\\
f^{1}(v) &= e^{-v}, && \text{for } \alpha=1.
\end{aligned}
$$

The corresponding hazard rates are

$$
\begin{aligned}
h^{\alpha}(v) &= \frac{1}{1+(1-\alpha)v} && \text{for } 0\le\alpha<1;\\
h^{1}(v) &= 1 && \text{for } \alpha=1;
\end{aligned}
$$

with virtual valuation

$$\varphi^{\alpha}(v) = \alpha v - 1 \quad \text{for } 0 \le \alpha \le 1.$$

A quick calculation shows that

$$
(\mathcal{F}^{\alpha})^{-1}(q) = \begin{cases} \frac{1}{1-\alpha}\left[\left(\frac{1}{q}\right)^{1-\alpha} - 1\right] & \text{if } 0\le\alpha<1\\ \ln\frac{1}{q} & \text{if } \alpha=1. \end{cases} \tag{4}
$$

The Construction We define a distribution over distributions. Each bidder i is either type A or type B (50/50 and independently). For a type B bidder i, we draw q uniformly from the interval $[\frac{\delta}{2k}, \frac{\delta}{k}]$ and set $H_i = (F^{\alpha})^{-1}(1-q)$. We then define bidder i's distribution F_i as equal to F^{α} on $[0, H_i)$ with a point mass with the remaining probability $1 - F^{\alpha}(H_i)$ at H_i. For a type A bidder we proceed similarly except that H_i is set to $(F^{\alpha})^{-1}(1-\frac{\delta}{2k})$. These distributions are always α-strongly regular. An important point is that the virtual valuation of these bidders is given by

$$\varphi(v) = \begin{cases} \alpha v - 1 & \text{if } v < H_i \\ H_i & \text{if } v = H_i. \end{cases} \tag{5}$$

Ultimately, it is the gap in virtual valuation between these two cases that is responsible for the lower bound in Theorem 3.2.

A Preliminary Lemma Let q_α^A denote the monopoly price in an auction with a single type A bidder. We define $q_0^A = \frac{\delta}{2k}$, as this is the largest quantile q for which $\varphi(q) \geq 0$ when $\alpha = 0$. For $\alpha > 0$ we begin by determining the monopoly price q_α for F^α. We note that $\varphi(v) = 0$ when $v = \frac{1}{\alpha}$. From (4), we deduce this occurs at $q_\alpha = \alpha^{1/(1-\alpha)}$ for $0 < \alpha < 1$, and at $q_1 = \frac{1}{e}$ for $\alpha = 1$. Thus, for bidder i, for $0 < \alpha < 1$, $q_\alpha^A = \max\{\frac{\delta}{2k}, \alpha^{1/(1-\alpha)}\}$, and for $\alpha = 1$, $q_1^A = \max\{\frac{\delta}{2k}, \frac{1}{e}\} = \frac{1}{e}$ (as we can assume that $k \geq 2$, $\delta \leq 1$).

Let $F^{A,\alpha}$ denote the distribution of a type A bidder. Now let $v^* = (F^{A,\alpha})^{-1}(\max\{1 - q_\alpha^A, \frac{k-1}{k}\})$, the value corresponding to quantile $\min\{q_\alpha^A, \frac{1}{k}\}$ in $F^{A,\alpha}$, and let $R^* = \min\{kq_\alpha^A, 1\} \cdot v^*$, k times the revenue at this quantile. From (4), we obtain that

$$v^* = \begin{cases} \frac{1}{1-\alpha}\left[\max\left\{\frac{1}{q_\alpha^A}, k\right\}^{1-\alpha} - 1\right] & \text{if } 0 \leq \alpha < 1 \\ \ln\max\left\{\frac{1}{q_\alpha^A}, k\right\} = \ln\max\{e, k\} & \text{if } \alpha = 1. \end{cases} \tag{6}$$

Lemma 5.1. (Upper Bound on Optimal Revenue). *The expected revenue (over* $\mathbf{v}$*) of the optimal auction (with respect to the* H_i*'s) is at most* R^*.

Proof. First, the expected revenue of the optimal auction is upper bounded by that of the optimal auction for the case where all H_i's are $(F^\alpha)^{-1}(1 - \frac{\delta}{2k})$ — i.e., where $F_i = F^{A,\alpha}$ for every i. Second, by symmetry, when bidders valuations are i.i.d. draws from $F^{A,\alpha}$, every bidder has the same purchase probability q in the (symmetric) optimal auction, and since there is only one item, this purchase probability q is at most $\frac{1}{k}$; it is also at most q_α^A. Third, we obtain an upper bound by dropping the constraint of selling only one item and instead optimally selling to each bidder with probability q. Fourth, this is precisely k times the revenue of selling to a single bidder with valuation from $F^{A,\alpha}$ using the posted price $(F^{A,\alpha})^{-1}(1 - q)$. Fifth, by regularity, selling to a single bidder with posted price $(F^{A,\alpha})^{-1}(1 - q)$ with $q \leq \min\{q_\alpha^A, \frac{1}{k}\}$ is no better than selling with posted price $v^* = (F^{A,\alpha})^{-1}(\max\{1 - q_\alpha^A, \frac{k-1}{k}\})$. The expected revenue from any one bidder is therefore at most the sale probability times v^*, namely $\min\{q_\alpha^A, \frac{1}{k}\} \cdot v^*$. The overall revenue, with k bidders, is thus at most $k \times \min\{q_\alpha^A, \frac{1}{k}\} \cdot v^* = R^*$, as claimed. □

Overview of Proof The high-level plan is the following. Fix an arbitrary auction strategy. Think of the random choices as occurring in three stages: in the first stage, the F_i's are chosen; in the second stage, m sample valuation profiles $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(m)}$ are chosen (i.i.d. from $F_1 \times \cdots \times F_k$); in the final stage, the input $\mathbf{v}$ is chosen (independently from $F_1 \times \cdots \times F_k$). We prove that the expected revenue of the auction strategy (with respect to all three stages of randomness) is at most $1 - \epsilon(\alpha, \delta)$ times that of the optimal auction (with respect to all three stages or, equivalently, the first and third stages only).[11] Again, $1 - \epsilon(\alpha, \delta) < 1$ will be independent of k. This implies that, for every auction strategy, there exists a choice of $F_1, \ldots, F_k$ such that the expected revenue of the auction strategy is at most $1 - \epsilon(\alpha, \delta)$ times the expected revenue of the optimal auction for the distributions $F_1, \ldots, F_k$.

[11] We prove this statement about the expected virtual welfare, which is equivalent by Theorem 2.1.

By Lemma 5.1, R^* is an upper bound on the optimal auction's expected revenue (equivalently, expected virtual welfare) for every choice of $F_1, \ldots, F_k$. The main argument is the following: there is an event $\mathcal{E}$ such that, for every auction strategy:

(i) the probability of $\mathcal{E}$ (over all three stages of randomness) is lower bounded by a function $\gamma(\delta)$ of δ (and independent of k and α);

(ii) given $\mathcal{E}$, the expected virtual welfare of the auction strategy is at least $\epsilon(\alpha, \delta)R^*$ smaller than that of the optimal auction, where $\epsilon(\alpha, \delta) > 0$ is a function of α and δ only.

Since by (2), for each set of bids, the virtual welfare earned by the optimal auction is always at least that of the auction strategy, (i)–(ii) imply that the expected virtual welfare (and hence revenue) of the optimal auction exceeds that of the auction strategy by $\epsilon(\alpha, \delta)R^*$ for some constant $\epsilon(\alpha, \delta)$ depending on α and δ. Lemma 5.1 then implies that the auction strategy's expected revenue is at most $1 - \epsilon(\alpha, \delta)$ times optimal.

The Main Argument To define the event $\mathcal{E}$, we use the principle of deferred decisions. We can flip the second- and third-stage coins before those of the first stage by sampling quantiles — $(m+1)k$ i.i.d. draws $\{q_i^{(j)}\}$ from the uniform distribution on [0,1]. (Once the distributions are chosen in the first stage, the valuation $v_i^{(j)}$ is just $F_i^{-1}(1 - q_i^{(j)})$.) We further break the first-stage coin flips into two substages; in the first, we determine bidder types (A and B); in the second, we choose H_i's for the type-B bidders. The event $\mathcal{E}$ is defined as the set of coin flips (across all stages) that meet the following criteria:

(P1) There are exactly two quantiles of the form $q_i^{(m+1)}$ that are at most $\frac{\delta}{k}$, say of bidders j and ℓ;

(P2) $q_j^{(m+1)}$ and $q_\ell^{(m+1)}$ are greater than $\frac{\delta}{2k}$;

(P3) for $i = 1, 2, \ldots, m$, $q_j^{(i)}$ and $q_\ell^{(i)}$ are greater than $\frac{\delta}{k}$;

(P4) one of the bidders j, ℓ is type A, the other is type B (we leave random which is which);

(P5) the type B bidder (from among j, ℓ) has valuation equal to the maximum valuation from its distribution.

The next lemma corresponds to step (i) in the proof approach above.

Lemma 5.2. *The probability of $\mathcal{E}$ (over all three stages of randomness) is lower bounded by $\frac{\delta^2}{32e^3}$.*

Proof. We first sample the k quantiles corresponding to the third stage. Elementary computations show that property (P1) holds with probability at least $\frac{1}{2e}\delta^2$ (independent of α and k). Conditioned on (P1) holding, (P2) holds with probability $\frac{1}{4}$. (P3) is independent of the first two properties, as it depends only on the second-stage randomness, and it holds with constant probability of at least $\frac{1}{e^2}$ (independent of α, k). Proceeding to the first stage, (P4) is independent of the first three properties and holds with 50% probability. Conditioned on (P1), (P2), and (P4) (as (P3) is irrelevant), the probability of (P5) equals the probability that a uniform draw from $[\frac{\delta}{2k}, \frac{\delta}{k}]$ (used to determine the H-value) is at least the q-value of the type B bidder, which is conditionally distributed uniformly on $(\frac{\delta}{2k}, \frac{\delta}{k}]$. This happens with probability $\frac{1}{2}$. We conclude that all of (P1)–(P5) hold with a positive probability, namely

$$\gamma(\delta) = \frac{\delta^2}{32e^3}.$$

□

To work toward statement (ii), we next prove that, for every auction strategy, conditioned on $\mathcal{E}$, the strategy fails to allocate the item to the optimal bidder — the type-B bidder with its maximum-possible valuation — with constant probability. It suffices to analyze the auction strategy that, conditioned on $\mathcal{E}$, maximizes the probability (over the remaining randomness) of allocating to the optimal bidder — of guessing, from among the two bidders j, ℓ that in $\mathbf{v}^{(m+1)}$ have valuation at least $(F^\alpha)^{-1}(1-\frac{\delta}{k})$, which one is type A and which one is type B. Since the two bidders were symmetric ex ante, Bayes' rule implies that the probability of guessing correctly (given $\mathcal{E}$) is maximized by, for every $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(m+1)}$, choosing the scenario that maximizes the likelihood of the valuation profiles $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(m+1)}$ (given $\mathcal{E}$).

