

Optimal Federated Learning for Functional Mean Estimation under Heterogeneous Privacy Constraints

Tony Cai, Abhinav Chakraborty, Lasse Vuursteen

February 10, 2025

Abstract

Federated learning (FL) is a distributed machine learning technique designed to preserve data privacy and security, and it has gained significant importance due to its broad range of applications. This paper addresses the problem of optimal functional mean estimation from discretely sampled data in a federated setting.

We consider a heterogeneous framework where the number of individuals, measurements per individual, and privacy parameters vary across one or more servers, under both common and independent design settings. In the common design setting, the same design points are measured for each individual, whereas in the independent design, each individual has their own random collection of design points. Within this framework, we establish minimax upper and lower bounds for the estimation error of the underlying mean function, highlighting the nuanced differences between common and independent designs under distributed privacy constraints.

We propose algorithms that achieve the optimal trade-off between privacy and accuracy and provide optimality results that quantify the fundamental limits of private functional mean estimation across diverse distributed settings. These results characterize the cost of privacy and offer practical insights into the potential for privacy-preserving statistical analysis in federated environments.

Keywords: Differential Privacy, Distributed Learning, Federated learning, Functional Mean Estimation, Heterogeneous Data, Minimax Rates, Privacy-Accuracy Trade-off

1 Introduction

In the era of big data, sensitive information is often distributed across multiple sources in fields such as healthcare and financial services. Privacy concerns prevent direct data pooling, making it essential to develop efficient statistical inference methods that preserve privacy while leveraging the collective power of distributed data. Federated learning addresses this challenge by enabling organizations or groups to collaboratively train and improve a shared global machine learning model without sharing raw data externally, ensuring that privacy is maintained at each data source. Differential privacy (DP) [14] has emerged as the leading framework for providing

rigorous privacy guarantees, offering a principled approach to quantifying the extent to which individual privacy is protected within a dataset.

The problem of estimating the mean of random functions based on discretely sampled data arises naturally in functional data analysis and has significant applications in areas such as signal processing, biomedical imaging, environmental monitoring, and financial modeling [27, 16]. Classical methods have explored this problem extensively, with adaptive algorithms and theoretical results establishing optimal rates of convergence under varying sampling schemes [9, 20, 32, 28]. These methods emphasize a critical balance: increasing the number of subjects reduces variability across the population, while increasing the number of measurements per subject improves the precision of individual function estimates. This trade-off is fundamental to designing efficient sampling strategies.

Incorporating differential privacy into functional mean estimation introduces unique challenges. DP requires a trade-off between privacy and accuracy by adding noise to the data or the estimation process. For functional data, this trade-off is further complicated by the relationship between sampling intensity and privacy risk – taking more measurements per individual increases privacy risks, even though it enhances estimation accuracy. Developing DP-compliant estimators requires careful consideration of both privacy constraints and the intrinsic properties of functional data, diverging significantly from traditional non-private methods.

The federated learning paradigm introduces additional complexities to functional mean estimation under privacy constraints. In federated settings, data is distributed across multiple servers or entities, such as hospitals, mobile devices, or autonomous vehicles, each with distinct privacy requirements, see e.g. [6] and references therein for an overview. These settings often involve heterogeneous data characterized by varying numbers of individuals, measurements, and privacy budgets across servers. This heterogeneity exacerbates the challenges of balancing privacy and statistical accuracy.

This paper investigates the cost of privacy in functional mean estimation within the framework of Federated Differential Privacy (FDP). FDP extends DP principles to distributed environments, enabling privacy-preserving statistical analysis at various levels, including individual data holders, servers, or aggregated data. Private functional data analysis, especially when the private outputs are functions, has received recent attention, primarily within the central DP framework [18, 25].

Our work makes several key contributions. First, we establish minimax upper and lower bounds for the estimation error of the mean function under both common and independent design settings, highlighting key differences under FDP constraints. Second, we propose novel algorithms that achieve the optimal trade-off between privacy and accuracy, accommodating heterogeneity in the number of individuals, measurements, and privacy budgets across servers. Third, we extend lower bound techniques by introducing innovative data processing methods that address information bottlenecks arising from server heterogeneity. Lastly, we quantify the fundamental limits of private functional mean estimation across various distributed environ-

ments, offering practical insights for privacy-preserving statistical analysis of functional data under diverse conditions, from local to central DP frameworks.

1.1 Problem Formulation

We consider the problem of distributed functional mean estimation under privacy constraints in a setting where data is distributed across multiple servers. Let $N = \sum_{s=1}^S n_s$ denote the total number of observations, with server $s = 1, \dots, S$ holding n_s independent observations.

Let $X(\cdot)$ be a random function defined on $[0, 1]$, and for $s = 1, \dots, S$, let $X_1^{(s)}, \dots, X_{n_s}^{(s)}$ be independent copies of X . The goal is to estimate the mean function $f(\cdot) = \mathbb{E}(X(\cdot))$ based on noisy observations from discrete locations on these curves distributed across S servers:

$$Y_{ij}^{(s)} = X_i(\zeta_{ij}^{(s)}) + \xi_{ij}^{(s)}, i = 1, \dots, n_s \text{ and } j = 1, \dots, m_s,$$

where $\zeta_{ij}^{(s)}$ are the sampling points and $\xi_{ij}^{(s)} \sim N(0, \sigma^2)$ (without loss of generality we assume $\sigma = 1$). The sample path of X is assumed to be smooth in that it belongs to $\mathcal{H}^\alpha(R)$ which also implies that the mean $f \in \mathcal{H}^\alpha(R)$ (where $\mathbb{E}X_i(\cdot) = f(\cdot)$). Here, $\mathcal{H}^\alpha(R)$ denotes the Hölder ball of radius R , within the class of α -Hölder continuous functions. In the classical setting without privacy constraints and distributed data, this problem has been extensively studied in the statistical literature (e.g., [16, 9]). In this article, we focus on quantifying the cost of privacy in a distributed setup for this problem.

So far, we have not discussed the design of the experiment. We will consider two cases: the common design and the independent design. In the *common design* case, all individuals are observed on the same locations (grid points) i.e all $s \in [S]$ and $i \in [n_s]$, $\zeta_{ij}^{(s)} = \zeta_j$ for all $j \in [m_s]$. Note that for the common design the design points are not considered to be private since it is shared among all the individuals. In the *independent design* case, the grid points are random and distinct for each individual. In this case, the design points are considered to be private. This reflects real-world situations where individuals could be identified based on specific measurement time points. We make the simplifying assumption that they are i.i.d. draws from a uniform distribution on $[0, 1]$. This can be relaxed to non-uniform distributions that are e.g. bounded above and below by a constant, and the estimator presented in this paper can be easily adapted to non-uniform designs using the privacy preserving technique presented in the Supplementary Material of [8].

To account for privacy constraints, we adopt the general framework of distributed estimation under privacy constraints introduced in [8]. Let $Z^{(s)} = \{Z_i^{(s)}\}_{i=1}^{n_s}$ denote the dataset held by server s , with $Z_i^{(s)} = \{Y_{ij}^{(s)}, \zeta_{ij}^{(s)}\}_{j=1}^{m_s}$. Each server s outputs a (randomized) transcript $T^{(s)}$ based on $Z^{(s)}$, with the law of the transcript determined conditionally on $Z^{(s)}$, i.e., $\mathbb{P}(\cdot|Z^{(s)})$. The collection of transcripts $T = (T^{(s)})_{s=1}^S$ must satisfy $(\epsilon, \delta) = (\epsilon_s, \delta_s)_{s=1}^S$ -federated differential privacy (FDP), defined as follows:

Definition 1. The transcript $T = (T^{(s)})_{s=1}^S$ is (ϵ, δ) -federated differentially private (FDP) if for all $s \in [S]$, $A \in \mathcal{T}$ and $z, z' \in \mathcal{Z}^{n_j}$ differing in one individual datum it holds that

$$\mathbb{P}\left(T^{(s)} \in A | Z^{(s)} = z\right) \leq e^{\epsilon_s} \mathbb{P}\left(T^{(s)} \in A | Z^{(s)} = z'\right) + \delta_s.$$

In the above definition, “differing in one datum” refers to being Hamming distance “neighbors.” Specifically, local datasets $Z^{(s)}$ and $\tilde{Z}^{(s)}$ are *neighboring* if their Hamming distance is at most 1, calculated over $\mathcal{Z}^{n_s} \times \mathcal{Z}^{n_s}$. In other words, $\tilde{Z}^{(s)}$ can be derived from $Z^{(s)}$ by modifying at most one observation among $Z_1^{(s)}, \dots, Z_{n_s}^{(s)}$. In the common design setup, where the design points are shared among all individuals, neighboring datasets share the same design points, and only the measurements Y differ. The smaller the values of ϵ_s and δ_s , the stricter the privacy constraint. We consider $\epsilon_s \leq C_\epsilon$ for $s = 1, \dots, S$, with a fixed constant $C_\epsilon > 0$ that does not affect the derived rates.

The Federated Differential Privacy (FDP) framework addresses scenarios where sensitive data is distributed across multiple parties, each generating outputs while maintaining differential privacy. In this distributed protocol, each server’s transcript depends solely on its local data, with no direct communication or data exchange between servers. This framework is particularly relevant in settings like multi-site studies or trials conducted on the same population, where institutions (e.g., hospitals) aim to collaborate without sharing raw data due to privacy concerns. The FDP framework generalizes commonly studied settings, including the local DP model ($n_s = 1$), where privacy mechanisms operate at the individual level, and the central DP model ($S = 1$), where data is aggregated at a single location.

Each server transmits its transcript to a central server. Using all transcripts $T = (T^{(1)}, \dots, T^{(S)})$, the central server computes an estimator $\hat{f} : \mathcal{T}^S \rightarrow \mathcal{F}$. We refer to the pair $(\hat{f}, \{\mathbb{P}(\cdot | Z^{(s)})\}_{s=1}^S)$ as a distributed estimation protocol. The class of distributed estimation protocols satisfying Definition 1 is denoted by $\mathcal{M}(\epsilon, \delta)$. We let $\mathbb{P}_f$ denote the joint law of transcripts and the $N = \sum_{s=1}^S n_s$ observations. We let $\mathbb{E}_f$ denote the expectation corresponding to $\mathbb{P}_f$.

The aim is to estimate the function f based on the distributed data. The difficulty of this estimation task arises from both the distributed nature of the data and privacy constraints that limit the sharing of information between servers. As in the conventional decision-theoretical framework, the estimation accuracy of a distributed estimator $\hat{f} \equiv \hat{f}(T)$ is measured by the integrated mean squared error (IMSE), $\mathbb{E}_f \|\hat{f} - f\|_2^2$, where the expectation is taken over the randomness in both the data and construction of the transcripts, and the quantity of particular interest is the *minimax risk* for the distributed private protocols over function class $\mathcal{F}$ which we take to be the Hölder class $\mathcal{H}^\alpha(R)$ with $\alpha > 1/2$,

$$\inf_{\hat{f} \in \mathcal{M}(\epsilon, \delta)} \sup_{f \in \mathcal{F}} \mathbb{E}_f \|\hat{f} - f\|_2^2. \quad (1)$$

This ‘global’ risk characterizes the difficulty of the federated learning problem over the function class when trying to infer the entire function underlying the data whilst adhering to the

heterogeneous privacy constraints.

1.2 Related work

Federated Differential Privacy (FDP), as considered in this paper, applies DP at the level of local samples consisting of multiple observations. Optimal statistical inference under this framework has been studied in the context of nonparametric models for estimation [8] and goodness-of-fit testing [7], with [8] addressing server heterogeneity. The homogeneous setting has been explored for discrete distributions [24, 3] and parametric mean estimation [23, 26]. Additionally, [10] examines discrete distribution testing in a two-server setting ($m = 2$) with differing DP constraints.

Upon completion of this work, we came across the paper [33], which considers a setting similar to ours and claims some partially overlapping results. There are important differences between the papers, however. Foremost, both their general upper- and lower-bounds in [33] exhibit a polynomial gap and our lower bound matches our upper bound up to polylogarithmic factors. Another important difference is that in [33], only the random design case is considered, which means the contrast with the fixed design setup are not treated in their work. Furthermore, the methods considered in the two papers differ: the method in [33] requires a sequential federated learning algorithm, whereas our method is a one-shot algorithm.

In terms of lower bound techniques, several approaches have been specifically developed for private settings. [22] and [4] explore private adaptations of general methods, such as Fano, Assouad, and Le Cam techniques, for establishing lower bounds in the central DP setting. In contrast, [5] develops Van Trees-based lower bounds tailored to local DP, which do not directly extend to central DP. Similarly, the techniques in [1, 2] are designed for the local DP setting and lack straightforward extensions to central DP.

Building on these foundations, [8] extends the Van Trees-based lower bound techniques of [5] to the federated DP setting, addressing the complexities introduced by distributed data and federated privacy constraints. In this paper, we further advance these methodologies by applying them to the functional mean estimation problem. We introduce novel data processing techniques that account for the diverse information bottlenecks arising from user, measurement, and server heterogeneity, enabling the derivation of optimal lower bounds for heterogeneous configurations.

1.3 Organization of the paper

The rest of the paper is organized as follows. In Section 2, we present our main results, which summarize the fundamental performance limits of private functional mean estimation in the common and independent design settings. In Section 3, we describe the methods for differentially private functional mean estimation in the common and independent design settings. In Section 4, we present the lower bound for the common design setting. Proofs of the main results are provided in the Supplementary Material.

1.4 Notation, definitions and assumptions

Throughout the paper, we shall write $N := \sum_{s=1}^S m_s n_s$ and consider asymptotics in S , m_s , the n_s 's and the privacy budget $(\epsilon, \delta) := \{\epsilon_s, \delta_s\}_{s=1}^S$. We assume that $N \rightarrow \infty$ and $\max_s \epsilon_s = O(1)$. For two positive sequences a_k, b_k we write $a_k \lesssim b_k$ if the inequality $a_k \leq C b_k$ holds for some universal positive constant C . Similarly, we write $a_k \asymp b_k$ if $a_k \lesssim b_k$ and $b_k \lesssim a_k$ hold simultaneously and let $a_k \ll b_k$ denote that $a_k/b_k = o(1)$.

We use $a \vee b$ and $a \wedge b$ for the maximum and minimum, respectively, between a and b . For $k \in \mathbb{N}$, $[k]$ shall denote the set $\{1, \dots, k\}$. Throughout the paper c and C denote universal constants whose value can differ from line to line. The Euclidean norm of a vector $v \in \mathbb{R}^d$ is denoted by $\|v\|_2$. For a matrix $M \in \mathbb{R}^{d \times d}$, the norm $M \mapsto \|M\|$ is the spectral norm and $\text{Tr}(M)$ is its trace. Furthermore, we let I_d denote the $d \times d$ identity matrix.

2 Main results

In most practical settings, different servers often contain differing numbers of samples in terms of individuals, measurements per individual, and privacy parameters. We consider the most general setting where each server $s = 1, \dots, S$ can have a different number of individuals n_s , a different number of measurements m_s for each individual in the server, and varying privacy parameters ϵ_s and δ_s . In the common design setting, our theory allows for heterogeneity in all of the aforementioned characteristics except for the design points.

Our main results are summarized in Theorem 1 and 2 below, which capture the fundamental performance limits of private functional mean estimation in the common and independent design settings, respectively.

2.1 Optimal performance when the design is common

In this section, we consider the case of common design, where $m_s = m$ for all $s = 1, \dots, S$. To maintain simplicity of the theoretical statement here, we will assume the design to be a equispaced for the main result of this section; $\zeta_j = j/m$. We note that our upper bounds for the common design setting hold for more general design point setups (see Section 3.2).

Theorem 1 below describes the minimax risk for differentially private estimators in the common design setting. The rate-determining quantity in this setting is governed by the parameter $D_* \geq 1$, which is the solution to the following equation:

$$D^{2\alpha} = m^{2\alpha} \bigwedge \sum_{s=1}^S \min \{n_s, D^{-1} n_s^2 \epsilon_s^2\}. \quad (2)$$

The existence and uniqueness of the solution $D_* > 0$ to equation (2) can be established under standard assumptions on the parameters m, n_s, ϵ_s , and α . Intuitively, D_* represents the balance between the design complexity, captured by $m^{2\alpha}$, and the effective sample size across all samples.

The left-hand side of the equation, $D^{2\alpha}$, increases monotonically with D , while the right-hand side combines terms that are bounded and decrease monotonically with D_* , guaranteeing a unique solution.

Theorem 1. *Suppose that for $\delta' = \min_s \delta_s$, it holds that $\delta' \log(1/\delta') \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2/N$. Let D_* solve (2) then,*

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \asymp \eta_N D_*^{-2\alpha}.$$

where η_N is a polylogarithmic factor in $N = \sum_{s=1}^S n_s$ and $1/\delta'$.

It is instructive to consider two specific limiting cases to better understand the implications of this result. First, in the *non-private case*, where the privacy constraints are removed by setting $\varepsilon_s \rightarrow \infty$ for all $s \in [S]$, the rate simplifies to the known minimax rates established in [9]. This scenario corresponds to the classical non-private setting where privacy considerations do not impose any restrictions. The effective rate is determined solely by the design complexity and sample sizes, without the additional penalties introduced by privacy constraints. Second, in the *large-sample limit*, where $n_s \rightarrow \infty$ for all $s \in [S]$, the rate reduces to $m^{-2\alpha}$. Here, the noise becomes negligible due to the abundance of samples, and the discretization error from the common design dominates. This reflects a regime where the design complexity m dictates the achievable performance, independent of the privacy constraints or noise levels. In order to gain more fine grained insight it would be helpful to look at the homogeneous setting which we do section 2.3.

2.2 Optimal performance when the design is independent

The convergence rate in the independent design setting is, just as in the common design setting, given by the solution to an equation. Let $D_* \geq 1$ be the largest number such that

$$D_*^{2\alpha} = \inf_{1 \leq D \leq D_*} \sum_{s=1}^S \min \{D^{-1} n_s m_s, D^{-2} m_s n_s^2 \varepsilon_s^2, D^{2\alpha} n_s, D^{2\alpha-1} n_s^2 \varepsilon_s^2\}. \quad (3)$$

The rate of convergence for differentially private estimators in the independent design setting is given by the following theorem.

Theorem 2. *Suppose that for $\delta' = \min_s \delta_s$, it holds that $\delta' \log(1/\delta') \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2/N$. Let D_* solve (3).*

Then,

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \asymp \eta_N D_*^{-2\alpha}.$$

where η_N is a polylogarithmic factor in $N = \sum_{s=1}^S n_s$ and $1/\delta'$.

The optimization problem in (3) can be interpreted minimum the four terms, which correspond to the sampling error, the measurement error, and their respective private counterparts.

Our upper bounds rely on solving the optimization problem (3), which determines, roughly speaking, the effective dimension of the problem. The effective dimension D_* balances the trade-offs between design complexity, sample sizes, noise levels, and the smoothness of the underlying function class.

In the *non-private case*, where the privacy constraints are removed by setting $\varepsilon_s \rightarrow \infty$ for all $s \in [S]$, the rate simplifies to the classical non-private minimax rate. Notably, minimax optimal rates in the heterogeneous setup, where design points differ across samples, were not well understood prior to this work. The optimization problem in this case reveals how differences in sample sizes and design complexities across datasets interact with the smoothness of the underlying function to determine the rate.

2.3 Comparison of the common and independent design settings

In this section, we compare the minimax rates of convergence for the common and independent design settings under the homogeneous case, where $n_s = n$, $m_s = m$, $\varepsilon_s = \varepsilon$, and $\delta_s = \delta$. While in practice, settings are more often heterogeneous than not, the homogenous case allows for a more straightforward formulation minimax rate than the implicit ones of Theorems 1 and 2. In this simplified formulation, the rates between the independent and common designs are more easily compared. The comparison highlights the fundamental differences in statistical efficiency and the impact of differential privacy constraints across these two design settings.

We begin by stating the minimax risks for the homogeneous setup under the two design settings, starting with the independent design setting.

Corollary 1. *Suppose that it holds that $\delta \log(1/\delta) \lesssim (\frac{n}{m} \wedge \sqrt{\frac{n}{m}}) \varepsilon^2/N$.*

It holds that

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \asymp \eta_N \left((Sn)^{-1} + (Smn)^{-\frac{2\alpha}{2\alpha+1}} + (Smn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+2}} + (Sn^2\varepsilon^2)^{-1} \right), \quad (4)$$

where η_N is a polylogarithmic factor in $N = Sn$ and $1/\delta$.

For the common design setting, the minimax rate of convergence in the homogeneous setup is given by the following corollary.

Corollary 2. *Suppose that it holds that $\delta \log(1/\delta) \lesssim (\frac{n}{m} \wedge \sqrt{\frac{n}{m}}) \varepsilon^2/N$.*

It holds that

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \asymp \eta_N \left((Sn)^{-1} + m^{-2\alpha} + (Sn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+1}} \right), \quad (5)$$

where η_N is a polylogarithmic factor in $N = Sn$ and $1/\delta$.

