

L_q Lower Bounds on Distributed Estimation via Fisher Information

Wei-Ning Chen
wnchen@stanford.edu

Ayfer Özgür
aozgur@stanford.edu

Stanford University

Abstract

Van Trees inequality, also known as the Bayesian Cramér-Rao lower bound, is a powerful tool for establishing lower bounds for minimax estimation through Fisher information. It easily adapts to different statistical models and often yields tight bounds. Recently, its application has been extended to distributed estimation with privacy and communication constraints where it yields order-wise optimal minimax lower bounds for various parametric tasks under squared L_2 loss.

However, a widely perceived drawback of the van Trees inequality is that it is limited to squared L_2 loss. The goal of this paper is to dispel that perception by introducing a strengthened version of the van Trees inequality that applies to general L_q loss functions by building on the Efroimovich's inequality – a lesser-known entropic inequality dating back to the 1970s. We then apply the generalized van Trees inequality to lower bound L_q loss in distributed minimax estimation under communication and local differential privacy constraints. This leads to lower bounds for L_q loss that apply to sequentially interactive and blackboard communication protocols. Additionally, we show how the generalized van Trees inequality can be used to obtain *local* and *non-asymptotic* minimax results that capture the hardness of estimating each instance at finite sample sizes.

Contents

1	Introduction	2
2	Preliminaries	3
3	Distributed Estimation under Information Constraints	5
3.1	Problem Setup	6
3.2	Global Minimax Lower Bounds	6
3.3	Achievability	7
4	Local Minimax Bounds on Discrete Distribution Estimation	7
4.1	b -bit communication constraint	8
4.2	ϵ -local differential privacy	9
5	Conclusion and Future Works	9
A	Global Minimax Lower Bounds under LDP	13
B	Missing Proofs	13
B.1	Proof of Theorem 2.4	13
B.2	Proof of Theorem 4.2	14
B.3	Proof of Corollary 3.2 and Corollary A.1	16
B.4	Proof of Theorem 4.4	16

C Proof of Lemmas	17
C.1 Proof of Lemma B.1	17
C.2 Proof of Lemma B.3	17

1 Introduction

For a real-valued parameter $\theta \in \Theta$ and an observation $X \sim P_\theta$, the basic question of parametric statistics is how well one can estimate θ from X under a given loss function $\ell(\hat{\theta} - \theta)$. The Fisher information $I_X(\theta)$ plays a crucial role in this context by capturing the *local* hardness of estimating θ from X , with implications both asymptotically (Hájek, 1961; Van der Vaart, 2000) and non-asymptotically (Cramér, 1999; Rao, 1992; Van Trees, 2004; Gill and Levit, 1995; Efroimovich, 1980; Aras et al., 2019). For example, the famous Cramér-Rao lower bound (Cramér, 1999; Rao, 1992) states that the squared L_2 estimation error (a.k.a. the mean squared error) of any unbiased estimator is lower bounded by the inverse of the Fisher information. The unbiasedness assumption on the estimator can be removed by assuming that the parameter θ is distributed according to a prior on Θ . This is known as the Bayesian Cramér-Rao lower bound or the van Trees Inequality (Van Trees, 2004; Gill and Levit, 1995). By strategically constructing the prior, this approach can also be used to prove lower bounds for minimax estimation, where the goal is to minimize the worst-case error of the estimator over the parameter space Θ .

The classic parametric estimation tasks have gained renewed popularity over recent years, driven by the prevalence of modern datasets often generated and stored on local devices, such as in federated learning (FL) and analytics (FA) (McMahan et al., 2017a,b; Kairouz et al., 2019). In these settings, the collection and utilization of decentralized data encounter various resource constraints, including communication and privacy considerations. These constraints have raised the question of how well one can estimate the unknown parameter $\theta \in \Theta$ from a processed observation Y , which corresponds to the output of a privatization and/or compression mechanism applied to $X \sim P_\theta$. In Barnes et al. (2019a,b, 2020), the classical Fisher information framework has been extended to the case of privacy and communication (or compression) constraints. This approach first upper bounds the Fisher information from a differentially private and/or compressed sample Y and then uses the van Trees inequality to lower bound the minimax squared L_2 error of distributed estimation. This, in a unified fashion, leads to order-optimal lower bounds for various parametric tasks under privacy and/or communication constraints, including two tasks that have been of significant interest in the recent literature: distributed mean estimation (Duchi et al., 2013; Gandikota et al., 2021; Suresh et al., 2017; Zhang et al., 2012; Duchi and Ruan, 2018; Duchi and Rogers, 2019; Asi et al., 2022; Agarwal et al., 2018; Chen et al., 2020; Feldman and Talwar, 2021; Shah et al., 2022; Asi et al., 2023; Isik et al., 2023, see) and discrete distribution (or frequency) estimation (Han et al., 2018a,b; Úlfar Erlingsson et al., 2014; Kairouz et al., 2016; Wang et al., 2017; Ye and Barg, 2017; Acharya et al., 2019, 2020a; Chen et al., 2020; Feldman and Talwar, 2021; Feldman et al., 2022, see). While order-optimal minimax lower bounds for distributed estimation under information constraints can also be derived by using other techniques, e.g., leveraging strong data-processing inequalities in Duchi et al. (2013); Zhang et al. (2013); Garg et al. (2014); Braverman et al. (2016); Duchi and Rogers (2019); Asodeh and Zhang (2022) or employing methods based on Le Cam, Fano, or Assaud (Acharya et al., 2020a,b, 2023), the constrained Fisher information approach has several distinct advantages:

- It is relatively straightforward and applies to various parametric models in a unified fashion, e.g., mean estimation or distribution estimation. In contrast, methods based on Le Cam, Fano, or Assaud are more versatile in that they can be potentially adopted to diverse statistical problems beyond parametric settings, e.g., testing, but the construction of worst-case instances tends to be intricate and is heavily tied to the specific problem structure.
- The Fisher information approach elucidates a clear relationship between the tail behavior of the score function of the parametric model and error rates under privacy/communication constraints.
- Leveraging the chain rule of Fisher information, the Fisher information approach naturally extends to sequentially interactive communication protocols or even fully interactive blackboard protocols (Kushilevitz, 1997).

- Finally, Fisher information is inherently a local measure of hardness. For example, the influential Hájek-Le Cam’s local asymptotic minimax (LAM) theorem asserts that, for a general ball-shaped loss function $\ell(\cdot)$, the *asymptotic* estimation error at $\theta \in \Theta$ is at least $\mathbb{E}[\ell(Z)]$, where $Z \sim \mathcal{N}(0, I_X^{-1}(\theta))$ (Hájek, 1961; Van der Vaart, 2000). Using the van Trees inequality with a carefully constructed prior, Chen et al. (2021b) shows that Fisher information can be used to prove *local* and *non-asymptotic* minimax lower bounds that capture the hardness of estimating each instance of the problem at finite sample size. To the best of our knowledge, this is the only example of *non-asymptotic local* lower bounds for distributed estimation, combining the power of Hájek-Le Cam type results to capture *local* hardness with the *non-asymptotic* nature of global minimax lower bounds that have been almost exclusively the focus of the existing distributed estimation literature.