Lemma 5.3. *Every auction strategy, conditioned on $\mathcal{E}$, allocates to a non-optimal bidder with probability at least $\frac{1}{3}$.*

Proof. The only valuations that affect the relative likelihoods of the two scenarios are $v_j^{(m+1)}$ and $v_\ell^{(m+1)}$. We already know the optimal bidder is either j or ℓ. Property (P3) of event $\mathcal{E}$ implies that the m sample valuations from j and ℓ are equally likely to be generated under the two scenarios — the distributions of type-A and type-B bidders differ only for quantiles in $[0, \frac{\delta}{k}]$.

Now, given $v_j^{(m+1)}$ and $v_\ell^{(m+1)}$, the posterior probabilities of the two scenarios are *not* equal. The reason is that, conditioned on $\mathcal{E}$, the type-A bidder's valuation is distributed according $(F^\alpha)^{-1}(q)$ where q is uniform in $[\frac{\delta}{2k}, \frac{\delta}{k}]$, while the type-B bidder's valuation is distributed according to the smaller of two i.i.d. such samples.[12] Thus, assigning the item to the bidder of j, ℓ with the lower valuation (in $\mathbf{v}^{(m+1)}$) maximizes the probability of allocating to the optimal (type-B) bidder. The probability that this allocation rule erroneously allocates the item to the type-A bidder is the probability that a sample for a distribution (the type-A bidder) is smaller than the minimum of two other samples from the same distribution (the type-B bidder), which is precisely $\frac{1}{3}$. □

The following lemma completes the proof of Theorem 3.2.

Lemma 5.4. *The revenue of every auction strategy is at most the following fraction of an optimal auction's revenue:*

$$
\begin{array}{ll}
1 - \dfrac{1-\ln 2}{96e^3} \dfrac{1}{\ln \max\{e, k\}} \min\{1, \dfrac{k}{e}\} \delta^2 & \textit{if } \alpha = 1 \\
1 - (1 - \alpha 2^{1-\alpha}) \dfrac{1}{96e^3} \delta^{1+\alpha} & \textit{if } 0 < \alpha < 1 \textit{ and } \dfrac{1}{k} \le q_\alpha = \alpha^{1/(1-\alpha)} \\
1 - (1 - \alpha 2^{1-\alpha}) \dfrac{1}{96e^3} \delta^{1+\alpha} \dfrac{1}{(k\alpha^{1/(1-\alpha)})^\alpha} & \textit{if } 0 < \alpha < 1 \textit{ and } \dfrac{\delta}{2k} < q_\alpha = \alpha^{1/(1-\alpha)} < \dfrac{1}{k} \\
1 - (1 - \alpha 2^{1-\alpha}) \dfrac{2^\alpha}{96e^3} \delta & \textit{if } 0 < \alpha < 1 \textit{ and } q_\alpha = \alpha^{1/(1-\alpha)} \le \dfrac{\delta}{2k} \\
1 - \dfrac{1}{96e^3} \delta & \textit{if } \alpha = 0.
\end{array}
$$

[12]In more detail, consider a type-B bidder i and condition on the event that its quantile $q_i = 1 - F_i(v_i)$ is in $[\frac{\delta}{2k}, \frac{\delta}{k}]$ and that its valuation is its maximum possible, which is equivalent to the condition that its fictitious quantile q_i' that generates its threshold H_i lies in $[q_i, \frac{\delta}{k}]$. The joint distribution of (q_i, q_i') is the same as the process that generates two i.i.d. draws from $[\frac{\delta}{2k}, \frac{\delta}{k}]$ and assigns q_i and q_i' to the smaller and larger one, respectively. Note that the valuation of the bidder is, by definition, $(F^\alpha)^{-1}(1 - q_i')$.

Proof. Condition on the event $\mathcal{E}$. By (5), the virtual value φ_B of the type B bidder i equals $H_i \geq (F^\alpha)^{-1}(\frac{\delta}{k})$; substituting $q = \frac{\delta}{k}$ in (4) yields a lower bound on H_i which implies that

$$\varphi_B \;\geq\; \begin{cases} \frac{1}{1-\alpha}\left[\left(\frac{k}{\delta}\right)^{1-\alpha} - 1\right] & \text{if } \alpha < 1 \\ \ln\frac{k}{\delta} & \text{if } \alpha = 1. \end{cases}$$

For the type A bidder, by (5), the virtual value φ_A is at most $\alpha \cdot (F^\alpha)^{-1}(\frac{\delta}{2k}) - 1$; using (4) again, this implies

$$\varphi_A \;\leq\; \begin{cases} \alpha\left[\frac{1}{1-\alpha}\left[\left(\frac{2k}{\delta}\right)^{1-\alpha} - 1\right]\right] - 1 \;=\; \frac{\alpha\cdot 2^{1-\alpha}}{1-\alpha}\left(\frac{k}{\delta}\right)^{1-\alpha} - \frac{1}{1-\alpha} & \text{if } \alpha < 1 \\ \ln\frac{2k}{\delta} - 1 & \text{if } \alpha = 1. \end{cases}$$

Thus, still conditioned on $\mathcal{E}$,

$$\varphi_B - \varphi_A \;\geq\; \begin{cases} \frac{1}{1-\alpha}\left(\frac{k}{\delta}\right)^{1-\alpha}\left(1 - \alpha\cdot 2^{1-\alpha}\right) & \text{if } \alpha < 1 \\ 1 - \ln 2 & \text{if } \alpha = 1. \end{cases}$$

We now bound the fractional loss of revenue. By Lemma 5.2, $\mathcal{E}$ occurs with probability at least $\delta^2/(32e^3)$. By Lemma 5.3, conditioned on $\mathcal{E}$, a type A rather than a type B bidder is wrongly allocated the item with probability $\frac{1}{3}$. Thus the expected loss of revenue is at least

$$\frac{1}{3}\frac{\delta^2}{32e^3}(\varphi_B - \varphi_A).$$

Recall from Lemma 5.1 that the optimal revenue is bounded above by $R^* = k \times \min\{q^A_\alpha, \frac{1}{k}\} \cdot v^*$, where $v^* = (F^{A,\alpha})^{-1}(\max\{1 - q^A_\alpha, \frac{k-1}{k}\})$ is the value corresponding to quantile $\min\{q^A_\alpha, \frac{1}{k}\}$ in $F^{A,\alpha}$. Recalling (6), we can lower bound the fractional loss of revenue at follows.

If $\alpha \in [0, 1)$ and $q^A_\alpha \geq \frac{1}{k}$, then the fractional loss of revenue is at least

$$\frac{\delta^2}{3\cdot 32e^3}\frac{\varphi_B - \varphi_A}{R^*} \;=\; \frac{1}{3\cdot 32e^3}(1 - \alpha 2^{1-\alpha})\frac{\frac{1}{1-\alpha}\left[\left(\frac{k}{\delta}\right)^{1-\alpha}\right]\delta^2}{\frac{1}{1-\alpha}(k^{1-\alpha} - 1)} \;\geq\; \frac{1 - \alpha 2^{1-\alpha}}{96e^3}\delta^{1+\alpha}.$$

If $\alpha \in [0, 1)$ and $q^A_\alpha \leq \frac{1}{k}$, then the fractional loss of revenue is at least

$$\begin{aligned} \frac{1}{3\cdot 32e^3}(1 - \alpha 2^{1-\alpha})\frac{\frac{1}{1-\alpha}\left[\left(\frac{k}{\delta}\right)^{1-\alpha}\right]\delta^2}{\frac{1}{1-\alpha}kq^A_\alpha\left[\left(\frac{1}{q^A_\alpha}\right)^{1-\alpha} - 1\right]} \;&\geq\; \frac{(1 - \alpha 2^{1-\alpha})\left(\frac{k}{\delta}\right)^{1-\alpha}\delta^2}{96e^3k(q^A_\alpha)^\alpha} \\ &=\; \frac{(1 - \alpha 2^{1-\alpha})\delta^{1+\alpha}}{96e^3(q^A_\alpha k)^\alpha}. \end{aligned}$$

If $q^A_\alpha = \alpha^{1/(1-\alpha)}$, this becomes

$$\frac{(1 - \alpha 2^{1-\alpha})\delta^{1+\alpha}}{96e^3(\alpha^{1/(1-\alpha)}k)^\alpha}$$

and if $q^A_\alpha = \frac{\delta}{2k}$ this simplifies to

$$\frac{2^\alpha(1 - \alpha 2^{1-\alpha})\delta}{96e^3}.$$

Finally, if $\alpha = 1$, as $q^A_\alpha = \frac{1}{e}$, the fractional loss of revenue is at least

$$\frac{1 - \ln 2}{96e^3\min\{1, \frac{k}{e}\}\ln\max\{e, k\}}\delta^2.$$

□

6 The Upper Bound

Section 6.1 describes in detail the empirical Myerson auction, the auction strategy for which the guarantee in Theorem 3.1 holds.

6.1 The Empirical Myerson Auction

In the *empirical Myerson auction*, we assume we are given m independent samples from each distribution F_i. The gist is to treat the resulting empirical distribution as the actual distribution in a Myerson auction (Section 2), though some additional technical details are required. In our variant, a number of the samples with the highest values are discarded, and there is a further detail regarding how to handle any high bids that occur in the auction (i.e. bids larger than the largest non-discarded sample).

In detail, for each bidder i, we use the samples from F_i to construct an "empirical revenue curve" as follows (see also Figure 2):

1. Suppose that the m independent samples drawn from F_i have values $v_{i1} \geq v_{i2} \geq \ldots \geq v_{im}$. Define the "empirical quantile" of v_{ij} as $\frac{2j-1}{2m}$.

2. Discard the $\lfloor \hat{\xi} m \rfloor - 1$ largest samples, for a suitable $\hat{\xi} > 0$.[13] Let S denote the remaining samples.

