

Efficient Estimation of a Gaussian Mean with Local Differential Privacy

Nikita Kalinin

Institute of Science and Technology Austria (ISTA)

NIKITA.KALININ@IST.AC.AT

Lukas Steinberger

University of Vienna

LUKAS.STEINBERGER@UNIVIE.AC.AT

Abstract

In this paper we study the problem of estimating the unknown mean θ of a unit variance Gaussian distribution in a locally differentially private (LDP) way. In the high-privacy regime ($\epsilon \leq 0.67$), we identify the exact optimal privacy mechanism that minimizes the variance of the estimator asymptotically. It turns out to be the extraordinarily simple sign mechanism that applies randomized response to the sign of $X_i - \theta$. However, since this optimal mechanism depends on the unknown mean θ , we employ a two-stage LDP parameter estimation procedure which requires splitting agents into two groups. The first n_1 observations are used to consistently but not necessarily efficiently estimate the parameter θ by $\tilde{\theta}_{n_1}$. Then this estimate is updated by applying the sign mechanism with $\tilde{\theta}_{n_1}$ instead of θ to the remaining $n - n_1$ observations, to obtain an LDP and efficient estimator of the unknown mean.

Keywords: local differential privacy, statistical efficiency, Gaussian mean estimation

1. Introduction

We consider the problem of estimating the unknown mean of a unit variance normal distribution in a locally differentially private and statistically efficient way. In this scenario, we have n agents or data-holders, each of which owns data X_i sampled independently from the probability distribution $P_\theta = N(\theta, 1)$ which depends on the unknown mean parameter $\theta \in \mathbb{R}$. Denote the class of all those potential data generating distributions by $\mathcal{P} = \{P_\theta : \theta \in \mathbb{R}\}$. We protect the privacy of the data owners by local differential privacy. This means that each agent generates a sanitized version Z_i of their original data X_i independently of everybody else by applying a privacy mechanism Q , which is a Markov kernel or a conditional distribution of Z_i given $X_i = x$. In other words, $Q(A|x) = P(Z_i \in A | X_i = x)$. From the privacy mechanism Q we can generate random variates in some arbitrary space \mathcal{Z} endowed with a sigma algebra \mathcal{G} . We say that Q satisfies the ϵ -local differential privacy property if

$$Q(A|x) \leq e^\epsilon Q(A|x')$$

for all $x, x' \in \mathbb{R}$ and all events $A \in \mathcal{G}$. We denote by $\mathcal{Q}_\epsilon = \mathcal{Q}_\epsilon(\mathbb{R})$ the set of all possible such privacy mechanisms taking inputs from \mathbb{R} .

After each agent applied the privacy mechanism Q we obtain iid samples Z_1, \dots, Z_n from the distribution QP_θ , which formally is the probability measure obtained by integrating Q with respect to $P_\theta(dx)$. We can only use these samples to estimate the unknown parameter θ and our goal is to do this with the smallest amount of estimation variance possible, that is, to estimate θ in a statistically efficient way. Standard asymptotic theory (cf. [van der Vaart, 2007](#), Chapter 8) shows

that a maximum likelihood estimator $\hat{\theta}_n^{(mle)}$ based on Z_1, \dots, Z_n achieves an asymptotic normal distribution

$$\sqrt{n}(\hat{\theta}_n^{(mle)} - \theta) \xrightarrow[n \rightarrow \infty]{d} N(0, I_\theta(Q\mathcal{P})^{-1})$$

where the inverse of the Fisher-Information

$$I_\theta(Q\mathcal{P}) := \mathbb{E}_{Z \sim QP_\theta} \left[\left(\frac{\partial \log QP_\theta(Z)}{\partial \theta} \right)^2 \right]$$

is the smallest variance achievable by regular estimators (van der Vaart, 2007, Theorem 8.8). Thus, to minimize asymptotic variance, we should solve the optimization problem

$$\max_Q I_\theta(Q\mathcal{P}) \quad \text{subject to} \quad Q \in \mathcal{Q}_\epsilon. \quad (1)$$

This is a very challenging optimization problem because the set of Markov kernels \mathcal{Q}_ϵ is a huge, non-parametric set and we have to even search over all possible output spaces $(\mathcal{Z}, \mathcal{G})$ from which Q can draw Z_i . Furthermore, it turns out that for fixed output space $(\mathcal{Z}, \mathcal{G})$ the mapping $Q \mapsto I_\theta(Q\mathcal{P})$ is convex, but we are trying to maximize this objective function. Hence, we may be confronted with numerous local and global optima at the boundary of the feasible set \mathcal{Q}_ϵ . Steinberger (2024) provided a numerical scheme for approximately solving (1) based on a linear program representation of Kairouz et al. (2016). However, the runtime of the linear program scales exponentially in the accuracy of approximation. Steinberger (2024) also showed that indeed, $[\sup_{Q \in \mathcal{Q}_\epsilon} I_\theta(Q\mathcal{P})]^{-1}$ is the smallest possible asymptotic variance among all (sequentially interactive) locally differentially private and regular estimators. Here, regular means that the scaled estimation error converges in distribution to a limiting distribution for all local parameters $\theta_n = \theta + h/\sqrt{n}$ (see Theorem 2 below).

Our main contribution is to show that the following simple binary privacy mechanism is an exact solution of (1) in the case where the privacy parameter ϵ is sufficiently small: Given that $X_i = x$, we generate Z_i by

$$Z_i = \begin{cases} \text{sgn}(x - \theta) & \text{with probability } p_\epsilon \\ -\text{sgn}(x - \theta) & \text{with probability } 1 - p_\epsilon, \end{cases}$$

where $p_\epsilon := \frac{e^\epsilon}{1+e^\epsilon}$ and $\text{sgn}(0) := 1$. In other words, for $z \in \{-1, 1\}$ and $x \in \mathbb{R}$, we set

$$Q_{\theta, \epsilon}^{\text{sgn}}(\{z\} | x) := P(Z_i = z | X_i = x) = \begin{cases} p_\epsilon, & \text{if } z = \text{sgn}(x - \theta), \\ 1 - p_\epsilon, & \text{if } z \neq \text{sgn}(x - \theta). \end{cases} \quad (2)$$

Notice that this mechanism can also be represented as a composition $Q_{\theta, \epsilon}^{\text{sgn}} = Q_\epsilon^{RR} T_\theta$, where $T_\theta(x) := \text{sgn}(x - \theta)$ and Q_ϵ^{RR} is the randomized response mechanism of Warner (1965) that flips the sign of the input with probability $1 - p_\epsilon$. In particular, we easily conclude that $Q_{\theta, \epsilon}^{\text{sgn}} \in \mathcal{Q}_\epsilon$. We can now state and prove the following result.

Theorem 1 *If $\epsilon \leq 0.67$ then*

$$I_\theta(Q\mathcal{P}) \leq I_\theta(Q_{\theta, \epsilon}^{\text{sgn}}\mathcal{P}) = \frac{2(e^\epsilon - 1)^2}{\pi(e^\epsilon + 1)^2},$$

for all $\theta \in \mathbb{R}$ and all $Q \in \mathcal{Q}_\epsilon$. In particular, the sign mechanism solves (1).

