

Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues

Noga Alon* Moshe Babaioff† Yannai A. Gonczarowski‡ Yishay Mansour§
Shay Moran¶ Amir Yehudayoff||

May 23, 2017

Abstract

In this work we derive a variant of the classic Glivenko-Cantelli Theorem, which asserts uniform convergence of the empirical Cumulative Distribution Function (CDF) to the CDF of the underlying distribution. Our variant allows for tighter convergence bounds for extreme values of the CDF.

We apply our bound in the context of *revenue learning*, which is a well-studied problem in economics and algorithmic game theory. We derive sample-complexity bounds on the uniform convergence rate of the empirical revenues to the true revenues, assuming a bound on the k th moment of the valuations, for any (possibly fractional) $k > 1$.

For uniform convergence in the limit, we give a complete characterization and a zero-one law: if the first moment of the valuations is finite, then uniform convergence almost surely occurs; conversely, if the first moment is infinite, then uniform convergence almost never occurs.

1 Introduction

A basic task in machine learning is to learn an unknown distribution μ , given access to samples from it. A natural and widely studied criterion for learning a distribution is approximating its Cumulative Distribution Function (CDF). The seminal Glivenko-Cantelli Theorem (Glivenko, 1933; Cantelli, 1933) addresses this question when the distribution μ is over the real numbers. It determines the behavior of the empirical distribution function as the number of samples grows: let X_1, X_2, \dots be a sequence of i.i.d. random variables drawn from a distribution μ on \mathbb{R}

*Tel Aviv University, Israel and Microsoft Research, Herzliya, Israel. nogaa@tau.ac.il. Research supported in part by an ISF grant and by a GIF grant.

†Microsoft Research, Herzliya, Israel. moshe@microsoft.com.

‡The Hebrew University of Jerusalem, Israel and Microsoft Research, Herzliya, Israel. yannai@gonch.name. Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. His work is supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 740282).

§Tel Aviv University, Israel and Microsoft Research, Herzliya, Israel. mansour@tau.ac.il. This research was supported in part by The Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), by a grant from the Israel Science Foundation, and by a grant from United States-Israel Binational Science Foundation (BSF).

¶University of California, San Diego, Simons Institute for the Theory of Computing, Berkeley, and Max Planck Institute for Informatics, Saarbrücken, Germany. shaymoran1@gmail.com.

||Technion — Israel Institute of Technology, Israel. amir.yehudayoff@gmail.com. Research supported by ISF grant 1162/15.

with Cumulative Distribution Function (CDF) F , and let x_1, x_2, \dots be their realizations. The *empirical distribution* μ_n is

$$\mu_n \triangleq \frac{1}{n} \sum_{i=1}^n \delta_{x_i},$$

where δ_{x_i} is the constant distribution supported on x_i . Let F_n denote the CDF of μ_n , i.e., $F_n(t) \triangleq \frac{1}{n} \cdot |\{1 \leq i \leq n : x_i \leq t\}|$. The Glivenko-Cantelli Theorem formalizes the statement that μ_n converges to μ as n grows, by establishing that $F_n(t)$ converges to $F(t)$, uniformly over all $t \in \mathbb{R}$:

Theorem 1.1 (Glivenko-Cantelli Theorem, 1933). *Almost surely,*

$$\lim_{n \rightarrow \infty} \sup_t |F_n(t) - F(t)| = 0.$$

Some twenty years after Glivenko and Cantelli discovered this theorem, Dvoretzky, Kiefer, and Wolfowitz (DKW) strengthened this result by giving an almost¹ tight quantitative bound on the convergence rate. In 1990, Massart proved a tight inequality, confirming a conjecture due to Birnbaum and McCarty (1958):

Theorem 1.2 (Massart, 1990). $\Pr \left[\sup_t |F_n(t) - F(t)| > \epsilon \right] \leq 2 \exp(-2n\epsilon^2)$ for all $\epsilon > 0$, $n \in \mathbb{N}$.

The above theorems show that, with high probability, F and F_n are close up to some *additive* error. We would have liked to prove a stronger, *multiplicative* bound on the error:

$$\forall t : |F(t) - F_n(t)| \leq \epsilon \cdot F(t).$$

However, for some distributions, the above event has probability 0, no matter how large n is. For example, assume that μ satisfies $F(t) > 0$ for all t . Since the empirical measure μ_n has finite support, there is t with $F_n(t) = 0$; for such a value of t , such a multiplicative approximation fails to hold.

So, the above multiplicative requirement is too strong to hold in general. A natural compromise is to consider a *submultiplicative* bound:

$$\forall t : |F(t) - F_n(t)| \leq \epsilon \cdot F(t)^\alpha,$$

where $0 \leq \alpha < 1$. When $\alpha = 0$, this is the additive bound studied in the context of the Glivenko-Cantelli Theorem. When $\alpha = 1$, this is the unattainable multiplicative bound. Our first main result shows that the case of $\alpha < 1$ is attainable:

Theorem 1.3 (Submultiplicative Glivenko-Cantelli Theorem). *Let $\epsilon > 0$, $\delta > 0$ and $0 \leq \alpha < 1$. There exists $n_0(\epsilon, \delta, \alpha)$ such that for all $n > n_0$, with probability $1 - \delta$:*

$$\forall t : |F(t) - F_n(t)| \leq \epsilon \cdot F(t)^\alpha.$$

It is worth pointing out a central difference between Theorem 1.3 and other generalizations of the Glivenko-Cantelli Theorem: for example, the seminal work of Vapnik and Chervonenkis (1971) shows that for every class of events \mathcal{F} of VC dimension d , there is $n_0 = n_0(\epsilon, \delta, d)$ such that for every $n \geq n_0$, with probability $1 - \delta$ it holds that $\forall A \in \mathcal{F} : |p(A) - p_n(A)| \leq \epsilon$. This yields Glivenko-Cantelli by plugging $\mathcal{F} = \{(-\infty, t] : t \in \mathbb{R}\}$, which has VC dimension 1. In contrast, the submultiplicative bound from Theorem 1.3 does not even extend to the VC dimension 1 class $\mathcal{F} = \{\{t\} : t \in \mathbb{R}\}$. Indeed, pick any distribution p over \mathbb{R} such that $p(\{t\}) = 0$ for every t , and observe that for every sample x_1, \dots, x_n , it holds that $p_n(\{x_i\}) \geq 1/n$, however $p(\{x_i\}) = 0$, and therefore, as long as $\alpha > 0$, it is never the case that $|p(\{x_i\}) - p_n(\{x_i\})| \leq p(\{x_i\})^\alpha$. Theorem 1.3 is proven in Section 3, which also includes other extensions.

Our second main result gives an explicit upper bound on $n_0(\epsilon, \delta, \alpha)$:

¹The inequality due to Dvoretzky et al. (1956) has a larger constant C in front of the exponent on the right hand side.

Theorem 1.4 (Submultiplicative Glivenko-Cantelli Bound). *Let $\epsilon, \delta \leq 1/4$, and $\alpha < 1$. Then*

$$n_0(\epsilon, \delta, \alpha) \leq \max \left\{ \frac{\ln(6/\delta)}{2\epsilon^2} \left(\frac{\epsilon\delta}{3}\right)^{-\frac{4\alpha}{1-\alpha}}, \quad (D+1) \left(10 \cdot \ln \left(12 \cdot \frac{D+4}{\delta(1-\alpha)} \right) \right)^{\frac{4\alpha}{1-\alpha}} \right\},$$

where $D = \frac{\ln(6/\delta)}{2\epsilon^2} \left(\frac{\epsilon\delta}{6} \cdot \ln\left(\frac{1+\alpha}{2\alpha}\right) \right)^{-\frac{4\alpha}{1-\alpha}}$.