In the common design case, the rate of convergence in (5) is dominated by the privacy-related term $(Sn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+1}}$ when $\varepsilon \lesssim \min(S^{-1/2}n^{-1/2}m^{\alpha+1/2}, S^{1/4\alpha}n^{(1-2\alpha)/4\alpha})$. In contrast, under the

independent design setting, the rate is given by the combination $(Smn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+2}} + (Sn^2\varepsilon^2)^{-1}$ when privacy constraints dominate.

If the rate is dominated by the second term $(Sn^2\varepsilon^2)^{-1}$, it is strictly faster than the common design privacy rate. If the first term $(Smn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+2}}$ dominates, it is faster than the common design rate whenever

$$(Smn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+2}} \lesssim (Sn^2\varepsilon^2)^{-\frac{2\alpha}{2\alpha+1}} \quad \text{iff} \quad \varepsilon \lesssim S^{-1/2}n^{-1/2}m^{\alpha+1/2}.$$

Thus, whenever the “privacy-related” terms dominate, an independent design achieves a faster rate of convergence, meaning the cost of privacy is consistently lower in this setting.

What is particularly noteworthy is that under the independent design, when privacy constraints are binding, there are benefits to having both more individuals and more measurements per individual. In contrast, under the common design, increasing the number of measurements per individual offers no benefit when the privacy constraint is binding. This can be interpreted as a privacy-related penalty for the common design. This might seem counterintuitive at first, since in the common design setting, the design points do not need to be private, whereas in the independent design setting, they do. The key insight here is about information leakage: In the common design where all individuals share the same measurement points, adding more measurements doesn’t help because the privacy cost is already maximized - everyone’s data is perfectly correlated through the shared design points. In contrast, with independent designs, while more measurements do leak more information about each individual (requiring more noise for privacy), this is offset by the statistical benefits of having more diverse measurement points across the population.

3 Methods for Differentially Private Functional Mean Estimation

We now describe the methods for differentially private functional mean estimation in the independent and common design settings. As revealed by Theorems 1 and 2, these settings exhibit distinct principal phenomena and require different approaches. We present an optimal algorithm for each setting, both of which provably achieve the minimax rate of convergence. The theorems accompanying the algorithms in this section provide the upper bounds for the respective theorems mentioned earlier, starting with the independent design case.

3.1 An optimal algorithm for the independent design setting

In the independent design setting, we consider a scenario where all design points are distinct and treated as private information. The challenge in this setup, compared to e.g. private nonparametric regression as studied in [8], stems from the fact that the model itself is effectively subject to three different sources of noise: the per-subject $X_i^{(s)}$ -function noise, the noise stemming from

the randomness in the design points, and the measurement noise. An optimal privacy mechanism is calibrated to account for all three sources of noise, ensuring that the privacy budget for each server is allocated efficiently across the different components. This means that the procedure must differ from server to server, depending on the number of individuals, measurements per individual, and privacy parameters.

Our approach is based on a private, truncated projection estimator, where we project onto compactly supported orthonormal basis functions. This approach allows effective privatization of both the design points and the outcome values. As basis functions, we use $A > \alpha$ -regular wavelets, which form an orthonormal basis of $L_2[0, 1]$. Wavelet bases allow characterization of Hölder spaces, with α capturing the decay of wavelet coefficients. For further details, see e.g. [19].

For any $A \in \mathbb{N}$ one can follow Daubechies' construction of the father $\phi(\cdot)$ and mother $\psi(\cdot)$ wavelets with A vanishing moments and bounded support on $[0, 2A - 1]$ and $[-A + 1, A]$, respectively, for which we refer to [12]. The basis functions are then obtained as

$$\{\phi_{l_0+1,m}, \psi_{lk} : m \in \{0, \dots, 2^{l_0+1} - 1\}, \quad l \geq l_0 + 1, \quad k \in \{0, \dots, 2^l - 1\}\},$$

with $\psi_{lk}(x) = 2^{l/2}\psi(2^l x - k)$, for $k \in [A - 1, 2^l - A]$, and $\phi_{l_0+1,k}(x) = 2^{l_0+1}\phi(2^{l_0+1}x - m)$, for $m \in [0, 2^{l_0+1} - 2A]$, while for other values of k and m , the functions are specially constructed, to form a basis with the required smoothness property. A function $f \in L_2[0, 1]$ can then be expressed as

$$f = \sum_{l=l_0}^{\infty} \sum_{k=0}^{2^l-1} f_{lk} \psi_{lk},$$

where $f_{lk} = \int f \psi_{lk}$. A well known estimator for the wavelet coefficients f_{lk} – in the absence of privacy constraints – is given by

$$\hat{f}_{lk}^{(s)} = \frac{1}{n_s} \sum_{i=1}^{n_s} \frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}),$$

where $Y_{ij}^{(s)}$ are the observations and $\zeta_{ij}^{(s)}$ are the corresponding design points.

For notational simplicity, let us define the contribution of a single individual $i \in [n_s]$, on server $s \in [S]$ to the estimator as:

$$U_{i,lk}^{(s)} = \frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}).$$

Thus, the standard (unclipped) estimator can be rewritten as: $\hat{f}_{lk}^{(s)} = \frac{1}{n_s} \sum_{i=1}^{n_s} U_{i,lk}^{(s)}$. In order to make an informed decision on how to clip requires sharp concentration bounds on $U_{i,lk}^{(s)}$. The following lemma establishes the concentration of $U_{i,lk}^{(s)}$ for a fixed l and k , using Bernstein's inequality.

Lemma 1. Fix $c > 0$, l and k , the sequence $\{U_{i,lk}^{(s)}\}_{i \in [n_s]}$ satisfies:

$$\mathbb{P} \left(\forall i, |U_{i,lk}^{(s)}| < \tau_l^{(s)} \right) \geq 1 - \gamma,$$

with the threshold $\tau_l^{(s)}$ given by

$$\tau_l^{(s)} = 2(2c \log N)^{3/2} \left(\sqrt{\frac{1}{m_s}} + \frac{1}{3} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right) + R2^{-l(\alpha+1/2)}, \quad (6)$$

and $\gamma = N^{-c}$, where the constant $c > 0$ can be arbitrary.

A proof of the lemma is given in Section 6.1.2 of the Supplementary Material. We define what is effectively the local clipped estimator for the l, k -th wavelet coefficient as

$$\hat{f}_{lk}^{\tau,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} \left[U_{i,lk}^{(s)} \right]_{\tau_l^{(s)}}, \quad (7)$$

where $\left[U_{i,lk}^{(s)} \right]_{\tau_l^{(s)}}$ denotes the clipping of $U_{i,lk}^{(s)}$ to the range $[-\tau_l^{(s)}, \tau_l^{(s)}]$. Clipping at this magnitude ensures that the estimator has bounded sensitivity, enabling it to be privatized effectively.

To achieve differential privacy, we aim to employ the Gaussian mechanism. However, adding Gaussian noise to the clipped averages of (7) directly does not balance the amount of signal versus sensitivity present in each wavelet coefficient. Instead, we employ a modified Gaussian mechanism, akin to ideas used in anisotropic private mean estimation of e.g. [11].

The coordinates of $\hat{f}_{lk}^{\tau}$ exhibit varying greatly varying sensitivities across the various resolution levels, proportional to $\tau_l^{(s)}$. In order to obtain the right signal-sensitivity trade-off, we rescale each coordinate, yielding at the l -th resolution level:

$$\tilde{f}_{lk}^{\tau,s} = (\tau_l^{(s)} \sqrt{2^l \wedge m_s})^{-1} \hat{f}_{lk}^{\tau,s}.$$

The ℓ_2 sensitivity of the vector $\{\tilde{f}_{lk}^{\tau,s}, k = 0, \dots, 2^l - 1, l_0 \leq l \leq L\}$ is bounded as shown in Lemma 2.

Lemma 2. Let $Z^{(j)}$ and $\tilde{Z}^{(j)}$ be any neighboring datasets. Let $\tilde{f}^{\tau,s}(Z^{(j)})$, $\tilde{f}^{\tau,s}(\tilde{Z}^{(j)})$ denote the vectors of coefficient estimators $\{\tilde{f}_{lk}^{\tau,s}, k = 0, \dots, 2^l - 1, l_0 \leq l \leq L\}$ computed on the basis of $Z^{(j)}$ and $\tilde{Z}^{(j)}$ respectively. Then:

$$\left\| \tilde{f}^{\tau,s}(Z^{(j)}) - \tilde{f}^{\tau,s}(\tilde{Z}^{(j)}) \right\|_2 \leq c_A \frac{\sqrt{L}}{n_s},$$

where c_A is a computable constant depending only the choice of wavelet basis.

Using the Gaussian mechanism (for a details, see Appendix A of [15]), we obtain a $(\varepsilon_s, \delta_s)$ -differentially private estimator

$$\tilde{f}_{lk}^{P,s} = \tilde{f}_{lk}^{\tau,s} + \widetilde{W}_{lk}^{(s)}, \quad \widetilde{W}_{lk}^{(s)} \stackrel{i.i.d.}{\sim} \mathcal{N} \left(0, \frac{4c_A^2 L \log(2/\delta)}{n_s^2 \varepsilon_s^2} \right).$$

To construct the final private estimator $\widehat{f}_{lk}^{P,s}$ in the central server, we rescale $\tilde{f}_{lk}^{P,s}$ back to the original scale, yielding

$$\widehat{f}_{lk}^{P,s} = \widehat{f}_{lk}^{\tau,s} + W_{lk}^{(s)}, \quad W_{lk}^{(s)} \stackrel{i.i.d.}{\sim} \mathcal{N}\left(0, \frac{4c_A^2 L(2^l \wedge m_s)(\tau_l^{(s)})^2 \log(2/\delta_s)}{n_s^2 \varepsilon_s^2}\right).$$

Each server s then outputs a transcript

$$T_L^{(s)} = \{\widehat{f}_{lk}^{P,s} : k = 0, \dots, 2^l - 1, l = l_0, \dots, L\}.$$

The final estimator $\widehat{f}^P$ is obtained via a post-processing step, where the transcripts are aggregated using weights that account for the heterogeneity of the servers. These weights depend on the resolution level l , the number of local observations n_s , the number of design points m_s , the local privacy constraints ε_s , and the smoothness level α . Specifically, the final estimator is:

$$\widehat{f}_L^P = \sum_{l=l_0}^L \sum_{k=0}^{2^l-1} \widehat{f}_{lk}^P \psi_{lk} \quad \text{where} \quad \widehat{f}_{lk}^P = \sum_{s=1}^S w_l^s \widehat{f}_{lk}^{P,s} \quad (8)$$

with $\sum_{s=1}^S w_l^{(s)} = 1$ for all $l \geq l_0$. The weights $w_l^{(s)}$ are chosen to be inversely proportional to the variance of $\widehat{f}_{lk}^{P,s}$:

$$w_l^{(s)} = \frac{u_l^{(s)}}{\sum_{s=1}^S u_l^{(s)}}, \quad \text{where} \quad u_l^{(s)} = n_s m_s \wedge m_s 2^{l(2\alpha+1)} \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2l\alpha} n_s^2 \varepsilon_s^2. \quad (9)$$

Remark 1. The weights $w_l^{(s)}$ account for the heterogeneity in data distribution across servers, including variations in the number of observations (n_s), the number of design points (m_s), and the privacy budgets (ε_s). By being inversely proportional to the variance of $\widehat{f}_{lk}^{P,s}$, the weights prioritize servers that contribute more reliable estimates. The structure of $u_l^{(s)}$ incorporates terms that reflect the impact of resolution level l , smoothness parameter α , and privacy constraints.

We summarize our final estimator in Algorithm 1.

Theorem 3. Let $\widehat{f}_L^P$ be the estimator defined in Algorithm 1. Let D^* satisfy (3) and let $L = \lceil \log_2 D^* \rceil$ we obtain the following upper bound on the minimax rate,

$$\sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\widehat{f}_L^P - f\|_2^2 \lesssim (\log N)^5 \log(2/\delta') (D^*)^{-2\alpha}.$$

where $\delta' = \min_s \delta_s$.

Remark 2. The parameter D^* , which satisfies the rate-determining relation (3), can be interpreted as the effective dimension of the problem. It reflects the interplay between the resolution level L and the sample size $N = \sum_{s=1}^S n_s$. Larger D^* corresponds to higher resolution levels and greater flexibility in representing the function f , while also increasing the complexity of the estimation task. The optimal choice of D^* balances this trade-off to achieve the minimax rate.

Algorithm 1 Private Estimation via Truncated Projection in Independent Design

1: **Input:** Observations $Y_{ij}^{(s)}$, design points $\zeta_{ij}^{(s)}$, wavelet basis ψ_{lk} , truncation level L , privacy budgets ε_s, δ_s

2: **Output:** Final estimator $\hat{f}_L^P$

3: **Clipped Coefficients:** Each server computes:

$$\hat{f}_{lk}^{\tau,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} \left[\frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}) \right]_{\tau_l^{(s)}} \quad \forall l = l_0, \dots, L, k = 0, \dots, 2^l - 1,$$

where $\tau_l^{(s)}$ is given by (6).

4: **Privatized Coefficients:** Apply the modified Gaussian mechanism:

$$\hat{f}_{lk}^{P,s} = \hat{f}_{lk}^{\tau,s} + W_{lk}^{(s)}, \quad W_{lk}^{(s)} \stackrel{i.i.d.}{\sim} N \left(0, \frac{4c_A^2 L(2^l \wedge m)(\tau_l^{(s)})^2 \log(2/\delta)}{n_s^2 \varepsilon_s^2} \right).$$

5: **Weighted Aggregation:** Aggregate across servers:

$$\hat{f}_{lk}^P = \sum_{s=1}^S w_l^{(s)} \hat{f}_{lk}^{P,s}, \quad w_l^{(s)} = \frac{u_l^{(s)}}{\sum_{s=1}^S u_l^{(s)}}, \quad u_l^{(s)} = n_s m_s \wedge m_s 2^{l(2\alpha+1)} \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2l\alpha} n_s^2 \varepsilon_s^2.$$

6: **Final Estimator:** Combine wavelet coefficients:

$$\hat{f}_L^P(x) = \sum_{l=l_0}^L \sum_{k=0}^{2^l-1} \hat{f}_{lk}^P \psi_{lk}(x).$$

Remark 3. In the absence of privacy constraints ($\varepsilon_s \rightarrow \infty$), the result in Theorem 3 obtains the minimax rates for the setting where different individuals may have different numbers of design points m_s . To the best of our knowledge, these rates were not previously established in the literature.

3.2 An optimal algorithm for the common design setting

In the common design setup, the design points $\zeta_1, \zeta_2, \dots, \zeta_m$ are shared among all individuals and servers. These points are assumed to be publicly known, requiring no privatization. Here, the challenge lies in the fact that the design points are shared across all servers, which means that an optimal private estimator must balance the discretization error due to the shared design points with the privacy leakage.

Our estimation strategy follows two main steps:

1. First, we obtain privatized averages of observations at each design point, effectively creating

private, noisy estimates of the function at these specific points.

2. Second, if the number of design points is large, we introduce a bagging step that divides the design points into groups.
3. Third, these privatized averages present an inverse problem at the central server, which we solve using local polynomial regression. This regression effectively interpolates across the design points.

Private averages at the design points of first step effectively pose an inverse problem at the central server; solved by local polynomial regression in the third step. The noise in this inverse problem at the central server is does not only depend on measurement noise, but also on the noise added by the privacy mechanism. If the number of design points is large, the noise added by the privacy mechanism needs to increase to avoid the privacy leakage due to revealing more measurements per individual.

The second step employing bagging mitigates this by dividing the design points into groups and computes smoothed estimators that interpolate across groups. Intuitively, the purpose of the bagging step is used to effectively reduce the number of measurements m , in the scenario where the measurements reveal too much information about the individuals if they were all communicated to the central server. In terms of attaining the minimax rate, the bagging step is superfluous, in the sense that one could also effectively use a smaller number of measurements by taking a selection of design points. Bagging, however, is a more practical approach, as it a straightforward implementation that is more robust to the choice of design points.

Next, we describe the algorithm for the common design formally.

Privatized averages at the design points Given S servers and a set of observations from a function f at deterministic points $\zeta_1, \zeta_2, \dots, \zeta_m$, the privatized mean $\bar{Y}_j^{P,s}$ at design point ζ_j on server s is computed as:

$$\bar{Y}_j^{P,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} \left[Y_{i,j}^{(s)} \right]_{\tau} + W_j^{(s)},$$

where $\left[Y_{i,j}^{(s)} \right]_{\tau}$ represents the clipped observations with clipping threshold $\tau = \sqrt{2 \log N} + C_{\alpha,R}$ ($C_{\alpha,R}$ is as defined in Lemma 8), and $W_j^{(s)} \sim \mathcal{N}\left(0, \frac{4\tau^2 m \log(2/\delta_s)}{n_s^2 \varepsilon_s^2}\right)$ is a Gaussian noise term ensuring differential privacy. Here, N is the total sample size across all servers, n_s is the sample size on server s , and ε_s is the privacy budget for server s .

Each server outputs a transcript $T^{(s)} = (\bar{Y}_j^{P,s})_{j \in [m]}$. The central server aggregates these transcripts using inverse variance weighting:

$$\bar{Y}_j^P = \sum_{s=1}^S w_s \bar{Y}_j^{P,s},$$

where the weights w_s are designed to account for server heterogeneity:

$$w_s = \frac{u_s}{\sum_{s=1}^S u_s}, \quad u_s = D^{-1} n_s^2 \varepsilon_s^2 \wedge n_s.$$

where D solves (2).

Bagging The design points $\zeta_1, \zeta_2, \dots, \zeta_m$ are ordered and divided into B disjoint sets $\{G_b\}_{b \in [B]}$, each containing m_0 points, such that $\lfloor m/m_0 \rfloor \leq B \leq \lceil m/m_0 \rceil$. We choose groups $(G_b)_{b \in [B]}$ of sizes $m_0 = \lceil D \rceil$ where D solves (2). In case of equispaced design (i.e $\zeta_j = j/m$), this amounts to $G_b = \left\{ \frac{B(j-1)+b}{m} : j = 1, \dots, m_0 \right\}$.

For each group G_b , a local polynomial estimator $\hat{f}_b^P(x)$ is computed based on the privatized means $\bar{Y}_j^P$ corresponding to the design points in G_b . Finally, the overall estimator $\hat{f}_B^P(x)$ is obtained by averaging across all groups:

$$\hat{f}_B^P(x) = \frac{1}{B} \sum_{b=1}^B \hat{f}_b^P(x).$$

Function approximation via local polynomials The local polynomial estimator approximates the function $f \in \mathcal{H}^\alpha(R)$ by expanding it around x as:

$$f(z) = f(x) + f'(x)(z-x) + \dots + \frac{f^{(p)}(x)}{p!} (z-x)^p,$$

where $p = \lfloor \alpha \rfloor$. This expansion can be written compactly as:

$$f(z) = \Theta^\top(x) V \left(\frac{z-x}{h} \right),$$

where $V(u) = \left(1, u, \frac{u^2}{2!}, \dots, \frac{u^p}{p!} \right)^\top$ is the vector of scaled monomials, and $\Theta(x) = \left(f(x), f'(x)h, \dots, f^{(p)}(x)h^p \right)^\top$ represents the function's coefficients at x , scaled by powers of the bandwidth h .

To estimate $f(x)$, we solve a weighted least squares problem for each group:

$$\hat{\Theta}_b(x) = \arg \min_{\Theta \in \mathbb{R}^{l+1}} \sum_{j \in G_b} \left[\bar{Y}_j^P - \Theta^\top V \left(\frac{\zeta_j - x}{h} \right) \right]^2 K \left(\frac{\zeta_j - x}{h} \right).$$

where K can be any smooth enough kernel function satisfying

$$\text{supp}(K) \subseteq [-1, 1], \quad \|K\|_\infty < \infty, \quad \text{and} \quad \int K(u) du = 1.$$

The first component of $\hat{\Theta}_b(x)$, corresponding to $\hat{f}_b^P(x)$, is the estimate of $f(x)$ for group b .

The local polynomial estimator $\hat{f}_b^P(x)$ for the b -th group can be shown to be equal to:

$$\hat{f}_b(x) = \sum_{j \in G_b} \bar{Y}_j^P W_{b,j}^*(x),$$

where the weights $W_{b,j}^*(x)$ are determined by the kernel and the design points of the b th group:

$$W_{b,j}^*(x) = \frac{1}{m_0 h} V^\top(0) B_{b,x}^{-1} V \left(\frac{\zeta_j - x}{h} \right) K \left(\frac{\zeta_j - x}{h} \right). \quad (10)$$

The matrix $B_{b,x}$ is the “effective” weighted design matrix for group b induced by K , given by

$$B_{b,x} = \frac{1}{m_0 h} \sum_{j \in G_b} V \left(\frac{\zeta_j - x}{h} \right) V \left(\frac{\zeta_j - x}{h} \right)^\top K \left(\frac{\zeta_j - x}{h} \right).$$

For the subsequent analysis, we rely on the following assumptions, where the choice of kernel and groups allows for additional flexibility. We remark that both assumptions (LP1) and (LP2) hold in particular for regular equi-spaced design under mild conditions on the Kernel K . See for example Lemma 1.5 in [30].