Our proof is based on the idea of quantizing the normal distribution and solving the problem of maximizing the private Fisher Information for discrete distributions simultaneously for all quantization levels. We apply results from [Kairouz et al. \(2016\)](#) and a delicate duality argument to solve the discrete case. The main steps of the proof are collected in Subsection 2.1.

An obvious issue with the sign mechanism in (2) is that it depends on the unknown parameter θ which we want to estimate. By incorporating the sign mechanism into a two-stage scheme as in [Steinberger \(2024\)](#) we end up with the following procedure for asymptotically efficient parameter estimation: Split the agents into two groups of size n_1 and $n_2 := n - n_1$ such that $\lim \frac{n_1}{n} = 0$.¹ We use the first group to get a consistent but not necessarily efficient estimate of the unknown mean θ and the second group to update the first-stage estimate to make it asymptotically efficient. The two stages of the procedure are described below, where θ_0 is an arbitrary initial guess for the unknown mean θ .

1. • Apply the sign mechanism $Q_{\theta_0, \epsilon}^{\text{sgn}}$ at the initial value θ_0 to X_1, \dots, X_{n_1} to obtain iid sanitized data $Z_i \sim Q_{\theta_0, \epsilon}^{\text{sgn}} P_\theta$, $i = 1, \dots, n_1$.

- Compute

$$\tilde{\theta}_{n_1} = \begin{cases} \theta_0 - \Phi^{-1} \left(\frac{1}{2} - \frac{1}{2} \frac{e^\epsilon + 1}{e^\epsilon - 1} \bar{Z}_{n_1} \right), & \text{if } |\bar{Z}_{n_1}| < \frac{e^\epsilon - 1}{e^\epsilon + 1}, \\ \theta_0, & \text{else,} \end{cases} \quad (3)$$

where Φ is the cumulative distribution function of the standard normal distribution and $\bar{Z}_{n_1} = \frac{1}{n_1} \sum_{i=1}^{n_1} Z_i$.

2. • Apply the sign mechanism $Q_{\tilde{\theta}_{n_1}, \epsilon}^{\text{sgn}}$ at the first stage estimate $\tilde{\theta}_{n_1}$ to the data X_{n_1+1}, \dots, X_n in the second group to generate values $Z_i \sim Q_{\tilde{\theta}_{n_1}, \epsilon}^{\text{sgn}} P_\theta$, $i = n_1 + 1, \dots, n$.

- Compute

$$\hat{\theta}_n = \begin{cases} \tilde{\theta}_{n_1} - \Phi^{-1} \left(\frac{1}{2} - \frac{1}{2} \frac{e^\epsilon + 1}{e^\epsilon - 1} \bar{Z}_{n_2} \right), & \text{if } |\bar{Z}_{n_2}| < \frac{e^\epsilon - 1}{e^\epsilon + 1}, \\ \tilde{\theta}_{n_1}, & \text{else,} \end{cases} \quad (4)$$

where $\bar{Z}_{n_2} = \frac{1}{n_2} \sum_{i=n_1+1}^n Z_i$.

Relying on our main Theorem 1, we can show that the above two-stage estimation procedure is regular and asymptotically achieves the minimal variance. See Subsection 2.2 for the proof. Let R_θ denote the distribution of the full sanitized data Z_1, \dots, Z_n when the true unknown parameter is $\theta \in \mathbb{R}$.

Theorem 2 *If $\epsilon \leq 0.67$ then the two-stage locally private estimation procedure described above satisfies*

$$\sqrt{n} \left(\hat{\theta}_n - [\theta + h/\sqrt{n}] \right) \xrightarrow[n \rightarrow \infty]{R_{\theta+h/\sqrt{n}}} N \left(0, \left[\sup_{Q \in \mathcal{Q}_\epsilon} I_\theta(Q\mathcal{P}) \right]^{-1} \right), \quad \forall h \in \mathbb{R}.$$

1. We parametrize $n_1 = n_1(n)$ by n such that all limits are understood as $n \rightarrow \infty$.

Notice that the two-stage estimator $\hat{\theta}_n$ depends on the choice of initial value θ_0 and on the size n_1 of the first group. Hence, θ_0 and n_1 are tuning parameters in our estimation procedure. In Section 3, we investigate the impact of θ_0 and n_1 on the performance of our locally private estimator in a finite sample simulation study.

1.1. Related Literature

Local differential privacy originated about 20 years ago (cf. [Dinur and Nissim, 2003](#); [Dwork and Nissim, 2004](#); [Dwork et al., 2006](#); [Dwork, 2008](#); [Evmimievski et al., 2003](#)) and has since become an incredibly popular way of data privacy protection when there is no trusted third party available. Despite its popularity, some of the most basic learning problems, such as the one considered here, have not been fully solved. Much attention has recently been paid to discrete distribution estimation and deriving upper bounds on the estimation error. For instance, [Wang et al. \(2016\)](#) have suggested an optimal mechanism for mutual information maximization for discrete distributions under LDP. [Ye and Barg \(2018\)](#) are studying the locally private minimax risk for discrete distributions. [Nam and Lee \(2022\)](#) is perhaps closer to our work. They consider the problem of maximizing a Fisher-Information in the local privacy mechanism, but they restrict to the 1-bit communication constraint which, for our purpose, is an oversimplified case. [Barnes et al. \(2020\)](#) also investigate and bound Fisher-Information when original data are perturbed using a differentially private randomization mechanism. But they have not attempted to find an optimal randomization mechanism or an efficient estimation procedure. Gaussian mean estimation is considered by [Joseph et al. \(2019\)](#) and they also use a two-stage procedure for parameter estimation. They obtain high-probability bounds for the estimation error, which is of order $n^{-1/2}$ for the two-round protocol, but unlike our work, they don't obtain the optimal asymptotic constant. The conceptual foundation to the problem of finding the optimal Fisher-Information mechanism was laid by [Steinberger \(2024\)](#). However, while in that work the optimal mechanism and locally private estimation procedure have to be approximated by a computationally expensive numerical optimization routine, we here derive an exact and closed form privacy mechanism and estimator.

2. Proof of main results

2.1. Proof of Theorem 1

In this section, we prove Theorem 1. We first follow the approach of [Steinberger \(2024\)](#) and quantize the Fisher-Information to then solve a discrete version of (1). However, while [Steinberger \(2024\)](#) had to rely on a numerical algorithm for the discrete optimization, which has exponential runtime in the resolution of the discretization, we here solve all the discrete problems exactly by providing a closed form solution and show that it is the same for all resolution levels k of the approximation. We then conclude that the solution to the simplified discrete problems is also the global solution. Notice that all the regularity conditions imposed by [Steinberger \(2024\)](#) are satisfied for the Gaussian location model $\mathcal{P} = \{N(\theta, 1) : \theta \in \mathbb{R}\}$ that we consider here (cf. Section 5.2 in [Steinberger, 2024](#)). In particular, the Fisher-Information $I_\theta(Q\mathcal{P})$ is well defined and finite for any privacy mechanism $Q \in \mathcal{Q}_\epsilon$.

