
Learning discrete distributions with infinite support

Doron Cohen

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
doronv@post.bgu.ac.il

Aryeh Kontorovich

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
karyeh@cs.bgu.ac.il

Geoffrey Wolfer

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva, Israel
geoffrey@post.bgu.ac.il

Abstract

We present a novel approach to estimating discrete distributions with (potentially) infinite support in the total variation metric. In a departure from the established paradigm, we make no structural assumptions whatsoever on the sampling distribution. In such a setting, distribution-free risk bounds are impossible, and the best one could hope for is a fully empirical data-dependent bound. We derive precisely such bounds, and demonstrate that these are, in a well-defined sense, the best possible. Our main discovery is that the half-norm of the empirical distribution provides tight upper and lower estimates on the empirical risk. Furthermore, this quantity decays at a nearly optimal rate as a function of the true distribution. The optimality follows from a minimax result, of possible independent interest. Additional structural results are provided, including an exact Rademacher complexity calculation and apparently a first connection between the total variation risk and the missing mass.

1 Introduction

Estimating a discrete distribution in the total variation (TV) metric is a central problem in computer science and statistics (see, e.g., Han et al. [2015], Kamath et al. [2015], Orlitsky and Suresh [2015] and the references therein). The TV metric, which we use throughout the paper, is a natural and abundantly motivated choice [Devroye and Lugosi, 2001]. For support size d , a sample of size $\mathcal{O}(d/\varepsilon^2)$ suffices for the maximum-likelihood estimator (MLE) to be ε -close (with constant probability) to the unknown target distribution. A matching lower bound is known [Anthony and Bartlett, 1999], and has been computed down to the exact constants [Kamath et al., 2015].

Classic VC theory — and, in particular, the aforementioned results — imply that for infinite support, no distribution-free sample complexity bound is possible. If μ is the target distribution and $\hat{\mu}_m$ is its empirical (i.e., MLE) estimate based on m iid samples, then Berend and Kontorovich [2013] showed that

$$\frac{1}{4}\Lambda_m(\mu) - \frac{1}{4\sqrt{m}} \leq \mathbb{E}[\|\mu - \hat{\mu}_m\|_{\text{TV}}] \leq \Lambda_m(\mu), \quad m \geq 2, \quad (1)$$

where

$$\Lambda_m(\mu) = \sum_{j \in \mathbb{N}: \mu(j) < 1/m} \mu(j) + \frac{1}{2\sqrt{m}} \sum_{j \in \mathbb{N}: \mu(j) \geq 1/m} \sqrt{\mu(j)}. \quad (2)$$

The quantity $\Lambda_m(\boldsymbol{\mu})$ has the advantage of always being finite and of decaying to 0 as $m \rightarrow \infty$. The bound in (1) suggests that $\Lambda_m(\boldsymbol{\mu})$, or a closely related measure, controls the sample complexity for learning discrete distributions in TV. Further supporting the foregoing intuition is the observation that for finite support size d and $m \gg 1$, we have $\Lambda_m \lesssim \sqrt{d/m}$, recovering the known minimax rate. Additionally, a closely related measure turns out to control a minimax risk rate in a sense made precise in Theorem 2.5.

One shortcoming of (1) is that the lower bound only holds for the MLE, leaving the possibility that a different estimator could achieve significantly improved bounds. Another shortcoming of (1) and related estimates is that they are not *empirical*, in that they depend on the unknown quantity we are trying to estimate. A fully empirical bound, on the other hand, would give a high-probability estimate on $\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_m\|_{\text{TV}}$ solely in terms of observable quantities such as $\hat{\boldsymbol{\mu}}_m$. Of course, such a bound should also be non-trivial, in the sense of improving with growing sample size and approaching 0 as $m \rightarrow \infty$. A further desideratum might be something akin to *instance optimality*: We would like the rate at which the empirical bound decays to be “the best” possible for the given $\boldsymbol{\mu}$, in an appropriate sense. Our analogue of instance optimality is inspired by, but distinct from, that of Valiant and Valiant [2016], as discussed in detail in Related work below.

Our contributions. We address the shortcomings of existing estimators detailed above by providing a fully empirical bound on $\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_m\|_{\text{TV}}$. Our main discovery is that the quantity $\Phi_m(\hat{\boldsymbol{\mu}}_m) := \frac{1}{\sqrt{m}} \sum_{j \in \mathbb{N}} \sqrt{\hat{\boldsymbol{\mu}}_m(j)}$ satisfies all of the desiderata posed above for an empirical bound. As we show in Theorems 2.1 and 2.2, $\Phi_m(\hat{\boldsymbol{\mu}}_m)$ provides tight, high-probability upper and lower bounds on $\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_m\|_{\text{TV}}$. Further, Theorem 2.3 shows that $\mathbb{E}[\Phi_m(\hat{\boldsymbol{\mu}}_m)]$ behaves as $\Lambda_m(\boldsymbol{\mu})$ defined in (2). Finally, a result in the spirit of instance optimality, Theorem 2.4, shows that no other estimator-bound pair can improve upon $(\hat{\boldsymbol{\mu}}_m, \Phi_m)$, other than by small constants. The latter follows from a minimax bound of independent interest, Theorem 2.5. Additional structural results are provided, including an exact Rademacher complexity calculation and a connection (apparently the first) between the total variation risk and the missing mass.

Definitions, notation and setting. As we are dealing with discrete distributions, there is no loss of generality in taking our sample space to be the natural numbers $\mathbb{N} = \{1, 2, 3, \dots\}$. For $k \in \mathbb{N}$, we write $[k] := \{i \in \mathbb{N} : i \leq k\}$. The set of all distributions on \mathbb{N} will be denoted by $\Delta_{\mathbb{N}}$, which we enlarge to include the “deficient” distributions: $\Delta_{\mathbb{N}} \subset \Delta_{\mathbb{N}}^{\circ} := \{\boldsymbol{\mu} \in [0, 1]^{\mathbb{N}} : \sum_{i \in \mathbb{N}} \boldsymbol{\mu}(i) \leq 1\}$. For $d \in \mathbb{N}$, we write $\Delta_d \subset \Delta_{\mathbb{N}}$ to denote those $\boldsymbol{\mu}$ whose support is contained in $[d]$.

For $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}^{\circ}$ and $I \subseteq \mathbb{N}$, we write $\boldsymbol{\mu}(I) = \sum_{i \in I} \boldsymbol{\mu}(i)$. We define the *decreasing permutation* of $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}^{\circ}$, denoted by $\boldsymbol{\mu}^{\downarrow}$, to be the sequence $(\boldsymbol{\mu}(i))_{i \in \mathbb{N}}$ sorted in non-increasing order, achieved by a¹ permutation $\Pi_{\boldsymbol{\mu}}^{\downarrow} : \mathbb{N} \rightarrow \mathbb{N}$; thus, $\boldsymbol{\mu}^{\downarrow}(i) = \boldsymbol{\mu}(\Pi_{\boldsymbol{\mu}}^{\downarrow}(i))$. For $0 < \eta < 1$, define $T_{\boldsymbol{\mu}}(\eta) \in \mathbb{N}$ as the least t for which $\sum_{i > t} \boldsymbol{\mu}^{\downarrow}(i) < \eta$. This induces a truncation of $\boldsymbol{\mu}$, denoted by $\boldsymbol{\mu}[\eta] \in \Delta_{\mathbb{N}}^{\circ}$ and defined by $\boldsymbol{\mu}[\eta](i) = \mathbf{1}[\Pi_{\boldsymbol{\mu}}^{\downarrow}(i) \leq T_{\boldsymbol{\mu}}(\eta)]\boldsymbol{\mu}(i)$.

For $\boldsymbol{\mu}, \boldsymbol{\nu} \in \Delta_{\mathbb{N}}^{\circ}$, we define the *total variation distance* in terms of the ℓ_1 norm:

$$\|\boldsymbol{\mu} - \boldsymbol{\nu}\|_{\text{TV}} := \frac{1}{2} \|\boldsymbol{\mu} - \boldsymbol{\nu}\|_1 = \frac{1}{2} \sum_{i \in \mathbb{N}} |\boldsymbol{\mu}(i) - \boldsymbol{\nu}(i)|. \quad (3)$$

For $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}^{\circ}$, we also define the *half-norm*² as

$$\|\boldsymbol{\mu}\|_{1/2} := \left(\sum_{i \in \Omega} \sqrt{\boldsymbol{\mu}(i)} \right)^2; \quad (4)$$

note that while $\|\boldsymbol{\mu}\|_{1/2}$ may be infinite, we have $\|\boldsymbol{\mu}\|_{1/2} \leq \|\boldsymbol{\mu}\|_0$, where the latter denotes the support size.