3. For each remaining sample $v_{ij} \in S$, plot a point $(\frac{2j-1}{2m}, \frac{2j-1}{2m} v_{ij})$.

4. Add points at (0,0) and (1,0).

5. While only needed for the analysis, it will be helpful to define the "empirical revenue curve", $\overline{R}_i(\bar{q})$: this is the curve comprising straight-line segments joining the sequence of points specified in Steps 3 and 4 above.

6. Take the convex hull — the least concave upper bound — of this point set. Denote the resulting "ironed empirical revenue curve" by $\overline{\text{CR}}_i$. This curve has constant slope between any two consecutive empirical quantiles of points of S.

Now define empirical ironed virtual values as follows. For $v > v_{i,\hat{\xi}m}$, it is simply the value v. For $v \leq v_{i,\hat{\xi}m}$, identify the two samples $v_{ij}, v_{i(j+1)} \in S$ that "sandwich" v. The empirical ironed virtual value of v is defined as the slope of the revenue curve $\overline{\text{CR}}_i$ in the interval defined by the empirical quantiles of v_{ij} and $v_{i(j+1)}$.[14] Note that the empirical ironed virtual value of $v_{i,\hat{\xi}m}$ is also $v_{i,\hat{\xi}m}$, and that the empirical ironed virtual valuation is a nondecreasing function of v.

Finally we run Myerson's auction on these empirical ironed virtual valuations. That is, the item is awarded to the bidder, if any, with the highest non-negative virtual value (with ties broken arbitrarily). The winner's payment is the lowest bid needed to ensure a (tied) win.

6.2 Notation

We next specify notation so as to clearly distinguish parameters for Myerson's optimal auction from those for the empirical auction, as our analysis will be repeatedly comparing these two auctions. After a couple of simple results, Lemma 6.2 bounds the empirical quantiles as a

[13]The reason for discarding the largest samples is that if they were present there is a non-negligible probability that they would create a poor approximation at the high value end of the distribution, which is the end that matters the most. See also [18].

[14]If v is one of the points of S and there are multiple choices for this slope, we take the largest one.

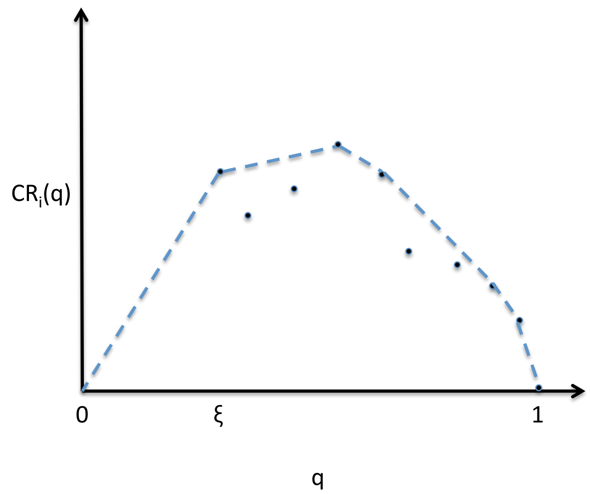

Figure 2: Construction of the ironed empirical revenue curve.

function of the actual quantiles, and vice versa (this is essentially Lemma 4.1 in [18]). Next, Lemma 6.4 relates the empirical and actual virtual values. With these in hand, in Section 6.5, we bound the expected revenue loss due to using the empirical auction as opposed to Myerson's optimal auction, assuming for the latter auction that the actual distributions were fully known.

Myerson's Auction Let MR (the "Myerson Revenue") denote the expected revenue earned by Myerson's auction. Let $x_i(q)$ denote the probability that bidder i wins in Myerson's auction with a bid that has quantile q in its value distribution. Recall that $v_i(q)$ denotes the value corresponding to quantile q and $\varphi_i(q)$ denotes the virtual value at quantile q. Let $\mathrm{MR}_i = E_{q_i}[\varphi_i \cdot x_i]$ denote the expected revenue provided by i in Myerson's auction (recall Theorem 2.1). Let $q_i(v)$ denote the minimum quantile for value v; sometimes it will be convenient to let q_v^i denote $q_i(v)$. Let r_i be the reserve price applied to i in Myerson's auction, namely the largest value for which $\varphi_i(q_i(v)) = 0$. Let q_{r_i} denote $q_i(r_i)$. Let $\mathrm{SR}_i = E[\varphi_i(q) \,|\, q \geq r_i] = q_{r_i} \cdot r_i$; note that SR_i is the expected revenue if i were the only participant in Myerson's auction (SR_i is short for "Single buyer Revenue"). Sometimes, to reduce clutter, we suppress the index i when it is clear from the context. The following claim is immediate from the definitions.

Claim 6.1. *i.* $\mathrm{MR} = \sum_{i=1}^{n} \mathrm{MR}_i$.

ii. $\mathrm{SR}_i \leq \mathrm{MR}$ *for all* i.

The Empirical Auction The empirical auction is defined in terms of the "empirical quantile" $\bar{q}$, but its analysis will entail considering its revenue as a function of the actual quantile q. We specify notation which will distinguish between these two parameters. For $v \leq v_{i,\hat{\xi}m}$, we define the empirical quantile $\bar{q}(v)$ as the solution to $v \cdot q = \overline{R}_i(q)$ in q. (If there are multiple solutions, we take the smallest one.) Going the other way, for an empirical quantile $\bar{q} \geq \bar{\xi}$, we define $\bar{v}(\bar{q})$ as the solution to $v \cdot \bar{q} = \overline{R}_i(\bar{q})$ in v. The empirical ironed virtual value $\bar{v}_i(\bar{q})$ of an empirical quantile $\bar{q} \geq \bar{\xi}$ is the slope of $\overline{\mathrm{CR}}_i$ at $\bar{q}$. Let $\bar{x}_i(\bar{q})$ denote the probability that bidder i wins in the empirical auction with the bid $\bar{v}_i(\bar{q})$. Let $\bar{r}_i$ denote the empirical reserve price, which is the minimum of $\bar{v}_i(\bar{\xi})$ and the largest value $\bar{v}$ for which $\bar{\varphi}_i(\bar{q}_i(\bar{v})) = 0$, and let $\bar{q}_{\bar{r}_i}$ denote the corresponding empirical quantile. Again, to reduce clutter, we sometimes suppress the index i when it is clear from the context.

The actual quantile q corresponding to empirical quantile $\bar{q}$ is defined by the relation $v_i(q) = \bar{v}_i(\bar{q})$; it is denoted by $q(\bar{v}_i(\bar{q}))$; we write it as q for short. Finally, we write the empirical probability of winning as $\tilde{x}_i(q) = \bar{x}_i(\bar{q})$.

As the auction may draw values $v_i > \bar{v}_i(\bar{\xi})$ it is convenient for the purposes of our analysis to define $\bar{q}$ for $\bar{q} < \bar{\xi}$. Let ξ_i be defined by $v_i(\xi_i) = \bar{v}_i(\bar{\xi})$. For $q_i < \xi_i$, we define the corresponding value of $\bar{q}$, as $\bar{q}(q) = \frac{\bar{\xi}}{\xi_i} q$. Then, for $\bar{q} < \bar{\xi}$, $\bar{\varphi}_i(\bar{q}) = v_i(q)$.

6.3 Relating the Actual and Empirical Quantiles

The following result is essentially Lemma 4.1 in [18].

Lemma 6.2. *Let F be a regular distribution. Suppose m independent samples with values $v_1 \geq v_2 \geq \ldots \geq v_m$ are drawn from F. Let $\gamma > 0$, $\hat{\xi} = \frac{a}{m} < 1$ for some integer $a > 0$ be given, and let ν be defined by $1 + \nu = (1+\gamma)^2$. Let $t_h = \frac{2h-1}{2m}$. Then, for all $v \leq v_{\hat{\xi}m}$,*

$$q(v) \in \left[\frac{\bar{q}(v)}{(1+\gamma)^2}, \bar{q}(v)(1+\gamma)^2\right] = \left[\frac{\bar{q}(v)}{(1+\nu)}, \bar{q}(v)(1+\nu)\right]$$

or equivalently

$$\bar{q}(v) \in \left[\frac{q(v)}{(1+\nu)}, q(v)(1+\nu)\right]$$

with probability at least $1-\delta$, if $\gamma\hat{\xi}m \geq 1$ and $m \geq \frac{6(1+\gamma)}{\gamma^2\hat{\xi}} \max\{\frac{\ln 3}{\gamma}, \ln\frac{3}{\delta}\}$.

Proof. We begin by identifying a subsequence of the samples, $v_{l_1}, v_{l_2}, \ldots, v_{l_s}$, with $l_1 \leq l_2 \leq \ldots \leq l_s$; we rename the sequence $u_1, u_2, \ldots, u_s$ for notational ease. It will be the case that $t_{l_{i+1}} \leq (1+\gamma)t_{l_i}$, for $1 \leq i < s$, and $t_{l_s}(1+\gamma) > 1$. We will show that

$$q(v_h) \in \left[\frac{t_h}{(1+\gamma)}, t_h(1+\gamma)\right], \qquad \text{for } v_h \in U = \{u_1, \ldots, u_s\}.$$

The claimed bound is then immediate as either each $v \leq v_{\hat{\xi}m}$ is sandwiched between two items in U, or it is at most u_s.

We define the l_i as follows: $l_1 = \hat{\xi}m$ and $l_{i+1} = \lfloor(1+\gamma)l_i\rfloor$ if $\lfloor(1+\gamma)l_i\rfloor \leq m$, and otherwise l_{i+1} is not defined (i.e. $i = s$). As $\gamma\hat{\xi}m \geq 1$, $\lfloor(1+\gamma)\hat{\xi}m\rfloor \geq \hat{\xi}m + \lfloor\gamma\hat{\xi}m\rfloor \geq l_1 + 1$, from which we conclude that the sequence is strictly increasing and hence well defined.

Next we bound the probability that $q(u_i) > (1+\gamma)t_{l_i}$. Now $q(u_i) > (1+\gamma)t_{l_i}$ only if fewer than $l_i = t_{l_i}m + \frac{1}{2}$ samples have q values that are at most $(1+\gamma)t_{l_i}$. As the expected number of such samples is $(1+\gamma)t_{l_i}m$, a Chernoff bound gives the following upper bound on the probability that $q(u_i) > (1+\gamma)t_{l_i}$ (cf. [31]):

$$\exp\{-\frac{\gamma^2 t_{l_i} m}{2(1+\gamma)}\}.$$

Similarly, the probability that $q(u_i) < t_{l_i}/(1+\gamma)$ is bounded by

$$\exp\{-\frac{\gamma^2 t_{l_i} m}{3(1+\gamma)}\}.$$

It will be helpful to bound both $t_{l_1}m$ and $[t_{l_{i+1}} - t_{l_i}]m$. As $\hat{\xi}m \geq 1$, $t_{l_1}m = (2\hat{\xi}m - 1)/2 \geq \frac{1}{2}\hat{\xi}m$. And as $\gamma\hat{\xi}m \geq 1$, $[t_{l_{i+1}} - t_{l_i}]m \geq \lfloor(1+\gamma)l_i\rfloor - l_i \geq \lfloor\gamma l_i\rfloor \geq \lfloor\gamma\hat{\xi}m\rfloor \geq \frac{1}{2}\gamma\hat{\xi}m$.