However, a widely perceived drawback of the Fisher information approach, both for classical parametric statistics as well as its extension to the distributed setting with privacy and communication constraints, is that its applicability is limited to squared L_2 error. For example, Tsybakov (2004) acknowledges several advantages of van Trees inequality, such as its relative simplicity in application and ability to establish sharp bounds, and also notes “*a limitation is that the van Trees inequality applies only to the squared loss function.*” The goal of this paper is to dispel this perception by showing how Fisher information can be used to prove lower bounds for any L_q loss for $1 \leq q < \infty$. We do this by building on a less-known entropic inequality called Efroimovich’s inequality (Efroimovich, 1980) highlighted in recent works Aras et al. (2019); Lee (2022). Our paper makes the following contributions:

- We prove a van Trees type inequality for L_q loss by combining Efroimovich’s inequality with a maximum entropy argument.¹
- We leverage this generalized van Trees inequality to establish global minimax lower bounds for various distributed estimation tasks under L_q loss. As an immediate consequence, this yields lower bounds for L_q loss applicable to sequentially interactive and blackboard communication protocols. Our approach not only recovers previous lower bounds presented in Acharya et al. (2023) for sequentially interactive models in a cleaner and more straightforward manner, but also extends to fully interactive blackboard communication protocols.
- We show how the generalized van Trees inequality can be used to derive *local non-asymptotic* minimax lower bounds for distribution estimation under L_q loss extending the approach of Chen et al. (2021b). This *local* lower bounds match the performance of a scheme previously developed in Chen et al. (2021b) and establish its instance-optimality under L_1 loss. The results extend to L_q loss.

Organization. The rest of the paper is organized as follows. We recap van Trees inequality and extend it to L_q loss via Efroimovich inequality in Section 2. In Section 3, we leverage the generalized van Trees inequality to establish global minimax lower bounds for various distributed estimation tasks under L_q loss. In Section 4, we extend our focus to derive *local* minimax lower bounds for distribution estimation tasks and provide schemes that achieve the lower bounds pointwisely. Finally, we summarize the work in Section 5.

2 Preliminaries

In this section, we recap van Trees inequality and Efroimovich inequality and refer interested readers to Aras et al. (2019) and Lee (2022) for more details and recent advances. Throughout the paper, we consider the following parametric model: let $\{P_\theta | \theta \in \Theta\}$ with $\Theta \subseteq \mathbb{R}^d$ be a family of probability measures over $\mathcal{X}$ with a dominating σ -finite measure λ such that the density:

$$dP_\theta(\cdot) = f(\cdot; \theta) d\lambda(\cdot)$$

exists. We make the following assumption on the model:

¹In Aras et al. (2019); Lee (2022), the authors comment that Efroimovich’s inequality can be potentially used to bound loss functions beyond L_2 . However, we have not been able to find an explicit result in this direction.

Assumption 2.1 *The density function $f(x; \theta)$ is differentiable for $x \in \mathcal{X}$ λ -a.e. In addition, f satisfies*

$$\int_{\mathcal{X}} \nabla_{\theta} f(x; \theta) d\lambda(x) = 0,$$

for all $\theta \in \Theta$.

Assumption 2.1 is a common regularity condition in the Cramér-Rao type bounds, allowing for exchanging the differentiation and integration. For a prior distribution π over Θ , the information theorists' Fisher information is defined as

$$J(\pi) \triangleq \int_{\mathbb{R}^d} \frac{|\nabla_{\theta} \pi(\theta)|^2}{\pi(\theta)}, \quad (1)$$

if it exists. On the other hand, the Fisher information matrix of $\{P_{\theta} | \theta \in \Theta\}$ is defined as

$$[I_X(\theta)]_{ij} \triangleq \int_{\mathcal{X}} \frac{\frac{\partial}{\partial \theta_i} f(x; \theta) \cdot \frac{\partial}{\partial \theta_j} f(x; \theta)}{f(x; \theta)} d\lambda(x). \quad (2)$$

Van Trees inequality states that the L_2 estimation error for estimating $\theta \in \Theta$, given a prior distribution π over Θ , is lower bound by the inverse Fisher information:

Theorem 2.2 (Van Trees (2004); Gill and Levit (1995)) *Let $X \sim P_{\theta}$ and $\theta \sim \pi$ for some prior distribution π . Let Assumption 2.1 hold. Then, for any prior distribution π on Θ for which the information theorist's Fisher information $J(\pi)$ exists, it holds that*

$$\mathbb{E} \left[\left\| \theta - \hat{\theta} \right\|_2^2 \right] \geq \frac{d}{\det(I_X(\theta) + J(\pi))^{1/d}}. \quad (3)$$

While van Trees inequality only applies to the L_2 error, the following Efroimovich inequality can be used to establish lower bounds for more general loss functions.

Theorem 2.3 (Efroimovich (1980)) *Under the assumptions of Theorem 2.2, we also have*

$$\frac{1}{2\pi e} e^{\frac{2}{d} h(\theta|X)} \geq \frac{1}{\det(I_X(\theta) + J(\pi))^{1/d}}, \quad (4)$$

where $h(\theta|X)$ is the conditional (differential) entropy.

To see how Theorem 2.3 implies the classical van Trees inequality, observe that

$$\begin{aligned} \frac{2}{d} h(\theta|X) &\leq \frac{2}{d} h(\theta|\hat{\theta}(X)) \leq \frac{2}{d} h(\theta - \hat{\theta}) \\ &\leq \log \left(2\pi e \cdot \det \left(\mathbb{E} \left[(\hat{\theta} - \theta)(\hat{\theta} - \theta)^{\top} \right] \right)^{1/d} \right) \\ &\leq \log \left(2\pi e \cdot \frac{1}{d} \mathbb{E} \left[\left\| \hat{\theta} - \theta \right\|_2^2 \right] \right). \end{aligned} \quad (5)$$

The last two inequalities follow from the fact that the Gaussian distribution maximizes entropy for fixed second moments and the AM-GM inequality. Rearranging yields Theorem 2.2.

Through Efroimovich's inequality, one can easily extend van Trees inequality to L_1 loss: following (5), we obtain

$$\begin{aligned} \frac{2}{d} h(\hat{\theta} - \theta) &\leq \frac{2}{d} \sum_i h(\hat{\theta}_i - \theta_i) \leq \frac{2}{d} \sum_i \log \left(2e \mathbb{E} \left[\left| \theta_i - \hat{\theta}_i \right| \right] \right) \\ &\leq 2 \log \left(\frac{2e}{d} \sum_i \mathbb{E} \left[\left| \theta_i - \hat{\theta}_i \right| \right] \right) \\ &= \log \left(\left(\frac{2e}{d} \mathbb{E} \left[\left\| \hat{\theta} - \theta \right\|_1 \right] \right)^2 \right). \end{aligned}$$

The last two inequalities follow from (1) the fact that the Laplace distribution maximizes entropy for a fixed mean absolute error, and (2) the AM-GM inequality. Applying Theorem 2.3, we immediately obtain the following lower bound for L_1 loss:

$$\mathbb{E} \left[\left\| \hat{\theta} - \theta \right\|_1 \right] \geq \sqrt{\frac{\pi}{2e}} \frac{d}{\det(I_X(\theta) + J(\pi))^{1/2d}}. \quad (6)$$

The above lower bound is order-wise optimal. To see this, suppose that there exists an order-optimal estimator $\hat{\theta}^*(X)$ such that

$$\mathbb{E} \left[\left\| \theta - \hat{\theta}^* \right\|_2^2 \right] = C \cdot \frac{d}{\det(I_X(\theta) + J(\pi))^{1/d}}.$$

Then, the Cauchy–Schwartz inequality yields

$$\mathbb{E} \left[\left\| \theta - \hat{\theta}^* \right\|_1 \right] \leq \sqrt{d \cdot \mathbb{E} \left[\left\| \theta - \hat{\theta}^* \right\|_2^2 \right]} = \sqrt{C} \cdot \frac{d}{\det(I_X(\theta) + J(\pi))^{1/2d}}.$$