For $m \in \mathbb{N}$ and $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, we write $\mathbf{X} = (X_1, \dots, X_m) \sim \boldsymbol{\mu}^m$ to mean that the components of the vector \mathbf{X} are drawn iid from from $\boldsymbol{\mu}$. We reserve $\hat{\boldsymbol{\mu}}_m \in \Delta_{\mathbb{N}}$ for the empirical measure induced by the sample \mathbf{X} , i.e. $\hat{\boldsymbol{\mu}}_m(i) := \frac{1}{m} \sum_{t \in [m]} \mathbf{1}[X_t = i]$; the term MLE will be used interchangeably.

¹While $\boldsymbol{\mu}^{\downarrow}$ is uniquely defined, $\Pi_{\boldsymbol{\mu}}^{\downarrow}$ is not. Uniqueness could be ensured by taking the lexicographically first permutation, but will not be needed for our results.

²The half-norm is not a proper vector-space norm, as it lacks sub-additivity.

For the class of boolean functions over the integers $\{f: \mathbb{N} \rightarrow \{0, 1\}\}$, which we denote by $\{0, 1\}^{\mathbb{N}}$, recall the definition of the *empirical Rademacher complexity* [Mohri et al., 2012, Definition 3.1] conditional on the sample \mathbf{X} :

$$\hat{\mathfrak{R}}_m(\mathbf{X}) := \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \{0, 1\}^{\mathbb{N}}} \frac{1}{m} \sum_{t=1}^m \sigma_t f(X_t) \right], \quad (5)$$

where $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_m) \sim \text{Uniform}(\{-1, 1\}^m)$. The expectation of the above random quantity is the *Rademacher complexity* [Mohri et al., 2012, Definition 3.2]:

$$\mathfrak{R}_m := \mathbb{E}_{\mathbf{X} \sim \mu^m} [\hat{\mathfrak{R}}(\mathbf{X})]. \quad (6)$$

Related work. Given the classical nature of the problem, a comprehensive literature survey is beyond our scope; the standard texts Devroye and Györfi [1985], Devroye and Lugosi [2001] provide much of the requisite background. Chapter 6.5 of the latter makes a compelling case for the TV metric used in this paper, but see Waggoner [2015] and the works cited therein for results on other ℓ_p norms. Though surveying all of the relevant literature is a formidable task, a relatively streamlined narrative may be distilled. Conceptually, the simplest case is that of $\|\mu\|_0 < \infty$ (i.e., finite support). Since learning a distribution over $[d]$ in TV is equivalent to agnostically learning the function class $\{0, 1\}^d$, standard VC theory [Anthony and Bartlett, 1999, Kontorovich and Pinelis, 2019] entails that the MLE achieves the minimax risk rate of $\sqrt{d/m}$ over all $\mu \in \Delta_{\mathbb{N}}$ with $\|\mu\|_0 \leq d$. An immediate consequence is that in order to obtain quantitative risk rates for the case of infinite support, one must assume some sort of structure [Diakonikolas, 2016]. One can, for example, obtain minimax rates for μ with bounded entropy [Han et al., 2015], or, say, bounded half-norm (as we do here). Alternatively, one can restrict one’s attention to a finite class $\mathcal{Q} \subset \Delta_{\mathbb{N}}$; here too, optimal results are known [Bousquet et al., 2019]. Berend and Kontorovich [2013] was one of the few works that made no assumptions on $\mu \in \Delta_{\mathbb{N}}$, but only gave non-empirical bounds.

Our work departs from the paradigm of a-priori constraints on the unknown sampling distribution. Instead, our estimates hold for all $\mu \in \Delta_{\mathbb{N}}$. Of course, this must come at a price: no a-priori sample complexity bounds are possible in this setting. Absent any prior knowledge regarding μ , one can only hope for sample-dependent *empirical* bounds, and we indeed obtain these. Further, our empirical bounds are essentially the best possible, as formalized in Theorem 2.4. The latter result may be thought of as a learning-theoretic analogue of being *instance-optimal*, as introduced by Valiant and Valiant [2017] in the testing framework. Instance optimality is a very natural notion in the context of testing whether an unknown sampling distribution μ is identical to or ε -far from a given reference one, μ_0 . For example, Valiant and Valiant discovered that a truncated $2/3$ -norm of μ_0 — i.e., a quantity closely related to $\|\mu_0\|_{2/3}$ — controls the complexity of the testing problem in TV distance. Instance optimality is more difficult to formalize for distribution learning, since for any given $\mu \in \Delta_{\mathbb{N}}$, there is a trivial “learner” with μ hard-coded inside. Valiant and Valiant [2016] defined this notion in terms of competing against an oracle who knows the distribution up to a permutation of the atoms, and did not provide empirical confidence intervals. We do derive fully empirical bounds, and further show that they are impossible to improve upon — by *any estimator* — other than by constants. Our results suggest that the half-norm $\|\mu\|_{1/2}$ plays a role in learning analogous to that of $\|\mu\|_{2/3}$ in testing. As an intriguing aside, we note that the half-norm corresponds to the Tsallis q -entropy with $q = 1/2$, which was shown to be an optimal regularizer in some stochastic and adversarial bandit settings [Zimmert and Seldin, 2019]. We leave the question of investigating a deeper connection between the two results for future work.

2 Main results

In this section, we formally state our main results. Recall from the Definitions that the sample $\mathbf{X} = (X_1, \dots, X_m) \sim \mu^m$ induces the empirical measure (MLE) $\hat{\mu}_m$, and that a key quantity in our bounds is

$$\Phi_m(\hat{\mu}_m) = \frac{1}{\sqrt{m}} \|\hat{\mu}_m\|_{1/2}^{1/2} = \frac{1}{\sqrt{m}} \sum_{j \in \mathbb{N}} \sqrt{\hat{\mu}_m(j)}. \quad (7)$$

Our first result is a fully empirical, high-probability upper bound on $\|\hat{\mu}_m - \mu\|_{\text{TV}}$ in terms of $\Phi_m(\hat{\mu}_m)$:

Theorem 2.1. For all $m \in \mathbb{N}$, $\delta \in (0, 1)$, and $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, we have that

$$\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \leq \Phi_m(\hat{\boldsymbol{\mu}}_m) + 3\sqrt{\frac{\log \frac{2}{\delta}}{2m}}$$

holds with probability at least $1 - \delta$. We also have

$$\mathbb{E}[\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}}] \leq \mathbb{E}[\Phi_m(\hat{\boldsymbol{\mu}}_m)].$$

Since $\|\hat{\boldsymbol{\mu}}_m\|_{1/2} \leq \|\hat{\boldsymbol{\mu}}_m\|_0 \leq \|\boldsymbol{\mu}\|_0$, this recovers the minimax rate of $\sqrt{d/m}$ for $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\mu}\|_0 \leq d$. We also provide a matching lower bound:

Theorem 2.2. For all $m \in \mathbb{N}$, $\delta \in (0, 1)$, and $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, we have that

$$\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \geq \frac{1}{4\sqrt{2}}\Phi_m(\hat{\boldsymbol{\mu}}_m) - 3\sqrt{\frac{\log \frac{2}{\delta}}{m}}$$

holds with probability at least $1 - \delta$.

Our empirical measure $\Phi_m(\hat{\boldsymbol{\mu}}_m)$ is never much worse than the non-empirical $\Lambda_m(\boldsymbol{\mu})$, defined in (2):

Theorem 2.3. For all $m \in \mathbb{N}$ and $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ we have

$$\mathbb{E}[\Phi_m(\hat{\boldsymbol{\mu}}_m)] \leq 2\Lambda_m(\boldsymbol{\mu})$$

and, with probability at least $1 - \delta$,

$$\Phi_m(\hat{\boldsymbol{\mu}}_m) \leq 2\Lambda_m(\boldsymbol{\mu}) + \sqrt{\log(1/\delta)/m}.$$

Furthermore, no other estimator-bound pair $(\tilde{\boldsymbol{\mu}}_m, \Psi_m)$ can improve upon $(\hat{\boldsymbol{\mu}}_m, \Phi_m)$, other than by a constant. This is the ‘‘instance optimality’’ result alluded to above:

Theorem 2.4. There exist universal constants $a, b > 0$ such that the following holds. For any estimator-bound pair $(\tilde{\boldsymbol{\mu}}_m, \Psi_m)$ and any continuous function $\theta: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$\mathbb{E}[\|\tilde{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}}] \leq \mathbb{E}[\Psi_m(\tilde{\boldsymbol{\mu}}_m)] \leq \theta\left(\mathbb{E}[\Phi_m(\hat{\boldsymbol{\mu}}_m)]\right)$$

holds for all $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, θ necessarily verifies

$$\inf_{0 < x < b} \frac{\theta(x)}{x} \geq \frac{1}{a}.$$

The next result, framed in the high-probability setting, draws a direct parallel between our characterization of the learning sample complexity via the half-norm and Valiant and Valiant [2017]’s characterization of the testing sample complexity via the 2/3-norm. The truncation is needed to ensure finiteness, since the $\|\boldsymbol{\mu}\|_{1/2} = \infty$ for heavy-tailed distributions (e.g. $\boldsymbol{\mu}(i) \propto 1/i^2$).