Now, by the union bound applied to all the $q(u_i)$, we obtain a failure probability of at most:

$$\sum_{i=1}^{l} \exp\{-\frac{\gamma^2 t_{l_i} m}{3(1+\gamma)}\} + \exp\{-\frac{\gamma^2 t_{l_i} m}{2(1+\gamma)}\}$$

$$\leq\ 2\sum_{i=1}^{l} \exp\{-\frac{\gamma^2 t_{l_i} m}{3(1+\gamma)}\}$$

$$\leq\ 2\sum_{i=0}^{l-1} \exp\left\{-\frac{\gamma^2[\hat{\xi}m + (i-1)\gamma\hat{\xi}m]}{6(1+\gamma)}\right\}$$

using the bounds on $t_{l_1}m$ and $[t_{l_{i+1}} - t_{l_i}]m$

$$\leq\ \frac{2\exp\{-\frac{\gamma^2\hat{\xi}m}{6(1+\gamma)}\}}{1-\exp\{-\frac{\gamma^3\hat{\xi}m}{6(1+\gamma)}\}}\ \leq\ 3\exp\left\{-\frac{\gamma^2\hat{\xi}m}{6(1+\gamma)}\right\}$$

$$\text{if } \exp\left\{-\frac{\gamma^3\hat{\xi}m}{6(1+\gamma)}\right\} \leq \frac{1}{3}.$$

We want the failure probability to be at most δ. So we need $\frac{\gamma^2\hat{\xi}m}{6(1+\gamma)} \geq \ln\frac{3}{\delta}$, i.e. $m \geq \frac{6(1+\gamma)}{\gamma^2\hat{\xi}}\ln\frac{3}{\delta}$. We also need $m \geq \frac{6(1+\gamma)}{\gamma^3\hat{\xi}}\ln 3$ to satisfy the condition in the final inequality. □

6.4 Relating the Actual and the Empirical Virtual Values

Let $\mathcal{E}_a$ be the event that the high probability outcome of Lemma 6.2 occurs, namely that for all $v \leq v_{\lfloor\hat{\xi}m\rfloor}$, $q(v) \in \left[\frac{\bar{q}(v)}{(1+\nu)}, \bar{q}(v)(1+\nu)\right]$. $\mathcal{E}_a$ occurs with probability at least $1-\delta$. It will also be helpful to express the bound on v as a bound on $\bar{q}$. To this end, we define $\bar{\xi} = t_1 = \frac{2\hat{\xi}m-1}{2m}$.

We will repeatedly encounter terms of the form $\varphi(\lambda q)$ with $\lambda > 1$; For $\lambda q > 1$, $\varphi(\lambda q)$ is interpreted to mean $\varphi(1)$; similarly for $\bar{\varphi}$.

Lemma 6.3. *Conditioned on $\mathcal{E}_a$, for all empirical quantiles $\bar{q} \geq \bar{\xi}$, $\overline{CR}(\bar{q}) \leq \bar{q}\cdot v(\frac{\bar{q}}{1+\nu})$, and for all $t_h = \frac{2j-1}{2m} \geq \bar{\xi}$, $\overline{CR}(t_h) \geq t_h \cdot v(\ t_h(1+\nu))$.*

This lemma is not as obvious as it may seem for it concerns points on the convex hull $\overline{\text{CR}}$ of the set of points $\overline{R}$ that are used to specify the empirical revenue.

Proof. By Lemma 6.2, as $\mathcal{E}_a$ holds, for all $t_h \geq \bar{\xi}$,

$$t_h \cdot v(t_h(1+\nu)) \leq \overline{R}(t_h) \leq t_h \cdot v(\frac{t_h}{1+\nu}).$$

We define $\overline{L}(\bar{q}) = \bar{q}\cdot v(\bar{q}(1+\nu))$ and $\overline{U}(\bar{q}) = \bar{q}\cdot v(\frac{\bar{q}}{1+\nu})$ for all $\bar{q}$.

Note that for any pair $q \neq q'$ of quantiles, the line joining the actual revenue $R(\frac{q}{1+\nu}) = \frac{q}{1+\nu}v(\frac{q}{1+\nu})$ to $R(\frac{q'}{1+\nu}) = \frac{q'}{1+\nu}v(\frac{q'}{1+\nu})$ is parallel to the line joining $\overline{U}(q)$ to $\overline{U}(q')$, for the latter line is obtained by expanding the former line by a factor $1+\nu$ in both the quantile and revenue dimensions. By the regularity of φ, the curve defined by R is convex, and consequently, the points $\overline{U}(\bar{q})$ all lie on their convex hull.

For $t_h \geq \bar{\xi}$, $\overline{U}(t_h)$ is an upper bound on $\overline{R}(t_h)$; it follows that the convex hull for the empirical revenue, for $\bar{q} \geq \bar{\xi}$, is enclosed by the convex hull $\overline{U}(\bar{q})$, and consequently $\overline{\text{CR}}(\bar{q}) \leq \overline{U}(\bar{q}) = \bar{q}\cdot v(\frac{\bar{q}}{1+\nu})$.

For the second result, the lower bound, we use a similar argument, but it will apply just to the empirical quantiles $t_h \geq \bar{\xi}$. Now, for any pair $q \neq q'$ of quantiles, the line joining the actual revenue $R(q(1+\nu)) = q(1+\nu)v(q(1+\nu))$ to $R(q'(1+\nu)) = q'(1+\nu) \cdot v(q'(1+\nu))$ is parallel to the line joining $\overline{L}(q)$ to $\overline{L}(q')$, and hence the points $\overline{L}(\bar{q})$ all lie on their convex hull. But, for $t_h \geq \bar{\xi}$, $\overline{L}(t_h) \leq \overline{R}(t_h)$, and consequently the values $\overline{R}(t_h)$ all lie on or above the curve $\overline{L}(t_h)$. □

The following lemma, which lies at the heart of out analysis, shows that with high probability$\varphi(q)$ is close to some value $\bar{\varphi}(\bar{q}\,')$ with $\bar{q}\,' \in [\frac{\bar{q}}{(1+\Delta)(1+\nu)^3}, \bar{q}(1+\Delta)(1+\nu)]$.

Lemma 6.4. *Let F be a regular distribution. Suppose that $(1+\Delta) \geq (1+\nu)^2$. Let $t_h = \frac{2h-1}{2m}$, for $1 \leq h \leq m$. Conditioned on $\mathcal{E}_a$, if $t_{h-1} < \bar{q} \leq t_h$, then*

i. for all $\bar{q}$ with $\bar{\xi}(1+\Delta)(1+\nu)^3 \leq \bar{q}$, $\varphi(q) \leq \varphi(\frac{t_h}{(1+\nu)^2}) \leq \bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3}) + 2\frac{\nu}{\Delta}(1+\Delta)(1+\nu)^3\frac{\mathrm{SR}}{\bar{q}}$, and

ii. for all $\bar{q}$ with $\bar{\xi} \leq \bar{q}$, $\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu)) \leq \varphi(t_h(1+\nu)) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}}{\bar{q}} \leq \varphi(q) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}}{\bar{q}}$.

Proof. The main part of the proof concerns the second inequality in (i) and the first one in (ii). We begin by proving the inequality in (i). First we give an upper bound on $\varphi(\frac{t_h}{(1+\nu)^2})$ and a lower bound on $\bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3})$.

As F is regular, R is convex; thus:

$$\begin{aligned}
\varphi(\frac{t_h}{(1+\nu)^2}) &\leq \frac{R(\frac{t_h}{(1+\nu)^2}) - R(\frac{t_h}{(1+\Delta)(1+\nu)^4})}{\frac{t_h}{(1+\nu)^2} - \frac{t_h}{(1+\Delta)(1+\nu)^4}} \\
&= \frac{\frac{t_h}{(1+\nu)^2} \cdot v(\frac{t_h}{(1+\nu)^2}) - \frac{t_h}{(1+\Delta)(1+\nu)^4} v(\frac{t_h}{(1+\Delta)(1+\nu)^4})}{\frac{t_h}{(1+\nu)^2} - \frac{t_h}{(1+\Delta)(1+\nu)^4}} \\
&= \frac{(1+\Delta)(1+\nu)^2 v(\frac{t_h}{(1+\nu)^2}) - v(\frac{t_h}{(1+\Delta)(1+\nu)^4})}{2\nu + \nu^2 + \Delta(1+\nu)^2}.
\end{aligned}$$

The following bound applies only when $\frac{t_h}{(1+\Delta)(1+\nu)^3} \geq \bar{\xi}$ for otherwise $\overline{\mathrm{CR}}(\frac{t_h}{(1+\Delta)(1+\nu)^3})$ is not defined; the constraint $\bar{q} \geq \bar{\xi}(1+\Delta)(1+\nu)^3$ suffices.

$$\begin{aligned}
\bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3}) &\geq \frac{\overline{\mathrm{CR}}(\frac{t_h}{(1+\nu)^3}) - \overline{\mathrm{CR}}(\frac{t_h}{(1+\Delta)(1+\nu)^3})}{\frac{t_h}{(1+\nu)^3} - \frac{t_h}{(1+\Delta)(1+\nu)^3}} \\
&\geq \frac{\frac{t_h}{(1+\nu)^3} \cdot v(\frac{t_h}{(1+\nu)^2}) - \frac{t_h}{(1+\Delta)(1+\nu)^3} v(\frac{t_h}{(1+\Delta)(1+\nu)^4})}{\frac{t_h}{(1+\nu)^3} - \frac{t_h}{(1+\Delta)(1+\nu)^3}} \quad \text{(by Lemma 6.3)} \\
&= \frac{(1+\Delta)v(\frac{t_h}{(1+\nu)^2}) - v(\frac{t_h}{(1+\Delta)(1+\nu)^4})}{\Delta}.
\end{aligned}$$

Now, we combine the bounds so as to eliminate the term $v(\frac{t_h}{(1+\Delta)(1+\nu)^4})$.