(LP1) $\lambda_{\min}(B_{b,x}) \geq \lambda_0$ for all $b \in [B]$, and $x \in [0, 1]$.

(LP2) For $b \in [B]$, for some constant a , and any measurable set A , we have

$$\frac{1}{m_0} \sum_{j \in G_b} \mathbb{I}(\zeta_j \in A) \leq a \max \left(\text{Leb}(A), \frac{1}{m_0} \right).$$

Assumption (LP1) ensures that the local design points in each group provide sufficient information for local polynomial estimation. Together with (LP2), these assumption ensures that within each group the design points $(\zeta_j)_{j \in G_b}$ are sufficiently spread over the interval $[0, 1]$.

Algorithm 2 summarizes the differentially private protocol. By averaging across all groups in (11), the bagging estimator as $\hat{f}_B^P(x)$ is computed in the central server. The following theorem establishes the convergence rate of this estimator.

Theorem 4. *Let $\hat{f}_B^P$ be the estimator defined in Algorithm 2. We choose groups $(G_b)_{b \in [B]}$ of sizes $m_0 = D$ where D solves (2) and set the bandwidth h to be $1/m_0$. We also assume that for each of the groups $(G_b)_{b \in [B]}$, the design points in the groups $(\zeta_j)_{j \in G_b}$ satisfy (LP1) and (LP2). Then,*

$$\sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \lesssim \log(N) \log(2/\delta') (D^{-2\alpha} \vee m^{-2\alpha}),$$

where D solves (2) $\delta' = \min_s \delta_s$.

Remark 4. The same minimax risk as in Theorem 4 can be achieved by using data from only one of the groups of size m_0 instead of aggregating across all B groups. This is because, when using a single group, the amount of noise added to ensure differential privacy scales with m_0 , which is smaller than or equal to m . Consequently, the effective privacy noise is reduced, and the overall estimation error remains comparable. While using all B groups can improve practical stability/performance by averaging across groups, it does not offer a theoretical advantage in terms of the minimax rate.

Algorithm 2 Private Estimation via Smoothing and Bagging for Common Design

- 1: **Input:** Design points $T_1, \dots, T_m$, S servers, observations $Y_{i,j}^{(s)}$, privacy budgets ε_s , groups $\{G_b\}_{b \in [B]}$ each of size m_0 , bandwidth h , kernel K
- 2: **Output:** Final estimator $\hat{f}_B^P(x)$
- 3: **Privatized Means:** Each server computes privatized means:

$$\bar{Y}_j^{P,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} \left[Y_{i,j}^{(s)} \right]_{\tau} + W_j^{(s)}, \quad W_j^{(s)} \sim \mathcal{N} \left(0, \frac{4\tau^2 m \log(2/\delta_s)}{n_s^2 \varepsilon_s^2} \right),$$

where $\tau = \sqrt{2 \log N} + C_{\alpha,R}$.

- 4: **Aggregate Means:** Central server aggregates:

$$\bar{Y}_j^P = \sum_{s=1}^S w_s \bar{Y}_j^{P,s}, \quad w_s = \frac{u_s}{\sum_{s=1}^S u_s}, \quad u_s = m_0^{-1} n_s^2 \varepsilon_s^2 \wedge n_s.$$

- 5: **Group-wise Estimation:** For each group G_b , compute local polynomial estimator:

$$\hat{f}_b^P(x) = \sum_{j \in G_b} \bar{Y}_j^P W_{b,j}^*(x),$$

where $W_{b,j}^*(x)$ are kernel-based weights given by (10).

- 6: **Final Estimator:** Aggregate across groups:

$$\hat{f}_B^P(x) = \frac{1}{B} \sum_{b=1}^B \hat{f}_b^P(x). \tag{11}$$

4 Lower bounds

As a complement to the upper bounds, we provide lower bounds, which together establish the minimax rates for the federated learning problem; resulting in Theorems 1 and 2. We first present the lower bound for the independent design setting, followed by the lower bound for the common design setting.

4.1 Lower bound for independent design

We first present the lower bound, before discussing the proof.

Theorem 5. Suppose that for $\delta' = \min_s \delta_s$, it holds that $\delta' \log(1/\delta') \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}} \right) \varepsilon_s^2 / N$. Let D_* solve (3). Then,

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon,\delta}} \sup_{f \in \mathcal{H}^{\alpha}(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \gtrsim D_*^{-2\alpha}. \tag{12}$$

Below, we briefly sketch the proof of Theorem 5. The challenge in proving the lower bound stems from the heterogeneity of the servers. Essentially, for each respective server, either the

sample size, the number of design points, or the privacy constraints can form the bottleneck.

The proof involves a careful construction of a data generating process that allows to capture these different bottlenecks for each server. From a technical standpoint, the heterogeneous nature of the servers is captured by considering the Fisher information of the model in an appropriate finite dimensional sub-model. The Fisher information tensorizes along the servers, which allows us provide separate information theoretic arguments for each server.

To provide a more detailed description of the argument, we introduce the aforementioned notions formally. Consider the 2^L -dimensional sub-model given by

$$\mathcal{H}_{R,L}^\alpha : \left\{ f \in \mathcal{H}_\alpha^R : f = \sum_{k=1}^{2^L} f_k \phi_{Lk}, f_k \in [-2^{-L(\alpha+1/2)-1}R, 2^{-L(\alpha+1/2)-1}R] \right\}, \quad (13)$$

where ϕ_k are the wavelet basis functions at resolution level L . Let μ and ν denote dominating measures for $\mathbb{P}_f^{Y^{(s)}}$ and $\mathbb{P}_f^{(Y^{(s)}X^{(s)})}$, respectively, for $f \in \mathcal{H}_{R,L}^\alpha$. Define $f_L = (f_k)_{k=1}^{2^L}$ to be the vector of coefficients of f in the wavelet basis. The transcript $T^{(s)}$ induces a ‘‘Fisher information’’, which can be computed to be equal

$$I_f^{Y^{(s)}|T^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[S_f^{Y^{(s)}} \middle| T^{(s)} \right] \mathbb{E}_f \left[S_f^{Y^{(s)}} \middle| T^{(s)} \right]^\top \quad (14)$$

where $S_f^{Y^{(s)}}$ is given by

$$S_f^{Y^{(s)}} = \sum_{j=1}^m \left(Y_{ij}^{(s)} - \sum_{k=1}^{2^L} f_k \phi_l(\zeta_{ij}^{(s)}) \right) \phi(\zeta_{ij}^{(s)}), \quad (15)$$

which can be seen as a ‘‘score function’’ of the observational model for $Y^{(s)}$ in the sub-model $\mathcal{H}_{R,L}^\alpha$, if $X_i = \mathbb{E}_f X_i$ almost everywhere. Let $I_f^{Y^{(s)}}$ denote the Fisher information of the observational model for $Y^{(s)}$ in the sub-model $\mathcal{H}_{R,L}^\alpha$, which equals $\mathbb{E}_f S_f^{Y^{(s)}} S_f^{Y^{(s)}}^\top$.

When m_s tends to infinity, the observational model for $Y^{(s)}$ effectively contains close to the amount of information as the observational model in which $X^{(s)}$ is directly observed. Following this intuition, the right information quantity to consider for servers with large m_s is the Fisher information of the observational model for $X^{(s)} = (X^{(s)})_{i \in [n]}$. Through data processing arguments, we can pass to the Fisher information of the observational model for $X^{(s)}$ given the transcript in the sub-model $\mathcal{H}_{R,L}^\alpha$. We consider the following dynamics for $X^{(s)}$:

$$X_i^{(s)} = f + R 2^{-L(\alpha+1/2)-1} \sum_{k=1}^{2^L} (1/2 - B_{ki}^{(s)}) \phi_{Lk}, \quad i \in [n], \quad (16)$$

for i.i.d. Beta(2,2) random variables $B_{ki}^{(s)}$. It is easy to see that each $X_i^{(s)}$ takes values in a Hölder ball of radius R and has mean f . Letting $S^{X^{(s)}}$ denote the score along the vector f_L ,

we let $I_f^{X^{(s)}|T^{(s)}}$ denote the Fisher information of the observational model for $X^{(s)}$ given the transcript in the sub-model $\mathcal{H}_{R,L}^\alpha$,

$$I_f^{X^{(s)}|T^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[S_f^{X^{(s)}} \middle| T^{(s)} \right] \mathbb{E}_f \left[S_f^{X^{(s)}} \middle| T^{(s)} \right]^\top. \quad (17)$$

The Fisher information of $X_i^{(s)}$ is defined as $I_f^{X^{(s)}} = \mathbb{E}_f S_f^{X^{(s)}} S_f^{X^{(s)}}^\top$.

The crux of the proof lies in the following lemma, which provides a lower bound on the minimax risk by expressing it in terms of a combination of Fisher information quantities—each capturing a different bottleneck (sample sizes, number of design points, privacy constraints).

Lemma 3. *Let $L \in \mathbb{N}$. The minimax risk $\inf_{\hat{f} \in \mathcal{M}_{\varepsilon,\delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2$ is lower bounded by*

$$\frac{2^{2L}}{\sup_{f \in \mathcal{H}_{R,L}^\alpha} \sum_{s=1}^S \min \left\{ \text{Tr}(I_f^{Y^{(s)}|T^{(s)}}), \text{Tr}(I_f^{Y^{(s)}}), \text{Tr}(I_f^{X^{(s)}|T^{(s)}}), \text{Tr}(I_f^{X^{(s)}}) \right\} + 2^{2L(\alpha+2)}}. \quad (18)$$

It can be seen as a generalization of the technique introduced by [8], which allows us to account for the situation where the observational model for $Y^{(s)}$ leads to close to equal performance for estimating f as one would have if $X^{(s)}$ were to be observed directly. Its proof relies on a multivariate version of the von Trees inequality of [17] and data processing arguments exploiting the Markov chain structure $f \rightarrow X^{(s)} \rightarrow Y^{(s)}$ of the model.

The minimum of the summand in (18) effectively correspond to four different scenarios that can arise for each server, depending on which information theoretic property forms the bottleneck.

- (i) In the first scenario we consider, both the privacy constraint, the limited number of observations as well the limited number of measurements affect the information that the transcript $T^{(s)}$ contains on f . This corresponds with the quantity $I_f^{Y^{(s)}|T^{(s)}}$, which we in turn bound by a constant multiple of $m_s n_s^2 \varepsilon_s^2$.
- (ii) In the second scenario, the privacy constraint is relatively lenient for the s -th server, in which case the Fisher information of its transcript satisfies $I_f^{Y^{(s)}|T^{(s)}} \asymp I_f^{Y^{(s)}}$, the latter of which can be computed to equal $2^L m_s n_s$.
- (iii) In the scenario that the s -th server contains relatively many measurements per individual, the sample size n_s and privacy constraint form the bottleneck for the s -th server. In this case, $I_f^{X^{(s)}|T^{(s)}}$ is the relevant quantity to consider, which we bound by a constant multiple of $2^{L(2\alpha+1)} n_s^2 \varepsilon_s^2$ using data processing arguments.
- (iv) In the last scenario, both measurements and privacy budget are no longer part of the bottleneck, and just the (number of) X_i observations locally determine information transmitted by machine s . In this case, the $I_f^{X^{(s)}|T^{(s)}} \asymp I_f^{X^{(s)}} \leq 2^{2L(\alpha+1)} n_s$.

The bounds on the traces of the various Fisher information are given in the Lemma 4 below, for which we give a proof in Section 8.1 of the Supplementary Material.

Lemma 4. *Suppose that for $\delta' = \min_s \delta_s$, it holds that $\delta' \log(1/\delta') \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2/N$.*

Then,

$$\min \left\{ \text{Tr}(I_f^{Y^{(s)}|T^{(s)}}), \text{Tr}(I_f^{Y^{(s)}}), \text{Tr}(I_f^{X^{(s)}|T^{(s)}}), \text{Tr}(I_f^{X^{(s)}}) \right\} \lesssim \min \left\{ m_s n_s^2 \varepsilon_s^2, 2^L m_s n_s, 2^{L(2\alpha+1)} n_s^2 \varepsilon_s^2, 2^{2L(\alpha+1)} n_s \right\}.$$

The proof of Theorem 5 follows a choice of L as a function of D_* , where D_* is the solution to equation (3).

4.2 Lower bound for common design

For the common design setting, we prove the following lower bound on the minimax risk.

Theorem 6. *Assume that, for $\delta' = \min_s \delta_s$, we have $\delta' \log(1/\delta') \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2/N$.*

Then, given a common equispaced design $(\zeta_j)_{j \in [m]}$ across the S servers, the minimax estimation risk satisfies

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \gtrsim D_*^{-2\alpha}, \quad (19)$$

where D_* solves (2).

Below, we provide sketch of the proof, highlighting the differences with the independent design lower bound. A detailed proof is given in Section 4.2 of the Supplementary Material. The lower bound relies on a similar construction as in the independent design case, but with a different data generating process. Under this common design setup, the minimax rate is either completely governed by the approximation error due to a limited number of design points, which is a cross-server phenomenon, or by the privacy constraints, which are server-specific.

Formally, (2) indicates the occurrence of two scenarios. When $D_* = m^{-2\alpha}$, the rate is governed by the deterministic approximation error (see e.g. [13]). Here, the privacy budget is of no concern. The challenging regime in terms of its proof strategy is when D_* takes its value in the second argument of the minimum of (2). The server specific cost is captured by a similar argument as that of Lemma 3, bounding the minimax risk by a combination of Fisher information quantities using a multivariate version of the von Trees inequality:

$$\frac{2^{2L}}{\sup_{f \in \mathcal{H}_{R,L}^\alpha} \sum_{s=1}^S \min \left\{ \text{Tr}(I_f^{X^{(s)}|T^{(s)}}), \text{Tr}(I_f^{X^{(s)}}) \right\} + 2^{2L(\alpha+2)}},$$

where the “sub-model Fisher information” matrices $I_f^{X^{(s)}|T^{(s)}}$ and $I_f^{X^{(s)}}$ relate to different data generating process that target specifically the cost due to limited local sample sizes and privacy constraints. Combining a geometric argument with a data processing argument, these matrices have trace bounded by constant multiples of $2^L n_s^2 \varepsilon_s^2$ and $n_s 2^{2L}$, respectively. Optimizing over the choice of L then yields the result.

5 Simulation Study

To highlight the differences between the independent and common design settings, we consider the setup of central differential privacy. In such a setting, where data is collected by a central server, the difference between common and independent design just boils down to whether each individual has their measurements taken at the same design points, or whether the design points are random independent across individuals.

In Section 5.1, we describe the simulation design.

5.1 Simulation Design

Specifically, the random curves X_i for each individual are generated as follows:

$$X(t) = R \cdot \sum_{l=0}^{L^*} \sum_{k=0}^{2^l-1} s_{lk} 2^{-l(\alpha+1/2)} \psi_{lk}(t), \quad \text{where } s_{lk} \stackrel{\text{i.i.d.}}{\sim} \text{Rad}(p). \quad (20)$$

Here, the wavelets ψ_{lk} are selected from the Daubechies' extremal phase wavelet family. The parameter R represents the radius of the Hölder ball in which the random function resides, and α denotes the smoothness parameter. The term $\text{Rad}(p)$ refers to the Rademacher distribution, which takes values $+1$ and -1 with probabilities p and $1-p$, respectively.

This construction ensures that the random functions X_i exhibit the desired smoothness properties and fall within the Hölder ball of radius R , while the use of Daubechies' wavelets allows for compact support and flexibility in approximation.

We are interested in estimating the mean of these curves given by $f(\cdot) = \mathbb{E}X(\cdot)$ which evaluates to

$$f(t) = R \cdot \sum_{l=0}^{L^*} \sum_{k=0}^{2^l-1} (2p-1) \cdot 2^{-l(\alpha+1/2)} \psi_{lk}(t). \quad (21)$$

For both of the simulation studies below we set $R = 2$, $L^* = 15$ and $p = 0.9$.

5.2 Independent Design

In the independent design setting, each individual has their sampling points ζ_{ij} drawn independently from a uniform distribution over $[0, 1]$. This accounts for greater variability in the observed data, as both the sampling locations and the random curves X_i vary across individuals. To estimate the mean function $f(t) = \mathbb{E}[X(t)]$, we employ the wavelet-based private estimator of Section 3.1 (Algorithm 1). We study the impact of varying n , the number of individuals, and m , the number of sampling points per individual, for different privacy budgets ϵ . Throughout, we set the privacy parameter $\delta = 1/n^2$. The performance is evaluated using the mean squared error (MSE), averaged over multiple runs, to capture the estimator's accuracy across varying settings. Figure 1 presents the MSE (log-scale) against the number of samples (n) and design points (m).




Figure 1: MSE vs n (with fixed $m = 64$) and MSE vs m (with fixed $n = 200$) for Independent design

The intuitively obvious finding that can be observed from both panels in Figure 1, is that estimation error decreases as the privacy budget (ϵ) increases and / or as the number of individuals in the sample increase n . Similarly, for an increase in the number of measurements, we see that an increase in the number of design points m leads to a decrease in the estimation error. However, for very large m , the error plateaus, indicating a phase transition (depending on ϵ) where the MSE becomes independent of m , which is in line with the main results (see Corollary 1).

5.3 Common Design

In the common design setting, all individuals share the same set of fixed sampling points, chosen as m equally spaced points in the interval $[0, 1]$, whilst otherwise the simulation setting remains the same: For each individual, we observe noisy evaluations of their random curve X_i at these fixed design points. To this end, we employ the privatized local polynomial regression estimator of Section 3.2 (Algorithm 2), where we choose K to be Gaussian kernel.

We study the effect of varying n , the number of individuals, and m , the number of design points, under different privacy budgets ϵ . Throughout, we set $\delta = 1/n^2$ as the privacy parameter for all simulations. The performance of the estimators is evaluated using the integrated mean squared error. Figure 2 presents the MSE (log-scale) against the number of samples (n) and design points (m), providing a comparison of the impact of n , m , and ϵ on estimation accuracy.

Similarly to the dependent design case, the estimation error decreases as the privacy budget (ϵ) increases and / or as n increases, with a plateau occurring for very large n and ϵ . This plateau corresponds to the phase transition where the estimation error is essentially fully determined by the approximation error due to the limited number of design points, and is not observed in the independent design case. This corroborates the main results, comparing Corollaries 1 and 2.




Figure 2: Log-scale MSE vs n (with fixed $m = 64$) and vs m (with fixed $n = 200$) for the common design setting.

Another stark difference compared to the independent design case, is that the estimation error plateaus much quicker as m increases, transitioning to a regime where the rate becomes independent of m . This is in line with the main results (see Corollary 2), which state that there is only a benefit of having more measurements per individual if the privacy constraints are not forming a bottleneck. In the regimes where the differential privacy constraint is sufficiently lenient, there is a benefit of additional measurements up to the point that the number of individuals forms the bottleneck, as can be seen in the case where $\varepsilon = \infty$ in Figure 2.

References

- [1] Jayadev Acharya, Clément L Canonne, Aditya Vikram Singh, and Himanshu Tyagi. Optimal rates for nonparametric density estimation under communication constraints. *IEEE Transactions on Information Theory*, 2023.
- [2] Jayadev Acharya, Clément L Canonne, Ziteng Sun, and Himanshu Tyagi. Unified lower bounds for interactive high-dimensional estimation under information constraints. *Advances in Neural Information Processing Systems*, 36, 2024.
- [3] Jayadev Acharya, Yuhua Liu, and Ziteng Sun. Discrete distribution estimation under user-level local differential privacy. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206, pages 8561–8585. PMLR, 25–27 Apr 2023.
- [4] Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private assouad, fano, and le cam. In *Algorithmic Learning Theory*, pages 48–78. PMLR, 2021.