Note that for fixed ϵ, δ , when $\alpha \rightarrow 0$ the above bound approaches the familiar $O\left(\frac{\ln(1/\delta)}{\epsilon^2}\right)$ bound by DKW and Massart for $\alpha = 0$. On the other hand, when $\alpha \rightarrow 1$ the above bound tends to ∞ , reflecting the fact that the multiplicative variant of Glivenko-Cantelli ($\alpha = 1$) does not hold. Theorem 1.4 is proven in Appendix A.

Note that the dependence on the confidence parameter δ is polynomial. This contrasts standard uniform convergence rates, which due to applications of concentration bounds such as Chernoff/Hoeffding, achieve logarithmic dependencies on δ . These concentration bounds are not applicable in our setting, when the CDF values are very small, and we use Markov's inequality instead.

The following example shows that a polynomial dependency on δ is indeed necessary and is not due to a limitation of our proof.

An example. Consider n independent samples x_1, \dots, x_n from the uniform distribution over $[0, 1]$, and set $\alpha = 1/2$ and $\epsilon = 1$. The probability of the event

$$“\exists i : x_i \leq 1/n^3”$$

is roughly $1/n^2$: indeed, the complementary event has probability $(1 - 1/n^3)^n \approx \exp(-1/n^2) \approx 1 - 1/n^2$. When this happens, we have: $F_n(1/n^3) \geq 1/n \gg 1/n^3 + 1/n^{3/2} = F(1/n^3) + [F(1/n^3)]^{1/2}$. Note that this happens with probability inverse polynomial in n (roughly $1/n^2$) and not inverse exponential.

An application to revenue learning. We demonstrate an application of our Submultiplicative Glivenko-Cantelli Theorem in the context of a widely studied problem in economics and algorithmic game theory: the problem of revenue learning. In the setting of this problem, a seller has to decide which price to post for a good she wishes to sell. Assume that each consumer draws her private valuation for the good from an unknown distribution μ . We envision that a consumer with valuation v will buy the good at any price $p \leq v$, but not at any higher price. This implies that the expected revenue at price p is simply $r(p) \triangleq p \cdot q(p)$, where $q(p) = \Pr_{V \sim \mu}[V \geq p]$.

In the language of machine learning, this problem can be phrased as follows: the examples domain $Z \triangleq \mathbb{R}^+$ is the set of all valuations v . The hypothesis space $H \triangleq \mathbb{R}^+$ is the set of all prices p . The revenue (which is a gain, rather than loss) of a price p on a valuation v is the function $p \cdot 1_{\{p \leq v\}}$.

The well-known *revenue maximization* problem is to find a price p^* that maximizes the expected revenue, given a sample of valuations drawn i.i.d. from μ . In this paper, we consider the more demanding *revenue estimation* problem: the problem of well-approximating $r(p)$, simultaneously for all prices p , from a given sample of valuations. (This clearly also implies a good estimation of the maximum revenue and of a price that yields it.) More specifically, we address the following question: when do the *empirical revenues*, $r_n(p) \triangleq p \cdot q_n(p)$, uniformly converge to the true revenues $r(p)$? More specifically, we would like to show that for some n_0 , for $n \geq n_0$ we have with probability $1 - \delta$ that

$$|r(p) - r_n(p)| \leq \epsilon.$$

The revenue estimation problem is a basic instance of the more general problem of uniform convergence of empirical estimates. The main challenge in this instance is that the prices are unbounded (and so are the private valuations that are drawn from the distribution μ).

Unfortunately, there is no (upper) bound on n_0 that is only a function of ϵ and δ . Moreover, even if we add the expectation of valuations, i.e., $\mathbb{E}[V]$ where V is distributed according to μ , still there is no bound on n_0 that is a function of only those three parameters (see Section 2.3 for an example). In contrast, when we consider higher moments of the distribution μ , we are able to derive bounds on the value of n_0 . These bounds are based on our Submultiplicative Glivenko-Cantelli Bound. Specifically, assume that $\mathbb{E}_{V \sim \mu}[V^{1+\theta}] \leq C$ for some $\theta > 0$ and $C \geq 1$. Then, we show that for any $\epsilon, \delta \in (0, 1)$, we have

$$\Pr \left[\exists v : |r(v) - r_n(v)| > \epsilon \right] \leq \Pr \left[\exists v : |q(v) - q_n(v)| > \frac{\epsilon}{C^{\frac{1}{1+\theta}}} q(v)^{\frac{1}{1+\theta}} \right].$$

This essentially reduces uniform convergence bounds to our Submultiplicative Glivenko-Cantelli variant. It then follows that there exists $n_0(C, \theta, \epsilon, \delta)$ such that for any $n \geq n_0$, with probability at least $1 - \delta$,

$$\forall v : |r_n(v) - r(v)| \leq \epsilon.$$

We remark that when θ is large, our bound yields $n_0 \approx O\left(\frac{\ln(1/\delta)}{\epsilon^2}\right)$, which recovers the standard sample complexity bounds obtainable via DKW and Massart.

When $\theta \rightarrow 0$, our bound diverges to infinity, reflecting the fact (discussed above) that there is no bound on n_0 that depends only on ϵ, δ , and $\mathbb{E}[V]$. Nevertheless, we find that $\mathbb{E}[V]$ qualitatively determines whether uniform convergence occurs in the limit. Namely, we show that

- If $\mathbb{E}_\mu[V] < \infty$, then almost surely $\lim_{n \rightarrow \infty} \sup_v |r(v) - r_n(v)| = 0$,
- Conversely, if $\mathbb{E}_\mu[V] = \infty$, then almost never $\lim_{n \rightarrow \infty} \sup_v |r(v) - r_n(v)| = 0$.

1.1 Related work

Generalizations of Glivenko-Cantelli. Various generalizations of the Glivenko-Cantelli Theorem were established. These include uniform convergence bounds for more general classes of functions as well as more general loss functions (for example, Vapnik and Chervonenkis, 1971; Vapnik, 1998; Koltchinskii and Panchenko, 2000; Bartlett and Mendelson, 2002). The results that concern unbounded loss functions are most relevant to this work (for example, Cortes et al., 2010, 2013; Vapnik, 1998). We next briefly discuss the relevant results from Cortes et al. (2013) in the context of this paper; more specifically, in the context of Theorem 1.3. To ease presentation, set α in this theorem to be $1/2$. Theorem 1.3 analyzes the event where the empirical quantile is bounded by

$$\begin{aligned} q_n(p) &\leq q(p) + \epsilon \sqrt{q(p)}, \\ q_n(p) &\geq q(p) - \epsilon \sqrt{q(p)}. \end{aligned}$$

whereas, Cortes et al. (2013) analyzes the event where it is bounded it by:

$$\begin{aligned} q_n(p) &\leq \tilde{O}(q(p) + \sqrt{q(p)/n} + 1/n), \\ q_n(p) &\geq \tilde{\Omega}(q(p) - \sqrt{q_n(p)/n} - 1/n) \end{aligned}$$

Thus, the main difference is the additive $1/n$ term in the bound from Cortes et al. (2013). In the context of uniform convergence of revenues, it is crucial to use the upper bound on the empirical quantile as we do, as it guarantees that large prices will not overfit, which is the main

challenge in proving uniform convergence in this context. In particular, the upper bound from Cortes et al. (2013) does not provide any guarantee on the revenues of prices $p \gg n$, as for such prices $p \cdot 1/n \gg 1$.