Theorem 2.5. There is a universal constant $C > 0$ such that for all $\Lambda \geq 2$ and $0 < \varepsilon, \delta < 1$, the MLE $\hat{\boldsymbol{\mu}}_m$ verifies the following optimality property: For all $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} \leq \Lambda$, we have

$$m \geq C\varepsilon^{-2} \max\{\Lambda, \log(1/\delta)\} \implies \mathbb{P}(\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} < \varepsilon) \geq 1 - \delta.$$

On the other hand, for any estimator $\tilde{\boldsymbol{\mu}}_m: \mathbb{N}^m \rightarrow \Delta_{\mathbb{N}}$ there is a $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ with $\max\{\|\boldsymbol{\mu}[\varepsilon/18]\|_{1/2}, \|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2}\} \leq \Lambda$ such that:

$$m < C\varepsilon^{-2} \min\{\Lambda, \log(1/\delta)\} \implies \mathbb{P}(\|\tilde{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \geq \varepsilon) \geq \min\{3/4, 1 - \delta\}.$$

The above is a simplified statement chosen for brevity; a considerably refined version is stated and proved in Theorem 3.1.

3 Proofs

3.1 Proof of Theorem 2.1

The proof consists of two parts. The first is contained in Lemma 3.1, which provides a high-probability empirical upper bound, and an expectation bound, similar to Theorem 2.1, but in terms of $\hat{\mathfrak{R}}_m(\mathbf{X})$ instead of $\Phi_m(\hat{\boldsymbol{\mu}}_m)$. The second part, contained in Lemma 3.2, provides an estimate of $\hat{\mathfrak{R}}_m(\mathbf{X})$ in terms of $\Phi_m(\hat{\boldsymbol{\mu}}_m)$.

Lemma 3.1. *For all $m \in \mathbb{N}$, $\delta \in (0, 1)$, and $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, we have that*

$$\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \leq 2\hat{\mathfrak{R}}_m(\mathbf{X}) + 3\sqrt{\frac{\log \frac{2}{\delta}}{2m}}$$

holds with probability at least $1 - \delta$. We also have,

$$\mathbb{E} [\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}}] \leq 2\mathfrak{R}_m. \quad (8)$$

Proof. The high-probability bound from the observation,

$$\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} := \sup_{A \subseteq \mathbb{N}} (\boldsymbol{\mu}(A) - \hat{\boldsymbol{\mu}}_m(A)) = \sup_{f \in \mathcal{F}} \left(\mathbb{E}_{X \sim \boldsymbol{\mu}} [f(X)] - \frac{1}{m} \sum_{i=1}^m f(X_i) \right) \quad (9)$$

where $\mathcal{F} := \{\mathbb{I}_A | A \subseteq \mathbb{N}\} = \{0, 1\}^{\mathbb{N}}$, combined with [Mohri et al., 2012, Theorem 3.3], which states: Let \mathcal{G} be a family of functions from \mathcal{Z} to $[0, 1]$ and let ν be a distribution supported on a subset of \mathcal{Z} . Then, for any $\delta > 0$, with probability at least $1 - \delta$ over $\mathbf{Z} = (Z_1, \dots, Z_m) \sim \nu^m$, the following holds:

$$\sup_{g \in \mathcal{G}} \left(\mathbb{E}_{Z \sim \nu} [g(Z)] - \frac{1}{m} \sum_{i=1}^m g(Z_i) \right) \leq 2\hat{\mathfrak{R}}_m(\mathbf{Z}) + 3\sqrt{\frac{\log \frac{2}{\delta}}{2m}}.$$

Plugging in \mathcal{F} for \mathcal{G} and $\boldsymbol{\mu}$ for ν in the above theorem completes the proof of the high-probability bound. The expectation bound (eq. (8)) follows from the observation at eq. (9) and a symmetrization argument [Mohri et al., 2012, eq. (3.8) to (3.13)]. \square

In order to complete the proof, we apply

Lemma 3.2 (Empirical Rademacher estimates). *Let $\mathbf{X} = (X_1, \dots, X_m)$ and let $\hat{\boldsymbol{\mu}}_m$ be the empirical measure constructed from the sample \mathbf{X} . Then,*

$$\frac{1}{2\sqrt{2}} \Phi_m(\hat{\boldsymbol{\mu}}_m) \leq \hat{\mathfrak{R}}_m(\mathbf{X}) \leq \frac{1}{2} \Phi_m(\hat{\boldsymbol{\mu}}_m).$$

Proof. The proof is based on an argument that was also developed in [Scott and Nowak, 2006, Section 7.1, Appendix E.] in the context of histograms and dyadic decision trees, and that was credited to Gilles Blanchard.

Let $\hat{S} = \{X_i | i \in [m]\}$ be the empirical support according to the sample $\mathbf{X} = (X_1, X_2, \dots, X_m)$. Then,

$$\begin{aligned} m\hat{\mathfrak{R}}_m(\mathbf{X}) &= \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{f \in \{0,1\}^{\mathbb{N}}} \sum_{i=1}^m \sigma_i f(X_i) \right] = \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{A \subseteq \hat{S}} \sum_{i=1}^m \sigma_i \mathbb{I}_A(X_i) \right] \\ &= \sum_{x \in \hat{S}} \mathbb{E}_{\boldsymbol{\sigma}} \left[\sup_{A \subseteq \{x\}} \sum_{i: X_i=x} \sigma_i \mathbb{I}_A(X_i) \right] = \sum_{x \in \hat{S}} \mathbb{E}_{\boldsymbol{\sigma}} \left[\left(\sum_{i: X_i=x} \sigma_i \right)_+ \right] = \sum_{x \in \hat{S}} \frac{1}{2} \mathbb{E}_{\boldsymbol{\sigma}} \left[\left| \sum_{i=1}^m \sigma_i \right| \right], \end{aligned}$$

where the last equality follows from counting $\{i : X_i = x\}$ and the symmetry of the random variable $\sum_{i=1}^m \sigma_i$ for all $n \in \mathbb{N}$. Now, by Khintchine's inequality, for $0 < p < \infty$ and $x_1, x_2, \dots, x_m \in \mathbb{C}$ we have

$$A_p \left(\sum_{i=1}^m |x_i|^2 \right)^{1/2} \leq \left(\mathbb{E}_{\boldsymbol{\sigma}} \left[\left| \sum_{i=1}^m x_i \sigma_i \right|^p \right] \right)^{1/p} \leq B_p \left(\sum_{i=1}^m |x_i|^2 \right)^{1/2},$$

where $A_p, B_p > 0$ are constants depending on p . Sharp values for A_p, B_p were found by Haagerup [1981]. In particular, for $p = 1$ he found that $A_1 = \frac{1}{\sqrt{2}}$ and $B_1 = 1$. By using Khintchine's inequality for each $\mathbb{E}_\sigma \left[\left| \sum_{i=1}^m \widehat{\mu}_m(x) \sigma_i \right| \right]$ with these constants, we get

$$\frac{1}{\sqrt{2}} \sqrt{m \widehat{\mu}_m(x)} \leq \mathbb{E}_\sigma \left[\left| \sum_{i=1}^m \widehat{\mu}_m(x) \sigma_i \right| \right] \leq \sqrt{m \widehat{\mu}_m(x)},$$

and hence

$$\frac{1}{2\sqrt{2}} \sum_{x \in \hat{S}} \sqrt{m \widehat{\mu}_m(x)} \leq m \hat{\mathfrak{R}}_m(\mathbf{X}) \leq \frac{1}{2} \sum_{x \in \hat{S}} \sqrt{m \widehat{\mu}_m(x)}.$$

Dividing by m completes the proof. \square

Remark: We also give an exact expression for $\hat{\mathfrak{R}}_m(\mathbf{X})$ in Lemma A.1, and show in Corollary A.1 with a more delicate analysis that

$$\frac{\|\widehat{\mu}_m\|_{1/2}^{1/2}}{\sqrt{2\pi m}} - \frac{3}{2} \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \|\widehat{\mu}_m^+\|_{-1/2}^{-1/2} \leq \hat{\mathfrak{R}}_m(\mathbf{X}) \leq \frac{\|\widehat{\mu}_m\|_{1/2}^{1/2}}{\sqrt{2\pi m}} + \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \|\widehat{\mu}_m^+\|_{-1/2}^{-1/2}.$$

3.2 Proof of Theorem 2.2

The proof follows from applying the lower bound of Lemma 3.2 to the following lemma:

Lemma 3.3 (lower bound by empirical Rademacher). *For all $m \in \mathbb{N}$, $\delta \in (0, 1)$, and $\mu \in \Delta_{\mathbb{N}}$, we have that*

$$\|\widehat{\mu}_m - \mu\|_{\text{TV}} \geq \frac{1}{2} \hat{\mathfrak{R}}_m(\mathbf{X}) - 3 \sqrt{\frac{\log \frac{2}{\delta}}{m}}$$

holds with probability at least $1 - \delta$.