- [5] 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.
- [6] Francoise Beaufays, Kanishka Rao, Rajiv Mathews, and Swaroop Ramaswamy. Federated learning for emoji prediction in a mobile keyboard. *arXiv:1906.04329*, 2019.
- [7] T. Tony Cai, Abhinav Chakraborty, and Lasse Vuursteen. Federated nonparametric hypothesis testing with differential privacy constraints: Optimal rates and adaptive tests. 2024.
- [8] T Tony Cai, Abhinav Chakraborty, and Lasse Vuursteen. Optimal federated learning for nonparametric regression with heterogeneous distributed differential privacy constraints. *arXiv preprint arXiv:2406.06755*, 2024.
- [9] T. Tony Cai and Ming Yuan. Optimal estimation of the mean function based on discretely sampled functional data: Phase transition. *The Annals of Statistics*, 39(5):2330 – 2355, 2011.
- [10] Clément L Canonne and Yucheng Sun. Private distribution testing with heterogeneous constraints: Your epsilon might not be mine. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, 2024.
- [11] Yuval Dagan, Michael I. Jordan, Xuelin Yang, Lydia Zakynthinou, and Nikita Zhivotovskiy. Dimension-free private mean estimation for anisotropic distributions. 2024.
- [12] Ingrid Daubechies. *Ten lectures on wavelets*. SIAM, 1992.
- [13] Ronald A DeVore and George G Lorentz. *Constructive Approximation*. Springer, 1993.
- [14] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.
- [15] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [16] Frédéric Ferraty and Philippe Vieu. *Nonparametric Functional Data Analysis: Theory and Practice*. Springer Science & Business Media, Berlin, Heidelberg, 2006.
- [17] Richard D Gill and Boris Y Levit. Applications of the van trees inequality: a bayesian cramér-rao bound. *Bernoulli*, pages 59–79, 1995.
- [18] Rob Hall, Alessandro Rinaldo, and Larry Wasserman. Differential privacy for functions and functional data. *The Journal of Machine Learning Research*, 14(1):703–727, 2013.

- [19] Wolfgang Härdle, Gerard Kerkycharian, Dominique Picard, and Alexander Tsybakov. *Wavelets, approximation, and statistical applications*, volume 129. Springer Science & Business Media, 2012.
- [20] Akira Horiguchi, Li Ma, and Botond T Szabó. Sampling depth trade-off in function estimation under a two-level design. *arXiv preprint arXiv:2310.02968*, 2023.
- [21] Iain M Johnstone. *Function Estimation and Gaussian Sequence Models*. Unpublished manuscript, 2019.
- [22] Clément Lalanne, Aurélien Garivier, and Rémi Gribonval. On the statistical complexity of estimation and testing under privacy constraints. *Transactions on Machine Learning Research Journal*, 2023.
- [23] Daniel Levy, Ziteng Sun, Kareem Amin, Satyen Kale, Alex Kulesza, Mehryar Mohri, and Ananda Theertha Suresh. Learning with user-level privacy. *Advances in Neural Information Processing Systems*, 34:12466–12479, 2021.
- [24] Yuhan Liu, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Michael Riley. Learning discrete distributions: user vs item-level privacy. *Advances in Neural Information Processing Systems*, 33:20965–20976, 2020.
- [25] Ardalan Mirshani, Matthew Reimherr, and Aleksandra Slavković. Formal privacy for functional data with gaussian perturbations. In *International Conference on Machine Learning*, pages 4595–4604. PMLR, 2019.
- [26] Shyam Narayanan, Vahab Mirrokni, and Hossein Esfandiari. Tight and robust private mean estimation with few users. In *International Conference on Machine Learning*, pages 16383–16412. PMLR, 2022.
- [27] JO Ramsay and BW Silverman. Principal components analysis for functional data. *Functional data analysis*, pages 147–172, 2005.
- [28] John A Rice and Bernard W Silverman. Estimating the mean and covariance structure nonparametrically when the data are curves. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(1):233–243, 1991.
- [29] A. Tsybakov. Introduction to nonparametric estimation. In *Springer Series in Statistics*, 2008.
- [30] Alexandre B. Tsybakov. *Introduction to nonparametric estimation*. Springer series in statistics. Springer, New York ; London, 2009. OCLC: ocn300399286.
- [31] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 1 edition, September 2018.

- [32] Jane-Ling Wang, Jeng-Min Chiou, and Hans-Georg Müller. Functional data analysis. *Annual Review of Statistics and its application*, 3(1):257–295, 2016.
- [33] Gengyu Xue, Zhenhua Lin, and Yi Yu. Optimal estimation in private distributed functional data analysis. 2024.
- [34] R. Zamir. A proof of the fisher information inequality via a data processing argument. *IEEE Transactions on Information Theory*, 44(3):1246–1250, 1998.

Supplementary Material

6 Proofs related to the upper bound theorems

6.1 Proof for the independent design upper bound

We first prove a upper bound on the ℓ_2 -sensitivity i.e Lemma 2.

6.1.1 Proof of Lemma 2

Proof. We aim to bound the ℓ_2 -sensitivity of the statistic $\tilde{f}_L^{\tau,s}$ under neighboring datasets $Z^{(s)}$ and $Z'^{(s)}$. Neighboring datasets differ in at most one individual:

$$Z_i^{(s)} = (Y_{ij}^{(s)}, \zeta_{ij}^{(s)})_{j \in [m_s]} \neq Z_i'^{(s)} = (Y_{ij}'^{(s)}, \zeta_{ij}'^{(s)})_{j \in [m_s]}.$$

Define the privatized coefficient for $Z'^{(s)}$:

$$\tilde{f}_{lk}'^{\tau,s} := \frac{1}{\sqrt{2^l \wedge m_s n_s}} \sum_{i=1}^{n_s} \left[\frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}'^{(s)} \psi_{lk}(T_{ij}'^{(s)}) \right]_{\tau_l^{(s)}}.$$

The ℓ_2 sensitivity is given by:

$$\left\| \tilde{f}_L^{\tau,s}(Z^{(s)}) - \tilde{f}_L^{\tau,s}(Z'^{(s)}) \right\|_2^2 = \sum_{l=l_0}^L \sum_{k=0}^{2^l-1} (\tilde{f}_{lk}^{\tau,s} - \tilde{f}_{lk}'^{\tau,s})^2.$$

Expanding $(\tilde{f}_{lk}^{\tau,s} - \tilde{f}_{lk}'^{\tau,s})^2$, it can be seen to equal

$$\frac{1}{(2^l \wedge m_s)(\tau_l^{(s)})^2 n_s^2} \left(\left[\frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}) \right]_{\tau_l^{(s)}} - \left[\frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}'^{(s)} \psi_{lk}(\zeta_{ij}'^{(s)}) \right]_{\tau_l^{(s)}} \right)^2.$$

Since the wavelet basis ψ_{lk} is compactly supported, with its support of the order 2^{-l} at resolution level l , at most a constant number of wavelet functions overlap at any design point at each

resolution level. Thus, the summation over k is bounded by the number of overlapping supports. Let c_A denote this constant, which gives the upper bound

$$\sum_{l=l_0}^L \sum_{k=0}^{2^l-1} \frac{1}{(2^l \wedge m_s) n_s^2} \mathbb{1} \left\{ \exists j : \{\zeta_{ij}^{(s)}, \zeta'_{ij}{}^{(s)}\} \cap \text{supp}(\psi_{lk}) \neq \emptyset \right\} \leq \sum_{l=l_0}^L \frac{c_A (2^l \wedge m_s)}{(2^l \wedge m_s) n_s^2}.$$

Combining terms, the ℓ_2 sensitivity satisfies:

$$\left\| \tilde{f}_L^{\tau,s}(Z^{(s)}) - \tilde{f}_L^{\tau,s}(Z'^{(s)}) \right\|_2^2 \leq \frac{c_A^2 L}{n_s^2},$$

for some constant $c_A > 0$. Taking the square root completes the proof. $\square$

6.1.2 Proof of upper bound for independent design

We next prove the upper bound i.e Theorem 3.

Proof of Theorem 3. Let $U_i^{(s)} = (U_{i,lk}^{(s)})_{k=0,\dots,2^l-1, l_0 \leq l \leq L}$, where:

$$U_{i,lk}^{(s)} := \frac{1}{m_s} \sum_{j=1}^{m_s} Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}). \quad (22)$$

The randomized $(\varepsilon_s, \delta_s)$ -DP algorithm outputs a noisy version of the clipped coefficients:

$$\hat{f}_{lk}^{P,s} = \hat{f}_{lk}^{\tau,s} + W_{lk}^{(s)}, \quad \hat{f}_{lk}^{\tau,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} [U_{i,lk}^{(s)}]_{\tau_l^{(s)}},$$

where $W_{lk}^{(s)} \stackrel{\text{i.i.d.}}{\sim} N\left(0, \frac{4L(2^l \wedge m_s)(\tau_l^{(s)})^2 \log(2/\delta_s)}{n_s^2 \varepsilon_s^2}\right)$.

Given that the wavelets form an orthonormal basis, Plancharel's theorem yields that the squared L_2 -risk is given by

$$\sum_{l \geq l_0} \sum_{k=1}^{2^l} \mathbb{E}_f (\hat{f}_{lk}^P - f_{lk})^2. \quad (23)$$

where $\hat{f}_{lk}^P = 0$ for $l > L$. For $l \leq L$, each summand equals

$$\mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \hat{f}_{lk}^{P,s} - f_{lk} \right)^2 = \underbrace{\mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \hat{f}_{lk}^{\tau,s} - f_{lk} \right)^2}_{\text{bias + sampling variance}} + \underbrace{\mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} W_{lk}^{(s)} \right)^2}_{\text{Privacy Noise}}.$$

Below, we bound each of these terms separately.

Bounding the sampling variance and bias Under the event $A_{lk}^{(s)} \equiv \{\forall i \in [n_s] : [U_{i,lk}^{(s)}]_{\tau_l^{(s)}} = U_{i,lk}^{(s)}\}$, the clipped estimator $\widehat{f}_{lk}^{\tau,s}$ equals the unbiased sample mean $\widehat{f}_{lk}^{(s)}$. By Lemma 1, this occurs with high probability:

$$\mathbb{P}(A_{lk}^{(s)}) \geq 1 - \frac{1}{N^c}.$$

Decomposing the expectation into two parts, we have

$$\begin{aligned} & \mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \widehat{f}_{lk}^{\tau,s} - f_{lk} \right)^2 \\ &= \mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \widehat{f}_{lk}^{\tau,s} - f_{lk} \right)^2 \mathbb{1}_{\{\cap_{s=1}^S A_{lk}^{(s)}\}} + \mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \widehat{f}_{lk}^{\tau,s} - f_{lk} \right)^2 \mathbb{1}_{\{\cap_{s=1}^S A_{lk}^{(s)}\}^c} \\ &\leq \mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} \widehat{f}_{lk}^{(s)} - f_{lk} \right)^2 \mathbb{1}_{\{\cap_{s=1}^S A_{lk}^{(s)}\}} + c_R \sum_{s=1}^S \mathbb{E}_f \mathbb{1}_{\{(A_{lk}^{(s)})^c\}} \\ &\leq \sum_{s=1}^S (w_l^{(s)})^2 \mathbb{E}_f (\widehat{f}_{lk}^{(s)} - f_{lk})^2 + c_R \frac{S}{N^c}. \end{aligned}$$

Here, $\widehat{f}_{lk}^{(s)} - f_{lk}$ has $\mathbb{E}_f$ -expectation zero and satisfies

$$\mathbb{E}_f (\widehat{f}_{lk}^{(s)} - f_{lk})^2 \lesssim \frac{1}{m_s n_s} + \frac{2^{-l(2\alpha+1)}}{n_s},$$

by Lemma 6. Combining these bounds, we have that

$$\mathbb{E} \left(\sum_{s=1}^S w_l^{(s)} \widehat{f}_{lk}^{\tau,s} - f_{lk} \right)^2 \lesssim \sum_{s=1}^S (w_l^{(s)})^2 \left(\frac{1}{m_s n_s} + \frac{2^{-l(2\alpha+1)}}{n_s} \right) + \frac{S}{N^c}.$$

Bounding the Privacy Noise We have

$$\mathbb{E}_f \left(\sum_{s=1}^S w_l^{(s)} W_{lk}^{(s)} \right)^2 = \sum_{s=1}^S (w_l^{(s)})^2 \cdot \frac{4L(2^l \wedge m_s)(\tau_l^{(s)})^2 \log(4/\delta)}{n_s^2 \varepsilon_s^2}.$$

Recalling the definition of the clipped estimator, we have that

$$\tau_l^{(s)} = 2(2c \log N)^{3/2} \left(\sqrt{\frac{1}{m_s}} + \frac{1}{3} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right) + R 2^{-l(\alpha+1/2)},$$

from which it follows that

$$(2^l \wedge m_s)(\tau_l^{(s)})^2 \lesssim (\log N)^3 \left(\frac{1}{m_s} + \frac{2^l}{m_s^2} \right) (2^l \wedge m_s) + 2^{-2l(\alpha+1/2)} (2^l \wedge m_s).$$

Substituting this in the earlier expression, we obtain that

$$\mathbb{E} \left(\sum_{s=1}^S w_l^{(s)} W_{lk}^{(s)} \right)^2 \lesssim \sum_{s=1}^S (w_l^{(s)})^2 \cdot \left(\frac{L 2^l (\log N)^3 \log(4/\delta)}{m_s n_s^2 \varepsilon_s^2} + \frac{L 2^{-2\alpha l} \log(4/\delta)}{n_s^2 \varepsilon_s^2} \right).$$

Combining the bounds Combining the previously derived bounds, we find that $\mathbb{E}_f(\widehat{f}_{lk}^P - f_{lk})^2$ is bounded by a constant multiple of

$$\sum_{s=1}^S (w_l^{(s)})^2 \left[\frac{1}{m_s n_s} + \frac{2^{-l(2\alpha+1)}}{n_s} + \frac{L 2^l (\log N)^3 \log(4/\delta)}{m_s n_s^2 \varepsilon_s^2} + \frac{L 2^{-2\alpha l} \log(4/\delta)}{n_s^2 \varepsilon_s^2} \right] + \frac{S}{N^c},$$

for each $l \leq L$.

The optimal choice of weights $w_l^{(s)}$ is revealed to be inversely proportional to the variance;

$$w_l^{(s)} = \frac{u_l^{(s)}}{\sum_{s=1}^S u_l^{(s)}}, \quad \text{where} \quad u_l^{(s)} = n_s m_s \wedge m_s 2^{l(2\alpha+1)} \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2.$$

e.g.

$$w_l^{(s)} \propto \frac{1}{m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2}.$$

This implies that $\mathbb{E}_f(\widehat{f}_{lk}^P - f_{lk})^2$ is bounded above by

$$\begin{aligned} & L \log(N)^3 \log(4/\delta) \sum_{s=1}^S (w_l^{(s)})^2 \left(\frac{1}{m_s n_s} + \frac{2^{-l(2\alpha+1)}}{n_s} + \frac{2^l}{m_s n_s^2 \varepsilon_s^2} + \frac{2^{-2\alpha l}}{n_s^2 \varepsilon_s^2} \right) + \frac{S}{N^c} \\ & \lesssim L \log(N)^3 \log(4/\delta) \sum_{s=1}^S (w_l^{(s)})^2 \left(\frac{1}{m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2} \right) + \frac{S}{N^c} \\ & = L \log(N)^3 \log(4/\delta) \frac{1}{\sum_{s=1}^S m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2} + \frac{S}{N^c} \end{aligned}$$

Expanding (23), we obtain

$$\mathbb{E}_f \|\widehat{f}_L^P - f\|_2^2 = \sum_{l=l_0}^L \sum_{k=0}^{2^l-1} \mathbb{E}_f (\widehat{f}_{lk}^P - f_{lk})^2 + \sum_{l>L} \sum_{k=0}^{2^l-1} f_{lk}^2.$$

Hence we have that the estimation risk $\mathbb{E}_f \|\widehat{f}_L^P - f\|_2^2$ is bounded as

$$\begin{aligned} & L \log(N)^3 \log(4/\delta) \sum_{l=l_0}^L \frac{2^l}{\sum_{s=1}^S m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2} + \frac{S 2^L}{N^c} + 2^{-2L\alpha} \\ & \lesssim L \log(N)^3 \log(4/\delta) \sum_{l=l_0}^L \frac{2^l}{\sum_{s=1}^S m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2} + 2^{-2L\alpha} \\ & \lesssim L^2 \log(N)^3 \log(4/\delta) \max_{l=l_0, \dots, L} \frac{2^l}{\sum_{s=1}^S m_s n_s \wedge 2^{l(2\alpha+1)} n_s \wedge 2^{-l} m_s n_s^2 \varepsilon_s^2 \wedge 2^{2\alpha l} n_s^2 \varepsilon_s^2} + 2^{-2L\alpha} \end{aligned}$$

The second last line follows from the fact that $S 2^L / N = O(1/N)$ (since $2^L < N$ and $c > 3$). By setting $d = 2^l$ and $D = 2^L$, we obtain

$$\mathbb{E}_f \|\widehat{f}_L^P - f\|_2^2 \lesssim \frac{(\log D)^2 (\log(N))^3 \log(4/\delta)}{\min_{d \leq D} \sum_{s=1}^S d^{-1} m_s n_s \wedge d^{2\alpha} n_s \wedge d^{-2} m_s n_s^2 \varepsilon_s^2 \wedge d^{2\alpha-1} n_s^2 \varepsilon_s^2} + D^{-2\alpha}$$

Using D^* which satisfies (3) and letting $L = \lceil \log_2 D^* \rceil$ we obtain the minimax rate of

$$(\log D^*)^2 \log(N)^3 \log(4/\delta) (D^*)^{-2\alpha}$$

. Since $D^* \leq N$ this implies the minimax risk is further upper bounded by

$$(\log(N)^5 \log(4/\delta) (D^*)^{-2\alpha}.$$

□

6.1.3 Proof of Lemma 1

Lemma 5. For a fixed l and k , $\{U_{i,lk}^{(s)}\}_{i \in [n]}$ as defined in (22) satisfies

$$\mathbb{P}_f(\forall i |U_{i,lk}^{(s)}| < \tau_l^{(s)}) \geq 1 - \gamma$$

with

$$\tau_l^{(s)} = 2(2c \log N)^{3/2} \left(\sqrt{\frac{1}{m_s}} + \frac{1}{3} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right) + R 2^{-l(\alpha+1/2)} \quad \text{and} \quad \gamma = (N)^{-c}. \quad (24)$$

Proof. By a union bound, it is enough to show that $\mathbb{P}_f(|U_{i,lk}^{(s)}| > \tau_l^{(s)}) \leq \gamma/n$. By the triangle inequality a union bound,

$$\begin{aligned} \mathbb{P}_f(|U_{i,lk}^{(s)}| > \tau_l^{(s)}) &\leq \mathbb{P}_f \left(|U_{i,lk}^{(s)} - \mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)})| > \sqrt{\frac{c_1 \log(N)}{m_s}} \right) \\ &\quad + \mathbb{P}_f \left(|\mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)})| > c_2 2^{-l(\alpha+1/2)} \right). \end{aligned} \quad (25)$$

We first bound the second term by observing that:

$$\mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)}) = \mathbb{E}(X_i^{(s)}(T_{ij}) \psi_{lk}(T_{ij}) | X_i^{(s)}) = \int X_i^{(s)}(t) \psi_{lk}(t) dt.$$

Since $X_i^{(s)}$ lies in the α -Hölder ball with radius R , we have:

$$|\mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)})| \leq R 2^{-l(\alpha+1/2)}.$$

For $c_2 \geq R$, this implies:

$$\mathbb{P}(|\mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)})| > c_2 2^{-l(\alpha+1/2)}) = 0.$$

We now turn to the first term in (25). We have

$$U_{i,lk}^{(s)} - \mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)}) = \frac{1}{m_s} \sum_{j=1}^{m_s} \left(Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}) - \mathbb{E}[Y_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}) | X_i^{(s)}] \right),$$

which decomposes further as

$$\frac{1}{m_s} \sum_{j=1}^{m_s} \left(X_i^{(s)}(\zeta_{ij}^{(s)}) \psi_{lk}(\zeta_{ij}^{(s)}) - \int X_i^{(s)}(t) \psi_{lk}(t) dt \right) + \frac{1}{m_s} \sum_{j=1}^{m_s} \xi_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}). \quad (26)$$

We apply Bernstein's inequality (Lemma 7) to the first summand. Let

$$V_i = X_i^{(s)}(\zeta_{ij}^{(s)}) \psi_{lk}(\zeta_{ij}^{(s)}) - \int X_i^{(s)}(t) \psi_{lk}(t) dt.$$

Then:

$$\text{Var}(V_i) \leq \|X\|_\infty^2 = C_{\alpha,R}^2, \quad \max_i \|V_i\|_\infty \leq C_{\alpha,R} \|\psi\|_\infty 2^{l/2},$$

where $C_{\alpha,R}$ is a constant as in Lemma 8. Set $\lambda = 2c \log N \left(\frac{C_{\alpha,R}}{\sqrt{m_s}} + \frac{1}{3} C_{\alpha,R} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right)$. By Lemma 7, it holds with probability $1 - \kappa_1 N^{-c}$ that

$$\frac{1}{m_s} \sum_{j=1}^{m_s} \left(X_i^{(s)}(\zeta_{ij}^{(s)}) \psi_{lk}(\zeta_{ij}^{(s)}) - \int X_i^{(s)}(t) \psi_{lk}(t) dt \right) \leq 2c \log N \left(\frac{C_{\alpha,R}}{\sqrt{m_s}} + \frac{1}{3} C_{\alpha,R} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right),$$

where $K(\alpha, R)$ is a suitable constant.