It is also worth pointing out that our lower bound on the empirical quantile implies that with high probability the quantile of the maximum sampled point is at least $1/n^2$ (or more generally, at least $1/n^{1/\alpha}$ when $\alpha \neq 1/2$), while the bound from Cortes et al. (2013) does not imply any non-trivial lower bound.

Another, more qualitative difference is that unlike the bounds in Cortes et al. (2013) that apply for general VC classes, our bound is tailored for the class of thresholds (corresponding to CDF/quantiles), and does not extend even to other classes of VC dimension 1 (see the discussion after Theorem 1.3).

Uniform convergence of revenues. The problem of *revenue maximization* is a central problem in economics and Algorithmic Game Theory (AGT). The seminal work of Myerson (1981) shows that given a valuation distribution for a single good, the revenue-maximizing selling mechanism for this good is a posted-price mechanism. In the recent years, there has been a growing interest in the case where the valuation distribution is unknown, but the seller observes samples drawn from it. Most papers in this direction assume that the distribution meets some tail condition that is considered “natural” within the algorithmic game theory community, such as boundedness (Morgenstern and Roughgarden, 2015; Roughgarden and Schrijvers, 2016; Morgenstern and Roughgarden, 2016; Balcan et al., 2016; Gonczarowski and Nisan, 2017; Devanur et al., 2016)², such as a condition known as Myerson-regularity (Dhangwatnotai et al., 2015; Huang et al., 2015; Cole and Roughgarden, 2014; Devanur et al., 2016), or such as a condition known as monotone hazard rate (Huang et al., 2015).³ These papers then go on to derive computation- or sample-complexity bounds on learning an optimal price (or an optimal selling mechanism from a given class) for a distribution that meets the assumed condition.

A recurring theme in statistical learning theory is that learnability guarantees are derived via a, sometimes implicit, uniform convergence bound. However, this has not been the case in the context of revenue learning. Indeed, while some papers that studied bounded distributions (Morgenstern and Roughgarden, 2015; Roughgarden and Schrijvers, 2016; Morgenstern and Roughgarden, 2016; Balcan et al., 2016) did use uniform convergence bounds as part of their analysis, other papers, in particular those that considered unbounded distributions, had to bypass the usage of uniform convergence by more specialized arguments. This is due to the fact that many unbounded distributions do not satisfy any uniform convergence bound. As a concrete example, the (unbounded, Myerson-regular) *equal revenue distribution*⁴ has an infinite expectation and therefore, by our Theorem 2.3, satisfies no uniform convergence, even in the limit. Thus, it turns out that the works that studied the popular class of Myerson-regular distributions (Dhangwatnotai et al., 2015; Huang et al., 2015; Cole and Roughgarden, 2014; Devanur et al., 2016) indeed could not have hoped to establish learnability via a uniform convergence argument. For instance, the way Dhangwatnotai et al. (2015) and Cole and Roughgarden (2014) establish learnability for Myerson-regular distributions is by considering the guarded ERM algorithm (an algorithm that chooses an empirical revenue maximizing price that is smaller than, say, the \sqrt{n} th largest sampled price), and proving a uniform convergence bound, not for all

²The analysis of Balcan et al. (2016) assumes a bound on the realized revenue (from any possible valuation profile) of any mechanism/auction in the class that they consider. For the class of posted-price mechanisms, this is equivalent to assuming a bound on the support of the valuation distribution. Indeed, for any valuation v , pricing at v gives realized revenue v (from the valuation v), and so unbounded valuations (together with the ability to post unbounded prices) imply unbounded realized revenues.

³Both Myerson-regularity and monotone hazard rate are conditions on the second derivative of the revenue as a function of the quantile of the underlying distribution. In particular, they impose restrictions on the tail of the distribution.

⁴This is a distribution that satisfies the special property that all values have the same expected revenue.

prices, but only for prices that are, say, smaller than the \sqrt{n} th largest sampled price, and then arguing that larger prices are likely to have a small empirical revenue, compared to the guarded empirical revenue maximizer. This means that the guarded ERM will output a good price, but it does not (and cannot) imply uniform convergence for all prices.

We complement the extensive literature surveyed above in a few ways. The first is generalizing the revenue maximization problem to a revenue estimation problem, where the goal is to uniformly estimate the revenue of all possible prices, when no bound on the possible valuations is given (or even exists). The problem of revenue estimation arises naturally when the seller has additional considerations when pricing her good, such as regulations that limit the price choice, bad publicity if the price is too high (or, conversely, damage to prestige if the price is too low), or willingness to suffer some revenue loss for better market penetration (which may translate to more revenue in the future). In such a case, the seller may wish to estimate the revenue loss due to posting a discounted (or inflated) price.

The second, and most important, contribution to the above literature is that we consider arbitrary distributions rather than very specific and limited classes of distributions (e.g., bounded, Myerson-regular, monotone hazard rate, etc.). Third, we derive finite sample bounds in the case that the expected valuation is bounded for some moment larger than 1. We further derive a zero-one law for uniform convergence in the limit that depends on the finiteness of the first moment. Technically, our bounds are based on an additive error rather than multiplicative ones, which are popular in the AGT community.

1.2 Paper organization

The rest of the paper is organized as follows. We begin by presenting the application of our Submultiplicative Glivenko-Cantelli to revenue estimation in Section 2. In Section 3, we prove the Submultiplicative Glivenko-Cantelli variant, and discuss some extensions of it. Section 4 contains a discussion and possible directions of future work. Some of the proofs are deferred to the appendices.

2 Uniform Convergence of Empirical Revenues

In this section we demonstrate an application of our Submultiplicative Glivenko-Cantelli variant by establishing uniform convergence bounds for a family of unbounded random variables in the context of revenue estimation.

2.1 Model

Consider a good g that we wish to post a price for. Let V be a random variable that models the valuation of a random consumer for g . Technically, it is assumed that V is a nonnegative random variable, and we denote by μ its induced distribution over \mathbb{R}^+ . A consumer who values g at a valuation v is willing to buy the good at any price $p \leq v$, but not at any higher price. This implies that the realized revenue to the seller from a (posted) price p is the random variable $p \cdot 1_{\{p \leq V\}}$. The *quantile* of a value $v \in \mathbb{R}^+$ is

$$q(v) = q(v; \mu) \triangleq \mu(\{x : x \geq v\}).$$

This models the fraction of the consumers in the population that are willing to purchase the good if priced at v . The expected revenue from a (posted) price $p \in \mathbb{R}^+$ is

$$r(p) = r(p; \mu) \triangleq \mathbb{E}_{\mu}[p \cdot 1_{\{p \leq V\}}] = p \cdot q(p).$$

Let V_1, V_2, \dots be a sequence of i.i.d. valuations drawn from μ , and let v_1, v_2, \dots be their realizations. The *empirical quantile* of a value $v \in \mathbb{R}^+$ is

$$q_n(v) = q(v; \mu_n) \triangleq \frac{1}{n} \cdot |\{1 \leq i \leq n : v_i \geq v\}|.$$

The *empirical revenue* from a price $p \in \mathbb{R}^+$ is

$$r_n(p) = r(p; \mu_n) \triangleq \mathbb{E}_{\mu_n} [p \cdot \mathbf{1}_{\{p \leq V\}}] = p \cdot q_n(p).$$

The revenue estimation error for a given sample of size n is

$$\epsilon_n \triangleq \sup_p |r_n(p) - r(p)|.$$

It is worth highlighting the difference between revenue estimation and revenue maximization. Let p^* be a price that maximizes the revenue, i.e., $p^* \in \arg \sup_p r(p)$. The maximum revenue is $r^* = r(p^*)$. The goal in many works in revenue maximization is to find a price \hat{p} such that $r^* - r(\hat{p}) \leq \epsilon$, or alternatively, to bound $r^*/r(\hat{p})$.