$$
\begin{aligned}
&\frac{2\nu+\nu^2+\Delta(1+\nu)^2}{\Delta}\varphi(\frac{t_h}{(1+\nu)^2})-\bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3})\\
&\quad\le\ \frac{(1+\Delta)[(1+\nu)^2-(1+\Delta)]v(\frac{t_h}{(1+\nu)^2})}{\Delta}\\
&\quad\le\ \frac{(1+\Delta)(2\nu+\nu^2)}{\Delta}\frac{\mathrm{SR}(1+\nu)^2}{t_h}\qquad(\text{as SR}\ge\frac{t_h}{(1+\nu)^2}v\left(\frac{t_h}{(1+\nu)^2}\right)\ )\\
&\quad\le\ 2\frac{\nu}{\Delta}(1+\Delta)(1+\nu)^3\frac{\mathrm{SR}}{\bar{q}}.
\end{aligned}
$$

In other words,

$$
\left(1+\frac{\nu(1+\Delta)(2+\nu)}{\Delta}\right)\varphi(\frac{t_h}{(1+\nu)^2})-\bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3})\ \le\ 2\frac{\nu}{\Delta}(1+\Delta)(1+\nu)^3\frac{\mathrm{SR}}{\bar{q}}.
$$

Thus

$$
\varphi(\frac{t_h}{(1+\nu)^2})\le\bar{\varphi}(\frac{\bar{q}}{(1+\Delta)(1+\nu)^3})+2\frac{\nu}{\Delta}(1+\Delta)(1+\nu)^3\frac{\mathrm{SR}}{\bar{q}}.
$$

The second inequality in (ii) is shown similarly. We start with an upper bound on $\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu))$ and a lower bound on $\varphi(t_h(1+\nu))$. The first bound applies only when $t_h\ge\bar{\xi}$; here $\bar{q}\ge\bar{\xi}$ suffices.

$$
\begin{aligned}
\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu))\ \le\ \bar{\varphi}(t_h(1+\Delta))\ &\le\ \frac{\overline{\mathrm{CR}}(t_h(1+\Delta))-\overline{\mathrm{CR}}(t_h)}{(1+\Delta)t_h-t_h}\\
&\le\ \frac{t_h(1+\Delta)\cdot v(\frac{t_h(1+\Delta)}{1+\nu})-t_hv(t_h(1+\nu))}{\Delta t_h}\qquad(\text{by Lemma 6.3})\\
&=\ \frac{(1+\Delta)v(\frac{t_h(1+\Delta)}{1+\nu})-v(t_h(1+\nu))}{\Delta}.
\end{aligned}
$$

$$
\begin{aligned}
\varphi(t_h(1+\nu))\ \ge\ \frac{R(\frac{t_h(1+\Delta)}{(1+\nu)})-R(t_h(1+\nu))}{\frac{t_h(1+\Delta)}{(1+\nu)}-t_h(1+\nu)}\ &=\ \frac{\frac{t_h(1+\Delta)}{1+\nu}v(\frac{t_h(1+\Delta)}{1+\nu})-t_h(1+\nu)v(t_h(1+\nu))}{\frac{t_h(1+\Delta)}{1+\nu}-t_h(1+\nu)}\\
&=\ \frac{(1+\Delta)v(\frac{t_h(1+\Delta)}{1+\nu})-(1+\nu)^2v(t_h(1+\nu))}{\Delta-2\nu-\nu^2}.
\end{aligned}
$$

Again, we combine the bounds so as to eliminate the term $v(t_h(1+\nu))$.

$$
\begin{aligned}
&\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu))-\frac{\Delta-2\nu-\nu^2}{\Delta(1+\nu)^2}\varphi(t_h(1+\nu))\\
&\quad\le\ \frac{(1+\Delta)[(\Delta-2\nu-\nu^2)-\frac{\Delta-2\nu-\nu^2}{(1+\nu)^2}]v(\frac{t_h(1+\Delta)}{1+\nu})}{\Delta(\Delta-2\nu-\nu^2)}\\
&\quad\le\ \frac{(1+\Delta)\nu(2+\nu)}{\Delta(1+\nu)^2}\frac{SR(1+\nu)}{(1+\Delta)t_h}\qquad(\text{as SR}\ge\frac{(1+\Delta)t_h}{1+\nu}v\left(\frac{(1+\Delta)t_h}{1+\nu}\right)\ )\\
&\quad\le\ 2\frac{\nu}{\Delta}\frac{\mathrm{SR}}{\bar{q}}.
\end{aligned}
$$

In other words,

$$\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu)) - \left(1 - \frac{\nu(1+\Delta)(2+\nu)}{\Delta(1+\nu)^2}\right)\varphi(t_h(1+\nu)) \ \le \ 2\frac{\nu}{\Delta}\frac{\mathrm{SR}}{\bar{q}}.$$

Thus

$$\bar{\varphi}(\bar{q}(1+\Delta)(1+\nu)) \le \varphi(t_h(1+\nu)) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}}{\bar{q}}.$$

We now show the remaining inequalities. To obtain the first inequality in (i), we note that by Lemma 6.2 and $\mathcal{E}_a$, $q \ge \frac{\bar{q}}{(1+\nu)} > \frac{t_h}{(1+\nu)^2}$, from which the result follows. Similarly, for the second inequality in (ii), $t_h(1+\nu) \ge \bar{q}(1+\nu) \ge q$, and again the result follows. □

6.5 Bounding the Expected Revenue Loss

Finally, we consider an auction with k bidders, where the valuation for the ith bidder comes from regular distribution F_i. For brevity, bidder i is referred to as i.

We define $\mathcal{E}_b$ to be the event that $\mathcal{E}_a$ holds for every distribution F_i.

Let $\mathrm{Shtf} = \sum_i E[\varphi_i \cdot x_i] - \sum_i E[\varphi_i \cdot \tilde{x}_i]$. In other words, $\overline{\mathrm{MR}} + \mathrm{Shtf} = \mathrm{MR}$, so it suffices to show that $\mathrm{Shtf} \le \epsilon \mathrm{MR}$. Recall that q_{r_i} denotes the quantile of F_i corresponding to the reserve price for i in the Myerson auction and $q_{\bar{r}_i}$ denotes the quantile corresponding to the reserve price in the empirical auction. Also, we let $\bar{q}_i$ be a quantile for i in the empirical auction, and we let q_i denote the corresponding quantile in F_i. $\bar{q}_j$ and q_j are defined similarly with respect to j. In addition, to reduce clutter, we let $\beta = (1+\Delta)(1+\nu)^3 - 1$.

The next lemma provides an upper bound on Shtf as the sum of several terms which we will bound in turn.

Lemma 6.5. *Conditioned on $\mathcal{E}_b$,*

$$\begin{aligned}
\mathrm{Shtf} &= \sum_i \left[\int_{q_i \le q_{r_i}} \varphi_i(q_i)\cdot x_i(q_i)\ dq_i - \int_{q_i \le q_{\bar{r}_i}} \varphi_i(q_i)\cdot \tilde{x}_i(q_i)\ dq_i\right] \\
&\le \sum_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)\cdot [x_i(q_i) - \tilde{x}_i(\frac{q_i}{1+\beta})]\ dq_i \qquad (7)\\
&\qquad + \beta \sum_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)\ dq_i \qquad (8)\\
&\qquad + \sum_i \int_{q_{r_i} \le q_i \le q_{\bar{r}_i}} [-\varphi_i(q_i)]\ dq_i. \qquad (9)
\end{aligned}$$

Proof. We upper bound the second (negative) term in the expression for Shtf.

$$\begin{aligned}
-\int_{q_i \le q_{\bar{r}_i}} \varphi_i(q_i)\cdot\tilde{x}_i(q_i)\ dq_i &= -\int_{q_i \le q_{r_i}} \varphi_i(q_i)\cdot\tilde{x}_i(q_i)\ dq_i + \int_{q_{r_i} \le q_i \le q_{\bar{r}_i}} [-\varphi_i(q_i)]\cdot\tilde{x}_i(q_i)\ dq_i \\
&\le -\int_{q_i \le q_{r_i}/(1+\beta)} \varphi_i(q_i)\cdot\tilde{x}_i(q_i)\ dq_i + \int_{q_{r_i} \le q_i \le q_{\bar{r}_i}} [-\varphi_i(q_i)]\ dq_i
\end{aligned}$$

and

$$
\begin{aligned}
-&\int_{q_i \le q_{r_i}/(1+\beta)} \varphi_i(q_i) \cdot \tilde{x}_i(q_i)\ dq_i \\
&= -(1+\beta) \int_{q_i \le q_{r_i}/(1+\beta)} \varphi_i(q_i) \cdot \tilde{x}_i(q_i)\ dq_i + \beta \int_{q_i \le q_{r_i}/(1+\beta)} \varphi_i(q_i) \cdot \tilde{x}_i(q_i)\ dq_i \\
&\le -\int_{q_i \le q_{r_i}} \varphi_i(\frac{q_i}{1+\beta}) \cdot \tilde{x}_i(\frac{q_i}{1+\beta})\ dq_i + \beta \int_{q_i \le q_{r_i}} \varphi_i(q_i) \cdot \tilde{x}_i(q_i)\ dq_i \\
&\le -\int_{q_i \le q_{r_i}} \varphi_i(q_i) \cdot \tilde{x}_i(\frac{q_i}{1+\beta})\ dq_i + \beta \int_{q_i \le q_{r_i}} \varphi_i(q_i)\ dq_i.
\end{aligned}
$$

Substituting in the expression for Shtf yields the result. □

The bound on (8) is simply

$$\beta \sum_i \mathrm{SR}_i = k\beta \cdot \mathrm{MR}. \tag{10}$$

In the following lemmas we bound the terms (9) and (7). To bound (9) we partition the integral into two intervals. The intervals are the ranges $q_{r_i} \le q_i \le \max\{\xi_i, q_{r_i}\}$ and $\max\{\xi_i, q_{r_i}\} \le q_i \le q_{\bar{r}_i}$, respectively, where ξ_i is the quantile of F_i corresponding to empirical quantile $\bar{\xi}$.