Next, we consider the second sum in (26). Write $\tilde{V}_i := \xi_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)})$. Conditioning on $\xi_i^{(s)} := \{\xi_{ij}^{(s)}\}_{j \in [m_s]}$, we again apply Bernstein's inequality: We have

$$\text{Var}(\tilde{V}_i \mid \xi_i^{(s)}) \leq \|\xi_i^{(s)}\|_\infty^2, \quad \max_i \|\tilde{V}_i\|_\infty \leq \|\xi_i^{(s)}\|_\infty 2^{l/2} \|\psi\|_\infty,$$

so by Lemma 7, with probability $1 - 1/N^c$, it holds that

$$\frac{1}{m_s} \sum_{j=1}^{m_s} \xi_{ij}^{(s)} \psi_{lk}(\zeta_{ij}^{(s)}) \leq 2c \|\xi_i^{(s)}\|_\infty \log N \left(\frac{1}{\sqrt{m_s}} + \frac{1}{3} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right).$$

By standard calculations, $\|\xi_i^{(s)}\|_\infty \leq \sqrt{2c \log N}$ with probability at least $1 - 2/N^c$. Putting everything together, we obtain that

$$U_{i,lk}^{(s)} - \mathbb{E}(U_{i,lk}^{(s)} \mid X_i^{(s)}) \leq 2(2c \log N)^{3/2} \left(\frac{1}{\sqrt{m_s}} + \frac{1}{3} \|\psi\|_\infty \frac{2^{l/2}}{m_s} \right),$$

with probability $1 - \kappa/N^c$ with $\kappa > 0$. □

Lemma 6. Consider $\{U_{i,lk}^{(s)}\}_{i \in [n]}$ as defined in (22) and let $\hat{f}_{lk}^{(s)}$ denote the quantity $\frac{1}{n} \sum_{i=1}^n U_{i,lk}^{(s)}$. It holds that

$$\mathbb{E}(\hat{f}_{lk}^{(s)} - f_{lk})^2 \lesssim \frac{1}{m_s n_s} + \frac{2^{-l(2\alpha+1)}}{n_s}. \quad (27)$$

Proof of Lemma 6. Note $\mathbb{E}(\hat{f}_{lk}^{(s)} - f_{lk})^2 = \text{Var}(\hat{f}_{lk}^{(s)}) = \frac{1}{n_s} \text{Var}(U_{1,lk}^{(s)})$. Next using properties of conditional expectations and variance we have

$$\text{Var}(U_{1,lk}^{(s)}) = \text{Var}(\mathbb{E}(U_{1,lk}^{(s)} | X_i^{(s)})) + \mathbb{E}(\text{Var}(U_{1,lk}^{(s)} | X_i^{(s)}))$$

Note that since $X \in \mathcal{H}_\alpha(R)$ we have $\mathbb{E}(U_{i,lk}^{(s)} | X_i^{(s)}) = \mathbb{E}(X_i^{(s)}(\zeta_{ij}^{(s)})\psi_{lk}(\zeta_{ij}^{(s)}) | X_i^{(s)}) = \int X_i^{(s)}(t)\psi_{lk}(t)dt \lesssim 2^{-l(\alpha+1/2)}$ which implies $\text{Var}(\mathbb{E}(U_{1,lk}^{(s)} | X_i^{(s)})) \lesssim 2^{-l(2\alpha+1)}$.

Also note that $\text{Var}(U_{1,lk}^{(s)} | X_i^{(s)}) = \frac{1}{m_s} \text{Var}(Y_{ij}^{(s)}\psi_{lk}(\zeta_{ij}^{(s)}) | X_i^{(s)})$ using the fact that X_i is uniformly bounded by a constant (which implies that the mean of $Y_{ij}^{(s)}$ is also uniformly bounded) we have that $\mathbb{E}(\text{Var}(U_{1,lk}^{(s)} | X_i^{(s)})) \lesssim 1/m_s$. Putting everything together we have the desired bound. $\square$

6.2 Auxiliary lemmas

The following lemma is a well-known form of Bernstein's inequality for bounded random variables.

Lemma 7. *If $V_1, \dots, V_n$ are independent bounded random variables such that $\mathbb{E}[V_i] = 0$, $\mathbb{E}[V_i^2] \leq 1$, and $\max_i \sup |V_i| \leq M$, for some constant $M > 0$.*

Then,

$$\mathbb{P}\left(\left|n^{-1} \sum_{i=1}^n V_i\right| > \lambda\right) \leq 2 \exp\left(-\frac{n\lambda^2}{2\left(1 + \frac{M\lambda}{3}\right)}\right).$$

The following lemma is a standard result, see e.g. Chapter 9 in [21].

Lemma 8. *There exists a constant $C_{\alpha,R}$ such that $\|f\|_\infty \leq C_{\alpha,R}$ for all $f \in \mathcal{H}^\alpha(R)$ with $\alpha > 1/2$.*

7 Proofs for the common design upper bound

This appendix provides the detailed proofs, assumptions, and auxiliary results that support the analysis of the local polynomial estimator used in the main text.

7.1 Proof of Theorem 4

7.1.1 Aggregated mean estimator

For the s th server, $s \in [S]$, $j \in [m]$ the privatized mean $\bar{Y}_j^{P,s}$ at each point ζ_j is given by:

$$\bar{Y}_j^{P,s} = \frac{1}{n} \sum_{i=1}^n \left[Y_{i,j}^{(s)} \right]_\tau + W_j^{(s)},$$

where $W_j^{(s)} \sim \mathcal{N}\left(0, \frac{4\tau^2 m \log(2/\delta)}{n_s^2 \varepsilon_s^2}\right)$ is a noise term added to ensure differential privacy, and the clipping parameter $R = \sqrt{2 \log n_s}$ is used to bound the data.

The aggregated estimator is given by

$$\bar{Y}_j^P = \sum_{s=1}^S w_s \bar{Y}_j^{P,s}$$

The weights are chosen as

$$w_s = \frac{u_s}{\sum_{s=1}^S u_s} \quad \text{where} \quad u_s = m_0^{-1} n_s^2 \varepsilon_s^2 \wedge n_s.$$

7.1.2 Bias and Variance Analysis of the Local Polynomial Estimator

Our estimator $\hat{f}^P(x)$ is given by an average of B many local polynomial estimator $\{\hat{f}_b^P(x)\}_{b \in [B]}$. Our data is divided based on $\{\zeta_j\}_{j \in [m]}$ into B many disjoint groups $\{G_b\}_{b \in [B]}$ each of size m_0 . The local polynomial estimator $\hat{f}_b^P(x)$ for the b th group is defined as:

$$\hat{f}_b^P(x) = \sum_{j \in G_b} \bar{Y}_j^P W_{b,j}^*(x),$$

where $W_{b,j}^*(x)$ are the weights, given by:

$$W_{b,j}^*(x) = \frac{1}{m_0 h} V^\top(0) B_{b,x}^{-1} V \left(\frac{\zeta_j - x}{h} \right) K \left(\frac{\zeta_j - x}{h} \right).$$

and

$$B_{b,x} = \frac{1}{m_0 h} \sum_{j \in G_b} V \left(\frac{\zeta_j - x}{h} \right) V \left(\frac{\zeta_j - x}{h} \right)^\top K \left(\frac{\zeta_j - x}{h} \right)$$

To analyze the performance of $\hat{f}^P(x)$, we decompose the estimation error into bias and variance components. The choice of bandwidth h affects both components, allowing us to achieve the minimax risk rate by balancing these terms.

Bias The bias of the estimator $\hat{f}^P(x)$ is given by:

$$bias(x) = \mathbb{E} \hat{f}^P(x) - f(x) = \frac{1}{B} \sum_{b=1}^B bias_b(x)$$

where $bias_b(x) = \mathbb{E} \hat{f}_b^P(x) - f(x)$.

$$bias_b(x) = \mathbb{E} \hat{f}_b^P(x) - f(x) = \sum_{j \in G_b} f(\zeta_j) W_{b,j}^*(x) - f(x) + \sum_{j \in G_b} (\mathbb{E}(\bar{Y}_j^P) - f(\zeta_j)) W_{b,j}^*(x). \quad (28)$$

Lets look at the first term, using a Taylor expansion of f around x , we have:

$$\begin{aligned} \sum_{j \in G_b} f(\zeta_j) W_{b,j}^*(x) - f(x) &= \sum_{j \in G_b} (f(\zeta_j) - f(x)) W_{b,j}^*(x) \\ &= \sum_{j \in G_b} \frac{1}{p!} \left(f^{(p)}(x + \xi_j(\zeta_j - x)) - f^{(p)}(x) \right) W_{b,j}^*(x) \end{aligned}$$

where $\xi_j \in [0, 1]$, $p = \lfloor \alpha \rfloor$ where we have used the fact that $\sum_{j=1}^{m_0} W_{l,j}^*(x) = 1$ and $\sum_{j \in G_b} (\zeta_j - x)^k W_{l,j}^*(x) = 0$ for $k \in [p]$ from Proposition 11. Under Assumption (LP1) and Lemma 10 we can bound the above term as:

$$\begin{aligned} \sum_{j \in G_b} f(\zeta_j) W_{b,j}^*(x) - f(x) &\leq \sum_{j \in G_b} \frac{1}{p!} C_R |\zeta_j - x|^\alpha |W_{b,j}^*(x)| \\ &= \sum_{j \in G_b} \frac{1}{p!} C_R |\zeta_j - x|^\alpha |W_{b,j}^*(x)| \mathbb{1}(|\zeta_j - x| \leq h) \\ &\leq C_R \sum_{j \in G_l} \frac{1}{p!} h^\alpha |W_{b,j}^*(x)| \leq \frac{C_R C_f h^\alpha}{p!} \end{aligned}$$

where C_R is the lipchitz constant for the p th derivative of f and depends on the holder radius R .

We next look at the second term in (28)

$$\mathbb{E} \bar{Y}_j^P - f(\zeta_j) = \sum_{s=1}^S w_s (\mathbb{E} \bar{Y}_j^{P,s} - f(\zeta_j)) \quad (29)$$

Using the fact that noise W 's are mean zero we have that

$$\mathbb{E} \bar{Y}_j^{P,s} - f(\zeta_j) = \mathbb{E}[Y_{1j}^{(s)}]_\tau - f(\zeta_j)$$

We can decompose the above bias as

$$\begin{aligned} \mathbb{E}[Y_{1j}^{(s)}]_\tau - f(\zeta_j) &= \mathbb{E} \left(\mathbb{E} \left[[Y_{1j}^{(s)}]_\tau \mid X_1^{(s)} \right] - X_1^{(s)}(T_{1j}) \right) \\ &\leq \mathbb{E} \left| \mathbb{E} \left[[Y_{1j}^{(s)}]_\tau \mid X_1^{(s)} \right] - X_1^{(s)}(T_{1j}) \right| \end{aligned}$$

Next conditional of $X_1^{(s)}$ we have that $Y_{1j}^{(s)}$ is gaussian with mean $X_1^{(s)}(T_{1j})$ and variance 1. Hence using Lemma 12 we have that

$$\left| \mathbb{E} \left([Y_{1j}^{(s)}]_\tau \mid X_1^{(s)} \right) - X_1^{(s)}(\zeta_j) \right| \leq 4 |X_1^{(s)}(T_{1j})| e^{\frac{1}{2}(\tau - |X_1^{(s)}(T_{1j})|)^2}$$

Using the fact $X \in \mathcal{H}^\alpha(R)$ we have that $|X_1^{(s)}(T_{1j})| \leq \|X_1^{(s)}\|_\infty \leq C_{\alpha,R}$ and our choice of $\tau = \sqrt{2 \log \bar{N}} + C_{\alpha,R}$ we have that

$$\left| \mathbb{E} \left([Y_{1j}^{(s)}]_\tau \mid X_1^{(s)} \right) - X_1^{(s)}(\zeta_j) \right| \leq 4 C_{\alpha,R} \frac{1}{\bar{N}}$$

which implies $\mathbb{E}[Y_{1j}^{(s)}]_\tau - f(\zeta_j) = O(1/N)$. Using (29) we have that $|\mathbb{E} \bar{Y}_j^P - f(\zeta_j)| = O(1/N)$. Thus the second term in (28) satisfies,

$$\sum_{j \in G_b} |W_{b,j}^*(x)| \cdot O\left(\frac{1}{N}\right) \leq O\left(\frac{1}{N}\right),$$

which can be safely ignored since $O\left(\frac{1}{n}\right)$ is smaller than the order of the variance as we shall see later. Hence putting everything together

$$bias(x) = O\left(h^\alpha + \frac{1}{N}\right).$$

Variance The variance of $\widehat{f}_B^P(x)$, $\sigma^2(x)$ is given by:

$$\begin{aligned}\text{Var}\left(\widehat{f}_B^P(x)\right) &= \text{Var}\left(\frac{1}{B} \sum_{b=1}^B \widehat{f}_b^P(x)\right) \\ &= \text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} \bar{Y}_j^P W_{b,j}^*(x)\right)\end{aligned}$$

We can rewrite $\bar{Y}_j^P$ as $\bar{Y}_j^P = \bar{Y}_j^\tau + W_j$ where

$$\text{Var}\left(\widehat{f}_B^P(x)\right) = \bar{Y}_j^\tau = \sum_{s=1}^S w_s \bar{Y}_j^{\tau,s} \text{ with } \bar{Y}_j^{\tau,s} = \frac{1}{n_s} \sum_{i=1}^{n_s} [Y_{i,j}^{(s)}]_\tau, \quad \text{and} \quad W_j = \sum_{s=1}^S w_s W_j^{(s)}.$$

Hence

$$\text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} \bar{Y}_j^\tau W_{b,j}^*(x)\right) + \text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} W_j W_{b,j}^*(x)\right) \quad (30)$$

We begin by bounding the first term (30). Suppose we have an upper bound on $\text{Var}(\bar{Y}_j^\tau) \leq \tilde{\sigma}^2$. Set $A_b = \sum_{j \in G_b} \bar{Y}_j^\tau W_{b,j}^*(x)$ and use Lemma 9 to conclude

$$\begin{aligned}\text{Var}(A_b) &= \text{Var}\left(\sum_{j \in G_b} \bar{Y}_j^\tau W_{b,j}^*(x)\right) \\ &\leq \left(\sum_{j \in G_b} \sqrt{\text{Var}(\bar{Y}_j^\tau W_{b,j}^*(x))}\right)^2 \\ &\leq \tilde{\sigma}^2 \left(\sum_{j \in G_b} |W_{b,j}^*(x)|\right)^2 \leq \tilde{\sigma}^2 C_f^2\end{aligned}$$

the last line follows from Lemma 10. Hence again using Lemma 9

$$\begin{aligned}\text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} \bar{Y}_j^\tau W_{b,j}^*(x)\right) &= \text{Var}\left(\frac{1}{B} \sum_{b=1}^B A_b\right) \\ &\leq \frac{1}{B^2} \left(\sum_{b=1}^B \sqrt{\tilde{\sigma}^2 C_f^2}\right)^2 = \tilde{\sigma}^2 C_f^2\end{aligned}$$

Next we need to find $\tilde{\sigma}^2$ such that $\text{Var}(\bar{Y}_j^\tau) \leq \tilde{\sigma}^2$. In that direction we proceed as follows

$$\begin{aligned}\text{Var}(\bar{Y}_j^\tau) &= \text{Var}\left(\sum_{s=1}^S w_s \bar{Y}_j^{\tau,s}\right) \\ &= \sum_{s=1}^S w_s^2 \text{Var}\left(\bar{Y}_j^{\tau,s}\right) \\ &= \sum_{s=1}^S w_s^2 \frac{1}{n_s} \text{Var}\left([Y_{1,j}^{(s)}]_\tau\right)\end{aligned}$$

Next we use the fact that for any random variable V , $\text{Var}([V]_\tau) \leq \text{Var}(V)$ to obtain

$$\begin{aligned} \text{Var}(\bar{Y}_j^\tau) &\leq \sum_{s=1}^S w_s^2 \frac{1}{n_s} \text{Var}\left(Y_{1j}^{(s)}\right) \\ &= \sum_{s=1}^S w_s^2 \frac{1}{n_s} \left(\text{Var}\left(X_1^{(s)}(T_{1j})\right) + \text{Var}\left(\varepsilon_{1j}^{(s)}\right) \right) \\ &\leq \sum_{s=1}^S w_s^2 \frac{1}{n_s} (C_{\alpha,R}^2 + 1) \end{aligned}$$

Putting everything together we have that

$$\text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} \bar{Y}_j^\tau W_{l,j}^*(x)\right) \leq (C_{\alpha,R}^2 + 1) C_f^2 \sum_{s=1}^S \frac{w_s^2}{n_s}$$

Next we bound the second term in (30)

$$\begin{aligned} \text{Var}\left(\frac{1}{B} \sum_{b=1}^B \sum_{j \in G_b} W_j W_{b,j}^*(x)\right) &= \frac{1}{B^2} \sum_{b=1}^B \sum_{j \in G_b} \text{Var}(W_j) (W_{b,j}^*(x))^2 \\ &= \text{Var}(W_1) \frac{1}{B^2} \sum_{b=1}^B \sum_{j \in G_b} (W_{b,j}^*(x))^2 \\ &= \text{Var}(W_1) \frac{1}{B} \frac{C_f^2}{m_0 h} \end{aligned}$$

where we used $\sum_{j \in G_b} (W_{b,j}^*(x))^2 \leq \max_j |W_{b,j}^*(x)| \sum_{j \in G_b} |W_{b,j}^*(x)|$ and Lemma 10 in the last line. Next by independence of the noise terms we have that $\text{Var}(W_1) = \sum_{s=1}^S w_j^2 \text{Var}(W_1^{(s)})$. Set $\delta' = \min_s \delta_s$ and $h \asymp m_0^{-1}$. Using the properties of the noise term $W_1^{(s)}$ added for privacy, we have:

$$\sigma^2(x) \leq (C_{\alpha,R}^2 + 1) C_f^2 \sum_{s=1}^S \frac{w_s^2}{n_s} + \frac{C_f^2}{B m_0 h} \sum_{s=1}^S w_s^2 \frac{4\tau^2 m \log(2/\delta)}{n_s^2 \varepsilon_s^2}$$

where the first term represents the intrinsic variance and the second term represents the additional variance due to privacy noise.