Since it is well-known that the L_2 lower bound can be achieved, this argument implies the L_1 lower bound is also tight up to a constant. More generally, we can prove the following van Trees inequality for any L_q loss:

Theorem 2.4 (Generalized van Trees Inequality) *Let $q \geq 1$ and let $C_{\text{ME}}(q)$ be the partition function of the L_q max-entropy distribution, formally defined as²*

$$C_{\text{ME}}(q) \triangleq 2e^{\frac{1}{q}} \Gamma\left(\frac{1}{q}\right) q^{\frac{1}{q}-1}.$$

Then, it holds that

$$\mathbb{E}_{\theta, X} \left[\left\| \hat{\theta}(X) - \theta \right\|_q^q \right] \geq \left(\frac{\sqrt{2\pi e}}{C_{\text{ME}}(q)} \right)^q \frac{d}{\det(I_X(\theta) + J(\pi))^{\frac{q}{2d}}}.$$

We can simplify the lower bound with the following AM-GM inequality

$$\det(A)^{1/d} \leq \frac{\text{Tr}(A)}{d} \text{ for any PSD } A \in \mathbb{R}^{d \times d},$$

obtaining the following simpler but slightly weaker form:

Proposition 2.5 *Let $q \geq 1$ and $C_{\text{ME}}(q)$ be defined as in Theorem 2.4. Then, it holds that*

$$\mathbb{E}_{\theta, X} \left[\left\| \hat{\theta}(X) - \theta \right\|_q^q \right] \geq \left(\frac{\sqrt{2\pi e}}{C_{\text{ME}}(q)} \right)^q \cdot \frac{d^{1+\frac{q}{2}}}{\text{Tr}(I_X(\theta) + J(\pi))^{\frac{q}{2}}}.$$

3 Distributed Estimation under Information Constraints

The generalized van Trees inequality (in the form stated in Proposition 2.5), can be combined with an upper bound on the trace of the Fisher information matrix to obtain a lower bound on the minimax L_q loss achievable in a distributed setting under privacy and communication constraints. This approach is rather straightforward and is outlined in Barnes et al. (2019b, 2020, 2019a), where authors also develop the necessary upper bounds on the trace of the Fisher information matrix under privacy and communication constraints. In this section, we overview how privacy and communication constraints are mathematically modeled in a distributed estimation setting and state the corresponding lower bounds under L_q loss.

²One can verify that $C_{\text{ME}}(2) = \sqrt{2\pi e}$ and $C_{\text{ME}}(1) = \sqrt{2e}$.

3.1 Problem Setup

The general distributed statistical task we consider in this paper can be formulated as follows. Each one of the n clients observes a local sample $X_i \sim P_\theta$, processes it via a local channel (i.e., a randomized mapping), and then sends a message $Y_i \in \mathcal{Y}$ to the server, which, upon receiving Y^n , aims to estimate the unknown parameter θ .

At client i , the message Y_i is generated via a *sequential* communication protocol; that is, samples are communicated sequentially by broadcasting the communication to all nodes in the system, including the server. Therefore, the encoding function W_i of the i -th client can depend on all previous messages $Y_1, \dots, Y_{i-1} \in \mathcal{Y}$. Formally, it can be written as a randomized mapping (possibly using shared randomness across participating clients and the server) of the form $Y_i \sim W_i(\cdot | X_i, Y^{i-1})$.

b -bit communication constraint The b -bit communication constraint restricts $|\mathcal{Y}| \leq 2^b$, ensuring the local message can be described in b bits.

Remark 3.1 While we only overview the b -bit sequentially interactive communication model, our results extend to the b -bit blackboard communication protocol Kushilevitz (1997), where each node is allowed to write b -bits in total on a publicly seen blackboard in a randomized order that can depend on the samples. The blackboard model allows for much more interaction between the nodes as compared to the sequential model (e.g. the protocol can start with one of the nodes writing a single bit on the blackboard and the second node to write a bit can depend on the value of the first written bit etc.). The results in Corollary 3.2 trivially apply to these more powerful communication protocols simply because the Fisher information bounds in Barnes et al. (2019b, 2020) are proven under this more general model.

ε -local differential constraint The ε -LDP constraint requires that, for any $i \in [n]$, $x_i, x'_i \in \mathcal{X}$, $y^{i-1} \in \mathcal{Y}^{i-1}$, $\mathcal{S} \in \sigma(Y_i)$, it holds that

$$\frac{W_i(\mathcal{S} | x_i, y^{i-1})}{W_i(\mathcal{S} | x'_i, y^{i-1})} \leq e^\varepsilon. \quad (7)$$

It ensures that any adversary who observes the output Y_i and the context Y^{i-1} cannot infer the local sample X_i .

As a special case, when W_i depends only on X_i and is independent of the other messages Y^{i-1} for all i (i.e. $W_i(\cdot | X_i, Y^{i-1}) = W_i(\cdot | X_i)$), we say the corresponding protocol is *non-interactive*. Finally, we call the tuple $(W^n, \hat{\theta}(Y^n))$ an estimation scheme, where $\hat{\theta}(Y^n)$ is an estimator of θ . We use Π_{seq} and Π_{ind} to denote the collections of all sequentially interactive and non-interactive schemes, respectively. The goal is to lower bound the global minimax L_q risk:

$$R(L_q, \Theta) \triangleq \inf_{(W^n, \hat{\theta}) \in \Pi_{\text{seq}}} \sup_{\theta \in \Theta} \mathbb{E} \left[\left\| \theta - \hat{\theta}(Y^n) \right\|_q^q \right]$$

and construct schemes in Π_{seq} that match these lower bounds.

3.2 Global Minimax Lower Bounds

Combining Proposition 2.5 with the Fisher information bounds in Barnes et al. (2019b), we obtain the following lower bounds on the L_q risk of common statistical models, including the discrete distribution estimation and the Gaussian mean estimation.

Corollary 3.2 (Estimation under b -bit constraint) Let $\kappa(q) \triangleq \left(\frac{\sqrt{2\pi e}}{C_{\text{ME}}(q)} \right)^q$. Then the following lower bounds hold:

- *Gaussian location model:* let $X \sim \mathcal{N}(\theta, \sigma^2 I_d)$ with $[-B, B] \subset \Theta$. For $nB^2 \min(b, d) \geq d\sigma^2$, we have

$$R(L_q, \Theta) \gtrsim d\kappa(q) \max \left\{ \left(\frac{d\sigma^2}{nb} \right)^{\frac{q}{2}}, \left(\frac{\sigma^2}{n} \right)^{\frac{q}{2}} \right\}. \quad (8)$$

- *Gaussian covariance estimation:* Suppose that $X \sim \mathcal{N}(0, \text{diag}(\theta_1, \dots, \theta_d))$ with $[\sigma_{\min}^2, \sigma_{\max}^2] \subset \Theta$. Then for $n(\sigma_{\max}^2 - \sigma_{\min}^2)^2 \min(b^2, d) \geq d\sigma_{\min}^4$, we have

$$R(L_q, \Theta) \gtrsim d\kappa(q) \max \left\{ \left(\frac{d\sigma_{\min}^4}{nb^2} \right)^{\frac{q}{2}}, \left(\frac{\sigma_{\min}^4}{n} \right)^{\frac{q}{2}} \right\}. \quad (9)$$

- *Distribution estimation:* Suppose that $\mathcal{X} = \{1, 2, \dots, d\}$ and that $f(x|\theta) = \theta_x$. Let $\Theta = \Delta_d$ be the d -dim probability simplex. For $n \min(2^b, d) \geq d^2$, we have

$$R(L_q, \Theta) \gtrsim d\kappa(q) \max \left\{ \left(\frac{1}{n2^b} \right)^{\frac{q}{2}}, \left(\frac{1}{nd} \right)^{\frac{q}{2}} \right\}. \quad (10)$$