2.1.1. DISCRETE APPROXIMATION

We begin by defining what is called a consistent quantizer in [Steinberger \(2024, cf. Definition 3 and Section 5 in that reference\)](#). For an even positive integer k and for $j = 1, \dots, k-1$, let $B_j := (x_{j-1}, x_j]$ and $B_k := (x_{k-1}, \infty)$, where $x_j := \Phi^{-1}(j/k)$ and Φ is the cdf of the standard normal distribution. Now set $T_{k,\theta}(x) := \sum_{j=1}^k j \mathbb{1}_{B_j}(x - \theta)$. Notice that $T_{k,\theta}$ maps the real line \mathbb{R} into the discrete set $[k] := \{1, \dots, k\}$, hence, it is called a quantizer. We write $T_{k,\theta_0} \mathcal{P} := \{r_{\theta,\theta_0} : \theta \in \mathbb{R}\}$ for the resulting quantized model, where $r_{\theta,\theta_0}(j) := P_\theta(T_{k,\theta_0}^{-1}(\{j\})) = P_\theta(B_j + \theta_0) = \Phi(x_j + \theta_0 - \theta) - \Phi(x_{j-1} + \theta_0 - \theta)$ is the probability mass function of the quantized data $T_{k,\theta_0}(X_i)$ and we set $\Phi(-\infty) := 0$, $\Phi(\infty) := 1$, $\phi(x) := \Phi'(x)$ and $\phi(\pm\infty) := 0$. Lemmas 4.7, 4.9, 5.1 and 5.2 in [Steinberger \(2024\)](#) show that for every even k and every $\theta \in \mathbb{R}$

$$\sup_{Q \in \mathcal{Q}_\epsilon(\mathbb{R})} I_\theta(Q\mathcal{P}) \leq \sup_{Q \in \mathcal{Q}_\epsilon([k])} I_\theta(QT_{k,\theta}\mathcal{P}) + \Delta_k = \sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_\theta(QT_{k,\theta}\mathcal{P}) + \Delta_k, \quad (5)$$

where $\Delta_k \rightarrow 0$ as $k \rightarrow \infty$, $\mathcal{Q}_\epsilon([k])$ is the set of ϵ -private mechanisms that take inputs from $[k]$ and produce random outputs from some arbitrary measurable space $(\mathcal{Z}, \mathcal{G})$, and $\mathcal{Q}_\epsilon([k] \rightarrow [k])$ are the ϵ -private mechanisms whose inputs and outputs both take values in $[k]$. Thus, the elements of $\mathcal{Q}_\epsilon([k] \rightarrow [k])$ can be represented as $k \times k$ column stochastic matrices with the property that $Q_{ij} \leq e^\epsilon Q_{ij'}$, for all $i, j, j' \in [k]$. Since obviously $\sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_\theta(QT_{k,\theta}\mathcal{P}) \leq \sup_{Q \in \mathcal{Q}_\epsilon} I_\theta(Q\mathcal{P})$, (5) yields

$$\sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_\theta(QT_{k,\theta}\mathcal{P}) \rightarrow \sup_{Q \in \mathcal{Q}_\epsilon} I_\theta(Q\mathcal{P}), \quad \forall \theta \in \mathbb{R},$$

as $k \rightarrow \infty$, k even. In the remainder of the proof, we will show that $\sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_\theta(QT_{k,\theta}\mathcal{P}) = I_\theta(Q_{\theta,\epsilon}^{\text{sgn}}\mathcal{P})$ for every even k . Thus, the sign mechanism in (2) must be a solution of (1).

2.1.2. EVALUATING THE QUANTIZED OBJECTIVE FUNCTION

Next, we provide an explicit expression for the quantized private Fisher-Information $I_\theta(QT_{k,\theta_0}\mathcal{P})$, for arbitrary $Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])$. The quantized and Q -privatized model $QT_{k,\theta_0}\mathcal{P} = \{m_{\theta,\theta_0} : \theta \in \mathbb{R}\}$ can be expressed via its probability mass functions $m_{\theta,\theta_0}(i) := \sum_{j=1}^k Q_{ij} r_{\theta,\theta_0}(j)$, $i \in [k]$. Thus, we have

$$I_\theta(QT_{k,\theta_0}\mathcal{P}) = \mathbb{E}_{Z \sim QT_{k,\theta_0}\mathcal{P}} \left[\left(\frac{\partial \log m_{\theta,\theta_0}(Z)}{\partial \theta} \right)^2 \right] = \sum_{i=1}^k \frac{\dot{m}_{\theta,\theta_0}(i)^2}{m_{\theta,\theta_0}(i)},$$

where $\dot{m}_{\theta,\theta_0}(i) := \frac{\partial}{\partial \theta} m_{\theta,\theta_0}(i) = \sum_{j=1}^k Q_{ij} [\phi(x_{j-1} + \theta_0 - \theta) - \phi(x_j + \theta_0 - \theta)]$ and the ratio is to be understood as equal to zero if the denominator $m_{\theta,\theta_0}(i)$ is zero. Abbreviating $y_j := \phi(x_{j-1}) - \phi(x_j)$, we arrive at

$$I_\theta(QT_{k,\theta}\mathcal{P}) = \sum_{i=1}^k \frac{\left(\sum_{j=1}^k Q_{ij} y_j \right)^2}{\frac{1}{k} \sum_{j=1}^k Q_{ij}} = \sum_{i=1}^k \mu(Q_i^T), \quad (6)$$

for $\mu(v) := k \frac{(v^T y)^2}{v^T \mathbf{1}}$, $v \in \mathcal{C}_k := \{u \in \mathbb{R}_+^k : u_j \leq e^\epsilon u_{j'}, \forall j, j' \in [k]\}$, and $\mu(0) := 0$. In particular, we see that $I_\theta(QT_{k,\theta}\mathcal{P}) = I_0(QT_{k,0}\mathcal{P})$ for all $\theta \in \mathbb{R}$. For later use we note that $y_j = -y_{k-j+1}$, $\sum_{j=1}^{k/2} y_j = -\phi(0) = -\sum_{j=k/2+1}^k y_j$ and $\sum_{j=1}^k y_j = 0$.