Proof. The proof is closely based on [Wainwright, 2019, Proposition 4.12], which states: Let $\mathbf{Y} = (Y_1, \dots, Y_m) \sim \nu^m$ for some distribution ν on \mathcal{Z} , let $\mathcal{G} \subseteq [-b, b]^{\mathcal{Z}}$ be a function class, and let $\sigma = (\sigma_1, \dots, \sigma_m) \sim \text{Uniform}(\{-1, 1\}^m)$. Then

$$\sup_{g \in \mathcal{G}} \left| \mathbb{E}_{Y \sim \nu} [g(Y)] - \frac{1}{m} \sum_{i=1}^m g(Y_i) \right| \geq \frac{1}{2} \mathbb{E}_{\sigma, \mathbf{Y}} \left[\sup_{g \in \mathcal{G}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i g(Y_i) \right| \right] - \frac{\sup_{g \in \mathcal{G}} |\mathbb{E}_{Y \sim \nu} [g(Y)]|}{2\sqrt{m}} - \delta \quad (10)$$

holds with probability at least $1 - e^{-\frac{n\delta^2}{2b^2}}$. Plugging in \mathbf{X} for \mathbf{Y} , μ for ν , \mathbb{N} for \mathcal{Z} , 1 for b , and $\mathcal{F} := \{\mathbb{I}_A | A \subseteq \mathbb{N}\} = \{0, 1\}^{\mathbb{N}}$ for \mathcal{G} in (10) together with observing that

$$\begin{aligned} \|\widehat{\mu}_m - \mu\|_{\text{TV}} &:= \sup_{A \subseteq \mathbb{N}} (\mu(A) - \widehat{\mu}_m(A)) = \sup_{f \in \mathcal{F}} \left| \mathbb{E}_{X \sim \mu} [f(X)] - \frac{1}{m} \sum_{i=1}^m f(X_i) \right|, \\ \mathbb{E}_{\sigma, \mathbf{X}} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i f(X_i) \right| \right] &\geq \mathfrak{R}_m, \quad \text{and} \quad \sup_{f \in \mathcal{F}} \left| \mathbb{E}_{X \sim \mu} [f(X)] \right| = 1, \end{aligned}$$

followed by some algebraic manipulation we get

$$\|\widehat{\mu}_m - \mu\|_{\text{TV}} \geq \frac{1}{2} \mathfrak{R}_m - \frac{1}{2\sqrt{m}} - \sqrt{\frac{2 \log \frac{2}{\delta}}{m}} \quad (11)$$

with probability at least $1 - \delta/2$. Applying McDiarmid's inequality to the $1/m$ -bounded-differences function $\hat{\mathfrak{R}}_m(\mathbf{X})$ (similar to [Mohri et al., 2012, Eq. (3.14)]) we get:

$$\frac{1}{2} \mathfrak{R}_m \geq \frac{1}{2} \hat{\mathfrak{R}}_m(\mathbf{X}) - \frac{1}{2} \sqrt{\frac{\log \frac{2}{\delta}}{2m}} \quad (12)$$

with probability at least $1 - \delta/2$. To conclude the proof, combine (11) and (12) with the union bound to get:

$$\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \geq \frac{1}{2} \hat{\mathfrak{R}}_m(\mathbf{X}) - \frac{1}{2\sqrt{m}} - \frac{1}{2} \sqrt{\frac{\log \frac{2}{\delta}}{2m}} - \sqrt{\frac{2 \log \frac{2}{\delta}}{m}}$$

with probability at least $1 - \delta$, and use the fact $-\frac{1}{2\sqrt{m}} - \frac{1}{2} \sqrt{\frac{\log \frac{2}{\delta}}{2m}} - \sqrt{\frac{2 \log \frac{2}{\delta}}{m}} \geq -3\sqrt{\frac{\log \frac{2}{\delta}}{m}}$ for all $m \in \mathbb{N}, \delta \in (0, 1)$. \square

Remark 3.1. We note that by using a more careful analysis, the constants of Theorem 2.2 can be improved to yield, under the same assumptions, $\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} \geq \frac{1}{2} \hat{\mathfrak{R}}_m(\mathbf{X}) - \frac{1}{4\sqrt{m}} - \frac{3}{2} \sqrt{\frac{\log \frac{2}{\delta}}{2m}}$ with probability at least $1 - \delta$.

3.3 Proof of Theorem 2.3

Invoking Fubini's theorem, we write

$$\frac{1}{\sqrt{m}} \mathbb{E} \left[\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \right] = \frac{1}{m} \sum_{i=1}^{\infty} \mathbb{E}_{X \sim \text{Bin}(m, \boldsymbol{\mu}(i))} \left[\sqrt{X} \right].$$

Since $X \in \{0, 1, 2, \dots\}$, we have $\sqrt{X} \leq X$ and hence $\mathbb{E} \left[\sqrt{X} \right] \leq \mathbb{E} [X]$. On the other hand, Jensen's inequality implies $\mathbb{E} \left[\sqrt{X} \right] \leq \sqrt{\mathbb{E} [X]}$, whence

$$\frac{1}{\sqrt{m}} \mathbb{E} \left[\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \right] \leq \frac{1}{m} \sum_{i=1}^{\infty} \min\{\sqrt{m\boldsymbol{\mu}(i)}, m\boldsymbol{\mu}(i)\} \quad (13)$$

$$= \sum_{i: \boldsymbol{\mu}(i) \leq 1/m} \boldsymbol{\mu}(i) + \frac{1}{\sqrt{m}} \sum_{i: \boldsymbol{\mu}(i) > 1/m} \sqrt{\boldsymbol{\mu}(i)} \leq 2\Lambda_m(\boldsymbol{\mu}). \quad (14)$$

The high-probability bound follows from applying McDiarmid's inequality to the $2/m$ -bounded-differences function: for all $\delta \in (0, 1)$, we have

$$\Phi_m(\hat{\boldsymbol{\mu}}_m) \leq \mathbb{E} [\Phi_m(\hat{\boldsymbol{\mu}}_m)] + \sqrt{\log(1/\delta)/m}.$$

\square

3.4 Statement and proof of the refined version of Theorem 2.5

Theorem 3.1. *There is a universal constant $C > 0$ such that for all $\Lambda \geq 2$ and $0 < \varepsilon, \delta < 1$, the MLE verifies the following optimality property: For all $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} \leq \Lambda$, if $(X_1, \dots, X_m) \sim \boldsymbol{\mu}^m$ and $m \geq \frac{C}{\varepsilon^2} \max\{\Lambda, \ln \delta^{-1}\}$, then $\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} < \varepsilon$ holds with probability at least $1 - \delta$.*

On the other hand, for all $\Lambda \geq 2$ and $0 < \varepsilon < 1/16, 0 < \delta < 1$, for any estimator $\bar{\boldsymbol{\mu}}: \mathbb{N}^m \rightarrow \Delta_{\mathbb{N}}$ there is a $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\mu}[\varepsilon/18]\|_{1/2} \leq \Lambda$ such that $\bar{\boldsymbol{\mu}}$ must require at least $m \geq \frac{C}{\varepsilon^2} \Lambda$ samples in order for $\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} < \varepsilon$ to hold with probability at least $3/4$, and for any estimator $\bar{\boldsymbol{\nu}}: \mathbb{N}^m \rightarrow \Delta_{\mathbb{N}}$ there is a $\boldsymbol{\nu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\nu}[2\varepsilon\delta/9]\|_{1/2} \leq \Lambda$, such that $\bar{\boldsymbol{\nu}}$ must require at least $m \geq \frac{C}{\varepsilon^2} \ln \frac{1}{\delta}$ samples in order for $\|\bar{\boldsymbol{\nu}} - \boldsymbol{\nu}\|_{\text{TV}} < \varepsilon$ to hold with probability at least $1 - \delta$.