Given a revenue-estimation error ϵ_n , one can clearly maximize the revenue within an additive error of $2\epsilon_n$ by simply posting a price $p_n^* \in \arg \max_p r_n(p)$, thereby attaining revenue $r_n^* = r(p_n^*)$. This follows since

$$r_n^* = r(p_n^*) \geq r_n(p_n^*) - \epsilon_n \geq r_n(p^*) - \epsilon_n \geq r(p^*) - 2\epsilon_n = r^* - 2\epsilon_n.$$

Therefore, good revenue estimation implies good revenue maximization.

We note that the converse does not hold. Namely, there are distributions for which revenue maximization is trivial but revenue estimation is impossible. One such case is the *equal revenue distribution*, where all values in the support of μ have the same expected revenue. For such distributions, the problem of revenue maximization becomes trivial, since any posted price is optimal. However, as follows from Theorem 2.3, since the expected revenue of such distributions is infinite, almost never do the empirical revenues uniformly converge to the true revenues.

2.2 Quantitative bounds on the uniform convergence rate

Recall that we are interested in deriving sample bounds that would guarantee uniform convergence for the revenue estimation problem. We will show that given an upper bound on the k th moment of V for some $k > 1$, we can derive a finite sample bound. To this end we utilize our Submultiplicative Glivenko-Cantelli Bound (Theorem 1.4).

We also consider the case of $k = 1$, namely that $\mathbb{E}[V]$ is bounded, and show that in this case there is still uniform convergence in the limit, but that there cannot be any guarantees on the convergence rate. Interestingly, it turns out that $\mathbb{E}[V] < \infty$ is not only sufficient but also necessary so that in the limit, the empirical revenues uniformly converge to the true revenues (see Section 2.3).

We begin by showing that bounds on the k th moment for $k > 1$ yield explicit bounds on the convergence rate. It is convenient to parametrize by setting $k = 1 + \theta$, where $\theta > 0$.

Theorem 2.1. *Let $\mathbb{E}_{V \sim \mu}[V^{1+\theta}] \leq C$ for some $\theta > 0$ and $C \geq 1$, and let $\epsilon, \delta \in (0, 1)$. Set⁵*

$$n_0 = \tilde{O} \left(\frac{\ln(1/\delta)}{\epsilon^2} C^{\frac{2}{1+\theta}} \left(\frac{6 \cdot C^{\frac{1}{1+\theta}}}{\epsilon \delta \ln(1 + \theta/2)} \right)^{4/\theta} \right). \quad (1)$$

For any $n \geq n_0$, with probability at least $1 - \delta$,

$$\forall v : |r_n(v) - r(v)| \leq \epsilon.$$

⁵The \tilde{O} conceals low order terms.

Note that when θ is large, this bound approaches the standard $O\left(\frac{\ln(1/\delta)}{\epsilon^2}\right)$ sample complexity bound of the additive Glivenko-Cantelli. For example, if all moments are uniformly bounded, then the convergence is roughly as fast as in standard uniform convergence settings (e.g., VC-dimension based bounds).

The proof of Theorem 2.1 follows from Theorem 1.4 and the next proposition, which reduces bounds on the uniform convergence rate of the empirical revenues to our Submultiplicative Glivenko-Cantelli.

Proposition 2.2. *Let $\mathbb{E}_{V \sim \mu}[V^{1+\theta}] \leq C$ for some $\theta > 0$ and $C \geq 1$, and let $\epsilon, \delta \in (0, 1)$. Then,*

$$\Pr\left[\exists v : |r(v) - r_n(v)| > \epsilon\right] \leq \Pr\left[\exists v : |q(v) - q_n(v)| > \frac{\epsilon}{C^{\frac{1}{1+\theta}}} q(v)^{\frac{1}{1+\theta}}\right].$$

Thus, Theorem 2.1 follows by plugging $\epsilon \leftarrow \frac{\epsilon}{C^{\frac{1}{1+\theta}}}$ and $\alpha \leftarrow \frac{1}{1+\theta}$ in Theorem 1.4.

Proof of Proposition 2.2. By Markov's inequality:

$$q(v) = \Pr[V \geq v] = \Pr[V^{1+\theta} \geq v^{1+\theta}] \leq \frac{C}{v^{1+\theta}}. \quad (2)$$

Now,

$$\begin{aligned} \Pr\left[\exists v : |r(v) - r_n(v)| > \epsilon\right] &= \Pr\left[\exists v : |v \cdot q(v) - v \cdot q_n(v)| > \epsilon\right] \\ &= \Pr\left[\exists v : |v \cdot q(v) - v \cdot q_n(v)| > \frac{\epsilon}{(v^{1+\theta} \cdot q(v))^{\frac{1}{1+\theta}}} (v^{1+\theta} \cdot q(v))^{\frac{1}{1+\theta}}\right] \\ &\leq \Pr\left[\exists v : |v \cdot q(v) - v \cdot q_n(v)| > \frac{\epsilon}{C^{\frac{1}{1+\theta}}} (v^{1+\theta} \cdot q(v))^{\frac{1}{1+\theta}}\right] \\ &= \Pr\left[\exists v : |q(v) - q_n(v)| > \frac{\epsilon}{C^{\frac{1}{1+\theta}}} q(v)^{\frac{1}{1+\theta}}\right]. \end{aligned}$$

where the inequality follows from Equation (2). \square

2.3 A qualitative characterization of uniform convergence

The sample complexity bounds in Theorem 2.1 are meaningful as long as $\theta > 0$, but deteriorate drastically as $\theta \rightarrow 0$. Indeed, as the following example shows, there is no bound on the uniform convergence sample complexity that depends only on the first moment of V , i.e., its expectation.

Consider a distribution η_p so that with probability p we have $V = 1/p$ and otherwise $V = 0$. Clearly, $\mathbb{E}[V] = 1$. However, we need to sample $m_p = O(1/p)$ valuations to see a single nonzero value. Therefore, there is no bound on the sample size m_p as a function of the expectation, which is simply 1.

We can now consider the higher moments of η_p . Consider the k th moment, for $k = 1 + \theta$ and $\theta > 0$, so $k > 1$. For this moment, we have $A_{p,\theta} = \mathbb{E}[V^{1+\theta}] = p^{\theta/(1+\theta)}$, which implies that $m_p = O(1/(A_{p,\theta})^{(1+\theta)/\theta})$. This does allow us to bound m_p as a function of θ and $\mathbb{E}[V^{1+\theta}]$, but for small θ we have a huge exponent of approximately $1/\theta$.

While the above examples show that there cannot be a bound on the sample size as a function of the expectation of the value, it turns out that there is a very tight connection between the first moment and uniform convergence:

Theorem 2.3. *The following dichotomy holds for a distribution μ on \mathbb{R}^+ :*

1. *If $\mathbb{E}_\mu[V] < \infty$, then almost surely $\lim_{n \rightarrow \infty} \sup_v |r(v) - r_n(v)| = 0$.*
2. *If $\mathbb{E}_\mu[V] = \infty$, then almost never $\lim_{n \rightarrow \infty} \sup_v |r(v) - r_n(v)| = 0$.*

That is, the empirical revenues uniformly converge to the true revenues if and only if $\mathbb{E}_\mu[V] < \infty$.