Lemma 6.6. *Conditioned on $\mathcal{E}_b$,*

$$\sum_i \int_{q_{r_i} \le q_i \le \max\{\xi_i, q_{r_i}\}} [-\varphi_i(q_i)]\ dq_i \le k\bar{\xi}(1+\nu) \cdot \mathrm{MR}.$$

Proof. If $\xi_i \le q_{r_i}$ the integral is zero and the result is immediate. So we can assume that $\xi_i \ge q_{r_i}$. Note that $-\varphi_i(q_i)$ is a non-decreasing function of q_i; thus its smallest values in the range $q_i \ge q_{r_i}$ occur in the integral we are seeking to bound. It follows that

$$
\begin{aligned}
\int_{q_{r_i} \le q_i \le \xi_i} [-\varphi_i(q_i)]\ dq_i &\le \frac{\xi_i - q_{r_i}}{1 - q_{r_i}} \int_{q_{r_i} \le q_i} [-\varphi_i(q_i)]\ dq_i \\
&\le \xi_i \int_{q_{r_i} \le q_i} [-\varphi_i(q_i)]\ dq_i \\
&\le \xi_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)\ dq_i \qquad\qquad (\text{as } \int_{0 \le q_i \le 1} \varphi_i(q_i)\ dq_i = 0) \\
&= \xi_i \cdot \mathrm{SR}_i \le \xi_i \cdot \mathrm{MR} \le \bar{\xi}(1+\nu) \cdot \mathrm{MR}.
\end{aligned}
$$

The last two inequalities follow from Claim 6.1(ii) and Lemma 6.2, respectively. The result follows on summing over i. □

Lemma 6.7. *Conditioned on $\mathcal{E}_b$,*

$$E\left[\sum_i \int_{\max\{\xi_i, q_{r_i}\} \le q_i \le q_{\bar{r}_i}} [-\varphi_i(q_i)]\ dq_i\right] \le 2\nu \sum_i \mathrm{SR}_i \le 2k\nu \cdot \mathrm{MR}.$$

Proof. let $\chi_i = \max\{\xi_i, q_{r_i}\}$ and let $\overline{\chi}_i = \max\{\bar{\xi}, \bar{q}_{r_i}\}$ be the corresponding empirical quantile. Again, if $\chi_i \ge q_{\bar{r}_i}$ the integral is zero and the result is immediate. So we can assume that $\chi_i < q_{\bar{r}_i}$. The derivation below uses Lemma 6.2 to justify the first and third inequalities; the second inequality follows from the definition of $\overline{r}_i$ as the empirical reserve price since $\overline{\chi}_i \geq \bar{\xi}$. Conditioned on $\mathcal{E}_b$,

$$q_{\overline{r}_i} \cdot \overline{r}_i \;\geq\; \frac{\bar{q}_{\overline{r}_i} \cdot \overline{r}_i}{1+\nu} \;\geq\; \frac{\overline{\chi}_i \cdot \bar{v}_i(\overline{\chi}_i)}{1+\nu} \;\geq\; \frac{\chi_i \cdot v_i(\chi_i)}{(1+\nu)^2}. \quad (11)$$

Thus

$$\int_{\max\{\xi_i, q_{r_i}\} \leq q_i \leq q_{\overline{r}_i}} [-\varphi_i(q_i)]\, dq_i \;=\; \chi_i \cdot v_i(\chi_i) - q_{\overline{r}_i} \cdot \overline{r}_i \;\leq\; \chi_i \cdot v_i(\chi_i) \left[1 - \frac{1}{(1+\nu)^2}\right] \quad \text{(by (11))}$$
$$\leq \; \frac{\nu(2+\nu)}{(1+\nu)^2} \mathrm{SR}_i \leq 2\nu \cdot \mathrm{SR}_i \leq 2\nu \cdot \mathrm{MR}.$$

□

It remains to bound term (7).

Before proceeding to the next lemma we need some additional terminology, namely the notions of i-safety, and of large and small amounts.

Definition 6.1. *The vector of quantiles $q = (q_1, q_2, \cdots, q_k)$ is said to be i-*safe *if $\bar{q}_h \geq \bar{\xi}$ for all $h \neq i$.*

We will also write $\mathbf{q} = (q_i, q_j, q_{-ij})$ and $\mathbf{q} = (q_i, q_{-i})$, when we want to focus on just two or one coordinates of the quantile vector.

We define large and small amounts with respect to $\varphi_i(q_i)$ and $\varphi_j(q_j)$ as follows.

Definition 6.2. *Let ρ and ρ' be defined as in Lemma 6.2. Suppose that $\varphi_i(q_i) \geq \varphi_j(q_j)$. $\varphi_i(q_i)$ is said to exceed $\varphi_i(q_j)$ by a* large amount *in the following cases:*

i. For $\bar{q}_i \geq \bar{\xi}(1+\beta)$,

$$\varphi_i(q_i) - \varphi_i(q_j) \geq 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i} + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j},$$

ii. and for $\bar{q}_i < \bar{\xi}(1+\beta)$,

$$\varphi_i(q_i) - \varphi_i(q_j) \geq 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}.$$

Otherwise, $\varphi_i(q_i)$ is said to exceed $\varphi_i(q_j)$ by a small amount*.*

The following lemma bounds the probability that i wins by a large amount over j in the Myerson auction at quantile $\mathbf{q}$, while j wins in the empirical auction at quantile $\mathbf{q}/(1+\beta)$.

Lemma 6.8. *Conditioned on $\mathcal{E}_b$, for any pair i and j, and for any q_i, if $\bar{q}_j \geq \bar{\xi}$, then the probability of the following event is bounded by $(1+\beta)^2 - 1$:*

> *i wins in the Myerson auction by a large amount over j at quantile $\mathbf{q}$, and j wins in the empirical auction at quantile $\bar{\mathbf{q}}/(1+\beta)$, where $\mathbf{q} = (q_i, q_{-i})$ is i-safe.*

Proof. We begin with the case that $\bar{q}_i \geq \bar{\xi}(1+\Delta)$. Given $\mathcal{E}_b$, by Lemma 6.4, for $\bar{\xi}(1+\beta) \leq \bar{q}_i$, $\varphi_i(q_i) \leq \bar{\varphi}_i(\frac{\bar{q}}{1+\beta}) + 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i}$ and for $\bar{\xi} \leq \bar{q}_j$, $\bar{\varphi}_j(\bar{q}_j(1+\Delta)(1+\nu)) \leq \varphi_j(q_j) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}$. Thus,

$$\begin{aligned}
\bar{\varphi}_j(\bar{q}_j(1+\Delta)(1+\nu)) &\leq \varphi_j(q_j) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} \\
&< \varphi_i(q_i) - 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i} \quad \text{(from Definition 6.2(i))} \\
&\leq \bar{\varphi}_i(\frac{\bar{q}_i}{1+\beta}) \\
&< \bar{\varphi}_j(\frac{\bar{q}_j}{1+\beta}).
\end{aligned}$$

Thus we have a lower bound of $\bar{\varphi}_j(\bar{q}_j(1+\Delta)(1+\nu))$ and an upper bound of $\bar{\varphi}_j(\frac{\bar{q}_j}{1+\beta})$ on the remaining terms. Clearly these can both hold only for a limited range of $\bar{q}_j$ and hence of q_j, which we bound as follows. Define $\hat{\bar{q}}_j = \arg\inf_{\bar{q}_j}\{\bar{\varphi}_j(\bar{q}_j(1+\Delta)(1+\nu)) \le \bar{\varphi}_i(\bar{q}_i/[1+\beta])\}$. Then these bounds can hold at most for $\bar{q}_j$ satisfying $\hat{\bar{q}}_j \le \bar{q}_j < \hat{\bar{q}}_j(1+\Delta)(1+\nu)(1+\beta)$. To obtain a probability bound, one needs to express the range in terms of the q_j quantile, namely ranging at most from $\hat{\bar{q}}_j/(1+\nu)$ to $\min\{1, \hat{\bar{q}}_j(1+\Delta)(1+\nu)^2(1+\beta)\}$, i.e. with probability at most $(1+\Delta)(1+\nu)^3(1+\beta) - 1 = (1+\beta)^2 - 1$.

When $\bar{q}_i < \bar{\xi}(1+\beta)$, we proceed similarly. (The third inequality below follows because for $\bar{q}_i \le \bar{\xi}$, $\bar{\varphi}_i(\bar{q}_i) = v_i(q_i)$, and the fourth inequality holds because $\frac{\bar{q}_i}{(1+\beta)} \le \bar{\xi}$ by assumption.)

$$\begin{aligned}
\bar{\varphi}_j(\bar{q}_j(1+\Delta)(1+\nu)) &\le \varphi_j(q_j) + 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} < \varphi_i(q_i) \\
&\le \max\{v_i(q_i), v_i(\xi_i)\} \\
&\le \bar{\varphi}_i(\min\{\bar{q}_i, \bar{\xi}\}) \\
&\le \bar{\varphi}_i(\frac{\bar{q}_i}{1+\beta}) < \bar{\varphi}_j(\frac{\bar{q}_j}{1+\beta}).
\end{aligned}$$

The rest of the argument is as for (i). □

Lemma 6.9. *Conditioned on $\mathcal{E}_b$,*

$$\begin{aligned}
&\sum_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)\cdot[x_i(q_i) - \tilde{x}_i(\frac{q_i}{1+\beta})]\, dq_i \\
&\quad \le \left[(k-1)\bar{\xi}(1+\nu) + k(k-1)[(1+\beta)^2-1] + (\rho+\rho') + 4k(1+\beta)(1+\nu)\frac{\nu}{\Delta}\cdot \ln\frac{1+\nu}{\bar{\xi}}\right]\mathrm{MR}.
\end{aligned}$$

Proof. We let $\mathcal{E}_i^x(q_i)$ be the event that quantile $\mathbf{q} = (q_i, q_{-i})$ is not i-safe. Clearly $\mathbf{Pr}[\mathcal{E}_i^x] \le \sum_{j\ne i}\xi_j \le (k-1)\bar{\xi}(1+\nu)$. Let $x_i^e(q_i)$ be the probability that i wins when the quantile $\mathbf{q} = (q_i, q_{-i})$ is not i-safe. Note that $x_i(q_i \,|\, \mathbf{q}$ is not i-safe$) \le x_i(q_i)$, as having some q_j be small only increases the probability that j wins. Thus

$$x_i^e(q_i) \le (k-1)\bar{\xi}(1+\nu)x_i(q_i). \tag{12}$$

We let x_i^b denote the probability that i wins both in Myerson's auction with i-safe $\mathbf{q} = (q_i, q_{-i})$ and in the empirical auction at quantile $\mathbf{q}/(1+\beta)$.