We conclude this subsection by computing the Fisher-Information for the sign mechanism in (2). Use the fact that $Q_{\theta,\epsilon}^{\text{sgn}} = Q_{\epsilon}^{RR} \circ T_{2,\theta} \in \mathcal{Q}_{\epsilon}([2] \rightarrow [2])$, where the randomized response mechanism can be represented in matrix form as

$$Q_{\epsilon}^{RR} = \begin{pmatrix} \frac{e^{\epsilon}}{1+e^{\epsilon}} & \frac{1}{1+e^{\epsilon}} \\ \frac{1}{1+e^{\epsilon}} & \frac{e^{\epsilon}}{1+e^{\epsilon}} \end{pmatrix}.$$

Thus,

$$\begin{aligned} I_{\theta}(Q_{\theta,\epsilon}^{\text{sgn}}\mathcal{P}) &= I_{\theta}(Q_{\epsilon}^{RR}T_{2,\theta}\mathcal{P}) = \sum_{i=1}^2 \frac{\left(\sum_{j=1}^2 [Q_{\epsilon}^{RR}]_{ij}y_j\right)^2}{\frac{1}{2}\sum_{j=1}^2 [Q_{\epsilon}^{RR}]_{ij}} \\ &= \frac{1}{1+e^{\epsilon}} \left(\frac{\phi(0)^2(1-e^{\epsilon})^2}{\frac{1}{2}(1+e^{\epsilon})} + \frac{\phi(0)^2(e^{\epsilon}-1)^2}{\frac{1}{2}(1+e^{\epsilon})} \right) = \frac{2}{\pi} \frac{(e^{\epsilon}-1)^2}{(e^{\epsilon}+1)^2}. \end{aligned}$$

2.1.3. REFORMULATION AS A LINEAR PROGRAM

From Lemma 4.5 of Steinberger (2024), $Q \mapsto I_0(QT_{k,0}\mathcal{P})$ is a continuous, sublinear and convex function on the compact set $\mathcal{Q}_{\epsilon}([k] \rightarrow [k]) \subseteq \mathbb{R}^{k \times k}$. Therefore, a maximizer exists and Theorems 2 and 4 of Kairouz et al. (2016) allow us to conclude that the maximization problem

$$\max_Q I_0(QT_{k,0}\mathcal{P}) \quad \text{subject to} \quad Q \in \mathcal{Q}_{\epsilon}([k] \rightarrow [k]) \quad (7)$$

has the same optimal value as the linear program

$$\begin{aligned} \max_{\alpha \in \mathbb{R}^{2^k}} \quad & \sum_{j=1}^{2^k} \mu(S_j^{(k)})\alpha_j \\ \text{s.t.} \quad & S^{(k)}\alpha = \mathbf{1} \\ & \alpha \geq 0, \end{aligned} \quad (8)$$

where $S^{(k)}$ is a staircase matrix defined as follows: For $0 \leq j \leq 2^k - 1$ consider its binary representation $b_j \in \{0, 1\}^k$ with $j = \sum_{i=1}^k b_{ij}2^{k-i}$. Then $(S^{(k)})_{i,j+1} := b_{ij}(e^{\epsilon} - 1) + 1$. For instance, with $k = 4$ we get the following matrix:

$$S^{(4)} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} \\ 1 & 1 & 1 & 1 & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & 1 & 1 & 1 & 1 & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} & e^{\epsilon} \\ 1 & 1 & e^{\epsilon} & e^{\epsilon} & 1 & 1 & e^{\epsilon} & e^{\epsilon} & 1 & 1 & e^{\epsilon} & e^{\epsilon} & 1 & 1 & e^{\epsilon} & e^{\epsilon} \\ 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} & 1 & e^{\epsilon} \end{bmatrix}.$$

Moreover, if α^* is a solution of (8), then $Q^* = [S^{(k)} \text{diag}(\alpha^*)]^T \in \mathbb{R}^{2^k \times k}$ is a solution of (7). From Kairouz et al. (2016, Theorem 2) we know that an optimal mechanism Q^* has at most k non-zero rows. Since zero rows do not contribute to the Fisher-Information (6), we can remove them from Q^* to obtain $Q^* \in \mathbb{R}^{k \times k}$.

Finally, notice that $\alpha_j^* := (1+e^{\epsilon})^{-1}$ for $j \in \{2^{k/2}, 2^k - 2^{k/2} + 1\}$ and $\alpha_j^* := 0$ else, is a feasible point of (8), because $b_{2^{k/2}} = \mathbf{1} - b_{2^k - 2^{k/2} + 1}$ and thus

$$S^{(k)}\alpha^* = \frac{1}{1+e^{\epsilon}} [b_{2^{k/2}}(e^{\epsilon} - 1) + \mathbf{1} + b_{2^k - 2^{k/2} + 1}(e^{\epsilon} - 1) + \mathbf{1}] = \mathbf{1}.$$

Furthermore, α^* achieves an objective function value of

$$\begin{aligned} & \frac{1}{1+e^\epsilon} \left[\mu(b_{2^{k/2}}(e^\epsilon - 1) + 1) + \mu(b_{2^{k-2^{k/2}+1}}(e^\epsilon - 1) + 1) \right] \\ &= \frac{2}{1+e^\epsilon} \left[\frac{\phi(0)^2(e^\epsilon - 1)^2}{e^\epsilon + 1} + \frac{\phi(0)^2(e^\epsilon - 1)^2}{e^\epsilon + 1} \right] = I_\theta(Q_{\theta, \epsilon}^{\text{sgn}} \mathcal{P}). \end{aligned}$$

Hence, we have $\sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_\theta(QT_k, \theta \mathcal{P}) = \sup_{Q \in \mathcal{Q}_\epsilon([k] \rightarrow [k])} I_0(QT_k, 0 \mathcal{P}) \geq I_\theta(Q_{\theta, \epsilon}^{\text{sgn}} \mathcal{P})$, for every even $k > 0$. It remains to establish an upper bound, which we do by a duality argument.

2.1.4. GUESSING A DUAL SOLUTION

In this section, we study the dual of the LP in (8). By weak duality, any feasible value of the dual provides an upper bound on the objective function of the primal problem. Hence, the challenge is to identify a feasible value of the dual that achieves an objective function equal to $I_\theta(Q_{\theta, \epsilon}^{\text{sgn}} \mathcal{P})$. The dual program reads as follows:

$$\begin{aligned} & \min_{\beta \in \mathbb{R}^k} \mathbf{1}^T \beta \\ & \text{s.t. } (S^{(k)})^T \beta \geq \boldsymbol{\mu} \end{aligned} \quad (9)$$

where $\boldsymbol{\mu} := (\mu_j)_{j=1}^{2^k} := (\mu(S_{\cdot j}^{(k)}))_{j=1}^{2^k} \in \mathbb{R}^{2^k}$. The proof of Theorem 1 is finished if we can identify a feasible point $\beta \in \mathbb{R}^k$ of (9) satisfying $\mathbf{1}^T \beta = \sum_{j=1}^k \beta_j = \frac{2}{\pi} t_\epsilon^2$, where $t_\epsilon := \frac{e^\epsilon - 1}{e^\epsilon + 1}$. Our guess for such a $\beta^* \in \mathbb{R}^k$ is

$$\beta_j^* := -\frac{2t_\epsilon^2}{\pi k} + |y_j| t_\epsilon^2 \sqrt{\frac{8}{\pi}}.$$

This clearly satisfies $\mathbf{1}^T \beta^* = \frac{2}{\pi} t_\epsilon^2$. It is the main technical challenge of the proof to show that β^* is a feasible point of (9).