We use the following basic fact in the Proof of Theorem 2.3:

Lemma 2.4. *Let X be a nonnegative random variable. Then*

$$\sum_{n=1}^{\infty} \Pr[X \geq n] \leq \mathbb{E}[X] \leq \sum_{n=0}^{\infty} \Pr[X \geq n].$$

Proof. Note that:

$$\sum_{n=1}^{\infty} \mathbf{1}_{\{X \geq n\}} = \lfloor X \rfloor \leq X \leq \lfloor X \rfloor + 1 = \sum_{n=0}^{\infty} \mathbf{1}_{\{X \geq n\}}.$$

The lemma follows by taking expectations. \square

Proof of Theorem 2.3. We start by proving item 2. Let μ be a distribution such that $\mathbb{E}_\mu[V] = \infty$. If $\sup_v v \cdot q(v) = \infty$ then for every realization v_1, \dots, v_n there is some $v \geq \max\{v_1, \dots, v_n\}$ such that $v \cdot q(v) \geq 1$, but $v \cdot q_n(v) = 0$. So, we may assume $\sup_v v \cdot q(v) < \infty$. Without loss of generality we may assume that $\sup_v v \cdot q(v) = 1/2$ by rescaling the distribution if needed. Consider the sequence of events E_1, E_2, \dots where E_n denotes the event that $V_n \geq n$. Since $\mathbb{E}_\mu[V] = \infty$, Lemma 2.4 implies that $\sum_{n=1}^{\infty} \Pr[E_n] = \infty$. Thus, since these events are independent, the second Borel-Cantelli Lemma (Borel, 1909; Cantelli, 1917) implies that almost surely, infinitely many of them occur and so infinitely often

$$V_n \cdot q_n(V_n) \geq 1 \geq V_n \cdot q(V_n) + \frac{1}{2}.$$

Therefore, the probability that $v \cdot q_n(v)$ uniformly converge to $v \cdot q(v)$ is 0.

Item 1 follows from the following monotone domination theorem:

Theorem 2.5. *Let \mathcal{F} be a family of nonnegative monotone functions, and let F be an upper envelope⁶ for \mathcal{F} . If $\mathbb{E}_\mu[F] < \infty$, then almost surely:*

$$\lim_{n \rightarrow \infty} \sup_{f \in \mathcal{F}} \left| \mathbb{E}[f] - \mathbb{E}[f] \right| = 0.$$

Indeed, item 1 follows by plugging $\mathcal{F} = \{v \cdot \mathbf{1}_{x \geq v} : v \in \mathbb{R}^+\}$, which is uniformly bounded by the identity function $F(x) = x$. Now, by assumption $\mathbb{E}_\mu[F] < \infty$, and therefore, almost surely

$$\lim_{n \rightarrow \infty} \sup_{v \in \mathbb{R}^+} |r(v) - r_n(v)| = \lim_{n \rightarrow \infty} \sup_{f \in \mathcal{F}} \left| \mathbb{E}[f] - \mathbb{E}[f] \right| = 0. \quad \square$$

Theorem 2.5 follows by known results in the theory of empirical processes (for example, with some work it can be proved using Theorem 2.4.3 from Vaart and Wellner (1996)). For completeness, we give a short and basic proof in Appendix B.

3 Submultiplicative Glivenko-Cantelli

In this section μ is a fixed but otherwise arbitrary distribution, with CDF F and empirical CDF F_n .

Theorem 1.3 is a corollary of the following lemma, which gives a quantitative bound on the confidence parameter δ .

⁶ F is an upper envelope for \mathcal{F} if $F(v) \geq f(v)$ for every $v \in V$ and $f \in \mathcal{F}$.

Lemma 3.1. Let $n \in \mathbb{N}$, $\epsilon > 0$ and $\alpha, p, q \in (0, 1)$. Assume that $n \geq \epsilon^{-\frac{1}{1-\alpha}}$ and $p \leq \min\{\epsilon^{\frac{1}{1-\alpha}}, 1/e\}$. Then,

$$\Pr\left[\exists t : |F(t) - F_n(t)| > \epsilon \cdot F(t)^\alpha\right] \leq q + \left\lceil \frac{\ln \ln(\frac{n}{q})}{\ln(\frac{1+\alpha}{2\alpha})} \right\rceil \frac{p^{\frac{1-\alpha}{2}}}{\epsilon} + 2 \exp(-2n(\epsilon p^\alpha)^2).$$

Note that p and q appear only on the right-hand side, and therefore can be “tuned” in order to minimize the upper bound. Our proof of Lemma 3.1 uses Theorem 1.2. To better understand the parameters, we state the following corollary (whose first item is stronger than Theorem 1.3).

Corollary 3.2. There are constants $c_1, c_2 > 0$ so that the following holds.

1. If $\alpha(n) \leq 1 - c_1 \cdot \frac{\ln \ln(n)}{\ln(n)}$, then the probability of the event

$$\forall t : |F(t) - F_n(t)| \leq \epsilon \cdot F(t)^{\alpha(n)}$$

tends to 1 as n tends to ∞ .

2. If $\alpha(n) \geq 1 - c_2 \cdot \frac{1}{\ln(n)}$ and μ is uniform over $[0, 1]$, then the probability of the event

$$\forall t : |F(t) - F_n(t)| \leq \frac{1}{10} \cdot F(t)^{\alpha(n)}$$

is at most $1/2$, for all $n \geq 2$.

We leave as an open question to determine the behavior of these probabilities when

$$\alpha(n) \in \left[1 - c_1 \cdot \frac{\ln \ln(n)}{\ln(n)}, 1 - c_2 \cdot \frac{1}{\ln(n)}\right].$$

Corollary 3.2 is proven in Appendix C.

Proof of Lemma 3.1. Let $\epsilon, \alpha, q, p, n$ be as in the statement of the lemma. We partition the event in question to three parts, depending on the value of t as follows. Partition \mathbb{R} to

$$I_{[0, q/n]} = \left\{t \in \mathbb{R} : 0 \leq F(t) \leq \frac{q}{n}\right\}, \quad I_{(q/n, p)} = \left\{t \in \mathbb{R} : \frac{q}{n} < F(t) \leq p\right\}$$

and

$$I_{[p, 1]} = \{t \in \mathbb{R} : p < F(t) \leq 1\}.$$

There are three corresponding events $E_{[0, q/n]}$, $E_{(q/n, p)}$ and $E_{[p, 1]}$; for example, $E_{[0, q/n]}$ is the event that $\exists t \in I_{[0, q/n]} : B(t) = 1$, where $B(t)$ is the indicator of $|F(t) - F_n(t)| > \epsilon \cdot F(t)^\alpha$.

The following three claims bound from above the probabilities of these three events. The three claims and the union bound complete the proof of the theorem.

Claim 3.3. $\Pr[E_{[0, q/n]}] \leq q$.

Proof. Let $t \in I_{[0, q/n]}$ be so that $B(t) = 1$. For any $t \in I_{[0, q/n]}$ we have that $F(t) \leq q/n \leq 1/n \leq \epsilon^{\frac{1}{1-\alpha}}$, where the last inequality is by our assumption on n and ϵ . This implies that $F(t) \leq \epsilon F(t)^\alpha$. Since $B(t) = 1$ it must be the case that $F_n(t) > F(t) + \epsilon(F(t))^\alpha \geq 0$, and therefore at least one sample x_i satisfies $x_i \leq t \leq q/n$. Now, by the union bound,

$$\Pr[E_{[0, q/n]}] \leq \Pr[\exists i \in [n] : x_i \in I_{[0, q/n]}] \leq n \cdot \frac{q}{n} = q. \quad \square$$

Claim 3.4. $\Pr[E_{(q/n, p)}] \leq \left\lceil \frac{\ln \ln(\frac{n}{q})}{\ln(\frac{1+\alpha}{2\alpha})} \right\rceil \frac{p^{\frac{1-\alpha}{2}}}{\epsilon}$.