- *Product Bernoulli model:* suppose that $X \sim \prod_{i=1}^d \text{Bern}(\theta_i)$. If $\Theta = [0, 1]^d$, then for $n \min\{b, d\} \geq d$, we have

$$R(L_q, \Theta) \gtrsim d\kappa(q) \max \left\{ \left(\frac{d}{nb} \right)^{\frac{q}{2}}, \left(\frac{1}{n} \right)^{\frac{q}{2}} \right\}. \quad (11)$$

for some universal constant C . On the other hand, $\Theta = \Delta_d$, then for $n \min\{2^b, d\} \geq d^2$, we get instead

$$R(L_q, \Theta) \gtrsim d\kappa(q) \max \left\{ \left(\frac{1}{n2^b} \right)^{\frac{q}{2}}, \left(\frac{1}{nd} \right)^{\frac{q}{2}} \right\}. \quad (12)$$

Similarly, we can prove lower bounds under the ε -local DP model by using the Fisher information upper bounds in Barnes et al. (2020). We defer this to Corollary A.1 in Appendix A. Corollary 3.2 and Corollary A.1 recover several existing lower bounds and extend them into the broader blackboard communication models. Specifically, (8) and (15) recover the non-sparse setting of in (Acharya et al., 2023, Theorem 4) for $q < \infty$; (10) and (17) recover (Acharya et al., 2023, Corollary 3); (11), (12), (17), and (18) recover (Acharya et al., 2023, Theorem 3) for $q < \infty$.

3.3 Achievability

We note that nearly all of the aforementioned lower bounds are order-wise tight³, meaning that they accurately characterize the correct dependence on parameters such as d , n , b , and ε .

Lemma 3.3 *Assume the same conditions in Corollary 3.2. Then (8), (10), (11), (12) from Corollary 3.2 are tight. The same results hold for Corollary A.1 in Appendix A.*

For $1 \leq q \leq 2$, the upper bounds readily follow from Hölder's inequality. For $q > 2$, the order-optimal estimation schemes can be constructed via a sample splitting trick in a fashion similar to Han et al. (2018a); Barnes et al. (2020), which yields optimal error performance. Notably, these upper bounds can be achieved through independent protocols, indicating that the global minimax error is not improved by interaction among clients.

4 Local Minimax Bounds on Discrete Distribution Estimation

In Section 3, we establish minimax bounds for various distributed statistical estimation models. However, these *global* minimax lower bounds tend to be too conservative and may not accurately reflect the difficulty of estimating each instance $\theta \in \Theta$. This is especially notable in the case of discrete distribution estimation. For example, we would expect the estimation problem to be inherently easier when the underlying distribution

³The only exception is the sub-exponential case (9), in which the tightness of the lower bound (even under L_2 loss) remains open.

we aim to estimate happens to be sparse. Ideally, it is desirable to have estimation schemes that can automatically adapt to the hardness of the underlying instance, i.e. achieve smaller error in easier instances of the problem, without knowing the instance ahead of time. In contrast, globally minimax optimal estimation schemes are typically designed and tuned for worst-case scenarios and can therefore be too pessimistic, i.e. not able to exploit structures that make the problem instance easier.

In this section, we focus on the discrete distribution estimation problem and derive *local* bounds with respect to general L_q loss. We consider the same setting as in Section 3; however to highlight that the parameter of interest is an instance of the d -dim probability simplex, we use $p = (p_1, p_2, \dots, p_d) \in \Delta_d$ to denote the unknown parameter ($\theta \in \Theta$). The goal here is to design a scheme $(W^n, \hat{p}(Y^n))$ to minimize the L_q local statistical risk:

$$R(L_q, p, (W^n, \hat{p})) \triangleq \mathbb{E} \left[\|p - \hat{p}(Y^n)\|_q^q \right] \text{ for all } p \in \Delta_d.$$

We mainly focus on the regime $1 \ll d \ll n$ and aim to characterize the *pointwise* statistical convergence rates when n is sufficiently large.

4.1 b -bit communication constraint

We start with the b -bit communication constraint. Chen et al. (2021b) provides a two-round scheme that achieves the following estimation error:

Theorem 4.1 (L_2 local minimax error Chen et al. (2021b)) *Let $b \leq \lfloor \log_2 d \rfloor$.*

1. *There exists a sequentially interactive scheme (with a single round of interaction) $(W^n, \check{p}) \in \Pi_{\text{seq}}$, such that for all $p \in \Delta_d$,*

$$R(L_2, p, (W^n, \check{p})) \asymp \frac{\|p\|_{\frac{1}{2}} + o_n(1)}{n2^b}. \quad (13)$$

2. *There exists a sequentially interactive scheme (with a single round of interaction) $(\tilde{W}^n, \check{p}) \in \Pi_{\text{seq}}$, such that for all $p \in \Delta_d$,*

$$R(L_1, p, (\tilde{W}^n, \check{p})) \lesssim \sqrt{\frac{\|p\|_{\frac{1}{3}} + o_n(1)}{n2^b}}. \quad (14)$$

Note that the upper bound on the risk is now a function of the unknown p . Chen et al. (2021a) develops the matching lower bound in (13) for L_2 loss. The optimality of the L_1 rate in (14) has remained open. In this paper, we close this gap by proving the following matching lower bound for general L_q loss.

Theorem 4.2 (b -bit L_q local minimax lower bounds) *Let $p \in \Delta'_d \triangleq \{p \in \Delta_d \mid \frac{1}{2} < p_1 < \frac{2}{3}\}$. Then for any $\delta > 0$, $B \geq \sqrt{\frac{\|p\|_{1/2}}{d2^b}}$, as long as $n = \Omega\left(\frac{d^3 \log d}{\|p\|_{1/2}}\right)$, it holds that⁴*

$$\begin{aligned} & \inf_{(W^n, \check{p}) \in \Pi_{\text{seq}}} \sup_{p': \|p' - p\|_\infty \leq \frac{B}{\sqrt{n}}} \mathbb{E}_{p'} \left[\|\hat{p}(W^n(X^n)) - p'\|_q^q \right] \\ & \gtrsim \max \left(\frac{C_\delta \|p\|_{\frac{q}{q+2} + \delta}}{(n2^b)^{\frac{q}{2}}}, \frac{\|p\|_{\frac{q}{q+2}}}{(n2^b)^{\frac{q}{2}} \log d}, \frac{C_\delta \|p\|_{\frac{q/2 + \delta}{q/2 + \delta}}}{n^{\frac{q}{2}}}, \frac{\|p\|_{\frac{q/2}{q/2}}}{n^{\frac{q}{2}} \log d} \right), \end{aligned}$$

where $C_\delta \triangleq (\delta/(1+\delta))^2$ is a δ -dependent constant.

⁴Indeed, the lower bound holds for blackboard interactive schemes Kushilevitz (1997), a more general class of interactive schemes than Π_{seq} . See Barnes et al. (2019b) for a discussion of blackboard schemes.

Sketch of the proof The proof is based on the framework introduced in Barnes et al. (2019b) (see also Chen et al. (2021b) for the analysis of ℓ_2 case), where a global upper bound on the quantized Fisher information is given and used to derive the minimax lower bound on the ℓ_2 error. We extend their results to the local regime and develop a *local* upper bound on the quantized Fisher information around a neighborhood of p .