Minimax risk. For any $\Lambda \in [2, \infty), 0 < \varepsilon, \delta < 1$, we define the minimax risk

$$\mathcal{R}_m(\Lambda, \varepsilon, \delta) := \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu}: \|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} < \Lambda} \mathbb{P}_{\mathbf{X} \sim \boldsymbol{\mu}^m} (\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon),$$

where the infimum is taken over all functions $\bar{\boldsymbol{\mu}}: \mathbb{N}^m \rightarrow \Delta_{\mathbb{N}}$, and the supremum is taken over the subset of distributions such that $\|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} < \Lambda$.

Upper bound. Let $\Lambda \in [2, \infty)$, $0 < \varepsilon, \delta < 1$, $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$, such that $\|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} \leq \Lambda$, $m \in \mathbb{N}$, $(X_1, \dots, X_m) \sim \boldsymbol{\mu}$ and let $\hat{\boldsymbol{\mu}}_m$ be the MLE. For $\eta > 0$, consider the two truncated distributions $\boldsymbol{\mu}[\eta]$ and $\hat{\boldsymbol{\mu}}'_m$, where we define the latter as

$$\hat{\boldsymbol{\mu}}'_m(i) := \hat{\boldsymbol{\mu}}_m(i) \mathbf{1}[\boldsymbol{\mu}[\eta](i) > 0], \quad i \in \mathbb{N}.$$

By the triangle inequality, $\mathbb{P}(\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon) \leq \mathbb{P}(\mathcal{E}_1 + \mathcal{E}_2 + \mathcal{E}_3 > \varepsilon)$, where

$$\mathcal{E}_1 := \|\hat{\boldsymbol{\mu}}_m - \hat{\boldsymbol{\mu}}'_m\|_{\text{TV}}, \quad \mathcal{E}_2 := \|\hat{\boldsymbol{\mu}}'_m - \boldsymbol{\mu}[2\varepsilon\delta/9]\|_{\text{TV}}, \quad \mathcal{E}_3 := \|\boldsymbol{\mu}[2\varepsilon\delta/9] - \boldsymbol{\mu}\|_{\text{TV}}.$$

By Markov's inequality,

$$\begin{aligned} \mathbb{P}\left(\mathcal{E}_1 > \frac{\varepsilon}{3}\right) &\leq \frac{3}{\varepsilon} \mathbb{E}[\|\hat{\boldsymbol{\mu}}_m - \hat{\boldsymbol{\mu}}'_m\|_{\text{TV}}] = \frac{3}{2\varepsilon} \mathbb{E}\left[\sum_{i=1}^{\infty} |\hat{\boldsymbol{\mu}}_m(i) - \hat{\boldsymbol{\mu}}'_m(i)|\right] \\ &= \frac{3}{2\varepsilon} \mathbb{E}\left[\frac{1}{m} \sum_{i \in \mathbb{N}: \Pi_{\boldsymbol{\mu}}(i) > T_{\boldsymbol{\mu}}(\eta)} \sum_{t=1}^m \mathbf{1}[X_t = i]\right] = \frac{3}{2\varepsilon} \mathbb{P}(\Pi_{\boldsymbol{\mu}}(X_t) > T_{\boldsymbol{\mu}}(\eta)) \leq \frac{\delta}{3}. \end{aligned}$$

Moreover, $\mathcal{E}_3 = \frac{1}{2} \sum_{i > T_{\boldsymbol{\mu}}(\eta)}^{\infty} \boldsymbol{\mu}^{\downarrow}(i) \leq \frac{\varepsilon\delta}{9} \leq \frac{\varepsilon}{3}$. In order to apply the union bound, it remains to handle $\mathbb{P}(\mathcal{E}_2 > \varepsilon/3)$. This is achieved in two standard steps. The first follows an argument similar to that of [Berend and Kontorovich, 2013, Lemma 5], that bounds from above the quantity in expectation using Jensen's inequality, $\mathbb{E}[\mathcal{E}_2] \leq \frac{\|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2}^{1/2}}{\sqrt{m}} \leq \sqrt{\frac{\Lambda}{m}}$. An application of McDiarmid's inequality controls the fluctuations around the expectation [Berend and Kontorovich, 2013, (7.5)] and concludes the proof. \square

Sample complexity lower bound $m = \Omega\left(\frac{\log \delta^{-1}}{\varepsilon^2}\right)$. See Lemma B.1.

Sample complexity lower bound $m = \Omega\left(\frac{\Lambda}{\varepsilon^2}\right)$. Let $\varepsilon \in (0, 1/16)$ and $\Lambda > 2$. First observe that $\Lambda/2 \leq 2\lfloor \Lambda/2 \rfloor \leq \Lambda$, and $2\lfloor \Lambda/2 \rfloor \in 2\mathbb{N}$. As a result,

$$\begin{aligned} \mathcal{R}_m(\Lambda, \varepsilon, \delta) &\stackrel{(i)}{\geq} \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu}: \|\boldsymbol{\mu}[2\varepsilon\delta/9]\|_{1/2} \leq 2\lfloor \Lambda/2 \rfloor} \mathbb{P}_{\mathbf{X} \sim \boldsymbol{\mu}^m}(\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon) \\ &\stackrel{(ii)}{\geq} \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu} \in \Delta_{2\lfloor \Lambda/2 \rfloor}} \mathbb{P}_{\mathbf{X} \sim \boldsymbol{\mu}^m}(\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon) \stackrel{(iii)}{\geq} \frac{1}{2} \left(1 - \frac{mC\varepsilon^2}{2\lfloor \Lambda/2 \rfloor}\right) \geq \frac{1}{2} \left(1 - \frac{2mC\varepsilon^2}{\Lambda}\right) \end{aligned}$$

where (i) and (ii) follow from taking the supremum over increasingly smaller sets, (iii) is Lemma B.2 invoked for $2\lfloor \Lambda/2 \rfloor \in \mathbb{N}$, and $C > 0$ is a universal constant. To conclude, $m \leq \frac{\Lambda}{4C\varepsilon^2} \implies \mathcal{R}_m(\Lambda, \varepsilon, \delta) \geq 1/4$, which yields the second lower bound. \square

Remark: The universal constant in the lower bound obtained by Tsybakov's method at Lemma B.2 is suboptimal, and we give a short proof in the appendix for completeness. We refer the reader to the more involved methods of Kamath et al. [2015] for obtaining tighter bounds.

3.5 Proof of Theorem 2.4

Let $d \in 2\mathbb{N}$ and $m \in \mathbb{N}$, and restrict the problem to $\boldsymbol{\mu} \in \Delta_d$. Let $\varepsilon \in (0, 1/16)$. By Lemma B.2, $\bar{\mathcal{R}}_m(d, \varepsilon) := \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu} \in \Delta_d} \mathbb{P}(\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon) \geq \frac{1}{2} \left(1 - \frac{Cm\varepsilon^2}{d}\right)$ for some $C > 0$, whence Markov's inequality yields

$$\frac{1}{2} \left(1 - \frac{Cm\varepsilon^2}{d}\right) \leq \frac{1}{\varepsilon} \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu} \in \Delta_d} \mathbb{E}[\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}}].$$

Restrict $m \geq \frac{d}{b^2}$, with $b := \sqrt{3C/16}$ and set $\varepsilon = \sqrt{\frac{d}{3Cm}}$, so that

$$\inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu} \in \Delta_d} \mathbb{E}[\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}}] \geq \frac{1}{a} \sqrt{\frac{d}{m}}, \quad \text{where } a := \sqrt{27C} \quad (15)$$

Suppose that $\theta(\sqrt{d/m}) < \frac{1}{a}\sqrt{\frac{d}{m}}$, then by hypothesis,

$$\inf_{\tilde{\mu}} \sup_{\mu \in \Delta_d} \mathbb{E} [\|\tilde{\mu} - \mu\|_{\text{TV}}] \leq \sup_{\mu \in \Delta_d} \mathbb{E} [\Psi_m(\tilde{\mu}_m)] \leq \sup_{\mu \in \Delta_d} \theta \left(\mathbb{E} [\Phi_m] \right).$$

For $\mu \in \Delta_d$, $\mathbb{E} \left[\sqrt{\frac{\|\tilde{\mu}_m\|_{1/2}}{m}} \right] \leq \sqrt{\frac{d}{m}}$. It follows that

$$\sup_{\mu \in \Delta_d} \theta \left(\mathbb{E} [\Phi_m] \right) \leq \theta \left(\sqrt{\frac{d}{m}} \right) < \frac{1}{a} \sqrt{\frac{d}{m}},$$

which contradicts (15). We have therefore established, for

$$r \in R := \left\{ \sqrt{d/m} : (m, d) \in \mathbb{N} \times 2\mathbb{N}, m \geq \frac{d}{b^2} \right\},$$

the lower bound $\theta(r) \geq r/a$. We extend the lower bound to the open interval $(0, b)$, by observing that R is dense in $(0, b)$ followed by a continuity argument. \square

Broader Impact

This work is of purely theoretical nature and does not present any foreseeable societal consequence.