Proof. If $q/n \geq p$ then this event is empty and its probability is 0. Therefore, assume that $q/n < p$, and that this event is not empty.

For all $t \in I_{(q/n,p)}$ since $F(t) \leq p \leq \epsilon^{-\frac{1}{1-\alpha}}$ we have $F(t) - \epsilon F(t)^\alpha \leq 0$. So, it suffices to consider the event

$$\exists t \in I_{(q/n,p)} : F_n(t) > F(t) + \epsilon \cdot F(t)^\alpha.$$

Consider the decreasing sequence of numbers p_0, p_1, \dots, p_m defined by

$$p_i = p \left(\frac{1+\alpha}{2\alpha} \right)^i,$$

where m is such that $p_m < q/n \leq p_{m-1}$. Since $p \leq 1/e$, we can bound $m \leq \left\lceil \frac{\ln \ln(\frac{n}{q})}{\ln(\frac{1+\alpha}{2\alpha})} \right\rceil$. Let

$$t_i = \inf\{t \in I_{(q/n,p)} : F(t) \geq p_i\}.$$

Let $F_n^-(t) = \mu_n(\{x : x < t\})$. We claim that

$$\exists t \in I_{(q/n,p)} : B(t) = 1 \implies \exists i < m : F_n^-(t_i) > \epsilon \cdot p_{i+1}^\alpha,$$

Indeed, assume that $t \in I_{(q/n,p)}$ satisfies $F_n(t) > F(t) + \epsilon \cdot F(t)^\alpha$. Since $p_m < t < p_0$, there is some $0 \leq i \leq m-1$ such that $p_{i+1} \leq F(t) < p_i$. Note that $t_{i+1} \leq t < t_i$. Indeed, $t_{i+1} \leq t$ follows since $p_{i+1} \leq F(t)$, and $t < t_i$ follows since $F(t_i) \geq p_i$ (which is implied by right continuity of F). Hence,

$$F_n^-(t_i) \geq F_n(t) > F(t) + \epsilon \cdot F(t)^\alpha \geq p_{i+1} + \epsilon \cdot p_{i+1}^\alpha \geq \epsilon \cdot p_{i+1}^\alpha.$$

It hence remains to upper bound the union of these events. Note that

$$\mathbb{E}[F_n^-(t_i)] = \mu(\{x : x < t_i\}) \leq p_i.$$

Therefore, by Markov's inequality:

$$\Pr[F_n^-(t_i) > \epsilon \cdot p_{i+1}^\alpha] \leq \frac{p_i}{\epsilon p_{i+1}^\alpha} = \frac{1}{\epsilon} p^{\left(\frac{1-\alpha}{2}\right) \left(\frac{1+\alpha}{2\alpha}\right)^i}.$$

By the union bound,

$$\Pr[\exists i < m : F_n(t_i) > p_{i+1} + \epsilon \cdot p_{i+1}^\alpha] \leq \frac{1}{\epsilon} \sum_{i=0}^{m-1} p^{\left(\frac{1-\alpha}{2}\right) \left(\frac{1+\alpha}{2\alpha}\right)^i} \leq \frac{m}{\epsilon} p^{\frac{1-\alpha}{2}} \leq \left\lceil \frac{\ln \ln(\frac{n}{q})}{\ln(\frac{1+\alpha}{2\alpha})} \right\rceil \frac{p^{\frac{1-\alpha}{2}}}{\epsilon}. \square$$

Claim 3.5. $\Pr[E_{[p,1]}] \leq 2 \exp(-2n(\epsilon p^\alpha)^2)$.

Proof. For all $t \in I_{[p,1]}$ we have $F(t)^\alpha \geq p^\alpha$. The claim follows by Theorem 1.2. \square

Lemma 3.1 follows from combining Claims 3.3, 3.4, and 3.5. \square

4 Discussion

Our main result is a submultiplicative variant of the Glivenko-Cantelli Theorem, which allows for tighter convergence bounds for extreme values of the CDF. We show that for the revenue learning setting our submultiplicative bound can be used to derive uniform convergence sample complexity bounds, assuming a finite bound on the k th moment of the valuations, for any (possibly fractional) $k > 1$. For uniform convergence in the limit, we give a complete characterization, where uniform convergence almost surely occurs if and only if the first moment is finite.

It would be interesting to find other applications of our submultiplicative bound in other settings. A potentially interesting direction is to consider unbounded loss functions (e.g., the squared-loss, or log-loss). Many works circumvent the unboundedness in such cases by ensuring (implicitly) that the losses are bounded, e.g., through restricting the inputs and the hypotheses. Our bound offers a different perspective of addressing this issue. In this paper we consider revenue learning, and replace the boundedness assumption by assuming bounds on higher moments. An interesting challenge is to prove uniform convergence bounds for other practically interesting settings. One such setting might be estimating the effect of outliers (which correspond to the extreme values of the loss).

In the context of revenue estimation, this work only considers the most naïve estimator, namely of estimating the revenues by the empirical revenues. One can envision other estimators, for example ones which regularize the extreme tail of the sample. Such estimators may have a potential of better guarantees or better convergence bounds. In the context of uniform convergence of selling mechanism revenues, this work only considers the basic class of posted-price mechanisms. While for one good and one valuation distribution, it is always possible to maximize revenue via a selling mechanism of this class, this is not the case in more complex auction environments. While in many more-complex environments, the revenue-maximizing mechanism/auction is still not understood well enough, for environments where it is understood (Cole and Roughgarden, 2014; Devanur et al., 2016; Gonczarowski and Nisan, 2017) (as well as for simple auction classes that do not necessarily contain a revenue-maximizing auction (Morgenstern and Roughgarden, 2016; Balcan et al., 2016)) it would also be interesting to study relaxations of the restrictive tail or boundedness assumptions currently common in the literature.

References

- M.-F. Balcan, T. Sandholm, and E. Vitercik. Sample complexity of automated mechanism design. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NIPS)*, pages 2083–2091, 2016.
- P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- Z. W. Birnbaum and R. C. McCarty. A distribution-free upper confidence bound for $\Pr\{Y < X\}$, based on independent samples of X and Y . *The Annals of Mathematical Statistics*, 29(2): 558–562, 1958.
- É. Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909.
- F. P. Cantelli. Sulla probabilità come limite della frequenza. *Atti Accad. Naz. Lincei*, 26(1): 39–45, 1917.
- F. P. Cantelli. Sulla determinazione empirica delle leggi di probabilità. *Giornale dell’Istituto Italiano degli Attuari*, 4:421–424, 1933.
- R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252, 2014.
- C. Cortes, Y. Mansour, and M. Mohri. Learning bounds for importance weighting. In *Proceedings of the 24th Conference on Neural Information Processing Systems (NIPS)*, pages 442–450, 2010.