For fixed $j \in [2^k]$, we have to verify $(S_{\cdot j}^{(k)})^T \beta^* \geq \mu(S_{\cdot j}^{(k)})$. It will be convenient to partition the index set $[k]$ as follows $[k] = I_1^0 \cup I_{e^\epsilon}^0 \cup I_1^1 \cup I_{e^\epsilon}^1$ where $I_1^0 = \{i \mid i \leq \frac{k}{2}, S_{i,j}^{(k)} = 1\}$, $I_{e^\epsilon}^0 = \{i \mid i \leq \frac{k}{2}, S_{i,j}^{(k)} = e^\epsilon\}$ and $I_1^1, I_{e^\epsilon}^1$ defined similarly for the second half of indices $i = \frac{k}{2} + 1, \dots, k$. Then we can rewrite the required inequality $(S_{\cdot j}^{(k)})^T \beta^* \geq \mu(S_{\cdot j}^{(k)})$ as follows:

$$\sum_{i \in I_1^0} \beta_i^* + e^\epsilon \sum_{i \in I_{e^\epsilon}^0} \beta_i^* + \sum_{i \in I_1^1} \beta_i^* + e^\epsilon \sum_{i \in I_{e^\epsilon}^1} \beta_i^* \geq \frac{\left(\sum_{i \in I_1^0} y_i + e^\epsilon \sum_{i \in I_{e^\epsilon}^0} y_i + \sum_{i \in I_1^1} y_i + e^\epsilon \sum_{i \in I_{e^\epsilon}^1} y_i \right)^2}{\frac{1}{k} (|I_1^0| + |I_1^1| + e^\epsilon |I_{e^\epsilon}^0| + e^\epsilon |I_{e^\epsilon}^1|)}.$$

For convenience let us denote $m_1 = |I_1^0|$, $m_2 = |I_{e^\epsilon}^1|$, $a_1 = \sqrt{2\pi} \sum_{i \in I_1^0} |y_i|$, $a_2 = \sqrt{2\pi} \sum_{i \in I_{e^\epsilon}^1} |y_i|$

such that $m_1, m_2 \leq \frac{k}{2}$ and $a_1, a_2 \in [0, 1]$. With this new notation, we can rewrite the sum on the left-hand-side as

$$(1 - e^\epsilon) \sum_{i \in I_1^0} \beta_i^* + (e^\epsilon - 1) \sum_{i \in I_{e^\epsilon}^1} \beta_i^* + e^\epsilon \sum_{i=1}^{k/2} \beta_i^* + \sum_{i=k/2+1}^k \beta_i^* \quad (10)$$

and the ratio on the right-hand-side as

$$\frac{\left(\frac{e^\epsilon-1}{\sqrt{2\pi}}a_1 + \frac{e^\epsilon-1}{\sqrt{2\pi}}a_2 - \frac{e^\epsilon}{\sqrt{2\pi}} + \frac{1}{\sqrt{2\pi}}\right)^2}{\frac{1}{k}(k/2 + m_1 - m_2 + e^\epsilon(k/2 + m_2 - m_1))}. \quad (11)$$

Substituting the values for β^* we simplify the expression in (10) to obtain

$$\begin{aligned} & (1 - e^\epsilon) \sum_{i \in I_1^0} \beta_i^* + (e^\epsilon - 1) \sum_{i \in I_{e^\epsilon}^1} \beta_i^* + e^\epsilon \sum_{i=1}^{k/2} \beta_i^* + \sum_{i=k/2+1}^k \beta_i^* \\ &= (1 - e^\epsilon) \sum_{i \in I_1^0} \left[-\frac{2t_\epsilon^2}{\pi k} + |y_i|t_\epsilon^2 \sqrt{\frac{8}{\pi}} \right] + (e^\epsilon - 1) \sum_{i \in I_{e^\epsilon}^1} \left[-\frac{2t_\epsilon^2}{\pi k} + |y_i|t_\epsilon^2 \sqrt{\frac{8}{\pi}} \right] + (1 + e^\epsilon) \frac{t_\epsilon^2}{\pi} \\ &= (1 - e^\epsilon) \frac{2t_\epsilon^2}{\pi} \left(-\frac{m_1}{k} + a_1 \right) + (e^\epsilon - 1) \frac{2t_\epsilon^2}{\pi} \left(-\frac{m_2}{k} + a_2 \right) + (1 + e^\epsilon) \frac{t_\epsilon^2}{\pi} \\ &= (e^\epsilon - 1) \frac{2t_\epsilon^2}{\pi} \left[\frac{m_1 - m_2}{k} + a_2 - a_1 \right] + (1 + e^\epsilon) \frac{t_\epsilon^2}{\pi}. \end{aligned}$$

Next, we simplify (11) to get

$$\begin{aligned} & \frac{\left(\frac{e^\epsilon-1}{\sqrt{2\pi}}a_1 + \frac{e^\epsilon-1}{\sqrt{2\pi}}a_2 - \frac{e^\epsilon}{\sqrt{2\pi}} + \frac{1}{\sqrt{2\pi}}\right)^2}{\frac{1}{k}\left(\frac{k}{2} + m_1 - m_2 + e^\epsilon\left(\frac{k}{2} + m_2 - m_1\right)\right)} = \frac{(a_1 + a_2 - 1)^2 (e^\epsilon - 1)^2}{\frac{2\pi}{k} \left(\frac{k}{2}(1 + e^\epsilon) + (m_2 - m_1)(e^\epsilon - 1)\right)} \\ &= \frac{(a_1 + a_2 - 1)^2 (e^\epsilon - 1)^2}{\pi(1 + e^\epsilon) \left(1 + 2t_\epsilon \frac{m_2 - m_1}{k}\right)}. \end{aligned}$$

After these simplifications, we see that the inequality we need to verify is given by

$$(e^\epsilon - 1) \frac{2t_\epsilon^2}{\pi} \left[\frac{m_1 - m_2}{k} + a_2 - a_1 \right] + (1 + e^\epsilon) \frac{t_\epsilon^2}{\pi} \geq \frac{(a_1 + a_2 - 1)^2 (e^\epsilon - 1)^2}{\pi(1 + e^\epsilon) \left(1 + 2t_\epsilon \frac{m_2 - m_1}{k}\right)}.$$

Dividing by $(1 + e^\epsilon)t_\epsilon^2/\pi$, this can be further simplified to

$$2t_\epsilon(a_2 - a_1) + 4t_\epsilon^2 \frac{m_2 - m_1}{k} (a_2 - a_1) - 4t_\epsilon^2 \left(\frac{m_2 - m_1}{k} \right)^2 - (a_1 + a_2)^2 + 2(a_1 + a_2) \geq 0. \quad (12)$$