Acknowledgments and Disclosure of Funding

We are thankful to Clayton Scott for the insightful conversations, and to the anonymous referees for their valuable comments. This research was partially supported by the Israel Science Foundation (grant No. 1602/19) and Google Research.

References

- M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, 1999. ISBN 0-521-57353-X. doi: 10.1017/CBO9780511624216. URL <http://dx.doi.org/10.1017/CBO9780511624216>.
- D. Berend and A. Kontorovich. The missing mass problem. *Statistics & Probability Letters*, 82(6):1102–1110, 2012. URL <https://EconPapers.repec.org/RePEc:eee:stapro:v:82:y:2012:i:6:p:1102-1110>.
- D. Berend and A. Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
- O. Bousquet, D. Kane, and S. Moran. The optimal approximation factor in density estimation. In A. Beygelzimer and D. Hsu, editors, *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pages 318–341. PMLR, 2019. URL <http://proceedings.mlr.press/v99/bousquet19a.html>.
- F. Deutsch and H. Hundal. Slow convergence of sequences of linear operators i: Almost arbitrarily slow convergence. *Journal of Approximation Theory*, 162(9):1701–1716, 2010a.
- F. Deutsch and H. Hundal. Slow convergence of sequences of linear operators ii: Arbitrarily slow convergence. *Journal of Approximation Theory*, 162(9):1717–1738, 2010b.
- L. Devroye and L. Györfi. *Nonparametric density estimation: the L_1 view*. Wiley Series in Probability and Mathematical Statistics: Tracts on Probability and Statistics. John Wiley & Sons, Inc., New York, 1985. ISBN 0-471-81646-9.
- L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics. Springer-Verlag, New York, 2001. ISBN 0-387-95117-2. doi: 10.1007/978-1-4613-0125-7. URL <http://dx.doi.org/10.1007/978-1-4613-0125-7>.

- I. Diakonikolas. Learning structured distributions. In P. Bühlmann, P. Drineas, M. J. Kane, and M. J. van der Laan, editors, *Handbook of Big Data*, pages 267–283. Chapman and Hall/CRC, 2016. URL <http://www.crcnetbase.com/doi/abs/10.1201/b19567-21>.
- S. R. Dunbar. Topics in probability theory and stochastic processes. *Department of Mathematics. University of Nebraska, Lincoln/NE*, 2009.
- U. Haagerup. The best constants in the khintchine inequality. *Studia Mathematica*, 70(3):231–283, 1981. URL <http://eudml.org/doc/218383>.
- Y. Han, J. Jiao, and T. Weissman. Minimax estimation of discrete distributions under ℓ_1 loss. *IEEE Transactions on Information Theory*, 61(11):6343–6354, Nov 2015. ISSN 0018-9448. doi: 10.1109/TIT.2015.2478816.
- M. B. Handelsman. Solution to problem 436. Distributing ‘heads’ minus ‘tails’. *The College Mathematics Journal*, 22(5):444–449, 1991. ISSN 07468342, 19311346. URL <http://www.jstor.org/stable/2686609>.
- S. Kamath, A. Orlitsky, D. Pichapati, and A. T. Suresh. On learning distributions from their samples. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, pages 1066–1100, 2015. URL <http://jmlr.org/proceedings/papers/v40/Kamath15.html>.
- A. Kontorovich and I. Pinelis. Exact lower bounds for the agnostic probably-approximately-correct (pac) machine learning model. *Ann. Statist.*, 47(5):2822–2854, 2019. ISSN 0090-5364. doi: 10.1214/18-AOS1766.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2012. ISBN 026201825X.
- A. Orlitsky and A. T. Suresh. Competitive distribution estimation: Why is good-turing good. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2143–2151, 2015. URL <http://papers.nips.cc/paper/5762-competitive-distribution-estimation-why-is-good-turing-good>.
- C. D. Scott and R. D. Nowak. Learning minimum volume sets. *J. Mach. Learn. Res.*, 7:665–704, Dec. 2006. ISSN 1532-4435.
- A. B. Tsybakov. Introduction to nonparametric estimation, 2009. URL <https://doi.org/10.1007/b13794>. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.
- G. Valiant and P. Valiant. Instance optimal learning of discrete distributions. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 142–155, 2016.
- G. Valiant and P. Valiant. An automatic inequality prover and instance optimal identity testing. *SIAM Journal on Computing*, 46(1):429–455, 2017.
- B. Waggoner. L_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 347–356, 2015. doi: 10.1145/2688073.2688095. URL <http://doi.acm.org/10.1145/2688073.2688095>.
- M. M. J. Wainwright. *High-dimensional statistics : a non-asymptotic viewpoint*. Cambridge series on statistical and probabilistic mathematics ; 48. Cambridge University Press, Cambridge, United Kingdom, 2019. ISBN 9781108498029.
- J. Zimmert and Y. Seldin. An optimal algorithm for stochastic and adversarial bandits. In K. Chaudhuri and M. Sugiyama, editors, *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pages 467–475. PMLR, 2019. URL <http://proceedings.mlr.press/v89/zimmert19a.html>.

A Analysis of the Empirical Rademacher complexity

From Lemma 3.2 (see also [Scott and Nowak, 2006, Section 7.1, Appendix E.]), we see that the Khintchine inequality already yields a control of $\hat{\mathfrak{R}}_m(\mathbf{X})$ by $\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}$ up to universal constants.

$$\frac{1}{2\sqrt{2}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \leq \hat{\mathfrak{R}}_m(\mathbf{X}) \leq \frac{1}{2} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}.$$

Furthermore, it is possible to derive an exact expression for it, from the expected absolute distance of a symmetric random walk:

Lemma A.1 (Empirical Rademacher complexity, exact expression). *Let $\mathbf{X} = (X_1, \dots, X_m)$ and let $\hat{\boldsymbol{\mu}}_m$ be the empirical measure constructed from the sample \mathbf{X} . Then,*

$$\hat{\mathfrak{R}}_m(\mathbf{X}) = \frac{1}{m} \sum_{x: \hat{\boldsymbol{\mu}}_m(x) > 0} \frac{1}{2^{m\hat{\boldsymbol{\mu}}_m(x)}} \left\lfloor \frac{m\hat{\boldsymbol{\mu}}_m(x)}{2} \right\rfloor \binom{m\hat{\boldsymbol{\mu}}_m(x)}{\lfloor m\hat{\boldsymbol{\mu}}_m(x)/2 \rfloor}.$$

Proof. Write $m\hat{\mathfrak{R}}_m(\mathbf{X}) = \sum_{x: \hat{\boldsymbol{\mu}}_m(x) > 0} \frac{1}{2} \mathbb{E}_{\sigma} \left[\left| \sum_{i=1}^{m\hat{\boldsymbol{\mu}}_m(x)} \sigma_i \right| \right]$ as in the proof of Lemma 3.2.

Now, observe that $\mathbb{E}_{\sigma} \left[\left| \sum_{i=1}^{m\hat{\boldsymbol{\mu}}_m(x)} \sigma_i \right| \right]$ is the expectation value of the absolute distance of a 1-dimensional symmetric random walk after $m\hat{\boldsymbol{\mu}}_m(x)$ steps, also known as the ‘‘heads minus tails’’ process [Handelsman, 1991]:

$$\mathbb{E}_{\sigma} \left[\left| \sum_{i=1}^{m\hat{\boldsymbol{\mu}}_m(x)} \sigma_i \right| \right] = \frac{1}{m2^{m\hat{\boldsymbol{\mu}}_m(x)}} \left\lfloor \frac{m\hat{\boldsymbol{\mu}}_m(x)}{2} \right\rfloor \binom{m\hat{\boldsymbol{\mu}}_m(x)}{\lfloor m\hat{\boldsymbol{\mu}}_m(x)/2 \rfloor}.$$

□

However, the above is inconvenient and involves the computation of factorials. Leveraging delicate bounds for the central binomial coefficient obtained with the Wallis product in Dunbar [2009], we derive the following corollary, that gives exact the first-order constant in terms of the *half-norm*, makes the *minus-half-norm* appear as a second dominant term, and that is easily computable.