We also introduce notation to measure the probability of wins by small and large amounts, for i-safe quantiles. We will be measuring the probability that i wins in the Myerson auction at quantile $\mathbf{q} = (q_i, q_j, q_{-ij})$ and j wins in the empirical auction at quantile $\mathbf{q}/(1+\beta)$, for some q_{-ij}. $x_{ij}^s(q_i, q_j)$ measures the probability of this event in the case that the win in the Myerson auction is by a small amount, and $x_{ij}^l(q_i, q_j)$ measures the probability of the event when the win margin is large.

Switching perspectives, we let $\tilde{x}_{ij}^s(q_i, q_j)$ denote the probability that i wins in the empirical auction at quantile $\mathbf{q}/(1+\beta)$ and j wins by a small amount over i in the Myerson auction at quantile $\mathbf{q}$, where $\mathbf{q} = (q_i, q_j, q_{-ij})$ is j-safe. Clearly, $\tilde{x}_{ji}^s(q_j, q_i) = x_{ij}^s(q_i, q_j)$.

We also note that

$$x_i(q_i) = x_i^e(q_i) + x_i^b(q_i) + \sum_{j\ne i}\int_{\bar{q}_j \ge \bar{\xi}} x_{ij}^s(q_i, q_j) + x_{ij}^l(q_i, q_j)\, dq_j, \tag{13}$$

and for $\bar{q}_i \ge \bar{\xi}$,

$$\tilde{x}_i(\frac{q_i}{1+\beta}) \ge x_i^b(q_i) + \sum_{i\ne j}\int_{\bar{q}_j \ge 0} \tilde{x}_{ij}^s(q_i, q_j)\, dq_j. \tag{14}$$

By Lemma 6.8,

$$\int_{q_j \geq \bar{\xi}} x^l_{ij}(q_i, q_j) \leq (1+\beta)^2 - 1. \tag{15}$$

We obtain:

$$\begin{aligned}
&\sum_i \int_{q_i \leq q_{r_i}} \varphi_i(q_i) \cdot [x_i(q_i) - \tilde{x}_i(\frac{q_i}{1+\beta})]\ dq_i \\
&\leq \sum_i \int_{q_i \leq q_{r_i}} \varphi_i(q_i) \left[\left(x^b_i(q_i) + x^e_i(q_i) + \sum_{j \neq i} \int_{\bar{q}_j \geq \bar{\xi}} x^s_{ij}(q_i, q_j) + x^l_{ij}(q_i, q_j) \right)\ dq_j \right. \\
&\qquad \left. - \left(x^b_i(q_i) + \sum_{j \neq i} \int_{\bar{q}_j \geq \bar{\xi}} \tilde{x}^s_{ij}(q_i, q_j) \right)\ dq_i \right]\ dq_i \quad \text{(using (13) and (14))} \\
&\leq \sum_i \int_{q_i \leq q_{r_i}} \varphi_i(q_i) \left[x^e_i(q_i) + \sum_{j \neq i} \int_{\bar{q}_j \geq \bar{\xi}} [x^l_i(q_i, q_j) + x^s_i(q_i, q_j)] \right]\ dq_i\ dq_j \\
&\quad - \sum_j \int_{q_j \leq q_{r_j}} \varphi_j(q_j) \cdot \sum_{i \neq j} \int_{\bar{q}_i \geq \bar{\xi}} \tilde{x}^s_{ji}(q_j, q_i)\ dq_j\ dq_i \qquad (16) \\
&\leq \sum_i \int_{q_i \leq q_{r_i}} (k-1)\bar{\xi}(1+\nu) \cdot \varphi_i(q_i) \cdot x_i(q_i)\ dq_i + (k-1)[(1+\beta)^2 - 1] \int_{q_i \leq q_{r_i}} \varphi_i(q_i)\ dq_i \\
&\qquad \text{(using (12) and (15))} \\
&\quad + \sum_i \sum_{j \neq i} \left[\int_{q_i \leq q_{r_i}, \bar{q}_j \geq \bar{\xi}} \varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j)\ dq_i\ dq_j \right. \\
&\qquad \left. - \int_{q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i)\ dq_i\ dq_j \right] \\
&\leq (k-1)\bar{\xi}(1+\nu) \sum_i \mathrm{MR}_i + (k-1)[(1+\beta)^2 - 1] \sum_i \mathrm{SR}_i \qquad (17) \\
&\quad + \sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} [\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i)]\ dq_i\ dq_j \qquad (18) \\
&\quad + \sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, q_j > q_{r_j}, \bar{q}_j \geq \bar{\xi}} \varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j)\ dq_i\ dq_j. \qquad (19)
\end{aligned}$$

We bound (17)–(19) in turn.

$$\begin{aligned}
&(k-1)\bar{\xi}(1+\nu) \sum_i \mathrm{MR}_i + (k-1)[(1+\beta)^2 - 1] \sum_i \mathrm{SR}_i \\
&\qquad \leq (k-1)\bar{\xi}(1+\nu)\mathrm{MR} + k(k-1)[(1+\beta)^2 - 1]\mathrm{MR}. \qquad (20)
\end{aligned}$$

We can deduce from (13) that

$$\sum_{j \neq i} \int_{\bar{q}_j \geq \bar{\xi}} x^s_{ij}(q_i, q_j)\ dq_j \leq x_i(q_i) \leq 1, \tag{21}$$

and from (14), for $\bar{q}_j \geq \bar{\xi}$,

$$\sum_{i \neq j} \int_{q_i \geq 0} x^s_{ij}(q_i, q_j)\ dq_i = \sum_{i \neq j} \int_{q_i \geq 0} \tilde{x}^s_{ji}(q_j, q_i)\ dq_i \leq \tilde{x}_j(\frac{q_j}{1+\beta}) \leq 1. \tag{22}$$

For (18), recall that $x^s_{ij}(q_i, q_j) = \tilde{x}^s_{ji}(q_j, q_i)$. Thus when $\bar{q}_j \geq \bar{\xi}_j$, if $\bar{q}_i \geq \bar{\xi}(1+\beta)$, $\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i) \leq x^s_{ij}[2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} + 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i}]$, and if $\bar{q}_i < \bar{\xi}(1+\beta)$, $\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i) \leq x^s_{ij}[2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}]$. Equivalently, if $\bar{q}_i \geq \bar{\xi}(1+\beta)$, $\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i) \leq \frac{x^s_{ij}(q_i,q_j)}{1+\rho'}[2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} + 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i}]$, and if $\bar{q}_i < \bar{\xi}(1+\beta)$, $\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i) \leq \frac{x^s_{ij}(q_i,q_j)}{1+\rho'}[2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}]$. Thus

$$\begin{aligned}
&\sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} [\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i)]\ dq_i\ dq_j \\
&\quad \leq \sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, \bar{q}_i \geq \bar{\xi}(1+\beta), q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} \frac{x^s_{ij}(q_i, q_j)}{1+\rho'} \cdot \left[2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} + 2\frac{\nu}{\Delta}(1+\beta) \cdot \frac{\mathrm{SR}_i}{\bar{q}_i}\right]\ dq_i\ dq_j \\
&\qquad + \sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, \bar{q}_i < \bar{\xi}(1+\beta), q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} \frac{x^s_{ij}(q_i, q_j)}{1+\rho'} \cdot 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}\ dq_i\ dq_j \\
&\quad \leq \sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, q_i \geq \frac{1+\beta}{1+\nu}\bar{\xi}, q_j \leq q_{r_j}, \bar{q}_j \geq \bar{\xi}} 2\frac{\nu}{\Delta}(1+\beta)(1+\nu) \cdot \frac{\mathrm{SR}_i}{q_i} \cdot x^s_{ij}(q_i, q_j)\ dq_i\ dq_j \\
&\qquad + \sum_j \int_{q_{r_j} \geq q_j \geq \frac{\bar{\xi}}{1+\nu}} 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j(1+\nu)}{q_j}\ dq_j \quad \text{(using (22))}.
\end{aligned} \tag{23}$$

For (19), we note that as $q_j > q_{r_j}$, $\varphi_j(q_j) \leq 0$, and then when there is a small margin win by i, by definition, if $\bar{\xi}(1+\beta) \leq \bar{q}_i$, $\varphi_i(q_i) \leq 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j} + 2\frac{\nu}{\Delta}(1+\beta)\frac{\mathrm{SR}_i}{\bar{q}_i}$, and if $\bar{\xi}(1+\beta) > \bar{q}_i$, $\varphi_i(q_i) < 2\frac{\nu}{\Delta}\frac{\mathrm{SR}_j}{\bar{q}_j}$. Also, the constraint $\bar{q}_j \geq \bar{\xi}$ implies $q_j \geq \frac{\bar{\xi}}{(1+\nu)}$. Thus

$$\begin{aligned}
&\sum_i \sum_{j \neq i} \int_{q_i \leq q_{r_i}, q_j > q_{r_j}, \bar{q}_j \geq \bar{\xi}} \varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j)\ dq_i\ dq_j \\
&\quad \leq \sum_j \int_{q_j > \max\{q_{r_j}, \frac{\bar{\xi}}{1+\nu}\}} 2\frac{\nu}{\Delta}(1+\nu)\frac{\mathrm{SR}_j}{q_j}\ dq_j \quad \text{(using (22))} \\
&\qquad + \sum_i \sum_{j \neq i} \int_{q_j \geq q_{r_j}, \bar{q}_j \geq \bar{\xi}, \bar{q}_i \geq \bar{\xi}(1+\beta)} 2\frac{\nu}{\Delta}(1+\beta)(1+\nu)\frac{\mathrm{SR}_i}{q_i} \cdot x^s_{ij}(q_i, q_j)\ dq_i\ dq_j
\end{aligned} \tag{24}$$

Combining (23) and (24) yields

$$\begin{aligned}
&\sum_i \sum_{j\neq i} \int_{q_i \le q_{r_i}, q_j \le q_{r_j}, \bar{q}_j \ge \bar{\xi}} [\varphi_i(q_i) \cdot x^s_{ij}(q_i, q_j) - \varphi_j(q_j) \cdot \tilde{x}^s_{ji}(q_j, q_i)]\ dq_i\ dq_j \\
&\quad + \sum_i \sum_{j\neq i} \int_{q_i \le q_{r_i}, q_j > q_{r_j}, \bar{q}_j \ge \bar{\xi}} \varphi_i(q_i) x^s_{ij}(q_i, q_j)\ dq_i\ dq_j \\
&\quad \le \sum_j \int_{q_j \ge \frac{\bar{\xi}}{1+\nu}} 2\frac{\nu}{\Delta}(1+\nu)\frac{\mathrm{SR}_j}{q_j}\ dq_j \\
&\qquad + \sum_i \int_{q_i \ge \bar{\xi}\frac{(1+\beta)}{1+\nu}} 2\frac{\nu}{\Delta}(1+\beta)(1+\nu)\frac{\mathrm{SR}_i}{q_i}\ dq_i \qquad \text{(using (21))} \\
&\quad \le \sum_i (\rho+\rho')\mathrm{MR}_i + \sum_j 2\frac{\nu}{\Delta}(1+\beta)(1+\nu)\mathrm{SR}_i \ln\frac{(1+\nu)}{\bar{\xi}(1+\beta)} \\
&\qquad + \sum_i 2\frac{\nu}{\Delta}\mathrm{SR}_j(1+\nu)\ln\frac{1+\nu}{\bar{\xi}} \\
&\quad \le 4k\frac{\nu}{\Delta}(1+\beta)(1+\nu)\ln\frac{1+\nu}{\bar{\xi}}\mathrm{MR}.
\end{aligned} \tag{25}$$

□

We are now ready to bound Shtf.