7.1.3 Proof of the Minimax Rate

Hence we can simplify the upper bound on variance as

$$\begin{aligned} \sigma^2(x) &\lesssim \log(N) \log(2/\delta') \sum_{s=1}^S w_s^2 \frac{1}{n_s \wedge m_0^{-1} n_s^2 \varepsilon_s^2} \\ &\lesssim \log(N) \log(2/\delta') \frac{1}{\sum_{s=1}^S n_s \wedge m_0^{-1} n_s^2 \varepsilon_s^2} \end{aligned}$$

where we used the fact that $m/B \leq m_0$ and our choice of weights. Set $m_0 = D \wedge m$ where D solves (2) to obtain that $\sigma^2(x) \lesssim \log(N) \log(2/\delta')(D \wedge m)^{-2\alpha}$, and we also know that $\text{bias}^2(x) = h^{2\alpha} = (m_0)^{2\alpha} = (D \wedge m)^{-2\alpha}$. Combining these two the minimax rate is given by

$$\log(N) \log(2/\delta') (D^{-2\alpha} \vee m^{-2\alpha}).$$

7.2 Auxiliary Lemmas and Propositions

Lemma 9 (Variance Identity).

$$\text{Var}\left(\sum_{l=1}^L A_l\right) = \sum_{l,l'} \text{Cov}(A_l, A_{l'}) \leq \sum_{l,l'} \sqrt{\text{Var}(A_l)} \sqrt{\text{Var}(A_{l'})} \leq \left(\sum_{l=1}^L \sqrt{\text{Var}(A_l)}\right)^2.$$

Lemma 10 (Properties of Weights $W_{m_0,j}^*(x)$). *Under Assumptions (LP1)–(LP2), $m_0 \geq \tilde{m}$, $h \geq 1/2m_0$, and $x \in [0, 1]$, the weights $W_{m_0,j}^*(x)$ satisfy:*

- (i) $\sup_{j: |\zeta_j - x| \leq h} |W_{m_0,j}^*(x)| \leq \frac{C_f}{m_0 h},$
- (ii) $\sum_{j=1}^{m_0} |W_{m_0,j}^*(x)| \leq C_f,$
- (iii) $W_{m_0,j}^*(x) = 0$ if $|\zeta_j - x| > h.$

Proof. For a proof see Lemma 1.3 of [29]. □

Lemma 11 (Polynomial Interpolation Property). *For any polynomial Q of order l , the weights $W_{m_0,j}^*(x)$ satisfy:*

$$\sum_{j=1}^{m_0} Q(\zeta_j) W_{m_0,j}^*(x) = Q(x).$$

in particular it implies that $\sum_{j=1}^{m_0} W_{m_0,j}^(x) = 1$ and $\sum_{j=1}^{m_0} (\zeta_j - x)^k W_{m_0,j}^*(x) = 0$ for $k \in [l]$.*

Proof. For a proof see Proposition 1.12 of [29]. □

Lemma 12 (Bias of truncated gaussians). *Let $W \sim N(\mu, \sigma^2)$, $\mu \in \mathbb{R}$ then the bias of $[W]_{-\tau}^\tau$ for $\tau > |\mu|$ is bounded by*

$$|\mathbb{E}[W]_{-\tau}^\tau - \mu| \leq 4|\mu| e^{-\frac{1}{2\sigma^2}(\tau - |\mu|)^2}$$

Proof. For simplicity assume $\sigma = 1$. First note that $\mathbb{E}[W]_{-\tau}^\tau - \mu = \mathbb{E}[Z]_{-\tau+\mu}^{\tau+\mu}$ where $Z \sim N(0, 1)$. Then, for $\mu > 0$ and $\tau > \mu$,

$$0 < \mathbb{E}[Z]_{-\tau+\mu}^{\tau+\mu} \leq \mathbb{E}[Z]_{-\tau+\mu}^{\tau-\mu} + (2\mu)\mathbb{P}(\tau - \mu \leq Z \leq \tau + \mu) \leq 4\mu e^{-\frac{1}{2}(\tau-\mu)^2}.$$

Similarly, for $\mu < 0$ and $\tau > \mu$, we obtain that

$$0 > \mathbb{E}[Z]_{-\tau+\mu}^{\tau+\mu} \geq 4\mu e^{-\frac{1}{2}(\tau+\mu)^2}.$$

Combining these two cases together, we have for $\mu \in \mathbb{R}$ and τ large that

$$\mathbb{E}[Z]_{-\tau+\mu}^{\tau+\mu} \leq 4|\mu| e^{-\frac{1}{2}(\tau-|\mu|)^2}.$$

□

8 Proofs related to the lower bound theorems

8.1 Proofs for the independent design lower bound (Theorem 5)

8.1.1 Proof of Lemma 3

As a first step, we enlarge the collection of differentially private protocols to selectively include protocols which rely on both $Y^{(s)}$ and $X^{(s)}$ for estimating f , for a selection of servers. This step allows us to circumvent technical data processing arguments later on in the proof. To that extent, consider $\mathcal{S} \subset [S]$, to be determined later. Write $\Delta X^{(s)} = (X_i^{(s)} - \mathbb{E}_f X_i^{(s)})_{i \in [n]}$. Let $\overline{\mathcal{M}_{\varepsilon, \delta}}$ denote the collection of all (ε, δ) -differentially private protocols where the s -th transcript $T^{(s)}$ is allowed to depend on $(X^{(s)}, Y^{(s)})$ if $s \in \mathcal{S}$, whereas $T^{(s)}$ is allowed to depend on $(Y^{(s)}, \Delta X^{(s)})$ if $s \notin \mathcal{S}$. Note that this does not impact our definition of the channel generating $T^{(s)}$ being differentially private otherwise, as the unit of privacy is the same for $(Y_i^{(s)}, X_i^{(s)}, T_i^{(s)})$ as it is for $(Y_i^{(s)}, T_i^{(s)})$. The class $\overline{\mathcal{M}_{\varepsilon, \delta}}$ is a superset of $\mathcal{M}_{\varepsilon, \delta}$, so

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon, \delta}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \geq \inf_{\hat{f} \in \overline{\mathcal{M}_{\varepsilon, \delta}}} \sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2.$$

Consider D_* to be the solution to the equation (3) and set $L = \lceil \log_2(D_*) \rceil$. Consider now the L -dimensional sub-model given by

$$\mathcal{H}_{R,L}^\alpha : \left\{ f \in \mathcal{H}_\alpha^R : f = \sum_{k=1}^{2^L} f_k \phi_{Lk}, f_k \in [-2^{-L(\alpha+1/2)}R, 2^{-L(\alpha+1/2)}R] \right\}, \quad (31)$$

where ϕ_k are the wavelet basis functions at resolution level L . Let μ and ν denote dominating measures for $\mathbb{P}_f^{Y^{(s)}}$, and $\mathbb{P}_f^{(Y^{(s)}X^{(s)})}$, respectively, for $f \in \mathcal{H}_{R,L}^\alpha$. Denote f_L by the vector $\{f_l : l = 1, \dots, L\}$. Let ∇_{f_L} denote the vector of partial derivatives with respect to the vector f_L .

Since $\mathcal{H}_{R,L}^\alpha$ is a subset of $\mathcal{H}_\alpha^R$, we have that

$$\sup_{f \in \mathcal{H}^\alpha(R)} \mathbb{E}_f \|\hat{f} - f\|_2^2 \geq \sup_{f \in \mathcal{H}_{R,L}^\alpha} \mathbb{E}_f \|\hat{f} - f\|_2^2,$$

where $f = \sum_{k=1}^{2^L} f_k \phi_{Lk}$. Furthermore, for any probability distribution π on $\mathcal{H}_{R,L}^\alpha$, the latter is further lower bounded by

$$\int_{f_L} \mathbb{E}_f \|\hat{f} - f\|_2^2 \pi(f_L) df_L.$$

Note that $\mathbb{E}_f \|\hat{f} - f\|_2^2 \geq \mathbb{E}_f \left(\sum_{k=1}^{2^L} (\hat{f}_k - f_k)^2 \right)$, where $\hat{f}_k = \int \hat{f} \phi_{Lk}$. For getting a lower bound on the minimax risk we will use a multivariate version of the Van-Trees inequality due to [17] (Theorem 1), which bounds the ℓ^2 -risk by

$$\int_{f_L} \mathbb{E} \left(\sum_{k=1}^{2^L} (\hat{f}_k - f_k)^2 \right) \lambda(f_L) df_L \geq \frac{2^{2L}}{\int \text{Tr}(I_{f_L}^T) \pi(f_L) df_L + J(\pi)},$$

where $I_{f_L}^T$ is the sub-model Fisher information associated with the transcript $T = (T^{(1)}, \dots, T^{(S)})$:

$$I_{f_L}^T = \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^T}{d\mu'} \right) \right] \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^T}{d\mu'} \right) \right]^\top,$$

and $\pi(f_L) = \prod_{k=1}^{2^L} \pi_k(f_k)$ is a prior for the parameter f_L and $J(\pi)$ is the Fisher information associated with the prior π :

$$J(\pi) = \sum_{k=1}^{2^L} \int \frac{\pi'_k(f_k)^2}{\pi_k(f_k)} df_{Lk}.$$

The independence of the data in the servers implies independence of the transcripts, which yields

$$\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^T}{d\mu'} \right) = \sum_{s=1}^S \nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{T^{(s)}}}{d\mu'} \right).$$

Let $Z^{(s)}$ denote either $(Y^{(s)}, \Delta X^{(s)})$ if $s \in \mathcal{S}^c$ or $(Y^{(s)}, X^{(s)})$ if $s \in \mathcal{S}$ and let η denote μ or ν accordingly. By Bayes rule, it follows that

$$\begin{aligned} \nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{T^{(s)}}}{d\mu'} \right) &= \frac{\int \frac{d\mathbb{P}^{T^{(s)}|Z^{(s)}}}{d\mu'} \nabla_{f_L} \frac{d\mathbb{P}_f^{Z^{(s)}}}{d\eta} d\eta}{\frac{d\mathbb{P}_f^{T^{(s)}}}{d\mu'}} \\ &= \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{Z^{(s)}}}{d\mu} \right) \middle| T^{(s)} \right]. \end{aligned}$$

The identity (17) follows. The data generating mechanism forms a Markov chain; $f \rightarrow X^{(s)} \rightarrow Y^{(s)}$, where the joint distribution degenerates as

$$d\mathbb{P}_f^{(X^{(s)}, Y^{(s)})} = d\mathbb{P}_f^{(X^{(s)})} d\mathbb{P}^{(Y^{(s)}|X^{(s)})}.$$

Hence,

$$\mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X^{(s)}, Y^{(s)})}}{d\mu \times \nu} \right) \middle| T^{(s)} \right] = \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X^{(s)})}}{d\nu} \right) \middle| T^{(s)} \right]$$

which implies that the Fisher information corresponding to the joint distribution $Z^{(s)} = (Y^{(s)}, X^{(s)})$ equals

$$I_f^{X^s|T^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X^{(s)})}}{d\nu} \right) \middle| T^{(s)} \right] \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X^{(s)})}}{d\nu} \right) \middle| T^{(s)} \right]^\top.$$

Let $X^{(s)}$ follow the dynamics of (16). By the orthogonality of the ϕ_{lk} 's, the wavelet transformation of $X_i^{(s)}$;

$$X_{ik}^{(s)} = \int X_i^{(s)}(t) \phi_{Lk}(t) dt \stackrel{d}{=} f_l + 2^{-L(\alpha+1/2)} (1/2 - B_{ik}^{(s)}) \quad \text{where} \quad B_{ik}^{(s)} \stackrel{iid}{\sim} \text{Beta}(2, 2), \quad (32)$$

with $k = 1, \dots, 2^L$, is sufficient for the observations $X^{(s)}$ and submodel $\mathcal{H}_{R,L}^\alpha$. By the Neyman-Pearson factorization lemma,

$$\mathbb{E} \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X^{(s)})}}{d\nu} \right) \middle| T^{(s)} \right] = \mathbb{E} \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(X_L^{(s)})}}{d\lambda} \right) \middle| T^{(s)} \right],$$

where $X_L^{(s)} = \{X_{ik}^{(s)} : k = 1, \dots, 2^L\}$ and λ is the Lebesgue measure on $\mathbb{R}^{2^L}$.

Similarly, since for the observations $(Y_{ij}^{(s)}, X_i^{(s)} - \mathbb{E}_f X_i^{(s)})$ the distribution of $X_i^{(s)} - \mathbb{E}_f X_i^{(s)}$ is stationary in f , the statistics $\tilde{Y}_{ij}^{(s)} = Y_{ij}^{(s)} - [X_i^{(s)} - \mathbb{E}_f X_i^{(s)}](\zeta_{ij}^{(s)})$ are sufficient for f ;

$$\tilde{Y}_{ij}^{(s)} = f(\zeta_{ij}^{(s)}) + \xi_{ij}^{(s)}, \quad (33)$$

and for $\tilde{Y}^{(s)} = (\tilde{Y}_{ij}^{(s)})_{i \in [n], j \in [m_s]}$, we obtain

$$I_f^{Y^{(s)}|T^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(\tilde{Y}^{(s)}|T^{(s)})}}{d\mu} \right) \middle| T^{(s)} \right] \mathbb{E}_f \left[\nabla_{f_L} \log \left(\frac{d\mathbb{P}_f^{(\tilde{Y}^{(s)}|T^{(s)})}}{d\mu} \right) \middle| T^{(s)} \right]^\top.$$

Let $\phi(\zeta_{ij}^{(s)}) = \left(\phi_l(\zeta_{ij}^{(s)}) \right)_{l \in [L]}$. The score function for $\tilde{Y}^{(s)}$ is given in (15). By linearity of the trace operation,

$$\int \text{Tr}(I_{T^{(s)}}(f)) d\pi(f) = \sum_{s \in \mathcal{S}} \int \text{Tr}(I_{T^{(s)}}(f)) d\pi(f) + \sum_{s \in \mathcal{S}^c} \int \text{Tr}(I_{T^{(s)}}(f_L)) d\pi(f).$$

Now, setting

$$\mathcal{S} = \left\{ s \in [S] : \text{Tr}(I^{X^{(s)}|T^{(s)}}) < \text{Tr}(I^{Y^{(s)}|T^{(s)}}) \right\},$$

and noting that by standard data processing arguments it holds that

$$\text{Tr}(I^{X^{(s)}|T^{(s)}}) \leq \text{Tr}(I^{X^{(s)}}) \quad \text{and} \quad \text{Tr}(I^{Y^{(s)}|T^{(s)}}) \leq \text{Tr}(I^{Y^{(s)}}),$$

(see e.g. [34]).

By taking as prior π_k a rescaled version of the density $t \mapsto \cos^2(\pi t/2) \mathbb{1}\{|t| \leq 1\}$ such that it is supported on $[-2^{-L(\alpha+1/2)}R, 2^{-L(\alpha+1/2)}R]$, we obtain that prior is supported on $\mathcal{H}_{R,L}^\alpha$. By a straightforward calculation for the Fisher information associated with the prior, we have that the minimax risk is lower bounded by

$$\inf_{\hat{f} \in \mathcal{M}_{\varepsilon,\delta}} \sup_{f \in \mathcal{H}_{R,L}^\alpha} \mathbb{E}_f \|\hat{f} - f\|_2^2 \geq \frac{2^{2L}}{\sup_{f_L} \text{Tr}(I_T(f_L)) + \frac{2^L \pi^2}{(2^{-L(\alpha+1/2)}R)^2}}.$$

Combining this with the earlier identities of Lemmas 3 and 4 yields the (further) lower bound

$$\frac{1}{C \left(\sum_{s=1}^S \min \{ 2^{-2L} m_s n_s^2 \varepsilon_s^2, 2^{-L} m_s n_s, 2^{L(2\alpha-1)} n_s^2 \varepsilon_s^2, 2^{2L\alpha} n_s \} + 2^{2\alpha L} \right)}$$

For $D_* \geq 1$ from (3) and $1 \leq \tilde{D} \leq D_*$ such that

$$D_*^{2\alpha} = \sum_{s=1}^S \min \left\{ \tilde{D}^{-1} n_s m_s, \tilde{D}^{-2} m_s n_s^2 \varepsilon_s^2, \tilde{D}^{2\alpha} n_s, \tilde{D}^{2\alpha-1} n_s^2 \varepsilon_s^2 \right\},$$

Set $2^L = \tilde{D}$ in the rate above to obtain the lower bound

$$\frac{1}{C \left(\sum_{s=1}^S \min \left\{ \tilde{D}^{-1} n_s m_s, \tilde{D}^{-2} m_s n_s^2 \varepsilon_s^2, \tilde{D}^{2\alpha} n_s, \tilde{D}^{2\alpha-1} n_s^2 \varepsilon_s^2 \right\} + \tilde{D}^{2\alpha} \right)} = \frac{C^{-1}}{D_*^{2\alpha} + \tilde{D}^{-2\alpha}} \gtrsim D_*^{-2\alpha}.$$

8.1.2 Proof of Lemma 4

We provide a bound for each of the four terms in the minimum separately. The data processing inequalities capturing the loss of information due to privacy constraints follow from Lemma 13 and Lemma 17. Recall that

$$I_f^{Y^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[S_f^{Y^{(s)}} | T^{(s)} \right] \mathbb{E}_f \left[S_f^{Y^{(s)}} | T^{(s)} \right]^\top,$$

where $S_f^{Y^{(s)}}$ is given by $S_f^{Y^{(s)}} = \sum_{i=1}^n S_f^{Y_i^{(s)}}$ and

$$S_f^{Y_i^{(s)}} = \sum_{j=1}^m \left(Y_{ij}^{(s)} - \sum_{k=1}^{2^L} f_k \phi_l(\zeta_{ij}^{(s)}) \right) \phi(\zeta_{ij}^{(s)}), \quad (34)$$

By a straightforward calculation,

$$\mathbb{E}_f [S_f^{Y^{(s)}}] [S_f^{Y^{(s)}}]^\top = m_s n_s I_{2^L},$$

so the bound $\lambda_{\max}(\mathbb{E}_f [S_f^{Y^{(s)}}] [S_f^{Y^{(s)}}]^\top) \leq m_s n_s$ in combination with Lemma 13 yields that $\text{Tr}(I_f^{Y^{(s)}}) \lesssim m_s n_s^2 \varepsilon_s^2$ and

$$\text{Tr}(I_f^{Y^{(s)}}) = 2^L m_s n_s.$$

The relationship (32) defines a regular location-shift model, which yields that the Fisher information for f in the observational model for $X^{(s)}$ is given by

$$I_f^{X^{(s)}} = \mathbb{E}_f \mathbb{E}_f \left[\nabla_f \log \left(\frac{d\mathbb{P}_f^{(X_L^{(s)})}}{d\nu} \right) \right] \mathbb{E} \left[\nabla_f \log \left(\frac{d\mathbb{P}_f^{(X_L^{(s)})}}{d\lambda} \right) \right]^\top = c n_s 2^{2L(\alpha+1/2)} I_{2^L},$$

for some constant $c > 0$. The lemma follows by combining the bounds for the four terms in the minimum.

8.2 Proof of the common design lower bound (Theorem 6)

The lower bound of the order $m^{-2\alpha}$ is the same as in the non-private case, see e.g. [9], Theorem 2.1. Hence, we are required to show to “private part” of the lower bound. Consider the $L + K$ -th resolution level of an $A > \alpha$ smooth orthonormal wavelet basis $\phi_{lk}(x) = 2^{l/2}\phi(2^l x - k)$ for $l \in \{l_0, l_0 + 1, \dots\}$, and $k = 1, \dots, 2^{L+K}\}$, with $K \in \mathbb{N}$ such that a subset of the wavelets $\{\phi_k : k = 1, \dots, 2^L\}$ satisfies $\phi_k = \phi_{(L+K)i_k}$ disjoint support.

Define f_μ with $\mu = (\mu_1, \dots, \mu_{2^L}) \in \mathbb{R}^{2^L}$ as

$$f_\mu = \sum_{k=1}^{2^L} \mu_k \phi_k.$$

Let $d \in \mathbb{N}$ such that $m = \lfloor d \cdot 2^L \rfloor$. Consider now the model

$$Y_{ij}^{(s)} = f_\mu(\zeta_j) + V_{ij} + \xi_{ij}^{(s)}, \quad V_i(t) = \sum_k^{2^L} Z_{ik}^{(s)} u_k(t) \quad \xi_{ij}^{(s)} \sim N(0, 1),$$

with $V_{ij} = V_i(\zeta_j)$, $u_k(t) = (2^{L/2} \|\phi\|_\infty)^{-1} \phi_k(t)$ and $Z_{ik}^{(s)}$ i.i.d. truncated standard normal $N_{[-1,1]}(0, 1)$.

Note that the function $X_i = f_\mu + V_i$ has mean f_μ and is in $\mathcal{H}^\alpha(R)$ for $R > 1$, for small enough $\|\mu\|_2$ and large enough m . Write $\phi_d := \left(\phi_1(\frac{j}{m})\right)_{j \in [d]}$ and $u = \frac{\phi_d}{2^{L/2} \|\phi\|_\infty}$. The vector $V_i = (V_{ij})_{i \in [n]}$ has covariance matrix proportional to $\Sigma := \text{diag}(uu^T, \dots, uu^T)$. We denote by $Y_i^{(s)}$ the vector $(Y_{ij}^{(s)})_{j \in [m]}$ and by $f_\mu(T)$ the vector $(f_\mu(\zeta_j))_{j \in [m]}$

Wherever it is non-zero, the likelihood for the i -th individual satisfies

$$p_\mu(Y_i^{(s)}) \propto \exp \left(-\frac{1}{2} \left(Y_i^{(s)} - f_\mu(T) \right)^T (\Sigma + I_m)^{-1} \left(Y_i^{(s)} - f_\mu(T) \right) \right),$$

where the normalizing constant can be seen to not depend on μ .

