Journal of Machine Learning Research ?? (2025) 1-69 Submitted 1/25; Revised ?/??; Published ?/??

Federated Nonparametric Hypothesis Testing with Differential Privacy Constraints: Optimal Rates and Adaptive Tests

T. Tony Cai TCAI@WHARTON.UPENN.EDU
Department of Statistics and Data Science
University of Pennsylvania
Philadelphia, PA 19104, USA

Abhinav Chakraborty ABCH@WHARTON.UPENN.EDU
Department of Statistics and Data Science
University of Pennsylvania
Philadelphia, PA 19104, USA

Lasse Vuursteen LASSEV@WHARTON.UPENN.EDU
Department of Statistics and Data Science
University of Pennsylvania
Philadelphia, PA 19104, USA

Editor: To be determined

Abstract

Nonparametric goodness-of-fit testing in the white-noise-with-drift model under federated differential privacy (FDP) constraints is studied. In this framework, data is distributed across multiple locations, with each submitting a differentially private summary to a central server–encompassing both local and central differential privacy.

We quantify the cost of privacy in the federated setting by establishing matching lower and upper bounds, up to a logarithmic factor, on the minimax separation rate. This optimal rate benchmarks the difficulty of the testing problem, factoring in model characteristics such as the number of observations, noise level, and regularity of the signal class, along with the strictness of the (ϵ, δ)-DP requirement and the degree to which the data is distributed.

Our results demonstrate interesting and novel phase transition phenomena: where the cost of DP is minimal for testing in more centralized settings, it can significantly affect distributed scenarios. Furthermore, it is revealed that distributed one-shot protocols with access to shared randomness outperform those without access to shared randomness. We also construct a data-driven testing procedure that can adapt to an unknown regularity parameter over a large collection of function classes with minimal additional cost, while adhering to the same set of DP constraints.

Keywords: Differential Privacy, Nonparametric Goodness-of-Fit Testing, Federated Learning, Distributed Inference, Minimax Theory

1 Introduction

Differential privacy (DP), introduced by Dwork et al. (2006), provides a rigorous mathematical guarantee that analyses can be made publicly available without revealing sensitive information about individuals. Many differentially private statistical methods have since

been developed. See, for example, Arachchige et al. (2019); Dwork and Smith (2010); Dwork et al. (2017). While several other privacy frameworks exist, DP holds a prominent position both theoretically and practically, finding application within industry giants like Google, Microsoft, Apple, as well as governmental entities such as the US Census Bureau, see e.g. Panavas et al. (2024) and references therein.

In parallel, the growing need to analyze data distributed across multiple entities has propelled the emergence of *federated learning*, a collaborative approach to distributed machine learning. Instead of pooling raw data, federated learning enables organizations or devices to train a shared model while keeping local data private. Applications span healthcare (where patient records reside in multiple hospitals), finance (with customer data spread across numerous branches), and modern technologies like smartphones and autonomous vehicles Beaufays et al. (2019).

Federated learning naturally raises the question of how to incorporate differential privacy in a distributed environment—leading to the notion of *federated differential privacy (FDP)*. Under FDP, each data holder produces a differentially private summary of its local dataset for aggregation. A special case is *local* DP, in which privacy protection is applied at the level of individual data entries. This is a stringent form of DP because each item of data is independently given privacy protection. In the other extreme, *central* DP, only the inference output needs to satisfy the DP constraint, meaning that if the output is a test, only the final decision needs to satisfy a DP constraint. The FDP framework offers flexibility in defining what constitutes private information (e.g., individual records versus institutional datasets), thereby unifying the central and local DP paradigms.

In this paper, we investigate nonparametric goodness-of-fit testing under FDP constraints. Goodness-of-fit testing is a fundamental statistical problem, aiming to determine whether observed data are consistent with a specified (null) distribution or should be rejected in favor of a composite alternative. Nonparametric goodness-of-fit testing has been extensively studied without privacy considerations, see Ingster and Suslina (2003a) for an overview. However, differential privacy introduces fundamental challenges: while the test ultimately produces only a binary decision (reject or fail to reject), achieving optimal power requires access to the complete dataset. When data are distributed across multiple servers, FDP's privacy requirements further restrict information sharing, making near-optimal performance even more challenging.