Without loss of generality, we can assume $a_1 + a_2 \leq 1$; otherwise, we could have chosen $a_1 = \sqrt{2\pi} \sum_{i \in I_1^1} |y_i|$ and $a_2 = \sqrt{2\pi} \sum_{i \in I_{e^\epsilon}^0} |y_i|$, maintaining the same inequality due to the symmetry of $\beta_i^* = \beta_{k+1-i}^*$ and antisymmetry in $y_i = -y_{k+1-i}$, but with the sum $a_1 + a_2 \leq 1$, because the total sum equals $\sqrt{2\pi} \sum_{i=1}^k |y_i| = 2$. Given this, we immediately get $-(a_1 + a_2)^2 + 2(a_1 + a_2) \geq a_1 + a_2$. Further more, it is obvious that

$$\begin{aligned} 2t_\epsilon(a_2 - a_1) &\geq -2t_\epsilon(a_1 + a_2), \quad -4t_\epsilon^2 \left(\frac{m_2 - m_1}{k} \right)^2 \geq -4t_\epsilon^2 \frac{m_1^2 + m_2^2}{k^2}, \\ 4t_\epsilon^2 \frac{m_2 - m_1}{k} (a_2 - a_1) &\geq -2t_\epsilon^2 \left| \frac{m_2 - m_1}{k/2} \right| \cdot |a_2 - a_1| \geq -2t_\epsilon^2 (a_1 + a_2). \end{aligned}$$

Now, combining these four inequalities, we see that if

$$-2t_\epsilon(a_2 + a_1) - 2t_\epsilon^2(a_1 + a_2) - 4t_\epsilon^2 \frac{m_1^2 + m_2^2}{k^2} + a_1 + a_2 \geq 0$$

holds true, then also (12) is satisfied. But this new inequality in the previous display is equivalent to

$$a_1 + a_2 \geq \frac{4t_\epsilon^2}{1 - 2t_\epsilon - 2t_\epsilon^2} \frac{m_1^2 + m_2^2}{k^2}, \quad (13)$$

provided that $1 - 2t_\epsilon - 2t_\epsilon^2 > 0$. To proceed, we establish two additional lemmas to get the lower bounds for a_1 and a_2 in terms of m_1 and m_2 respectively.

Lemma 3 *For any k , the sequence $y_j = \phi(\Phi^{-1}((j-1)/k)) - \phi(\Phi^{-1}(j/k))$, $j = 1, \dots, k$, is increasing.*

Proof For $x \in (1, k]$, define $g(x) := \phi(\Phi^{-1}((x-1)/k)) - \phi(\Phi^{-1}(x/k))$ and $g(1) := -\phi(\Phi^{-1}(1/k))$. Since g is continuous on $[1, k]$, it suffices to show that g is strictly increasing on $(1, k)$. Using the fact that $\phi'(x) = (-x)\phi(x)$ and the inverse function theorem we get

$$\begin{aligned} g'(x) &= \frac{1}{k} \left(\frac{\phi'}{\phi} \circ \Phi^{-1}((x-1)/k) - \frac{\phi'}{\phi} \circ \Phi^{-1}(x/k) \right) \\ &= \frac{1}{k} (\Phi^{-1}(x/k) - \Phi^{-1}((x-1)/k)) > 0. \end{aligned}$$

Thus, g is strictly increasing on $(1, k)$. ■

Lemma 4 *For $x \in [0, \frac{1}{2}]$ we have*

$$\phi(0) - \phi\left(\Phi^{-1}\left(\frac{1}{2} + x\right)\right) \geq \sqrt{\frac{\pi}{2}}x^2.$$

Proof For $x \in [0, \frac{1}{2}]$, define $g(x) := \phi(0) - \phi(\Phi^{-1}(\frac{1}{2} + x)) - \sqrt{\frac{\pi}{2}}x^2$ then $g(0) = 0$, $g(\frac{1}{2}) = +\infty$. Since g is continuous on $[0, \frac{1}{2}]$, it suffices to show that g is strictly increasing on $(0, \frac{1}{2})$. Using the fact that $\phi'(x) = (-x)\phi(x)$ and the inverse function theorem we get

$$\begin{aligned} g'(x) &= -\frac{\phi'}{\phi} \circ \Phi^{-1}\left(\frac{1}{2} + x\right) - \sqrt{2\pi}x = \Phi^{-1}\left(\frac{1}{2} + x\right) - \sqrt{2\pi}x, \\ g''(x) &= \frac{1}{\phi(\Phi^{-1}(\frac{1}{2} + x))} - \sqrt{2\pi} \geq 0. \end{aligned}$$

Since the first derivative is equal to zero at zero and non-decreasing, we see that the function g itself must be non-decreasing. ■

By combining these two lemmas, we can easily show that

$$a_2 = \sqrt{2\pi} \sum_{i \in I_{e^\epsilon}^1} y_i \geq \sqrt{2\pi} \sum_{i=k/2+1}^{k/2+m_2} y_i = \sqrt{2\pi} \left(\phi(0) - \phi \left(\Phi^{-1} \left(\frac{1}{2} + \frac{m_2}{k} \right) \right) \right) \geq \pi \frac{m_2^2}{k^2}.$$

A similar lower bound holds for a_1 as well, namely $a_1 \geq \pi \frac{m_1^2}{k^2}$. Therefore, we obtain the required inequality (13) provided that

$$a_1 + a_2 \geq \pi \frac{m_1^2 + m_2^2}{k^2} \geq \frac{4t_\epsilon^2}{1 - 2t_\epsilon - 2t_\epsilon^2} \frac{m_1^2 + m_2^2}{k^2},$$

which is true whenever

$$\frac{4t_\epsilon^2}{1 - 2t_\epsilon - 2t_\epsilon^2} \leq \pi.$$

But since $\epsilon \mapsto t_\epsilon$ is increasing this is seen to hold for all $0 \leq \epsilon \leq 0.67$, for which we also have $1 - 2t_\epsilon - 2t_\epsilon^2 > 0$. This finishes the proof.

2.2. Proof of Theorem 2

Rather than relying on Theorem 4.12 in Steinberger (2024) and checking all their assumptions, we here present a direct proof that also has the advantage of being self-contained. We begin by showing the consistency of first and second-stage estimators. For convenience, fix $h \in \mathbb{R}$ and set $\theta_n := \theta + h/\sqrt{n}$. All probabilities, expectations, variances and stochastic convergence results below are with respect to R_{θ_n} .