Lemma 6.10.

$$\begin{aligned}
\mathrm{Shtf} \le \mathrm{MR}&\left[k\delta + k^2\delta + k\beta + k\bar{\xi}(1+\nu) + 2k\nu + (k-1)\bar{\xi}(1+\nu) + k(k-1)[(1+\beta)^2 - 1]\right] \\
&+ \mathrm{MR}\left[4k(1+\beta)(1+\nu)\frac{\nu}{\Delta}\cdot\ln\frac{1+\nu}{\bar{\xi}}\right].
\end{aligned}$$

Proof. In the event that $\mathcal{E}_a$ does not hold for some F_i, which occurs with probability at most $k\delta$, the contribution to Shtf is at most

$$\begin{aligned}
&k\delta\left[\sum_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)x_i(q_i)\ dq_i - \int_{q_i \le \bar{q}_{r_i}} \varphi_i(q_i)\tilde{x}_i(q_i)\ dq_i\right] \\
&\le k\delta\mathrm{MR} + k\delta\sum_i \int_{q_{r_i} < q_i \le \bar{q}_{r_i}} [-\varphi_i(q_i)]\ dq_i \\
&\le k\delta\mathrm{MR} + k\delta\sum_i \int_{q_i \le q_{r_i}} \varphi_i(q_i)\ dq_i \quad \left(\text{as } \textstyle\int_{q_i} \varphi_i(q_i)\ dq_i = 0\right) \\
&\le k\delta\mathrm{MR} + k\delta\sum_i \mathrm{SR}_i \le k\delta\mathrm{MR} + k^2\delta\mathrm{MR}.
\end{aligned}$$

Otherwise, the contribution is given by summing the bounds from (10), Lemmas 6.6–6.7, and 6.9. □

Proof of Theorem 3.1: We first choose $\Delta, \nu, \bar{\xi} \le \frac{1}{12}$. It is easy to check that then $(1+\beta)^2 - 1 = (1+\Delta)^2(1+\nu)^6 - 1 \le 2\Delta(1+\Delta)(1+\nu)^6 + 6\nu(1+\Delta)^2(1+\nu)^5 \le (2\Delta + 6\nu)\left(\frac{13}{12}\right)^7 \le 4\Delta + 11\nu$. Similarly, $\beta \le 2\Delta + 4\nu$, and $4(1+\beta)(1+\nu) \le 4\left(\frac{13}{12}\right)^5 \le 4\cdot\frac{3}{2} = 6$. Consequently,

$$\mathrm{Shtf} \le \mathrm{MR}\left[k\delta + k^2\delta + 2k\Delta + 4k\nu + 3k\bar{\xi} + 2k\nu + 4k(k-1)\Delta + 11k(k-1)\nu + 6k\frac{\nu}{\Delta}\ln\frac{(1+\nu)^2}{\bar{\xi}}\right].$$

It suffices that Shtf $\le \epsilon$MR. To this end, we bound the right hand side of the above expression by ϵ. To achieve this it suffices to choose ν, $\bar{\xi}$, δ, and Δ as follows:

$$\begin{aligned} k(k+1)\delta &= \frac{1}{4}\epsilon \\ (4k(k-1)+2k)\Delta &= \frac{1}{4}\epsilon \\ 3k\bar{\xi} &= \frac{1}{4}\epsilon \\ 4k\nu + 2k\nu + 11k(k-1)\nu + 6k\frac{\nu}{\Delta}\ln\frac{(1+\nu)^2}{\bar{\xi}} &= \frac{1}{4}\epsilon. \end{aligned}$$

It suffices that

$$\begin{aligned} \delta &= \Theta\left(\frac{\epsilon}{k^2}\right) \\ \Delta &= \Theta\left(\frac{\epsilon}{k^2}\right) \\ \bar{\xi} &= \Theta\left(\frac{\epsilon}{k}\right) \\ \nu &= \Theta\left(\frac{\epsilon^2}{k^3}\frac{1}{\ln\frac{k}{\epsilon}}\right). \end{aligned}$$

One final detail is that we need to set $\hat{\xi}$ also, but it suffices to note that $\hat{\xi} = \bar{\xi} + \frac{1}{2m}$ and so $\hat{\xi} = \Theta(\frac{\epsilon}{k}) + \frac{1}{2m}$.

By Lemma 6.2, $m = \Omega(\frac{1}{\gamma^3\hat{\xi}} + \ln\frac{1}{\delta}\cdot\frac{1}{\gamma^2\hat{\xi}})$ suffices. Recalling that $1+\nu = (1+\gamma)^2$, so $\gamma = \Theta(\nu)$, we obtain that $m = \Omega(\frac{k^{10}}{\epsilon^7}\ln^3\frac{k}{\epsilon})$ suffices. □

7 Conclusions

This paper proposes a general model for learning a near-optimal auction from data, in the form of i.i.d. samples from unknown distributions. It provides upper and lower bounds on the sample complexity for the case of single-item auctions, and shows that the number of samples required to obtain a $(1-\epsilon)$-approximation of the optimal expected revenue scales polynomially with both the number k of bidders and with $\frac{1}{\epsilon}$. We conclude by listing some of the many directions in which this work could be extended.

1. Prove tight bounds (in terms of k and $\frac{1}{\epsilon}$) on how many samples are necessary and sufficient to achieve expected revenue at least $1-\epsilon$ times the maximum possible. (See [28, 32, 16] for recent progress.)

2. Prove good sample complexity upper bounds for settings other than single-item auctions. (See [32] for recent progress.)

3. Are there natural settings where the learning problem is information-theoretically easy (meaning polynomial sample complexity) yet computationally hard (under complexity assumptions)?

4. Identify multi-parameter problems, less general than those in [19], where a near-optimal mechanism can be learned from a polynomial number of samples.

5. For problems where a $(1-\epsilon)$-approximate mechanism cannot be learned from a polynomial number of samples, identify the best-possible approximation factor for which an approximately optimal mechanism can be efficiently learned.

6. If some of the bidders that contribute the samples are the same as the bidders that participate in the final auction, then these bidders need not bid truthfully. (Underbidding could result in lower payments in the future.) Is it still possible to learn a near-optimal auction in this setting?

References

[1] M. Anthony and P. Bartlett. *Neural Network Learning: Theoretical Foundations.* Cambridge University Press, 1999.

[2] P. D. Azar and R. D. Kleinberg abd S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete pages = 1358–1377, year = 2014.*

[3] P. D. Azar, C. Daskalakis, S. Micali, and S. M. Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 596–604, 2013.

[4] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Innovations in Computer Science (ICS).* Tsinghua University Press, January 2011.

[5] Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, 74(8):1245–1270, 2008.

[6] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003. Article 5.

[7] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

[8] N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.

[9] S. Chawla, J. D. Hartline, and R. D. Kleinberg. Algorithmic pricing via virtual valuations. In *EC*, pages 243–251, 2007.

[10] S. Chawla, J. D. Hartline, and D. Nekipelov. Mechanism design for data science. In *Proceedings of the Fifteenth ACM Conf. on Economics and Computation*, pages 711–712, 2014.

[11] Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. Sequential posted pricing and multi-parameter mechanism design. In *STOC*, 2010.

[12] A. Chiesa, S. Micali, and Z. A. Zhu. Mechanism design with approximate valuations. In *Proceedings of Innovations in Theoretical Computer Science*, pages 34–38, 2012.

[13] R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 243–252, 2014.

[14] Richard Cole and Shravas Rao. Applications of α-strongly regular distributions to bayesian auctions.

[15] Jacques Crémer and Richard P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361, 1985.

[16] N. Devanur, Z. Huang, and C. A. Psomas. Sample complexity of auctions with side information. Working paper, 2015.

[17] Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. Prior-independent multi-parameter mechanism design. In *Workshop on Internet and Network Economics (WINE)*, 2011.

[18] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *EC*, 2010.

[19] S. Dughmi, L. Han, and N. Nisan. Sampling and representation complexity of revenue maximization. In *Workshop on Internet and Network Economics (WINE)*, pages 277–291, 2014.

[20] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745, 2007.

[21] Hu Fu, Nima Haghpanah, Jason D. Hartline, and Robert Kleinberg. Optimal auctions for correlated buyers with sampling. In *EC*, pages 23–36, 2014.

[22] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[23] B. Q. Ha and J. D. Hartline. Mechanism design via consensus estimates, cross checking, and profit extraction. *ACM Transactions on Economics and Computation*, 1(2):8, 2013.

[24] J. D. Hartline. Mechanism design and approximation. Book draft. October, 2013.

[25] J. D. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.

[26] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *EC*, 2009.

[27] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*, WWW '08, pages 189–198, New York, NY, USA, 2008. ACM.

[28] Z. Huang, Y. Mansour, and T. Roughgarden. Making the most of your samples. In *16th Annual ACM Conference on Economics and Computation (EC)*, pages 45–60, 2015.

[29] Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*, pages 594–605, New York, New York, USA., 2003. IEEE Computer Society.

[30] Andres Munoz Medina and Mehryar Mohri. Learning theory and algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of The 31st Intl. Conf. on Machine Learning*, pages 262–270, 2014.

[31] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, New York, NY, USA, 2005.

[32] J. Morgenstern and T. Roughgarden. The psuedo-dimension of near-optimal auctions. In *28th Annual Conference on Neural Information Processing Systems (NIPS)*, 2015.

[33] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[34] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

[35] M. Ostrovsky and M. Schwarz. Reserve prices in internet advertising auctions: A field experiment. Working paper, December 2009.

[36] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *EC*, 2012.

[37] I. Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529, 2003.

[38] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.