Corollary A.1 (Empirical Rademacher complexity, first order bound). *Let $\mathbf{X} = (X_1, \dots, X_m)$ and let $\hat{\boldsymbol{\mu}}_m$ be the empirical measure constructed from the sample \mathbf{X} . Then writing*

$$\phi_m(\mathbf{X}) := \frac{\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}}{\sqrt{2\pi m}},$$

it holds that

$$\frac{\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}}{\sqrt{2\pi m}} - \frac{3}{2} \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \|\hat{\boldsymbol{\mu}}_m^+\|_{-1/2}^{-1/2} \leq \hat{\mathfrak{R}}_m(\mathbf{X}) \leq \frac{\|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}}{\sqrt{2\pi m}} + \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \|\hat{\boldsymbol{\mu}}_m^+\|_{-1/2}^{-1/2}.$$

Proof. Let $n \in \mathbb{N}$, if $n = 2k, k \geq 1$,

$$\frac{1}{2^n} \left\lfloor \frac{n}{2} \right\rfloor \binom{n}{\lfloor n/2 \rfloor} = \frac{1}{4^k} k \binom{2k}{k},$$

and if $n = 2k - 1, k \geq 1, \lceil n/2 \rceil = k$ such that similarly,

$$\begin{aligned} \frac{1}{2^n} \left\lfloor \frac{n}{2} \right\rfloor \binom{n}{\lfloor n/2 \rfloor} &= \frac{1}{2^{2k-1}} k \binom{2k-1}{k} = \frac{2}{4^k} k \frac{(2k-1)!}{k!(2k-k-1)!} \\ &= \frac{2}{4^k} k \frac{(2k)!(2k-k)}{(2k)k!(2k-k)!} = \frac{2}{4^k} k \frac{2k-k}{2k} \binom{2k}{k} = \frac{1}{4^k} k \binom{2k}{k}. \end{aligned}$$

Moreover, from Dunbar [2009, p.11], for $k \geq 1$, an application of the Wallis product yields,

$$\frac{k}{\sqrt{\pi/2}\sqrt{2k+1}} \left(1 - \frac{1}{2k}\right) \leq \frac{1}{4^k} k \binom{2k}{k} \leq \frac{k}{\sqrt{\pi/2}\sqrt{2k+1}} \left(1 + \frac{1}{2k}\right).$$

If follows that when $n = 2k$,

$$\sqrt{\frac{n}{2\pi}} \left\{ \sqrt{\frac{n}{n+1}} \left(1 - \frac{1}{n}\right) \right\} \leq \frac{1}{2^n} \binom{n}{\lfloor n/2 \rfloor} \leq \sqrt{\frac{n}{2\pi}} \left\{ \sqrt{\frac{n}{n+1}} \left(1 + \frac{1}{n}\right) \right\},$$

and for $n = 2k - 1$,

$$\sqrt{\frac{n}{2\pi}} \left\{ \frac{n+1}{\sqrt{n(n+2)}} \left(1 - \frac{1}{n+1}\right) \right\} \leq \frac{1}{2^n} \binom{n}{\lfloor n/2 \rfloor} \leq \sqrt{\frac{n}{2\pi}} \left\{ \frac{n+1}{\sqrt{n(n+2)}} \left(1 + \frac{1}{n+1}\right) \right\}$$

For all $n \in \mathbb{N}$,

$$\begin{aligned} \sqrt{\frac{n}{n+1}} \left(1 + \frac{1}{n}\right) &\leq \frac{n+1}{\sqrt{n(n+2)}} \left(1 + \frac{1}{n+1}\right) \leq 1 + \frac{1}{n}, \\ \sqrt{\frac{n}{n+1}} \left(1 - \frac{1}{n}\right) &\geq \frac{n+1}{\sqrt{n(n+2)}} \left(1 - \frac{1}{n+1}\right) \geq 1 - \frac{3}{2n}, \end{aligned}$$

such that

$$\begin{aligned} \hat{\mathfrak{R}}_m(\mathbf{X}) &\leq \sqrt{\frac{1}{2\pi m}} \sum_{x: \hat{\boldsymbol{\mu}}_m(x) > 0} \sqrt{\hat{\boldsymbol{\mu}}_m(x)} \left\{ 1 + \frac{1}{m \hat{\boldsymbol{\mu}}_m(x)} \right\} \\ &\leq \phi_m(\mathbf{X}) + \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \left\| \hat{\boldsymbol{\mu}}_m^+ \right\|_{-1/2}^{-1/2}, \end{aligned}$$

where we wrote

$$\left\| \hat{\boldsymbol{\mu}}_m^+ \right\|_{-1/2}^{-1/2} := \sum_{x \in \mathbb{N}} \frac{\mathbf{1}[\hat{\boldsymbol{\mu}}_m(x) > 0]}{\sqrt{\hat{\boldsymbol{\mu}}_m(x)}},$$

and conversely,

$$\hat{\mathfrak{R}}_m(\mathbf{X}) \geq \phi_m(\mathbf{X}) - \frac{3}{2} \sqrt{\frac{1}{2\pi}} \frac{1}{m^{3/2}} \left\| \hat{\boldsymbol{\mu}}_m^+ \right\|_{-1/2}^{-1/2}.$$

□

B Auxiliary lemmas for lower bounds

Lemma B.1 (Sample complexity lower bound $m = \Omega(\log \delta^{-1}/\varepsilon^2)$). *Let $\Lambda \geq 2$, $0 < \varepsilon < 1/2$ and $0 < \delta < 1$. For any estimator $\bar{\boldsymbol{\nu}}: \mathbb{N}^m \rightarrow \Delta_{\mathbb{N}}$ there is a $\boldsymbol{\nu} \in \Delta_{\mathbb{N}}$ with $\|\boldsymbol{\nu}[2\varepsilon\delta/9]\|_{1/2} \leq \Lambda$, such that $\bar{\boldsymbol{\nu}}$ must require at least $m = \Omega\left(\frac{\log \delta^{-1}}{\varepsilon^2}\right)$ samples in order for $\|\bar{\boldsymbol{\nu}} - \boldsymbol{\nu}\|_{\text{TV}} < \varepsilon$ to hold with probability at least $1 - \delta$.*

Proof. The proof is standard and consists of lower bounding the difficulty of learning a biased coin. Recall that for $\boldsymbol{\mu}_0 := (1/2, 1/2)$, $\boldsymbol{\mu}_\varepsilon := (1/2 - \varepsilon, 1/2 + \varepsilon)$, direct computations lead to $\|\boldsymbol{\mu}_0 - \boldsymbol{\mu}_\varepsilon\|_1 = 2\varepsilon$, and $D_{\text{KL}}(\boldsymbol{\mu}_\varepsilon \|\boldsymbol{\mu}_0) = (1/2 - \varepsilon) \ln \frac{1/2 - \varepsilon}{1/2} + (1/2 + \varepsilon) \ln \frac{1/2 + \varepsilon}{1/2} \leq 4\varepsilon^2$, where $D_{\text{KL}}(\boldsymbol{\mu}_\varepsilon \|\boldsymbol{\mu}_0)$ is the KL divergence between $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\mu}_0$. We also verify that $\|\boldsymbol{\mu}_\varepsilon\|_{1/2} \leq \|\boldsymbol{\mu}_0\|_{1/2} \leq 2 \leq \Lambda$, hence also for their truncated version. From an immediate corollary of LeCam's theorem [Tsybakov, 2009, Theorem 2.2, Lemma 2.6], $\mathcal{R}_m(\Lambda, \varepsilon, \delta) \geq \frac{1}{2} \exp(-m D_{\text{KL}}(\boldsymbol{\mu}_\varepsilon \|\boldsymbol{\mu}_0))$, whence $m \leq \frac{1}{4\varepsilon^2} \log \frac{\delta^{-1}}{2} \implies \mathcal{R}_m(\Lambda, \varepsilon, \delta) \geq \delta$. □

Lemma B.2. *Let $d \in 2\mathbb{N}$, $d \geq 16$, $m \in \mathbb{N}$, $\varepsilon \in (0, 1/16)$, and let*

$$\bar{\mathcal{R}}_m(d, \varepsilon) := \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu}: \boldsymbol{\mu} \in \Delta_d} \mathbb{P}_{\mathbf{X} \sim \boldsymbol{\mu}^m} (\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon),$$

where the infimum is taken over all $\bar{\boldsymbol{\mu}}: [d]^m \rightarrow \Delta_d$. Then there is a universal $C > 0$ such that

$$\bar{\mathcal{R}}_m(d, \varepsilon) \geq \frac{1}{2} \left(1 - \frac{Cm\varepsilon^2}{d}\right).$$

Proof. As is customary in Analysis, the universal constant $C > 0$ may change its value from expression to expression. Consider the family of distributions

$$\mathcal{D}(d) := \left\{ \boldsymbol{\mu}^{(\boldsymbol{\sigma})} \in \Delta_d, \boldsymbol{\sigma} \in \{0, 1\}^{d/2} \right\},$$

where

$$\boldsymbol{\mu}^{(\boldsymbol{\sigma})} := \frac{1}{d} (1 + 16\varepsilon\sigma_1, 1 - 16\varepsilon\sigma_1, 1 + 16\varepsilon\sigma_2, 1 - 16\varepsilon\sigma_2, \dots, 1 + 16\varepsilon\sigma_{d/2}, 1 - 16\varepsilon\sigma_{d/2}).$$