To obtain a local upper bound, we construct an h -dimensional parametric sub-model Θ_p^h that contains p and is a subset of $\mathcal{P}_d$, where $h \in [d]$ is a tuning parameter and will be determined later. By considering the sub-model Θ_p^h , we can control its Fisher information around p with a function of h and p . Optimizing over $h \in [d]$ yield an upper bound that depends on $\|p\|_{\frac{q}{q+2}}$. Finally, the local upper bound on the quantized Fisher information is translated to a local minimax lower bound on the L_q error via the generalized van Trees inequality (Theorem 2.4). We defer the detailed proof to Appendix B.2. $\square$

Remark 4.3 Note that Theorem 4.2 complements Theorem 4.1 and is nearly tight (up to a $\log d$ factor) for $q = 1$ and $q = 2$. We believe the upper bound can be easily adapted to match the lower bound for all $q > 1$.

4.2 ε -local differential privacy

Similarly, under the ε -local DP constraint, we prove the following local lower bound:

Theorem 4.4 (ε -LDP L_q local minimax lower bounds) Let $p \in \Delta'_d \triangleq \{p \in \Delta_d | \frac{1}{2} < p_1 < \frac{2}{3}\}$. Then for any $\delta > 0, B \geq \sqrt{\frac{\|p\|_{1/2}}{d \min(e^\varepsilon, (e^\varepsilon - 1)^2)}}$, as long as $n = \Omega\left(\frac{d^3 \log d}{\|p\|_{1/2}}\right)$, it holds that

$$\begin{aligned} & \inf_{(W^n, \hat{p}) \in \Pi_{\text{seq}} p': \|p' - p\|_\infty \leq \frac{B}{\sqrt{n}}} \sup \mathbb{E}_{p'} \left[\|\hat{p}(W^n(X^n)) - p'\|_q^q \right] \\ & \gtrsim \max \left(\frac{C_\delta \|p\|_{\frac{q}{q+2} + \delta}}{\left(n \min(e^\varepsilon, (e^\varepsilon - 1)^2)\right)^{\frac{q}{2}}}, \right. \\ & \quad \left. \frac{C_1 \|p\|_{\frac{q}{q+2}}}{\left(n \min(e^\varepsilon, (e^\varepsilon - 1)^2)\right)^{\frac{q}{2}} \log d}, \frac{C_\delta \cdot \|p\|_{\frac{q}{2} + \delta}^{q/2 + \delta}}{n^{\frac{q}{2}}}, \frac{C_2 \|p\|_{\frac{q}{2}}^{q/2}}{n^{\frac{q}{2}} \log d} \right), \end{aligned}$$

for some $C_\delta, C_1, C_2 > 0$.

However, unlike the b -bit communication constraint, the tightness of Theorem 4.4 is unknown, and a matching achievability scheme currently remains unsettled.

5 Conclusion and Future Works

In this work, we first point out that the popular van Trees inequality can be generalized to accommodate general loss functions through Efronovich's inequality. Subsequently, we extend the application of the generalized van Trees inequality to L_q loss in the context of distributed estimation under communication and local differential privacy constraints. Notably, combining with previous bounds on constrained Fisher information, our results offer a significantly simplified analysis over existing (global) lower bounds and, more importantly, can be used to derive local minimax results that capture the hardness of instances.

Acknowledgement

We thank the anonymous Reviewer 2 for pointing out a critical error in the proof of Theorem 4.2.

References

- Acharya, J., Canonne, C. L., Sun, Z., and Tyagi, H. (2023). Unified lower bounds for interactive high-dimensional estimation under information constraints. *Advances in Neural Information Processing Systems*, 36.
- Acharya, J., Canonne, C. L., and Tyagi, H. (2020a). Inference under information constraints i: Lower bounds from chi-square contraction. *IEEE Transactions on Information Theory*, 66(12):7835–7855.
- Acharya, J., Canonne, C. L., and Tyagi, H. (2020b). Inference under information constraints ii: Communication constraints and shared randomness. *IEEE Transactions on Information Theory*, 66(12):7856–7877.
- Acharya, J., Sun, Z., and Zhang, H. (2019). Hadamard response: Estimating distributions privately, efficiently, and with little communication. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1120–1129.
- Agarwal, N., Suresh, A. T., Yu, F. X. X., Kumar, S., and McMahan, B. (2018). cpsgd: Communication-efficient and differentially-private distributed sgd. In *Advances in Neural Information Processing Systems*, pages 7564–7575.
- Aras, E., Lee, K.-Y., Pananjady, A., and Courtade, T. A. (2019). A family of bayesian cramér-rao bounds, and consequences for log-concave priors. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2699–2703. IEEE.
- Asi, H., Feldman, V., Nelson, J., Nguyen, H., and Talwar, K. (2023). Fast optimal locally private mean estimation via random projections. In *Neural Information Processing Systems*, volume 37.
- Asi, H., Feldman, V., and Talwar, K. (2022). Optimal algorithms for mean estimation under local differential privacy. In *International Conference on Machine Learning*, pages 1046–1056. PMLR.
- Asoodeh, S. and Zhang, H. (2022). Contraction of locally differentially private mechanisms. *arXiv preprint arXiv:2210.13386*.
- Barnes, L. P., Chen, W.-N., and Özgür, A. (2020). Fisher information under local differential privacy. *IEEE Journal on Selected Areas in Information Theory*, 1(3):645–659.
- Barnes, L. P., Han, Y., and Özgür, A. (2019a). Fisher information for distributed estimation under a blackboard communication protocol. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 2704–2708. IEEE.
- Barnes, L. P., Han, Y., and Özgür, A. (2019b). Lower bounds for learning distributions under communication constraints via fisher information.
- Borovkov, A. A. (1998). *Mathematical Statistics*. Gordon and Breach Science Publishers.
- Braverman, M., Garg, A., Ma, T., Nguyen, H. L., and Woodruff, D. P. (2016). Communication lower bounds for statistical estimation problems via a distributed data processing inequality. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 1011–1020.
- Chen, W.-N., Kairouz, P., and Özgür, A. (2020). Breaking the communication-privacy-accuracy trilemma. *Advances in Neural Information Processing Systems*, 33.
- Chen, W.-N., Kairouz, P., and Ozgur, A. (2021a). Breaking the dimension dependence in sparse distribution estimation under communication constraints. In *Conference on Learning Theory*, pages 1028–1059. PMLR.
- Chen, W.-N., Kairouz, P., and Özgür, A. (2021b). Pointwise bounds for distribution estimation under communication constraints. *Advances in Neural Information Processing Systems*, 34.
- Cramér, H. (1999). *Mathematical methods of statistics*, volume 43. Princeton university press.