2.2.1. CONSISTENCY

Showing consistency of the first-stage estimator (3) is rather straight forward, because Z_1, \dots, Z_{n_1} are iid with values in $\{-1, 1\}$ and

$$\begin{aligned} \mathbb{E}[Z_1] &= P(Z_1 = 1) - P(Z_1 = -1) = 2P(Z_1 = 1) - 1 \\ &= 2 \frac{e^\epsilon}{1 + e^\epsilon} - 2 \frac{e^\epsilon - 1}{e^\epsilon + 1} \Phi(\theta_0 - \theta_n) - 1 = \frac{e^\epsilon - 1}{e^\epsilon + 1} (1 - 2\Phi(\theta_0 - \theta_n)). \end{aligned}$$

Thus, by Markov's inequality, \bar{Z}_{n_1} converges in probability to $\frac{e^\epsilon - 1}{e^\epsilon + 1} (1 - 2\Phi(\theta_0 - \theta)) \in (-\frac{e^\epsilon - 1}{e^\epsilon + 1}, \frac{e^\epsilon - 1}{e^\epsilon + 1})$, and thus $P(|\bar{Z}_{n_1}| < \frac{e^\epsilon - 1}{e^\epsilon + 1}) \rightarrow 1$ as $n \rightarrow \infty$ and $\tilde{\theta}_{n_1}$ converges to θ in probability, by the continuous mapping theorem.

Consistency of (4) can be shown in a similar way, using the conditional Markov inequality. In view of consistency of $\tilde{\theta}_{n_1}$ it suffices to show that $\bar{Z}_{n_2} = \frac{1}{n_2} \sum_{i=n_1+1}^n Z_i \rightarrow 0$ in probability as $n \rightarrow \infty$. Conditionally on $\tilde{\theta}_{n_1}$, the Z_{n_1+1}, \dots, Z_n are iid. The conditional expectation is computed in the same way as above, that is

$$\mathbb{E}[\bar{Z}_{n_2} | \tilde{\theta}_{n_1}] = \mathbb{E}[Z_n | \tilde{\theta}_{n_1}] = \frac{e^\epsilon - 1}{e^\epsilon + 1} (1 - 2\Phi(\tilde{\theta}_{n_1} - \theta_n)),$$

and it converges to zero in probability as $n \rightarrow \infty$. The Z_i 's only take values 1 or -1 and thus their conditional variance is bounded by 1. Now

$$\begin{aligned}
 P(|\bar{Z}_{n_2}| > \varepsilon) &\leq P(|\bar{Z}_{n_2} - \mathbb{E}[\bar{Z}_{n_2}|\tilde{\theta}_{n_1}]| + |\mathbb{E}[\bar{Z}_{n_2}|\tilde{\theta}_{n_1}]| > \varepsilon) \\
 &\leq \mathbb{E} \left[P \left(|\bar{Z}_{n_2} - \mathbb{E}[\bar{Z}_{n_2}|\tilde{\theta}_{n_1}]| > \frac{\varepsilon}{2} \middle| \tilde{\theta}_{n_1} \right) \right] + P \left(|\mathbb{E}[\bar{Z}_{n_2}|\tilde{\theta}_{n_1}]| > \frac{\varepsilon}{2} \right) \\
 &\leq \frac{4}{\varepsilon^2} \mathbb{E} \left[1 \wedge \text{Var}[\bar{Z}_{n_2}|\tilde{\theta}_{n_1}] \right] + P \left(|\mathbb{E}[Z_n|\tilde{\theta}_{n_1}]| > \frac{\varepsilon}{2} \right) \\
 &\leq \frac{4}{\varepsilon^2 n_2} + P \left(|\mathbb{E}[Z_n|\tilde{\theta}_{n_1}]| > \frac{\varepsilon}{2} \right) \xrightarrow{n \rightarrow \infty} 0.
 \end{aligned}$$

2.2.2. ASYMPTOTIC NORMALITY

For asymptotic normality, first notice that we already showed above that $|\bar{Z}_{n_1}|$ and $|\bar{Z}_{n_2}|$ are bounded by $\frac{e^\varepsilon - 1}{e^\varepsilon + 1}$ with asymptotic probability one. Hence, it suffices to proceed on the event where both of these absolute values are appropriately bounded. Write $G(p) := \Phi^{-1}(p)$ and consider the scaled estimation error of (4), that is

$$\begin{aligned}
 \sqrt{n}(\hat{\theta}_n - \theta_n) &= -\sqrt{n} \left(\Phi^{-1} \left(\frac{1}{2} - \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \bar{Z}_{n_2} \right) - (\tilde{\theta}_{n_1} - \theta_n) \right) \\
 &= -\sqrt{n} \left(G \left(\frac{1}{2} - \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \bar{Z}_{n_2} \right) - G(\Phi(\tilde{\theta}_{n_1} - \theta_n)) \right).
 \end{aligned}$$

By a first order Taylor expansion and writing $A_n := \frac{1}{2} - \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \bar{Z}_{n_2} = \frac{1}{n_2} \sum_{i=n_1+1}^n \left(\frac{1}{2} - \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} Z_i \right)$ and $B_n := \Phi(\tilde{\theta}_{n_1} - \theta_n)$, we get

$$\sqrt{n}(G(A_n) - G(B_n)) = G'(B_n) \sqrt{n}(A_n - B_n) + \frac{1}{2} \sqrt{n} \int_{B_n}^{A_n} G''(t)(A_n - t)^2 dt. \quad (14)$$

From our results of the previous subsection we see that $A_n \rightarrow \frac{1}{2}$ and $B_n \rightarrow \frac{1}{2}$ in probability, and since G' is continuous on $[0, 1]$ we have $G'(B_n) \rightarrow G'(\frac{1}{2}) = [\phi \circ \Phi^{-1}(\frac{1}{2})]^{-1} = \sqrt{2\pi}$, in probability. Next, we show that $\sqrt{n}(A_n - B_n) \rightarrow N(0, \frac{1}{4}(\frac{e^\varepsilon + 1}{e^\varepsilon - 1})^2)$ in distribution. For $i = n_1 + 1, \dots, n$, define $Y_i := \frac{1}{2} - \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} Z_i - \Phi(\tilde{\theta}_{n_1} - \theta_n)$ and note that conditional on $\tilde{\theta}_{n_1}$ they are iid with conditional mean zero and conditional variance

$$\text{Var}[Y_i|\tilde{\theta}_{n_1}] = \frac{1}{4} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^2 \left[1 - \left(\frac{e^\varepsilon - 1}{e^\varepsilon + 1} (1 - 2\Phi(\tilde{\theta}_{n_1} - \theta_n)) \right)^2 \right] \xrightarrow{n \rightarrow \infty} \frac{1}{4} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^2.$$