- C. Cortes, S. Greenberg, and M. Mohri. Relative deviation learning bounds and generalization with unbounded loss functions. *CoRR*, abs/1310.5796, 2013.
- N. R. Devanur, Z. Huang, and C.-A. Psomas. The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 426–439, 2016.
- P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.
- A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, 27(3):642–669, 1956.
- V. Glivenko. Sulla determinazione empirica delle leggi di probabilita. *Giornale dell’Istituto Italiano degli Attuari*, 4:92–99, 1933.
- Y. A. Gonczarowski and N. Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, 2017. To appear.
- Z. Huang, Y. Mansour, and T. Roughgarden. Making the most of your samples. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 45–60, 2015.
- V. Koltchinskii and D. Panchenko. *Rademacher Processes and Bounding the Risk of Function Learning*, pages 443–457. Birkhäuser Boston, Boston, MA, 2000.
- P. Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990.
- J. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NIPS)*, pages 136–144, 2015.
- J. Morgenstern and T. Roughgarden. Learning simple auctions. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, pages 1298–1318, 2016.
- R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- T. Roughgarden and O. Schrijvers. Ironing in the dark. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 1–18, 2016.
- A. W. v. d. Vaart and J. A. Wellner. *Weak convergence and empirical processes : with applications to statistics*. Springer series in statistics. Springer, New York, 1996. Réimpr. avec corrections 2000.
- V. Vapnik. *Statistical Learning Theory*. Wiley, 1998. ISBN 978-0-471-03003-4.
- V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16:264–280, 1971.

A Proof of Theorem 1.4

Proof. Let $\epsilon, \delta \leq 1/4$ and $\alpha < 1$. By Lemma 3.1,

$$\Pr \left[\exists t : |F(t) - F_n(t)| > \epsilon \cdot F(t)^\alpha \right] \leq q + \left[\frac{\ln \ln(\frac{n}{q})}{\ln(\frac{1+\alpha}{2\alpha})} \right] \frac{p^{\frac{1-\alpha}{2}}}{\epsilon} + 2 \exp(-2n(\epsilon p^\alpha)^2) \quad (3)$$

for every $q \leq 1$, $n \geq \epsilon^{-\frac{1}{1-\alpha}}$ and $p \leq \epsilon^{\frac{1}{1-\alpha}}$.

Set q, p so that each of the first two summands in Equation 3 is at most δ . Specifically, $q = \delta$, and

1. if $\frac{\ln \ln(\frac{n}{\delta})}{\ln(\frac{1+\alpha}{2\alpha})} \geq 1$ then set $p = \left(\epsilon \delta \cdot \frac{\ln(\frac{1+\alpha}{2\alpha})}{2 \ln \ln(\frac{n}{\delta})} \right)^{\frac{2}{1-\alpha}}$, and
2. if $\frac{\ln \ln(\frac{n}{\delta})}{\ln(\frac{1+\alpha}{2\alpha})} < 1$ then set $p = (\epsilon \delta)^{\frac{2}{1-\alpha}}$.

Note that indeed the requirements $n \geq \epsilon^{-\frac{1}{1-\alpha}}$ and $p \leq \epsilon^{\frac{1}{1-\alpha}}$ are satisfied by p and by the desired n (from the theorem statement).

Plugging these p and q in the third summand in Equation 3 yields:

$$2 \exp \left(-2n\epsilon^2 \left(\epsilon \delta \cdot \frac{\ln(\frac{1+\alpha}{2\alpha})}{2 \ln \ln(\frac{n}{\delta})} \right)^{\frac{4\alpha}{1-\alpha}} \right)$$

when $\frac{\ln \ln(\frac{n}{\delta})}{\ln(\frac{1+\alpha}{2\alpha})} \geq 1$, or

$$2 \exp \left(-2n\epsilon^2 (\epsilon \delta)^{\frac{4\alpha}{1-\alpha}} \right)$$

otherwise. In order for the above to be at most δ , it suffices that

$$2n\epsilon^2 \left(\epsilon \delta \cdot \frac{\ln(\frac{1+\alpha}{2\alpha})}{2 \ln \ln(\frac{n}{\delta})} \right)^{\frac{4\alpha}{1-\alpha}} \geq \ln(2/\delta)$$

when $\frac{\ln \ln(\frac{n}{\delta})}{\ln(\frac{1+\alpha}{2\alpha})} \geq 1$, or

$$2n\epsilon^2 (\epsilon \delta)^{\frac{4\alpha}{1-\alpha}} \geq \ln(2/\delta)$$

otherwise. The second case implies an explicit bound of

$$n \geq \frac{\ln(2/\delta)}{2\epsilon^2} (\epsilon \delta)^{-\frac{4\alpha}{1-\alpha}}. \quad (4)$$

To get an explicit bound on n in the first case, we need to solve a recursion of the following type: find a lower bound on n so that the following inequality holds:

$$n \geq D(\ln \ln(E \cdot n))^F,$$

where $D \geq 0$, $E \geq 4$, $F \geq 0$. (Here $D = \frac{\ln(2/\delta)}{2\epsilon^2} \left(\frac{\epsilon \delta}{2} \cdot \ln(\frac{1+\alpha}{2\alpha}) \right)^{-\frac{4\alpha}{1-\alpha}}$, $E = \frac{1}{\delta}$, $F = \frac{4\alpha}{1-\alpha}$.) Setting

$$n \geq (D+1) \left(10(\ln(D+4) + \ln(F+4) + \ln(E)) \right)^F = (D+1) \left(10 \cdot \ln \left(4 \cdot \frac{D+4}{\delta(1-\alpha)} \right) \right)^{\frac{4\alpha}{1-\alpha}} \quad (5)$$

suffices. Therefore, the probability (i.e., the sum of all three summands of Equation 3) is bounded by 3δ . Replacing δ by $\delta/3$ in Equations 4 and 5 (and in the definition of D) yields the desired bound on $n_0(\epsilon, \delta, \alpha)$. \square

B Proof of Theorem 2.5

Proof. Let $\epsilon > 0$. Having $\mathbb{E}_\mu[F] < \infty$ implies that there is $v_0 \in \mathbb{R}^+$ such that $\mathbb{E}_\mu[1_{\{V \geq v_0\}}F] < \epsilon$. Since we can write

$$\mathbb{E}[f] = \mathbb{E}_\mu[f \cdot 1_{\{V \leq v_0\}}] + \mathbb{E}_\mu[f \cdot 1_{\{V > v_0\}}]$$

it suffices to show that almost surely there exist n_1 such that

$$(\forall n \geq n_1) (\forall f \in \mathcal{F}) : \left| \mathbb{E}_\mu[f \cdot 1_{\{V \geq v_0\}}] - \mathbb{E}_{\mu_n}[f \cdot 1_{\{V \geq v_0\}}] \right| \leq 2\epsilon \quad (6)$$

and that almost surely there exist n_2 such that

$$(\forall n \geq n_2) (\forall f \in \mathcal{F}) : \left| \mathbb{E}_\mu[f \cdot 1_{\{V < v_0\}}] - \mathbb{E}_{\mu_n}[f \cdot 1_{\{V < v_0\}}] \right| \leq 3\epsilon. \quad (7)$$

We begin by showing Equation (6): the law of large numbers implies that almost surely, there exists n_1 such that $\mathbb{E}_{\mu_n}[1_{\{V \geq v_0\}}F] < 2\epsilon$, for every $n \geq n_1$. Since every $f \in \mathcal{F}$ satisfies $0 \leq f \leq F$, it follows that $0 \leq \mathbb{E}_\mu[1_{\{V \geq v_0\}}f] < \epsilon$, and $0 \leq \mathbb{E}_{\mu_n}[1_{\{V \geq v_0\}}f] < 2\epsilon$ for $n \geq n_1$. This implies Equation (6).