- Duchi, J. C., Jordan, M. I., and Wainwright, M. J. (2013). Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE.
- Duchi, J. C. and Rogers, R. (2019). Lower bounds for locally private estimation via communication complexity. In *Conference on Learning Theory*, pages 1161–1191. PMLR.
- Duchi, J. C. and Ruan, F. (2018). The right complexity measure in locally private estimation: It is not the fisher information. *arXiv preprint arXiv:1806.05756*.
- Efremovich, S. Y. (1980). Information contained in a sequence of observations. *Problems in Information Transmission*, 15:24–39.
- Feldman, V., Nelson, J., Nguyen, H., and Talwar, K. (2022). Private frequency estimation via projective geometry. In *International Conference on Machine Learning*, pages 6418–6433. PMLR.
- Feldman, V. and Talwar, K. (2021). Lossless compression of efficient private local randomizers. In *International Conference on Machine Learning*, pages 3208–3219. PMLR.
- Gandikota, V., Kane, D., Maity, R. K., and Mazumdar, A. (2021). vqsgd: Vector quantized stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pages 2197–2205. PMLR.
- Garg, A., Ma, T., and Nguyen, H. (2014). On communication cost of distributed statistical estimation and dimensionality. In *Advances in Neural Information Processing Systems*, pages 2726–2734.
- Gill, R. D. and Levit, B. Y. (1995). Applications of the van trees inequality: a bayesian cramér-rao bound. *Bernoulli*, 1(1-2):59–79.
- Hájek, J. (1961). Some extensions of the wald-wolfowitz-noether theorem. *The Annals of Mathematical Statistics*, pages 506–523.
- Han, Y., Mukherjee, P., Özgür, A., and Weissman, T. (2018a). Distributed statistical estimation of high-dimensional and nonparametric distributions. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 506–510. IEEE.
- Han, Y., Özgür, A., and Weissman, T. (2018b). Geometric lower bounds for distributed parameter estimation under communication constraints. In *Conference On Learning Theory*, pages 3163–3188. PMLR.
- Isik, B., Chen, W.-N., Özgür, A., Weissman, T., and No, A. (2023). Exact optimality of communication-privacy-utility tradeoffs in distributed mean estimation. *Advances in Neural Information Processing Systems*, 37.
- Kairouz, P., Bonawitz, K., and Ramage, D. (2016). Discrete distribution estimation under local privacy. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 2436–2444, New York, New York, USA.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2019). Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*.
- Kushilevitz, E. (1997). Communication complexity. In *Advances in Computers*, volume 44, pages 331–360. Elsevier.
- Lee, K.-Y. (2022). *New Information Inequalities with Applications to Statistics*. PhD thesis, UC Berkeley.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. (2017a). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR.
- McMahan, H. B., Ramage, D., Talwar, K., and Zhang, L. (2017b). Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*.

- Rao, C. R. (1992). Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in Statistics: Foundations and basic theory*, pages 235–247. Springer.
- Shah, A., Chen, W.-N., Balle, J., Kairouz, P., and Theis, L. (2022). Optimal compression of locally differentially private mechanisms. In *International Conference on Artificial Intelligence and Statistics*, pages 7680–7723. PMLR.
- Suresh, A. T., Yu, F. X., Kumar, S., and McMahan, H. B. (2017). Distributed mean estimation with limited communication. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, page 3329–3337. JMLR.org.
- Tsybakov, A. (2008). *Introduction to Nonparametric Estimation*. Springer-Verlag.
- Tsybakov, A. B. (2004). Introduction to nonparametric estimation, 2009. URL <https://doi.org/10.1007/b13794>. Revised and extended from the, 9(10).
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Van Trees, H. L. (2004). *Detection, estimation, and modulation theory, part I: detection, estimation, and linear modulation theory*. John Wiley & Sons.
- Wang, T., Blocki, J., Li, N., and Jha, S. (2017). Locally differentially private protocols for frequency estimation. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, pages 729–745.
- Ye, M. and Barg, A. (2017). Optimal schemes for discrete distribution estimation under local differential privacy. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 759–763.
- Zhang, Y., Duchi, J. C., Jordan, M. I., and Wainwright, M. J. (2013). Information-theoretic lower bounds for distributed statistical estimation with communication constraints. In *Advances in Neural Information Processing Systems*, pages 2328–2336.
- Zhang, Y., Wainwright, M. J., and Duchi, J. C. (2012). Communication-efficient algorithms for statistical optimization. *Advances in neural information processing systems*, 25.
- Úlfar Erlingsson, Pihur, V., and Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 21st ACM Conference on Computer and Communications Security*, Scottsdale, Arizona.

A Global Minimax Lower Bounds under LDP

Subject to ε -LDP constraints, the following lower bounds on the risk holds:

Corollary A.1 (Distributed estimation under ε -bit local DP) *Making the same assumptions as in Corollary 3.2*

- *Gaussian location model: let $X \sim \mathcal{N}(\theta, \sigma^2 I_d)$ with $[-B, B] \subset \Theta$. For $nB^2 \min(\varepsilon, \varepsilon^2, d) \geq d\sigma^2$, we have*

$$\mathbf{R}(L_q, \Theta) \gtrsim d\kappa(q) \left(\frac{d}{n \min(\varepsilon, \varepsilon^2)} \right)^{\frac{q}{2}}. \quad (15)$$

- *Distribution estimation: Suppose that $\mathcal{X} = \{1, 2, \dots, d\}$ and that $f(x|\theta) = \theta_x$. Let $\Theta = \Delta_d$ be the d -dim probability simplex. For $n \min(e^\varepsilon, (e^\varepsilon - 1)^2, d) \geq d^2$, we have*

$$\mathbf{R}(L_q, \Theta) \gtrsim d\kappa(q) \left(\frac{1}{n \min(e^\varepsilon, (e^\varepsilon - 1)^2, d)} \right)^{\frac{q}{2}}. \quad (16)$$

- *Product Bernoulli model: suppose that $X \sim \prod_{i=1}^d \text{Bern}(\theta_i)$. If $\Theta = [0, 1]^d$, then for $n \min\{b, d\} \geq d$, we have*

$$\mathbf{R}(L_q, \Theta) \gtrsim d\kappa(q) \left(\frac{d}{n(\varepsilon, \varepsilon^2, d)} \right)^{\frac{q}{2}}. \quad (17)$$

On the other hand, $\Theta = \Delta_d$, then for $n \min\{2^b, d\} \geq d^2$, we get instead

$$\mathbf{R}(L_q, \Theta) \gtrsim d\kappa(q) \left(\frac{1}{n \min(e^\varepsilon, (e^\varepsilon - 1)^2, d)} \right)^{\frac{q}{2}}. \quad (18)$$

B Missing Proofs

B.1 Proof of Theorem 2.4

We start with the following maximum-entropy lemma.

Lemma B.1 (Maximum entropy under p -th moment constraints.) *For any random variable $X \in \mathbb{R}$, it holds that*

$$h(X) \leq \log \left(2e^{\frac{1}{q}} \Gamma \left(\frac{1}{q} \right) q^{\frac{1}{q}-1} (\mathbb{E}[|X|^q])^{\frac{1}{q}} \right) \triangleq \log \left(C_{\text{ME}}(q) \cdot \mathbb{E}[|X|^q]^{\frac{1}{q}} \right).$$

Now, we proceed with (5):

$$\begin{aligned} \frac{2}{d} h(\hat{\theta} - \theta) &\leq \frac{2}{d} \sum_i h(\hat{\theta}_i - \theta_i) \leq \frac{2}{d} \sum_i \log \left(C_{\text{ME}}(q) \left(\mathbb{E}[|\hat{\theta}_i - \theta_i|^q] \right)^{\frac{1}{q}} \right) \\ &\leq 2 \log \left(C_{\text{ME}}(q) \left(\frac{1}{d} \sum_i \mathbb{E}[|\hat{\theta}_i - \theta_i|^q] \right)^{\frac{1}{q}} \right) \\ &= \log \left(\left(C_{\text{ME}}(q) \left(\frac{1}{d} \mathbb{E}[\|\hat{\theta} - \theta\|_q^q] \right)^{\frac{1}{q}} \right)^2 \right). \end{aligned}$$

Plugging into Theorem 2.3 and rearranging yield the desired result. $\square$

B.2 Proof of Theorem 4.2

Before entering the main proof, we first introduce some notation that will be used throughout this section. Let $(p_{(1)}, p_{(2)}, \dots, p_{(d)})$ be the sorted sequence of $p = (p_1, p_2, \dots, p_d)$ in the non-increasing order; that is, $p_{(i)} \geq p_{(j)}$ for all $i > j$. Denote $\pi : [d] \rightarrow [d]$ as the corresponding sorting function⁵, i.e. $p_{(i)} = p_{\pi(i)}$ for all $i \in [d]$.