Writing ζ for the vector $(\zeta_1, \dots, \zeta_m)$, $f_\mu(\zeta)$ can be written as

$$f_\mu(\zeta) = \begin{pmatrix} \sum_k \mu_k \phi_k(\zeta_1) \\ \vdots \\ \sum_k \mu_k \phi_k(\zeta_m) \end{pmatrix} = \begin{pmatrix} \phi_1(\zeta_1) & \cdots & \phi_N(\zeta_1) \\ \vdots & \ddots & \vdots \\ \phi_1(\zeta_m) & \cdots & \phi_N(\zeta_m) \end{pmatrix} \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_{2^L} \end{pmatrix} = \Phi^T \mu,$$

where Φ is an $2^L \times m$ matrix. Hence, its score in μ is given by

$$S_f^{Y_i^{(s)}} := \Phi(\Sigma + I_m)^{-1} (Y_i^{(s)} - f_\mu(\zeta)) \quad (35)$$

Hence, the Fisher information for the i -th individual $-\mathbb{E}_f \frac{\partial^2}{\partial \mu^2} \log p_\mu(Y_i^{(s)})$ is bounded above by

$$\Phi(\Sigma + I_m)^{-1} \Phi^\top.$$

By Lemma 14, we have that

$$\Phi(\Sigma + I_m)^{-1} \Phi^\top \preceq 2^L \|\phi\|_\infty^2 I_{2^L}.$$

Write $S_f^Y = \sum_{i=1}^{n_s} S_f^{Y_i^{(s)}}$ and

$$\mathbf{C}_f(T) = \mathbb{E}[S_f^Y | T] \mathbb{E}[S_f^Y | T]^T.$$

By combining the previous inequality with Lemma 15 and 13, we have that

$$\mathbb{E}\text{Tr}(\mathbf{C}_f(T)) \lesssim 2n_s \sqrt{2^L} \varepsilon_s \sqrt{\mathbb{E}\text{Tr}(\mathbf{C}_f(T))} + n_s \delta_s \cdot 2^L (\sqrt{m_s n_s} \vee m_s) \log(1/\delta_s) + n_s \delta_s. \quad (36)$$

If $\sqrt{\text{Tr}(\mathbb{E}(\mathbf{C}_f(T)))} \leq \sqrt{2^L} n_s \varepsilon_s$, then there is nothing to prove. So assume instead that $\sqrt{\text{Tr}(\mathbb{E}(\mathbf{C}_f(T)))} \geq \sqrt{2^L} n_s \varepsilon_s$. Combining the above display with (36), we get

$$\sqrt{\mathbb{E}\text{Tr}(\mathbf{C}_f(T))} \lesssim 2\sqrt{2^L} n_s \varepsilon_s + \delta_s \varepsilon_s^{-1} n_s^{-1} \sqrt{2^L} (\sqrt{m_s n_s} \vee m_s) \log(1/\delta_s) + \frac{\delta_s}{\sqrt{2^L} \varepsilon_s}.$$

The assumption of small enough δ' ($\delta_s \log(1/\delta_s) \lesssim \frac{n_s^2 \varepsilon_s^2}{\sqrt{m_s n_s} \vee m_s}$), implies that the last two terms are $o(\sqrt{m_s n_s} \varepsilon_s)$ which yields that

$$\text{Tr}(\mathbb{E}(\mathbf{C}_f(T))) \lesssim 2^L n_s^2 \varepsilon_s^2.$$

We also have $I_f(\mathbf{Z}) = \text{Var}(S_f(Z_1)) \lesssim 2^L I_{2^L}$ as already shown. Hence we have that $\text{Tr}(\mathbf{C}_f(T)) \leq (2^L)^2 n_s$. Using the von-Trees construction of Lemma 3, we have that

$$\sup_{f \in \mathcal{H}_\alpha^R} \mathbb{E}_f \|\hat{f} - f\|_2^2 \gtrsim \frac{2^{2L}}{\sum_{s=1}^S (2^L n_s^2 \varepsilon_s^2) \wedge (n_s 2^{2L}) + \pi^2 (2^L)^{2\alpha+2}}.$$

Optimizing over L , the result follows.

Lemma 13. Consider $\Sigma = \text{diag} \left(\underbrace{uu^T, \dots, uu^T}_{N \text{ times}} \right)$ where $u = c_L \phi_d$, $\phi_d := \left(\phi_1(\frac{j}{m}) \right)_{j \in [d]}$, $c_L = \frac{1}{\sqrt{2^L}} \|\phi\|_\infty^{-1}$. and define

$$\Phi = \begin{pmatrix} \phi_1(\zeta_1) & \cdots & \phi_{2^L}(\zeta_1) \\ \vdots & \ddots & \vdots \\ \phi_1(\zeta_m) & \cdots & \phi_{2^L}(\zeta_m) \end{pmatrix}.$$

It holds that

$$\Phi (\Sigma + I_m)^{-1} \Phi^T = a_L I_{2^L} \preceq 2^L \|\phi\|_\infty^2 I_{2^L}.$$

where $a_L = \frac{\|\phi_d\|_2^2}{1 + c_L^2 \|\phi_d\|_2^2}$.

Proof. We have

$$\begin{aligned} \Phi (\Sigma + I_m)^{-1} \Phi^T &= \Phi \left(\text{diag} (uu^T, \dots, uu^T) + I_m \right)^{-1} \Phi^T \\ &= \Phi \left(\text{diag} \left[\left(I_d - \frac{uu^T}{1 + u^T u} \right), \dots, \left(I_d - \frac{uu^T}{1 + u^T u} \right) \right] \right) \Phi^T \\ &= \Phi \Phi^T - \frac{1}{1 + u^T u} \Phi \text{diag} (uu^T, \dots, uu^T) \Phi^T. \end{aligned}$$

Leveraging the fact that the ϕ 's have disjoint support, we have that

$$\begin{aligned}\Phi\Phi^T &= \begin{bmatrix} \sum_{j=1}^m \phi_1^2(\zeta_j) & \sum_{j=1}^m \phi_1(\zeta_j)\phi_2(\zeta_j) & \cdots \\ \sum_{j=1}^m \phi_1(\zeta_j)\phi_2(\zeta_j) & \sum_{j=1}^m \phi_2^2(\zeta_j) & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix} \\ &= \begin{bmatrix} \sum_{j=1}^m \phi_1^2(\zeta_j) & 0 & \cdots \\ 0 & \sum_{j=1}^m \phi_2^2(\zeta_j) & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}.\end{aligned}$$

Recalling that $\phi_k(x) = 2^{L/2+K/2}\phi(2^{L+K}x - k)$ and $d \in \mathbb{N}$ such that $m = \lfloor d2^L \rfloor$, we have that $\sum_{j=1}^m \phi_k^2(\zeta_j) \leq 2 \sum_{j=1}^d \phi_1^2(\zeta_j)$. This implies that

$$\Phi\Phi^T = 2 \left(\sum_{j=1}^d \phi_1^2(\zeta_j) \right) I_{2L}.$$

Similarly, we have that

$$\Phi \operatorname{diag}(uu^T, \dots, uu^T) \Phi^T \preceq \begin{pmatrix} \phi_d^T uu^T \phi_d & 0 & 0 & \cdots \\ 0 & \phi_d^T uu^T \phi_d & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} = (u^\top \phi_d)^2 I_{2L}$$

We have that $u = c_L \phi_d$. Denote $c_L = (2^{L/2} \|\phi\|_\infty)^{-1}$ and note that $u^\top \phi_d = c_L \|\phi_d\|_2^2$ and $u^\top u = c_L^2 \|\phi_d\|_2^2$. The result now follows:

$$\begin{aligned}\Phi(\Sigma + I_m)^{-1} \Phi^T &\preceq \left(\|\phi_d\|_2^2 - \frac{c_L^2 \|\phi_d\|_2^4}{1 + c_L^2 \|\phi_d\|_2^2} \right) I_{2L} \\ &= \frac{\|\phi_d\|_2^2}{1 + c_L^2 \|\phi_d\|_2^2} I_{2L} \preceq \frac{1}{c_L^2} I_{2L} \preceq 2^L \|\phi\|_\infty^2 I_{2L}.\end{aligned}$$

□

Lemma 14. For $M \asymp 2^L(\sqrt{m_s n_s} \vee m_s) \log(1/\delta_s)$ we have that

$$2M\delta_s + \int_M^\infty \mathbb{P}(|G_1| \geq t) dt + \int_M^\infty \mathbb{P}(|\check{G}_1| \geq t) dt \lesssim \delta_s 2^L(\sqrt{m_s n_s} \vee m_s) \log(1/\delta_s) + \delta_s.$$

where $S_f^{Y_i} := \Phi(\Sigma + I_m)^{-1}(Y_i - f_\mu(T))$, $S_f^{Y^{(s)}} := \sum_{i=1}^n S_f^{Y_i}$ and $G_i = \langle \mathbb{E}[S_f^{Y_i}|T], S_f^{Y_i} \rangle$ and $\check{G}_i = \langle \mathbb{E}[S_f^{Y_i}|T], \check{S}_f^{Y_i^{(s)}} \rangle$ with $\check{S}_f^{Y_i^{(s)}}$ an independent copy $S_f^{Y_i^{(s)}}$.

Proof of Lemma 14. We proceed along the same line as in Lemma16. We first analyze the term concerning $G_i = \langle \mathbb{E}[S_f^{Y^{(s)}}|T], S_f^{Y^{(s)}} \rangle$. Using Jensen's and the law of total probability, we have that for $i = 1, \dots, n_s$,

$$\mathbb{E} e^{t|G_i|} \leq \mathbb{E} e^{t|\langle S_f^{Y^{(s)}}, S_f^{Y^{(s)}} \rangle|}.$$

Without loss of generality, consider the index i to be equal to 1. Note that we have that $\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \rangle = \sum_{i=1}^{n_s} \langle S_f^{Y_i^{(s)}}, S_f^{Y_1^{(s)}} \rangle$.

Using our model we have that

$$S_f^{Y_i^{(s)}} = \Phi(\Sigma + I_{m_s})^{-1}(\xi_i + V_i)$$

This implies that $S_f^{Y_i^{(s)}} \sim N(0, \Phi(\Sigma + I_{m_s})^{-1}\Phi^T) = N(0, a_{2L}I_d)$, where a_{2L} is given by Lemma 13. Using the linearity of expectation, we have that $\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \rangle = \sum_{i=1}^{n_s} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle$. We can split the sum into two terms

$$\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \rangle = \|S_f^{Y_1^{(s)}}\|_2^2 + \sum_{i=2}^{n_s} \sum_{j=1}^{m_s} \xi_{ij} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle$$

Hence we have that

$$\begin{aligned} \mathbb{E} e^{t|\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \rangle|} &\leq \mathbb{E} \left(e^{t\|S_f^{Y_1^{(s)}}\|_2^2} e^{t \left| \sum_{i=2}^{n_s} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle \right|} \right) \\ &= \mathbb{E} \left(e^{t\|S_f^{Y_1^{(s)}}\|_2^2} \mathbb{E} \left(e^{t \left| \sum_{i=2}^{n_s} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle \right|} \mid Z_1 \right) \right) \end{aligned}$$

Conditional on Z_1 , we have that $\sum_{i=2}^{n_s} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle$ behaves as follows

$$\sum_{i=2}^{n_s} \langle \mathbf{S}_f(Z_i), S_f^{Y_1^{(s)}} \rangle \mid Z_1 \sim N(0, (n_s - 1)a_{2L}\|S_f(Z_1)\|_2^2)$$

Hence, using the m.g.f of Gaussian, we have that

$$\begin{aligned} \mathbb{E} e^{t|\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \rangle|} &\leq 2\mathbb{E} \left(e^{t\|S_f^{Y_1^{(s)}}\|_2^2} e^{\frac{1}{2}t^2(n_s-1)a_{2L}\|S_f^{Y_1^{(s)}}\|^2} \right) \\ &= 2\mathbb{E} \left(e^{(t+\frac{1}{2}t^2(n_s-1)a_{2L})\|S_f(Z_1)\|^2} \right) \\ &= 2\mathbb{E} \left(e^{(t+\frac{1}{2}t^2(n_s-1)a_{2L})a_{2L}\sum_{j=1}^{m_s} U_j^2} \right) \\ &= 2\mathbb{E} \left(e^{(t+\frac{1}{2}t^2(n_s-1)a_{2L})a_{2L}m_s U_1^2} \right) \end{aligned}$$

where we have used the fact that $S_f^{Y_i^{(s)}} \sim N(0, a_{2L}I_d)$ hence $\|S_f^{Y_i^{(s)}}\|_2 = a_{2L} \sum_{i=1}^{m_s} U_i^2$ with U_i as i.i.d standard Gaussians. The last expectation is bounded by a constant if $(t + \frac{1}{2}t^2(n_s - 1)a_{2L})a_{2L}m_s \leq \frac{1}{2}$ (using the fact that U_1^2 is χ_1^2). Hence if $t \lesssim \frac{1}{m_s 2^L} \wedge \frac{1}{\sqrt{m_s n_s} 2^L}$ then $\mathbb{E} e^{t|G_1|} \leq c$. It follows that

$$\begin{aligned} \mathbb{P}(|G_1| > u) &\leq \mathbb{P}(e^{t|G_1|} \geq e^{tu}) \\ &\leq e^{-tu} \mathbb{E} e^{t|G_1|} \\ &\leq 2e^c e^{-u \left(\frac{1}{m_s 2^L} \wedge \frac{1}{\sqrt{m_s n_s} 2^L} \right)}. \end{aligned}$$

This means that for $M \gtrsim 2^L(\sqrt{m_s n_s} \vee m_s) \log(1/\delta_s)$, we obtain

$$\int_M^\infty \mathbb{P}(|G_i| \geq t) dt \lesssim 2e^{-\log(1/\delta_s)} = 2\delta_s.$$

In a similar fashion we can show that $\int_M^\infty \mathbb{P}(|\check{G}_i| \geq t) dt \lesssim 2\delta_s$. $\square$

8.3 Differential privacy data processing inequalities for the Fisher information

Lemma 15. *Let $\mathbf{C}_f(T)$ denote the matrix*

$$\mathbb{E}_f \left[S_f^{Y^{(s)}} \mid T \right] \mathbb{E}_f \left[S_f^{Y^{(s)}} \mid T \right]^T, \quad (37)$$

where $S_f^{Y^{(s)}} = \sum_{i=1}^{n_s} S_f^{Y_i^{(s)}}$ denotes the score function of f for the observational model $Y^{(s)}$ as defined in (34).

It holds that $\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T)) \leq 2^L m_s n_s$.

Furthermore, assuming δ_s is such that $\delta_s \log(1/\delta_s) \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2 L^{-1}$, it also holds that

$$\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T)) \lesssim m_s n_s^2 \varepsilon_s^2.$$

Proof of Lemma 15. The first assertion follows by standard data processing arguments for the Fisher information, see e.g. [34]. By Lemma 4.2 in [8], there exists a universal constant $C > 0$ such that

$$\begin{aligned} \mathbb{E}_f \text{Tr}(\mathbf{C}_f(T)) &\leq C n_s \varepsilon_s \sqrt{\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T))} \sqrt{\lambda_{\max}(\text{Var}(S_f^{Y_i^{(s)}}))} \\ &\quad + 2n_s M \delta_s + n_s \int_M^\infty \mathbb{P}(|G_i| \geq t) dt + n_s \int_M^\infty \mathbb{P}(|\check{G}_i| \geq t) dt, \end{aligned}$$

with

$$G_i = \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T], S_f^{Y_i^{(s)}} \right\rangle \text{ and } \check{G}_i = \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T], \check{S}_f^{Y_i^{(s)}} \right\rangle, \quad (38)$$

and $\check{S}_f^{Y_i^{(s)}}$ denotes the data set $S_f^{Y_i^{(s)}}$ with its i th data point $Y_i^{(s)}$ replaced by an independent copy $\check{Y}_i^{(s)}$.

Note that $\text{Var}(S_f^{Y_i^{(s)}}) = m_s I_L$,

$$\text{Var}_f(S_f^{Y^{(s)}}) = \sum_{j=1}^{m_s} \left(Y_{ij}^{(s)} - \sum_{k=1}^{2^L} f_k \phi_l(\zeta_{ij}^{(s)}) \right) \phi(\zeta_{ij}^{(s)})$$

hence $\lambda_{\max}(\text{Var}(S_f^{Y^{(s)}})) = m_s$. Lemma 16 yields that for $M \asymp 2^L(\sqrt{n_s} m_s^{3/2} \vee m_s^2) \log(1/\delta_s)$, we have that

$$2M\delta_s + \int_M^\infty \mathbb{P}(|G_1| \geq t) dt + \int_M^\infty \mathbb{P}(|\check{G}_1| \geq t) dt \lesssim \delta_s 2^L(\sqrt{n_s} m_s^{3/2} \vee m_s^2) \log(1/\delta_s) + \delta_s.$$

Using Lemma 16, we have that

$$\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T)) \lesssim 2n_s \sqrt{m_s} \varepsilon_s \sqrt{\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T))} + n_s \delta_s 2^L (\sqrt{n_s} m_s^{3/2} \vee m_s^2) \log(1/\delta_s) + n_s \delta_s. \quad (39)$$

If $\sqrt{\text{Tr}(\mathbb{E}_f(\mathbf{C}_f(T)))} \leq \sqrt{m_s} n_s \varepsilon_s$, then there is nothing to prove. So assume instead that $\sqrt{\text{Tr}(\mathbb{E}_f(\mathbf{C}_f(T)))} \geq \sqrt{m_s} n_s \varepsilon_s$. Combining the above display with (39), we get

$$\sqrt{\mathbb{E}_f \text{Tr}(\mathbf{C}_f(T))} \lesssim 2\sqrt{m_s} n_s \varepsilon_s + \delta_s \varepsilon_s^{-1} 2^L (\sqrt{n_s} m_s \vee m_s^{3/2}) \log(1/\delta_s) + \frac{\delta_s}{\sqrt{m_s} \varepsilon_s}.$$

The assumption of small enough δ_s ($\delta_s \log(1/\delta_s) \lesssim \left(\frac{n_s}{m_s} \wedge \sqrt{\frac{n_s}{m_s}}\right) \varepsilon_s^2 2^{-L}$) implies that the last two terms are $o(\sqrt{m_s} n_s \varepsilon_s)$ which yields the result. $\square$

Lemma 16. *Let*

$$G_i = \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T], S_f^{Y_i^{(s)}} \right\rangle, \quad \check{G}_i = \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T], \check{S}_f^{Y_i^{(s)}} \right\rangle,$$

and $\check{S}_f^{Y_i^{(s)}}$ denotes the data set $S_f^{Y_i^{(s)}}$ with its i th data point $Y_i^{(s)}$ replaced by an independent copy $\check{Y}_i^{(s)}$.

For $M \asymp 2^L (\sqrt{n_s} m_s^{3/2} \vee m_s^2) \log(1/\delta_s)$ we have that

$$2M\delta_s + \int_M^\infty \mathbb{P}_f(|G_1| \geq t) dt + \int_M^\infty \mathbb{P}_f(|\check{G}_1| \geq t) dt \lesssim \delta_s 2^L (\sqrt{n_s} m_s^{3/2} \vee m_s^2) \log(1/\delta_s) + \delta_s.$$

Proof of Lemma 16. We first analyze the term with G_i in the integrand. Using Jensen's inequality and the law of total probability, we have for $i = 1, \dots, n_s$,

$$\begin{aligned} \mathbb{E}_f \exp(t |G_i|) &= \mathbb{E}_f \exp\left(t \left| \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T], S_f^{Y_i^{(s)}} \right\rangle \right|\right) \\ &= \int \int \exp\left(t \left| \left\langle \mathbb{E}_f[S_f^{Y^{(s)}} \mid T = u], S_f^{Y_i^{(s)}}(y_i) \right\rangle \right|\right) dP^{T|Y_i^{(s)}=y_i}(u) dP^{Y_i^{(s)}}(y_i) \\ &\leq \int \int \mathbb{E}_f \left[\exp\left(t \left| \left\langle S_f^{Y^{(s)}}, S_f^{Y_i^{(s)}}(y_i) \right\rangle \right|\right) \mid T = u \right] dP^{T|Y_i^{(s)}=y_i}(u) dP^{Y_i^{(s)}}(y_i) \\ &= \mathbb{E}_f \exp\left(t \left| \left\langle S_f^{Y^{(s)}}, S_f^{Y_i^{(s)}} \right\rangle \right|\right). \end{aligned}$$

Without loss of generality, consider $i = 1$. we have

$$\left\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle = \sum_{i=1}^{n_s} \left\langle S_f^{Y_i^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle = \sum_{i=1}^{n_s} \sum_{j=1}^{m_s} \xi_{ij}^{(s)} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle,$$

where $\xi_{ij}^{(s)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. We have

$$\left\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle = \left\| S_f^{Y_1^{(s)}} \right\|_2^2 + \sum_{i=2}^{n_s} \sum_{j=1}^{m_s} \xi_{ij}^{(s)} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle.$$

Hence,

$$\mathbb{E}_f \exp \left(t \left| \left\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle \right| \right) \leq \mathbb{E}_f \left(\exp \left(t \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \exp \left(t \left| \sum_{j=2}^{n_s} \left\langle S_f^{Y_j^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle \right| \right) \right)$$

which in turn equals

$$\mathbb{E}_f \left[\exp \left(t \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \mathbb{E}_f \left(\exp \left(t \left| \sum_{j=2}^{n_s} \left\langle S_f^{Y_j^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle \right| \right) \middle| \{\zeta_{ij}^{(s)}\}, Y_1^{(s)} \right) \right].$$

Conditionally on $\{\zeta_{ij}^{(s)}\}$ and $Y_1^{(s)}$, the sum $\sum_{j=2}^{n_s} \left\langle S_f^{Y_j^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle$ is normal with mean zero, since

$$\sum_{j=2}^{n_s} \sum_{j=1}^{m_s} \xi_{ij}^{(s)} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle \middle| [\{\zeta_{ij}^{(s)}\}, Y_1^{(s)}] \sim \mathcal{N} \left(0, \sum_{j=2}^{n_s} \sum_{j=1}^{m_s} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle^2 \right).$$

Using the moment generating function of the normal distribution, we have

$$\mathbb{E}_f \exp \left(t \left| \left\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle \right| \right) \leq 2 \mathbb{E}_f \left[\exp \left(t \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \exp \left(\frac{1}{2} t^2 \sum_{j=2}^{n_s} \sum_{j=1}^{m_s} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle^2 \right) \right]. \quad (40)$$

Next,

$$\sum_{j=2}^{n_s} \sum_{j=1}^{m_s} \left\langle \phi(\zeta_{ij}^{(s)}), S_f^{Y_1^{(s)}} \right\rangle^2 \leq \sum_{j=2}^{n_s} \sum_{j=1}^{m_s} \left\| \phi(\zeta_{ij}^{(s)}) \right\|_2^2 \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \leq c_\phi^2 m_s (n_s - 1) 2^L \left\| S_f^{Y_1^{(s)}} \right\|_2^2.$$

Hence, as a consequence of (40),

$$\begin{aligned} \mathbb{E}_f \exp \left(t \left| \left\langle S_f^{Y^{(s)}}, S_f^{Y_1^{(s)}} \right\rangle \right| \right) &\leq 2 \mathbb{E}_f \left[\exp \left(t \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \exp \left(\frac{1}{2} t^2 c_\phi^2 m_s (n_s - 1) 2^L \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \right] \\ &= 2 \mathbb{E}_f \left[\exp \left(\left(t + \frac{1}{2} t^2 c_\phi^2 m_s (n_s - 1) 2^L \right) \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \right]. \end{aligned} \quad (41)$$

Setting $s = t + \frac{1}{2} t^2 c_\phi^2 m_s (n_s - 1) 2^L$, we proceed to bound $\mathbb{E} \exp \left(s \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right)$. Observe that

$$\left\| S_f^{Y_1^{(s)}} \right\|_2^2 = \sum_{k=1}^{2^L} \left(\sum_{j=1}^{m_s} \xi_{1j} \phi_{Lk}(\zeta_{1j}^{(s)}) \right)^2 \leq \sum_{k=1}^{2^L} \left(\sum_{j=1}^{m_s} \xi_{1j}^2 \right) \left(\sum_{j=1}^{m_s} \phi_{Lk}^2(\zeta_{1j}^{(s)}) \right) \leq c_\phi^2 m_s 2^L \sum_{j=1}^{m_s} \xi_{1j}^2.$$

Hence by independence and the m.g.f. bound for ξ_{1j}^2 ,

$$\mathbb{E}_f \exp \left(s \left\| S_f^{Y_1^{(s)}} \right\|_2^2 \right) \leq \mathbb{E}_f \exp \left(s c_\phi^2 m_s 2^L \sum_{j=1}^{m_s} \xi_{1j}^2 \right) = (\mathbb{E}_f \exp(s c_\phi^2 m_s 2^L \xi_{11}^2))^{m_s}.$$

One checks this is bounded by a constant as soon as $s \leq \frac{1}{m_s^2 2^L}$, which, recalling that $s = t + \frac{1}{2} t^2 c_\phi^2 m_s (n_s - 1) 2^L$, is the case when $t \lesssim \frac{1}{m_s^2 2^L} \wedge \frac{1}{\sqrt{n_s} m_s^{3/2} 2^L}$.