It remains to show Equation (7): set $\epsilon' = \frac{\epsilon}{F(v_0)+1}$. The Glivenko-Cantelli Theorem implies that almost surely there exists n_2 such that

$$(\forall n \geq n_2) (\forall v \in \mathbb{R}^+) : |q(v) - q_n(v)| \leq \epsilon'.$$

Let $f \in \mathcal{F}$. By monotonicity of f , it follows that there is a sequence $0 = a_0, a_1, \dots, a_N = v_0$ such that f does not change by more than ϵ within each interval $[a_i, a_{i+1})$, (i.e., $\sup_{x,y \in [a_i, a_{i+1})} |f(x) - f(y)| < \epsilon$). Consider the piecewise constant function

$$f_\epsilon = f(a_0) + \sum_i (f(a_{i+1}) - f(a_i)) 1_{\{V \geq a_i\}}.$$

Note that f_ϵ gets the value $f(a_i)$ on each interval $[a_i, a_{i+1})$. Thus, $|f(v) - f_\epsilon(v)| \leq \epsilon$ for every $v \leq v_0$. Therefore, $|\mathbb{E}_\mu[f 1_{\{V < v_0\}}] - \mathbb{E}_\mu[f_\epsilon 1_{\{V < v_0\}}]| \leq \epsilon$ and $|\mathbb{E}_{\mu_n}[f] - \mathbb{E}_{\mu_n}[f_\epsilon]| \leq \epsilon$. So, it suffices to show that

$$(\forall n \geq n_2) : \left| \mathbb{E}_\mu[f_\epsilon 1_{\{V < v_0\}}] - \mathbb{E}_{\mu_n}[f_\epsilon 1_{\{V < v_0\}}] \right| \leq \epsilon.$$

Indeed, for $n \geq n_2$:

$$\begin{aligned} \left| \mathbb{E}_\mu[f_\epsilon 1_{\{V < v_0\}}] - \mathbb{E}_{\mu_n}[f_\epsilon 1_{\{V < v_0\}}] \right| &\leq \sum_i (f(a_{i+1}) - f(a_i)) \cdot |q(a_i) - q_n(a_i)| \\ &\leq \sum_i (f(a_{i+1}) - f(a_i)) \cdot \epsilon' && \text{(by definition of } n_2) \\ &\leq f(v_0) \cdot \epsilon' \\ &\leq \epsilon. && \text{(by definition of } \epsilon') \end{aligned}$$

□

C Proof of Corollary 3.2

Proof. We begin with the first item. Let $\delta > 0$. It suffices to prove that

$$\Pr \left[\forall t : |F(t) - F_n(t)| \leq \epsilon \cdot F(t)^{\alpha(n)} \right] \leq 3\delta$$

for a large enough n . To this end, we set q, p so that each of the first two summands in Lemma 3.1 is at most δ . Specifically,

$$q = \delta$$

and

$$p = \left(\frac{\epsilon \delta \ln\left(\frac{1+\alpha}{2\alpha}\right)}{\ln \ln\left(\frac{n}{\delta}\right) + \ln\left(\frac{1+\alpha}{2\alpha}\right)} \right)^{\frac{2}{1-\alpha}}.$$

As required by the premise of Lemma 3.1, $p \leq \epsilon^{\frac{1}{1-\alpha}}$. (The other requirement, $n \geq \epsilon^{-\frac{1}{1-\alpha}}$, will be verified at the end of the proof.)

Plugging these values for p, q , the last summand becomes

$$2 \exp(-2n(\epsilon p^\alpha)^2) = 2 \exp\left(-2n\epsilon^2 \left(\frac{\epsilon \delta \ln\left(\frac{1+\alpha}{2\alpha}\right)}{\ln \ln\left(\frac{n}{\delta}\right) + \ln\left(\frac{1+\alpha}{2\alpha}\right)}\right)^{\frac{4\alpha}{1-\alpha}}\right).$$

We need to verify that the above expression becomes less than δ for large n . Equivalently, that

$$\lim_{n \rightarrow \infty} n \left(\frac{\epsilon \delta \ln\left(\frac{1+\alpha}{2\alpha}\right)}{\ln \ln\left(\frac{n}{\delta}\right) + \ln\left(\frac{1+\alpha}{2\alpha}\right)} \right)^{\frac{4\alpha}{1-\alpha}} = \infty.$$

Rewriting $\alpha = 1 - \beta$ gives

$$\lim_{n \rightarrow \infty} n \left(\frac{\epsilon \delta \ln\left(1 + \frac{\beta}{2-2\beta}\right)}{\ln \ln\left(\frac{n}{\delta}\right) + \ln\left(1 + \frac{\beta}{2-2\beta}\right)} \right)^{\frac{4-4\beta}{\beta}} = \infty.$$

Since we are focusing on small value of β we can assume that $\beta \geq 1/2$. Using that $x/2 \leq \ln(1+x) \leq x$ for $x \in [0, 1]$ and $\beta \leq 1/2$, it suffices that we show

$$\lim_{n \rightarrow \infty} n \left(\frac{\epsilon \delta \beta/2}{\ln \ln(n/\delta) + 1} \right)^{\frac{1}{\beta}} = \infty,$$

or, by taking “ln”, that

$$\lim_{n \rightarrow \infty} \left(\ln n - \frac{1}{\beta} \left(\ln(1/\epsilon) + \ln(1/\delta) + \ln(4/\beta) + \ln(\ln \ln(n/\delta) + 1) \right) \right) = \infty.$$

To this end, it suffices that $1/\beta \ln(1/\beta) \leq \ln(n)/2$, which holds for $\beta \geq c \cdot \ln \ln(n)/\ln(n)$, where c is a sufficiently large constant.

It remains to check that the condition stated in Lemma 3.1, that $n \geq \epsilon^{-\frac{1}{1-\alpha}} = \epsilon^{-\frac{1}{\beta}}$, is satisfied. Indeed, for a sufficiently large n

$$\epsilon^{-\frac{1}{\beta}} \leq 1/\epsilon^{c \cdot \ln(n)/\ln \ln(n)} = \exp(c \ln(1/\epsilon) \cdot \ln(n)/\ln \ln(n)) < \exp(\ln(n)) = n.$$

For the second item, let $Y_1 \leq Y_2 \leq \dots \leq Y_n$ denote the sequence obtained by sorting X_1, X_2, \dots, X_n . Note that it suffices to show that the probability that

$$Y_1 < \frac{1}{2n}$$

is at least $1/2$: indeed, this event implies that

$$\begin{aligned} F_n\left(\frac{1}{2n}\right) - F\left(\frac{1}{2n}\right) &\geq \frac{1}{n} - \frac{1}{2n} \\ &= \frac{1}{2n} \\ &\geq \frac{1}{10} \cdot \left(\frac{1}{2n}\right)^{1 - \frac{1}{2\ln n}} && \text{(since } n \geq 2\text{)} \\ &= \frac{1}{10} \cdot F\left(\frac{1}{2n}\right)^{1 - \frac{1}{2\ln n}}, \end{aligned}$$

which implies the conclusion with $c_2 = 1/2$.

Thus, it remains to show that with probability of at least $\frac{1}{2}$, we have $Y_1 < \frac{1}{2n}$:

$$\Pr\left[Y_1 \geq \frac{1}{2n}\right] = \Pr\left[\forall i \leq n : X_i \geq \frac{1}{2n}\right] = \left(1 - \frac{1}{2n}\right)^n \geq \frac{1}{2}. \quad \square$$