Constructing the sub-model Θ_p^h . We construct Θ_p^h by freezing the $(d-h)$ smallest coordinates of p and only letting the largest $(h-1)$ coordinates to be free parameters. Mathematically, let

$$\Theta_p^h \triangleq \{(\theta_2, \theta_3, \dots, \theta_h) \mid \pi^{-1}(\theta_1, \theta_2, \dots, \theta_h, p_{(h+1)}, p_{(h+2)}, \dots, p_{(d)}) \in \mathcal{P}_d\}, \quad (19)$$

where $\theta_1 = 1 - \sum_{i=2}^h \theta_i - \sum_{i=h+1}^d p_i$ is fixed when $(\theta_2, \dots, \theta_h)$ are determined. For instance, if $p = (\frac{1}{16}, \frac{1}{8}, \frac{1}{2}, \frac{1}{16}, \frac{1}{4})$ (so $d = 5$) and $h = 3$, then the corresponding sub-model is

$$\Theta_p^h \triangleq \left\{(\theta_2, \theta_3) \mid \left(\frac{1}{16}, \theta_3, \theta_1, \frac{1}{16}, \theta_2\right) \in \mathcal{P}_d\right\}.$$

Bounding the quantized Fisher information. Now recall that under this model, the score function

$$S_\theta(x) \triangleq (S_{\theta_2}(x), \dots, S_{\theta_h}(x)) \triangleq \left(\frac{\partial \log p(x|\theta)}{\partial \theta_2}, \dots, \frac{\partial \log p(x|\theta)}{\partial \theta_h}\right)$$

can be computed as

$$S_{\theta_i}(x) = \begin{cases} \frac{1}{\theta_i}, & \text{if } x = \pi(i), 2 \leq i \leq h \\ -\frac{1}{\theta_1}, & \text{if } x = \pi(1) \\ 0, & \text{otherwise} \end{cases}$$

The next lemma shows that to bound the quantized Fisher information, it suffices to control the variance of the score function.

Lemma B.2 (Theorem 1 in Barnes et al. (2019b)) *Let W be any b -bit quantization scheme and $I_W(\theta)$ is the Fisher information of Y at θ where $Y \sim W(\cdot|X)$ and $X \sim p_\theta$. Then for any $\theta \in \Theta \subseteq \mathbb{R}^h$,*

$$\text{Tr}(I_W(\theta)) \leq \min \left(\text{Tr}(I_X(\theta)), 2^b \max_{\|u\|_2 \leq 1} \text{Var}(\langle u, S_\theta(X) \rangle) \right).$$

Therefore, for any unit vector $u = (u_2, \dots, u_h)$ with $\|u\|_2 = 1$, we control the variance as follows:

$$\begin{aligned} \text{Var}(\langle u, S_\theta(X) \rangle) &\stackrel{(a)}{=} \sum_{i=1}^h \theta_i \left(\sum_{j=2}^h u_j S_{\theta_j}(\pi(i)) \right)^2 \\ &= \theta_1 \left(\sum_{j=2}^h u_j \right)^2 \left(\frac{1}{\theta_1} \right)^2 + \sum_{i=2}^h \theta_i u_i^2 \frac{1}{\theta_i^2} \\ &= \frac{\left(\sum_{j=2}^h u_j \right)^2}{\theta_1} + \sum_{j=2}^h \frac{u_j^2}{\theta_j} \\ &\leq \frac{h}{\theta_1} + \frac{1}{\min_{j \in \{2, \dots, h\}} \theta_j}, \end{aligned} \quad (20)$$

where (a) holds since the score function has zero mean. This allows us to upper bound $I_W(\theta)$ in a neighborhood around $\theta(p)$, where $\theta(p)$ is the location of p in the sub-model Θ_p^h , i.e.

$$\theta(p) = (\theta_2(p), \dots, \theta_h(p)) \triangleq (p_{(2)}, \dots, p_{(h)}).$$

⁵With a slight abuse of notation, we overload π so that $\pi((p_1, \dots, p_d)) \triangleq (p_{\pi(1)}, \dots, p_{\pi(d)})$

In particular, for any $0 < B \leq \frac{p(h)}{3}$ and $p \in \mathcal{P}'_d \triangleq \{p \in \mathcal{P}_d \mid \frac{1}{2} < p_1 < \frac{5}{6}\}$, the neighborhood $\mathcal{N}_{B,h}(p) \triangleq \theta(p) + [-B, B]^h$ must be contained in Θ_p^h . Next, we control the quantized Fisher information over $\mathcal{N}_{B,h}(p)$ as follows.

1. For any $\theta' \in \mathcal{N}_{B,h}(p)$, it holds that

$$\theta'_1 \geq \theta_1(p) - \frac{hp(h)}{3} \geq \frac{1}{6},$$

where the second inequality holds since 1) $\theta_1(p) = p_{(1)} \geq \frac{1}{2}$ by our definition of $\mathcal{P}'_d$, and 2) $\frac{hp(h)}{3} \leq \frac{\sum_{i=1}^h p_{(i)}}{3} \leq \frac{1}{3}$. We also have

$$\min_{j \in \{2, \dots, h\}} \theta'_j \geq \min_{j \in \{2, \dots, h\}} \theta_j(p) - \frac{p(h)}{3} \geq \frac{2p(h)}{3}.$$

Therefore (20) implies for any $\theta' \in \mathcal{N}_{B,h}(p)$,

$$\text{Var}(\langle u, S'_{\theta}(X) \rangle) \leq 6h + \frac{3}{2p(h)}. \quad (21)$$

2. By the definition of $S_{\theta_i}(x)$ and Fisher information matrix,

$$\text{Tr}(I_X(\theta)) = \sum_x \sum_{i \in [h]} S_{\theta_i}(x) S_{\theta_i}(x) p_{\theta}(x) = \sum_{i=1}^h \frac{1}{p_{(i)}}.$$

Maximizing $\text{Tr}(I_X(\theta))$ yields

$$\max_{\theta' \in \mathcal{N}_{B,h}(p)} \text{Tr}(I_X(\theta)) = \sum_{i=1}^h \frac{1}{p_{(i)} - \frac{1}{3}p_{(h)}} \leq \frac{3}{2} \sum_i \frac{1}{p_{(i)}} \leq \frac{3}{2} \frac{h}{p(h)}. \quad (22)$$

Combining (21) and (22) with Lemma B.2, we arrive at

$$\forall \theta' \in \mathcal{N}_{B,h}(p), \quad \text{Tr}(I_W(\theta')) \leq \min \left(2^b \left(6h + \frac{3}{2p(h)} \right), \frac{3h}{2p(h)} \right) \leq 2^b \left(6h + \frac{3}{2p(h)} \right) + \frac{3h}{2p(h)}.$$