Within this regime, $\mathbb{E} \exp(t |G_1|) \leq e^c$ for some constant $c > 0$. Consequently,

$$\begin{aligned} \mathbb{P}(|G_1| > u) &\leq \mathbb{P}(\exp(t |G_1|) \geq \exp(t u)) \\ &\leq \exp(-t u) \mathbb{E} \exp(t |G_1|) \\ &\leq 2 e^c \exp\left(-u \left(\frac{1}{\sqrt{n_s} m_s^{3/2} 2^L} \wedge \frac{1}{m_s^2 2^L}\right)\right). \end{aligned}$$

Thus, setting

$$M \gtrsim 2^L \left(\sqrt{n_s} m_s^{3/2} \vee m_s^2\right) \log\left(\frac{1}{\delta_s}\right),$$

one gets

$$\int_M^\infty \mathbb{P}(|G_1| \geq t) dt \lesssim 2 \exp\left(-\log\left(\frac{1}{\delta_s}\right)\right) = 2 \delta_s.$$

We now turn to the term $\int_M^\infty \mathbb{P}(|\check{G}_1| \geq t) dt$. By a similar argument, observe

$$\mathbb{E}_f \exp\left(t |\check{G}_i|\right) = \mathbb{E}_f \exp\left(t \left| \left\langle \mathbb{E}_f[S_f^{Y(s)} \mid T], \check{S}_f^{Y(s)} \right\rangle \right|\right) \leq \mathbb{E}_f \exp\left(t \left| \left\langle S_f^{Y(s)}, \check{S}_f^{Y(s)} \right\rangle \right|\right).$$

By symmetry,

$$\left\langle S_f^{Y(s)}, \check{S}_f^{Y_1(s)} \right\rangle = \sum_{j=1}^{n_s} \sum_{s=1}^{m_s} \xi_{ij} \left\langle \phi(Y_{js}^{(s)}), \check{S}_f^{Y_1(s)} \right\rangle.$$

The same argument as before yields

$$\mathbb{E}_f \exp\left(t \left| \left\langle S_f^{Y(s)}, \check{S}_f^{Y_1(s)} \right\rangle \right|\right) \leq 2 \mathbb{E}_f \left[\exp\left(\frac{1}{2} t^2 c_\phi^2 m_s n_s 2^L \left\| \check{S}_f^{Y_1(s)} \right\|_2^2\right) \right].$$

As before, this remains bounded when $t \lesssim \frac{1}{m_s^2 2^L} \wedge \frac{1}{\sqrt{n_s} m_s^{3/2} 2^L}$. Hence the same exponential tail bound applies;

$$\int_M^\infty \mathbb{P}(|\check{G}_1| \geq t) dt \lesssim 2 \exp\left(-\log\left(\frac{1}{\delta_s}\right)\right) = 2 \delta_s \quad \text{for } M \gtrsim 2^L \left(\sqrt{n_s} m_s^{3/2} \vee m_s^2\right) \log\left(\frac{1}{\delta_s}\right).$$

Putting all pieces together yields the statement of the lemma. $\square$

Lemma 17. Consider $X_{L;i}^{(s)} = (X_{1i}^{(s)}, \dots, X_{2^L i}^{(s)})$ where $X_{ki}^{(s)}$ is defined as in (32). It holds that $\text{Tr}(I_f^{X(s)}) \leq n_s 2^L$ and if $0 \leq \delta_s \leq \left((n_s 2^{-L} \wedge n_s^{1/2} 2^{-L/2}) \varepsilon_s^2\right)^{1+p}$ for some $p > 0$, we also have that

$$\text{Tr}(I_f^{X(s)}) \leq C n_s^2 \varepsilon_s^2$$

for a constant $C > 0$.

Proof. The bound $\text{Tr}(I_f^{X^{(s)}}) \leq n_s 2^L$ follows by the fact that conditional expectation contracts the L_2 -norm. For the second statement, we start introducing the notations $\overline{X^{(s)}} = n_s^{-1} \sum_{i=1}^{n_s} X_{L;i}^{(s)}$ and

$$G_i = \left\langle \mathbb{E}_0 \left[n_s \overline{X^{(s)}} \mid T^{(s)} \right], X_{L;i}^{(s)} \right\rangle.$$

For the remainder of the proof, consider versions of $X^{(s)}$ and $T^{(s)}$ defined on the same probability space, and we shall write as a shorthand

$$\mathbb{P}^s \equiv \mathbb{P}_0^{(X^{(s)}, T^{(s)})} \quad \text{and} \quad \mathbb{E}^s \equiv \mathbb{E}_0^{(X^{(s)}, T^{(s)})}.$$

For random variables V, W defined on the same probability space, it holds that

$$\mathbb{E}W\mathbb{E}[W \mid V] = \mathbb{E}\mathbb{E}[W \mid V]\mathbb{E}[W \mid V],$$

since $W - \mathbb{E}[W \mid V]$ is orthogonal to $\mathbb{E}[W \mid V]$. Combining this fact with the linearity of the inner product and conditional expectation, we see that

$$\text{Tr}(I_f^{X^{(s)}}) = \mathbb{E}_0^{T^{(s)}} \left\| \mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}] \right\|_2^2 = \sum_{i=1}^{n_s} \mathbb{E}^s G_i.$$

Define also

$$\check{G}_i = \left\langle \mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}], \check{X}_{L;i}^{(s)} \right\rangle,$$

where $\check{X}_{L;i}^{(s)}$ is an independent copy of $X_{L;i}^{(s)}$ (defined on the same, possibly enlarged probability space) and note that $\mathbb{E}^s \check{G}_i = 0$. Write $G_i^+ := 0 \vee G_i$ and $G_i^- = -(0 \wedge G_i)$. We have

$$\begin{aligned} \mathbb{E}^s G_i^+ &= \int_0^\infty \mathbb{P}^s(G_i^+ \geq t) dt = \int_0^T \mathbb{P}^s(G_i^+ \geq t) dt + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt \\ &\leq e^{\varepsilon_s} \int_0^T \mathbb{P}^s(\check{G}_i^+ \geq t) dt + T\delta_s + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt \\ &\leq \int_0^T \mathbb{P}^s(\check{G}_i^+ \geq t) dt + 2\varepsilon_s \int_0^T \mathbb{P}^s(\check{G}_i^+ \geq t) dt + T\delta_s + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt \\ &\leq \int_0^\infty \mathbb{P}_0^s(\check{G}_i^+ \geq t) dt + 2\varepsilon_s \int_0^\infty \mathbb{P}^s(\check{G}_i^+ \geq t) dt + T\delta_s + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt, \end{aligned}$$

Similarly, we obtain

$$\begin{aligned} \mathbb{E}^s G_i^- &\geq \int_0^T \mathbb{P}^s(G_i^- \geq t) dt \\ &\geq e^{-\varepsilon_s} \int_0^T \mathbb{P}^s(\check{G}_i^- \geq t) dt - T\delta_s \\ &\geq \int_0^T \mathbb{P}^s(\check{G}_i^- \geq t) dt - 2\varepsilon_s \int_0^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt - T\delta_s \\ &\geq \int_0^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt - 2\varepsilon_s \int_0^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt - T\delta_s - \int_T^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt. \end{aligned}$$

Putting these together with $G_i = G_i^+ - G_i^-$, we get

$$\begin{aligned}\mathbb{E}^s G_i &\leq \int_0^\infty \mathbb{P}^s(\check{G}_i^+ \geq t) dt - \int_0^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt + 2\varepsilon_s \int_0^\infty \mathbb{P}^s(|\check{G}_i| \geq t) dt \\ &\quad + 2T\delta_s + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt + \int_T^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt \\ &= \mathbb{E}^s \check{G}_i + 2\varepsilon_s \mathbb{E}^s |\check{G}_i| + 2T\delta_s + \int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt + \int_T^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt.\end{aligned}$$

The first term in the last display equals 0. For the second term, observe that

$$\check{G}_i \left| \left[T^{(s)}, X^{(s)} \right] \right. \sim N(0, \|\mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}]\|_2^2),$$

so

$$\mathbb{E}^s |\check{G}_i| = \mathbb{E}^{X^{(s)}, T^{(s)}} \mathbb{E}^{\check{X}^{(s)}} |\check{G}_i| = \mathbb{E} \|\mathbb{E}[n_s \overline{X^{(s)}} \mid T^{(s)}]\|_2 \leq \sqrt{\text{Tr}(I_f^{X^{(s)}})},$$

where the last inequality is Cauchy-Schwarz. To bound the terms

$$\int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt + \int_T^\infty \mathbb{P}^s(\check{G}_i^- \geq t) dt,$$

we shall employ tail bounds, which follow after showing that G_i is $\sqrt{2^L n_s}$ -sub-exponential. To see this, note that by applying Cauchy-Schwarz and Jensen's inequality followed by the law of total probability, we have that

$$\begin{aligned}\mathbb{E}^s e^{t|G_i|} &= \mathbb{E}^s e^{t|\langle \mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}], X_{L;i}^{(s)} \rangle|} \\ &\leq \mathbb{E}^s e^{\frac{t}{2} \left(\|\mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}]\|_2^2 + \|X_{L;i}^{(s)}\|_2^2 \right)} \\ &\leq \sqrt{\mathbb{E}_0 e^{t \|\mathbb{E}_0[n_s \overline{X^{(s)}} \mid T^{(s)}]\|_2^2}} \sqrt{\mathbb{E}_0 e^{t \|X_{L;i}^{(s)}\|_2^2}} \\ &\leq \mathbb{E}_0 e^{t|\langle n_s \overline{X^{(s)}}, X_{L;i}^{(s)} \rangle|},\end{aligned}$$

where the last equality follows from the fact that conditional expectation contracts the L_1 -norm.

Next, we bound $\mathbb{E}_0^{X^{(s)}} e^{t|\langle n_s \overline{X^{(s)}}, X_{L;i}^{(s)} \rangle|}$. By the triangle inequality and Cauchy-Schwarz,

$$\mathbb{E}_0^{X^{(s)}} e^{t|\langle n_s \overline{X^{(s)}}, X_{L;i}^{(s)} \rangle|} \leq \sqrt{\mathbb{E}_0^{X^{(s)}} e^{2t|\langle \sum_{k \neq i}^{n_s} X_{Lk}^{(s)}, X_{L;i}^{(s)} \rangle|}} \sqrt{\mathbb{E}_0^{X_{L;i}^{(s)}} e^{2t|\langle X_{L;i}^{(s)}, X_{L;i}^{(s)} \rangle|}}.$$

The random variable $\langle X_{L;i}^{(s)}, X_{L;i}^{(s)} \rangle$ is the inner product of centered beta(2, 2) distributed random variables, so

$$\mathbb{E}_0^{X_{L;i}^{(s)}} e^{2t|\langle X_{L;i}^{(s)}, X_{L;i}^{(s)} \rangle|} \leq e^{2^{L-1}t}.$$

By Lemma 19,

$$\mathbb{E}_0^{X^{(s)}} e^{2t|\langle \sum_{k \neq i}^{n_s} X_{Lk}^{(s)}, X_{L;i}^{(s)} \rangle|} \leq e^{ct^2(n_s-1)2^L},$$

for some $c > 0$. By the fact that $G_i^+ \leq |G_i|$ and Markov's inequality,

$$\mathbb{P}^s(G_i^+ > T) \leq \mathbb{P}^s(|G_i| > T) \leq e^{-tT} \mathbb{E}^s e^{t|G_i|}, \text{ for all } T, t > 0.$$

Combining this with the bound for the moment generating function derived above means that for $\delta_s = 0$, the result follows from letting $T \rightarrow \infty$. If $\delta_s > 0$, take $T = 32(2^L \vee \sqrt{2^L n_s}) \log(1/\delta_s)$ to obtain that

$$\int_T^\infty \mathbb{P}^s(G_i^+ \geq t) dt \leq e^{-\log(1/\delta_s)}.$$

It is easy to see that the same bound applies to $\int_T^\infty \mathbb{P}_0^s(\check{G}_i^- \geq t) dt$. We obtain that

$$\sum_{i=1}^{n_s} \mathbb{E}^s G_i \leq 2n_s \varepsilon_s \sqrt{\text{Tr}(I_f^{X(s)})} + 64\delta_s(2^L \vee \sqrt{2^L n_s}) \log(1/\delta_s) + 2n_s \delta_s.$$

If $\sqrt{\text{Tr}(I_f^{X(s)})} \leq n_s \varepsilon_s$, the lemma holds (there is nothing to prove). So assume instead that $\sqrt{\text{Tr}(I_f^{X(s)})} \geq n_s \varepsilon_s$. We find that

$$\sqrt{\text{Tr}(I_f^{X(s)})} \leq 2n_s \varepsilon_s + 64\delta_s \frac{2^L \vee \sqrt{2^L n_s}}{n_s \varepsilon_s} \log(1/\delta_s) + \frac{2}{\varepsilon_s} \delta_s.$$

Since $x^p \log(1/x)$ tends to 0 as $x \rightarrow 0$ for any $p > 0$, the result follows for $\delta_s \leq \left(\left(\frac{n_s}{2^L} \wedge \frac{n_s^{1/2}}{\sqrt{2^L}}\right) \varepsilon_s^2\right)^{1+p}$ for some $p > 0$ as this implies that the last two terms are $O(n_s \varepsilon_s)$. $\square$

8.3.1 Auxiliary lemmas

The following lemmas are folklore that involve straightforward calculations, but included for completeness.

Lemma 18. *Let Z be $N(0, 1)$, $0 \leq \lambda \leq 1/4$. Then,*

$$\mathbb{E} e^{\lambda(Z^2-1)} \leq e^{2\lambda^2}.$$

Proof. Using the change of variables $u = z\sqrt{1-2\lambda}$,

$$\begin{aligned} \mathbb{E} e^{\lambda(Z^2-1)} &= \frac{1}{\sqrt{2\pi}} \int e^{\lambda(z^2-1)} e^{-\frac{1}{2}z^2} dz \\ &= \frac{e^{-\lambda}}{\sqrt{2\pi(1-2\lambda)}} \int e^{-\frac{1}{2}z^2} dz = \frac{e^{-\lambda}}{\sqrt{(1-2\lambda)}}. \end{aligned}$$

The MacLaurin series of $-\frac{1}{2} \log(1-2\lambda)$ reads

$$\frac{1}{2} \sum_{k=1}^{\infty} \frac{(2\lambda)^k}{k},$$

which yields that the second last display equals

$$\exp\left(\frac{3}{2}\lambda^2 + \frac{1}{2}\sum_{k=3}^{\infty}\frac{(2\lambda)^k}{k}\right).$$

If $\lambda \leq 1/4$,

$$\sum_{k=3}^{\infty}\frac{(2\lambda)^k}{k} \leq \frac{(2\lambda)^3}{1-2\lambda} \leq \lambda^2,$$

from which the result follows. $\square$

Lemma 19. *Let $a \in \mathbb{R}$ and let $Z, Z' \stackrel{i.i.d.}{\sim} N(0, I_d)$ for $d \in \mathbb{N}$.*

Then, $a\langle Z, Z' \rangle$ is $Ca\sqrt{d}$ -sub-exponential for a universal constant $C > 0$ and

$$\mathbb{E}e^{t|a\langle Z, Z' \rangle|} \leq 2e^{t^2a^2d},$$

whenever $|t| \leq (2a^2)^{-1}$.

Proof. Since $\langle Z, Z' \rangle | Z' \sim N(0, \|Z'\|_2^2)$,

$$\mathbb{E}e^{ta\langle Z, Z' \rangle} = \mathbb{E}^{Z'} \mathbb{E}^{Z|Z'} e^{ta\langle Z, Z' \rangle} = \mathbb{E}^{Z'} e^{\frac{t^2a^2}{2}\|Z'\|_2^2}.$$

By Lemma 18, the latter is further bounded by

$$e^{\frac{t^2a^2d}{2} + \frac{t^4a^4d}{2}} \leq e^{t^2a^2d},$$

whenever $t^2a^2 \leq 1/2$. The conclusion then follows by e.g. Proposition 2.7.1 in [31], since $\langle Z, Z' \rangle$ is mean zero. For the last statement,

$$\langle Z, Z' \rangle | Z' \stackrel{d}{=} -\langle Z, Z' \rangle | Z'.$$

Consequently,

$$\begin{aligned} \mathbb{E}^{Z|Z'} e^{t|a\langle Z, Z' \rangle|} &= \mathbb{E}^{Z|Z'} \mathbb{1}_{\{\langle Z, Z' \rangle > 0\}} e^{ta\langle Z, Z' \rangle} + \mathbb{E}^{Z|Z'} \mathbb{1}_{\{\langle Z, Z' \rangle \leq 0\}} e^{-ta\langle Z, Z' \rangle} \\ &\leq 2\mathbb{E}^{Z|Z'} e^{ta\langle Z, Z' \rangle}, \end{aligned}$$

and the proof follows by what was shown above. $\square$

8.4 Lower Bound Common Design Auxilliary Lemmas

Lemma 20. *Let $B \sim \text{Beta}(2, 2)$. Define $U = \frac{1}{2} - B$. Let $\mu \in \mathbb{R}$ and consider the random variable*

$$X = \mu + U.$$

The Fisher information for the parameter μ is given by a positive constant independent of μ .

Proof. The probability density of B is given by

$$f_B(b) = 6b(1 - b), \quad b \in [0, 1].$$

Hence, the density f_U of U satisfies

$$f_U(u) = f_B\left(\frac{1}{2} - u\right) = 6\left(\frac{1}{2} - u\right)\left(\frac{1}{2} + u\right) = 6\left(\frac{1}{4} - u^2\right).$$

and we find that probability density function of X is

$$f_X(x; \mu) = f_U(x - \mu) = 6\left(\frac{1}{2} - (x - \mu)\right)\left(\frac{1}{2} + (x - \mu)\right).$$

The differentiating the log-likelihood, we have

$$\frac{\partial}{\partial \mu} \log f_X(x; \mu) = \frac{\partial}{\partial \mu} \log f_U(u) \Big|_{u=x-\mu} = -\frac{f'_U(u)}{f_U(u)},$$

where $u = x - \mu$. Hence, the Fisher information for μ is

$$I(\mu) = \mathbb{E} \left[\left(\frac{\partial}{\partial \mu} \log f_X(X; \mu) \right)^2 \right] = \int_{-1/2}^{1/2} \frac{f'_U(u)^2}{f_U(u)} du,$$

since $X - \mu \stackrel{d}{=} U$. We have $\frac{(f'_U(u))^2}{f_U(u)} = \frac{24u^2}{\frac{1}{4} - u^2}$, so the integral is positive and independent of μ . $\square$