From the Varshamov-Gilbert bound [Tsybakov, 2009, Lemma 2.9], there exists a $\tilde{\mathcal{D}}(d) \subsetneq \mathcal{D}(d)$ satisfying (a) $|\tilde{\mathcal{D}}(d)| > 2^{d/16}$, (b) for $\boldsymbol{\mu}^{(\boldsymbol{\sigma})}, \boldsymbol{\mu}^{(\boldsymbol{\sigma}')} \in \tilde{\mathcal{D}}(d)$, $\boldsymbol{\sigma} \neq \boldsymbol{\sigma}' \implies \left\| \boldsymbol{\mu}^{(\boldsymbol{\sigma})} - \boldsymbol{\mu}^{(\boldsymbol{\sigma}')} \right\|_{\text{TV}} \geq 2\varepsilon$, and (c) $\boldsymbol{\mu}^{(0)} \in \tilde{\mathcal{D}}(d)$. It is straightforward to verify that $D_{\text{KL}}(\boldsymbol{\mu}^{(\boldsymbol{\sigma})} \parallel \boldsymbol{\mu}^{(0)}) \leq C\varepsilon^2$. Applying Tsybakov's method [Tsybakov, 2009, Theorem 2.5],

$$\begin{aligned} \bar{\mathcal{R}}_m(d, \varepsilon) &\geq \inf_{\bar{\boldsymbol{\mu}}} \sup_{\boldsymbol{\mu} \in \tilde{\mathcal{D}}(d)} \mathbb{P}(\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\|_{\text{TV}} > \varepsilon) \\ &\geq \frac{1}{2} \left(1 - \frac{\frac{4m}{|\tilde{\mathcal{D}}(d)|} \sum_{\boldsymbol{\mu}^{(\boldsymbol{\sigma})} \in \tilde{\mathcal{D}}(d)} D_{\text{KL}}(\boldsymbol{\mu}^{(\boldsymbol{\sigma})} \parallel \boldsymbol{\mu}^{(0)})}{\ln |\tilde{\mathcal{D}}(d)|} \right) \end{aligned}$$

so that $\bar{\mathcal{R}}_m(d, \varepsilon) \geq \frac{1}{2} \left(1 - \frac{Cm\varepsilon^2}{d} \right)$. \square

C Convergence properties of the empirical bound

In this section, we briefly analyze convergence of $\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}$. In Proposition C.1 we confirm that the quantity converges *almost surely* and in L_1 , but with Proposition C.2, we show that this convergence can be *arbitrarily slow*.

Proposition C.1 (L_1 and almost sure convergence). *Let $\boldsymbol{\mu} \in \Delta_{\mathbb{N}}$ and let $\mathbf{X} := (X_1, \dots, X_m) \sim \boldsymbol{\mu}^m$. Then, $\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \xrightarrow{L_1} 0$ and $\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \xrightarrow{\text{a.s.}} 0$.*

Proof. For L_1 convergence the proof is as follows:

$$\begin{aligned} \lim_{m \rightarrow \infty} \mathbb{E} \left[\left| \frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} - 0 \right| \right] &= \lim_{m \rightarrow \infty} \mathbb{E} \left[\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \right] \\ &\leq \lim_{m \rightarrow \infty} 2\Lambda_m(\boldsymbol{\mu}) && \text{(Theorem 2.3)} \\ &= 0. && \text{([Berend and Kontorovich, 2013, Lemma 7])} \end{aligned}$$

Now, for almost sure convergence, recall that $\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2}$ satisfies $2/m$ -bounded-differences. By the L_1 convergence established above, we have that for all $\varepsilon > 0$ there is an $M_\varepsilon \in \mathbb{N}$ s.t. for all $m \geq M_\varepsilon$, we have $\mathbb{E} \left[\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \right] \leq \varepsilon/2$. Invoking McDiarmid's inequality, for every $m \geq M_\varepsilon$, we have

$$\mathbb{P} \left(\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \geq \varepsilon/2 \right) \leq \exp \left(-\frac{m\varepsilon^2}{2} \right).$$

Thus,

$$\sum_{m=1}^{\infty} \mathbb{P} \left(\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \geq \varepsilon/2 \right) \leq M_\varepsilon + \sum_{m=M_\varepsilon}^{\infty} \exp \left(-\frac{m\varepsilon^2}{2} \right) < \infty.$$

An application of the Borel-Cantelli lemma completes the proof:

$$\frac{1}{\sqrt{m}} \|\hat{\boldsymbol{\mu}}_m\|_{1/2}^{1/2} \xrightarrow{\text{a.s.}} 0.$$

\square

To formalize our idea of arbitrarily slow convergence, we adapt the terminology developed in Deutsch and Hundal [2010a,b]. We begin with the set of all $[0, 1]$ -valued sequences that converge to 0:

$$\mathcal{U} := \{U \in [0, 1]^{\mathbb{N}} : \lim_{m \rightarrow \infty} U(m) = 0\}.$$

Following Deutsch and Hundal [2010a, Definition 2.7], we will say that the statistic $\hat{\theta}_m : \mathbb{N}^m \rightarrow [0, 1]$ converges arbitrarily slowly to 0 in L_1 if

1. $\forall \mu \in \Delta_{\mathbb{N}}, \lim_{m \rightarrow \infty} \mathbb{E} [\hat{\theta}_m] = 0,$
2. $\forall U \in \mathcal{U}, \exists \mu \in \Delta_{\mathbb{N}}$ such that $\forall m \in \mathbb{N}, \mathbb{E} [\hat{\theta}_m] \geq U(m).$

It turns out [Deutsch and Hundal, 2010b, Remark 2.8, Theorem 2.9] that restricting the set \mathcal{U} to the decreasing sequences,

$$\mathcal{U}^\downarrow := \{U \in [0, 1]^{\mathbb{N}} : \sup_{m \in \mathbb{N}} U(m+1)/U(m) \leq 1, \lim_{m \rightarrow \infty} U(m) = 0\},$$

does not change the above definition of arbitrarily slow convergence.

Proposition C.2 (Arbitrary slow convergence in L_1). *For any sequence $1 > r_1 > r_2 > \dots$ decreasing to 0, there is a distribution $\mu \in \Delta_{\mathbb{N}}$ such that $2 \mathbb{E}_{\mathbf{X} \sim \mu^m} [\|\mu - \hat{\mu}_m\|_{\text{TV}}] > r_m$ for all $m \geq 1$.*

Proof.

$$\begin{aligned} 2 \mathbb{E} [\|\mu - \hat{\mu}_m\|_{\text{TV}}] &= \mathbb{E} [\|\mu - \hat{\mu}_m\|_1] \\ &= \sum_{i=1}^{\infty} \mathbb{E} [|\mu(i) - \hat{\mu}_m(i)|] \\ &= \sum_{i=1}^{\infty} \mathbb{E} [\mathbf{1}[\hat{\mu}_m(i) > 0] |\mu(i) - \hat{\mu}_m(i)| + \mathbf{1}[\hat{\mu}_m(i) = 0] |\mu(i) - \hat{\mu}_m(i)|] \\ &= \sum_{i=1}^{\infty} \mathbb{E} [\mathbf{1}[\hat{\mu}_m(i) > 0] |\mu(i) - \hat{\mu}_m(i)|] + \sum_{i=1}^{\infty} \mathbb{E} [\mathbf{1}[\hat{\mu}_m(i) = 0] \mu(i)] \\ &= \sum_{i=1}^{\infty} \mathbb{E} [\mathbf{1}[\hat{\mu}_m(i) > 0] |\mu(i) - \hat{\mu}_m(i)|] + \mathbb{E} \left[\sum_{i: \hat{\mu}_m(i)=0} \mu(i) \right] \\ &\geq \mathbb{E} \left[\sum_{i: \hat{\mu}_m(i)=0} \mu(i) \right] = \mathbb{E} [\mu(\mathbb{N} \setminus \{X_1, \dots, X_m\})] \\ &= \mathbb{E} [U_m], \end{aligned}$$

where $U_m := \mu(\mathbb{N} \setminus \{X_1, \dots, X_m\})$ is the missing mass random variable. From [Berend and Kontorovich, 2012, Proposition 4], we have that: For any sequence $1 > r_1 > r_2 > \dots$ decreasing to 0, there is a distribution $\mu \in \Delta_{\mathbb{N}}$ such that $\mathbb{E} [U_m] > r_m$ for all $m \geq 1$. \square

Remark C.1. *To our knowledge, the above result is the first to establish a connection between the TV risk $\|\mu - \hat{\mu}_m\|_{\text{TV}}$ and the missing mass U_m .*