Moreover, $\text{Var}[Y_i|\tilde{\theta}_{n_1}] \geq \frac{1}{4} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^2 [1 - (\frac{e^\varepsilon - 1}{e^\varepsilon + 1})^2]$ and $\mathbb{E}[|Y_i|^3|\tilde{\theta}_{n_1}] \leq \frac{3}{2} + \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1}$. Thus, from the Berry-Esseen bound (cf. [Berry, 1941](#), Theorem 1) applied to the conditional distribution, we get for every $t \in \mathbb{R}$,

$$\begin{aligned}
 &\left| P \left(\sqrt{n_2} \frac{A_n - B_n}{\sqrt{\text{Var}(Y_n|\tilde{\theta}_{n_1})}} \leq t \right) - \Phi(t) \right| = \left| P \left(\sqrt{n_2} \sum_{i=n_1+1}^n \frac{Y_i}{\sqrt{\text{Var}(Y_i|\tilde{\theta}_{n_1})}} \leq t \right) - \Phi(t) \right| \\
 &\leq \mathbb{E} \left| P \left(\sqrt{n_2} \sum_{i=n_1+1}^n \frac{Y_i}{\sqrt{\text{Var}(Y_i|\tilde{\theta}_{n_1})}} \leq t \middle| \tilde{\theta}_{n_1} \right) - \Phi(t) \right| \leq \frac{1.88 \left(\frac{3}{2} + \frac{1}{2} \frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)}{\frac{1}{8} \left(\frac{e^\varepsilon + 1}{e^\varepsilon - 1} \right)^3 [1 - (\frac{e^\varepsilon - 1}{e^\varepsilon + 1})^2]^{\frac{3}{2}} \sqrt{n_2}}.
 \end{aligned}$$

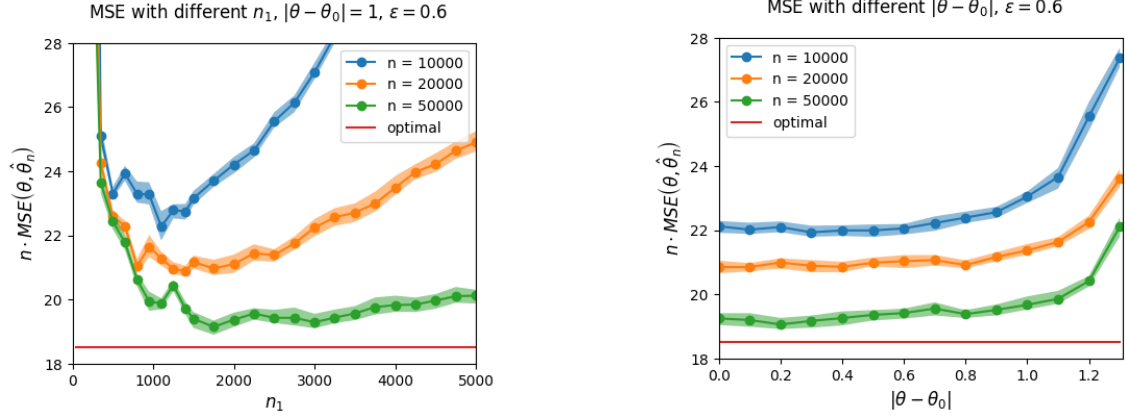


Figure 1: Left panel: Scaled MSE of the private estimator as a function of the number n_1 of first-stage samples. Right panel: Scaled MSE of the private estimator as a function of the initial guess θ_0 .

Applying Slutski's theorem and noting that $n/n_2 \rightarrow 1$, we get the desired convergence of $\sqrt{n}(A_n - B_n) \rightarrow N(0, \frac{1}{4}(\frac{e^\epsilon + 1}{e^\epsilon - 1})^2)$ in distribution. Hence, $G'(B_n)\sqrt{n}(A_n - B_n) \rightarrow N(0, \frac{\pi}{2}(\frac{e^\epsilon + 1}{e^\epsilon - 1})^2)$. From Theorem 1 we see that this is the optimal asymptotic variance. Thus, the proof is finished if we can show that the remainder term in (14) converges to zero in probability. But this is easily seen, because on the event where $A_n, B_n \in [\frac{1}{4}, \frac{3}{4}]$, which has asymptotic probability one, the remainder term is bounded in absolute value by

$$\frac{1}{2}|\sqrt{n}(A_n - B_n)||A_n - B_n|^2 \sup_{t \in [\frac{1}{4}, \frac{3}{4}]} |G''(t)| \xrightarrow[n \rightarrow \infty]{i.p.} 0.$$

3. Experiments

In this section, we provide numerical experiments to investigate the finite sample performance of our two-stage locally private estimation procedure described in (3) and (4) in dependence on the tuning parameters $\theta_0 \in \mathbb{R}$ and $n_1 \in \mathbb{N}$, for different sample sizes n and $\epsilon = 0.6$. For the plots, we generated 200000 Montecarlo samples and computed the 95% confidence intervals via bootstrapping. In the left panel of Figure 1, we plot the scaled mean squared error (MSE) of $\hat{\theta}_n$ as a function of the number n_1 of samples in the first stage of our private estimation procedure. From the plots, we see that a proper choice of n_1 is crucial to get a small MSE. Especially too small values of n_1 must be avoided.

In the right panel of Figure 1, we plot the scaled MSE as a function of the difference $\theta_0 - \theta$, where n_1 was chosen such that the corresponding MSE function in the left panel is minimized. The results are not surprising. The MSE increases as the initial guess moves away from the true value. Notice that our results do not provide uniform convergence in θ , and in fact it is known that with differential privacy it is impossible to achieve uniformity. Thus, extreme values of the parameter θ , in the sense that $|\theta_0 - \theta|$ is large, will be harder to estimate.

Acknowledgments

This research was supported by the Austrian Science Fund (FWF): I 5484-N, as part of the Research Unit 5381 of the German Research Foundation.

References

- Leighton Pate Barnes, Wei-Ning Chen, and Ayfer Özgür. Fisher information under local differential privacy. *IEEE Journal on Selected Areas in Information Theory*, 1(3):645–659, 2020.
- Andrew C Berry. The accuracy of the Gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136, 1941.
- Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 202–210. ACM, 2003.
- Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008. doi: 10.1007/978-3-540-79228-4.
- Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In *Annual International Cryptology Conference*, pages 528–544. Springer, 2004.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, Lecture Notes in Computer Science, pages 265–284. Springer, 2006. doi: 10.1007/11681878.
- Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 211–222. ACM, 2003. doi: 10.1145/773153.773174.
- Matthew Joseph, Janardhan Kulkarni, Jieming Mao, and Steven Z Wu. Locally private Gaussian estimation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. *The Journal of Machine Learning Research*, 17(1):492–542, 2016.
- Seung-Hyun Nam and Si-Hyeon Lee. A tighter converse for the locally differentially private discrete distribution estimation under the one-bit communication constraint. *IEEE Signal Processing Letters*, 29:1923–1927, 2022. doi: 10.1109/LSP.2022.3205276.
- Lukas Steinberger. Efficiency in local differential privacy. *arXiv preprint arXiv:2301.10600*, 2024.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, 8th edition, 2007.
- Shaowei Wang, Liusheng Huang, Pengzhan Wang, Yiwen Nie, Hongli Xu, Wei Yang, Xiang-Yang Li, and Chunming Qiao. Mutual information optimally local private discrete distribution estimation. *arXiv preprint arXiv:1607.08025*, 2016.

Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965. doi: 10.1080/01621459.1965.10480775.

Min Ye and Alexander Barg. Optimal schemes for discrete distribution estimation under locally differential privacy. *IEEE Transactions on Information Theory*, 64(8):5662–5676, 2018.