Bounding the L_q error. We choose the prior μ on $\mathcal{N}_{B,h}(p)$ to minimize the Fisher information, i.e., $\mu = \mu^{\otimes h}$ where μ is the cosine prior introduced in Borovkov (1998); Tsybakov (2008). It is straightforward to see that $\text{Tr}(I(\pi) = \frac{h\pi^2}{B^2})$. Then, applying Proposition 2.5 on $\mathcal{N}_{B,h}(p)$ yields

$$\begin{aligned} \sup_{\theta' \in \mathcal{N}_{B,h}(p)} \mathbb{E} \left[\left\| \hat{\theta} - \theta' \right\|_q^q \right] &\geq \frac{h^{1+\frac{q}{2}}}{\left(n2^b \left(6h + \frac{3}{2p(h)} \right) + \frac{h\pi^2}{B^2} + \frac{3nh}{2p(h)} \right)^{q/2}} \\ &\geq \frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{\left(n2^b \left(6hp(h) + \frac{3}{2} \right) + \frac{h\pi^2 p_{(h)}}{B^2} + \frac{3nh}{2} \right)^{q/2}} \\ &\stackrel{(a)}{\geq} \frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{\left(10n2^b + \frac{10hp(h)}{B^2} + 4nh \right)^{q/2}}. \end{aligned} \quad (23)$$

where the (a) is due to $hp(h) \leq 1$. Now, if we pick $B \geq \sqrt{\frac{hp(h)}{n2^b}}$, then (23) becomes

$$(23) \geq \frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{(20n2^b + 4nh)^{q/2}} \geq \min \left(\frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{(40n2^b)^{q/2}}, \frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{(8nh)^{q/2}} \right) = \min \left(\frac{h^{1+\frac{q}{2}} p_{(h)}^{\frac{q}{2}}}{(40n2^b)^{q/2}}, \frac{hp_{(h)}^{\frac{q}{2}}}{(8n)^{q/2}} \right). \quad (24)$$

Notice that in order to satisfy the condition $\mathcal{N}_{B,h}(p) \subseteq \Theta_p^h$, B must be at most $\frac{p(h)}{3}$, so we have an implicit sample size requirement: n must be at least $\frac{3h}{2^b p(h)}$.

Optimizing over h . Finally, we maximize (24) over $h \in [d]$ to obtain the best lower bound, via the following simple but crucial lemma.

Lemma B.3 *For any $p \in \mathcal{P}_d$ and $\delta > 0$, it holds that*

- $$\max_{h \in [d]} h^{1+\frac{q}{2}} \cdot p_{(h)}^{\frac{q}{2}} \geq \max \left(C_\delta \|p\|_{\frac{q}{q+2}+\delta}, \frac{C \|p\|_{\frac{q}{q+2}}}{\log d} \right),$$

for $C_\delta \triangleq \left(\frac{\delta}{1+\delta} \right)^{\frac{2}{1+\delta}}$ and a universal constant C small enough.

- $$\max_{h \in [d]} h \cdot p_{(h)}^{q/2} \geq \max \left(C_\delta \cdot \|p\|_{q/2+\delta}, \frac{C \|p\|_{q/2}^{q/2}}{\log d} \right),$$

for $C_\delta \triangleq \left(\frac{\delta}{1+\delta} \right)^{\frac{2}{1+\delta}}$ and a universal constant C small enough.

Picking $h^* = \operatorname{argmax}_{h \in [d]} h^2 p_{(h)}$ and by Lemma B.3 and (24), we obtain that for all $\hat{\theta}$

$$\sup_{\theta' \in \mathcal{N}_{B_n, h^*}(p)} \mathbb{E} \left[\left\| \hat{\theta} - \theta' \right\|_q^q \right] \geq \max \left(\frac{C_\delta \|p\|_{\frac{q}{q+2}+\delta}}{(n2^b)^{\frac{q}{2}}}, \frac{C_1 \|p\|_{\frac{q}{q+2}}}{(n2^b)^{\frac{q}{2}} \log d}, \frac{C_\delta \cdot \|p\|_{q/2+\delta}^{q/2}}{n^{\frac{q}{2}}}, \frac{C_2 \|p\|_{q/2}^{q/2}}{n^{\frac{q}{2}} \log d} \right),$$

as long as $p \in \mathcal{P}'_d$ and $B_n = \sqrt{\frac{h^* p_{(h^*)}}{n2^b}} \stackrel{(a)}{\leq} \sqrt{\frac{d}{n2^b \|p\|_{\frac{1}{2}}}}$, where (a) holds due to the second result of Lemma B.3 and $h^* \leq d$. In addition, the sample size constraint that n must be larger than $\frac{3h^*}{2^b p_{(h^*)}}$ can be satisfied if $n = \Omega \left(\frac{d^3 \log d}{2^b \|p\|_{\frac{1}{2}}} \right)$ since $\frac{h^*}{p_{(h^*)}} \leq \frac{(h^*)^3 \log d}{C \|p\|_{\frac{1}{2}}} \leq \frac{d^3 \log d}{C \|p\|_{\frac{1}{2}}}$, where the first inequality is due to Lemma B.3 and the second one is due to $h^* \leq d$. The proof is complete by observing that

$$\inf_{(W^n, \hat{p})} \sup_{p': \|p' - p\|_\infty \leq B_n} \mathbb{E} \left[\left\| \hat{p} - p' \right\|_q^q \right] \geq \inf_{(W^n, \hat{\theta})} \sup_{\theta' \in \mathcal{N}_{B_n, h^*}(p)} \mathbb{E} \left[\left\| \hat{\theta} - \theta' \right\|_q^q \right].$$

□

B.3 Proof of Corollary 3.2 and Corollary A.1

The corollaries follow from a direct application of Theorem 2.4, along with (Barnes et al., 2019b, Corollary 1-4) and (Barnes et al., 2020, Corollary 1-3).

B.4 Proof of Theorem 4.4

The proof follows exactly the same as the Section 4.2, except for replacing Lemma B.2 with the following:

Lemma B.4 (Proposition 2 and 3 and in Barnes et al. (2020)) *Let W be any ε -local DP scheme and $I_W(\theta)$ is the Fisher information of Y at θ where $Y \sim W(\cdot|X)$ and $X \sim p_\theta$. Then for any $\theta \in \Theta \subseteq \mathbb{R}^h$,*

$$\operatorname{Tr}(I_W(\theta)) \leq \min \left\{ \operatorname{Tr}(I_X(\theta)), e^\varepsilon \max_{\|u\|_2 \leq 1} \operatorname{Var}(\langle u, S_\theta(X) \rangle), (e^\varepsilon - 1)^2 \max_{\|u\|_2 \leq 1} \operatorname{Var}(\langle u, S_\theta(X) \rangle) \right\}.$$

C Proof of Lemmas

C.1 Proof of Lemma B.1

Let X satisfy the q -th moment constraint $\mathbb{E}[|X|^q] \leq \Delta$. Then, by the maximum entropy principle, the density of X that maximizes the entropy must take the following form:

$$f(x; \beta) = \frac{\exp(-|x|^q/\beta)}{\lambda(q, \beta)},$$

where

$$\lambda(q, \beta) \triangleq \int_{\mathbb{R}} \exp(-|x|^q/\beta) dx = \frac{2\beta^{1/q}}{q} \Gamma\left(\frac{1}{q}\right).$$

Simple calculations yield

$$\mathbb{E}_{X \sim f(x; \beta)}[|X|^q] = \frac{2}{\lambda(q, \beta)} \int_0^\infty x^q \exp(-x^q/\beta) dx = \frac{2 \frac{\beta^{1/q+1}}{q} \Gamma\left(1 + \frac{1}{q}\right)}{\lambda(q, \beta)} = \frac{\beta \Gamma\left(1 + \frac{1}{q}\right)}{\Gamma\left(\frac{1}{q}\right)} = \frac{\beta}{q} \leq \Delta.$$

On the other hand, the differential entropy can be calculated as

$$h(X) = \frac{1}{q} - \log\left(\frac{q}{2\beta^{1/q}\gamma\left(\frac{1}{q}\right)}\right) \leq \frac{1}{q} - \log\left(\frac{p^{1-1/q}}{2\Delta^{1/q}\Gamma\left(\frac{1}{q}\right)}\right),$$

which establishes the lemma. □

C.2 Proof of Lemma B.3

The proof follows from the same arguments as in (Chen et al., 2021b, Lemma 6.2) and is omitted here. □