We quantify the cost of privacy in the canonical nonparametric goodness-of-fit testing setting, establishing the theoretical performance limits under FDP constraints and optimal private testing procedures in an oracle scenario, assuming known regularity parameters. Specifically, we derive the minimax separation rate, which serves as a benchmark for the difficulty of the testing problem. Comparing these results with prior findings for federated nonparametric estimation Cai et al. (2023a) reveals notable differences between testing and estimation under differential privacy.

In addition to oracle rates – derived under regularity parameters that are rarely known in practice – we ask: Without the knowledge of the regularity parameters, is it possible to construct a test that is as good as when the parameters are known? This inquiry concerns *adaptation*, a major goal in nonparametric statistics. To address it, we propose a *data-driven test* that, while adhering the same FDP constraints, *adapts* to unknown regularity

parameters with only minimal additional cost compared to the performance of the oracle procedure.

1.1 Testing under federated differential privacy

We begin by formally introducing the general framework of federated inference under DP constraints. Consider a family of probability measures $\{P_f\}_{f\in\mathcal{F}}$ on the measurable space $(\mathcal{X}, \mathscr{X})$, parameterized by $f \in \mathcal{F}$. We consider a setting where $N = mn$ i.i.d. observations are drawn from a distribution P_f and distributed across m servers. Each server $j = 1, \ldots, m$ holding an equal amount (n many) observations.

Let us denote by $X^{(j)} = (X_i^{(j)})_{i=1}^n$ the n realizations from P_f on the j-th server. Based on $X^{(j)}$, each server outputs a (randomized) transcript $Y^{(j)}$ to the central server that satisfies the privacy constraint. The central server, utilizing all transcripts $Y := (Y^{(1)}, \ldots, Y^{(m)})$, decides between a null hypothesis and an alternative hypothesis, through means of a test $T = T(Y)$. Since we are concerned with testing between a null and alternative hypothesis, we shall consider the decision space $\{0, 1\}$, where 0 corresponds to DO NOT REJECT and 1 with REJECT. A *test* is then simply to be understood as a statistic taking values in $\{0, 1\}$. Figure 1 gives an illustration of a federated (ϵ, δ)-DP-constrained testing procedure.

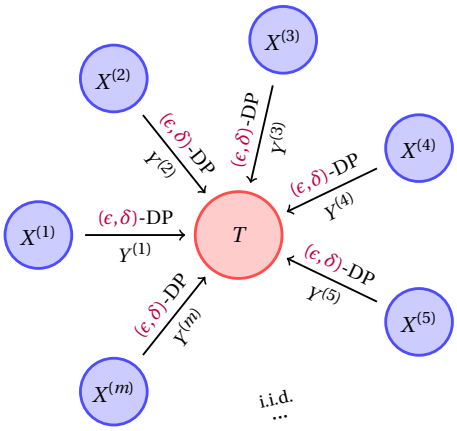

Figure 1: Illustration of the (ϵ, δ)-FDP-constrained testing.

The transcript $Y^{(j)}$ satisfies an (ϵ, δ)-DP constraint, which, loosely speaking, means that the transcript $Y^{(j)}$ cannot differ too much depending on whether a specific individual is in the data set or not. This is achieved through randomization, independent of the data.

We will consider two types of sources for randomization; independently among the servers or through a shared source of randomness U (e.g., the same random seed). Formally, a shared source of randomness means that the law of the transcript is given by a distribution conditionally on $X^{(j)}$ and U, $A \mapsto \mathbb{P}(Y^{(j)} \in A | X^{(j)}, U)$, defined on a measurable space $(\mathcal{Y}, \mathscr{Y})$. The presence of shared randomness is a slight, but important extension of protocols where $Y^{(j)}$ is allowed to be random only through locally generated randomness. To ensure that the source of shared randomness does not erode the notion of privacy, only the local source of randomness is used in the privacy mechanism, i.e. to guarantee privacy. We formalize this as follows.

Definition 1. *The transcript $Y^{(j)}$ is (ϵ, δ)-differentially private ((ϵ, δ)-DP) if for all $A \in \mathscr{Y}^{(j)}$, $u \in \mathcal{U}$ and data sets $x, x' \in \mathcal{X}^n$ differing in at most one individual datum*[1] *it holds that*

$$\mathbb{P}\left(Y^{(j)} \in A | X^{(j)} = x, U = u\right) \leqslant e^{\epsilon}\mathbb{P}\left(Y^{(j)} \in A | X^{(j)} = x', U = u\right) + \delta. \tag{1}$$

If the transcripts $Y^{(j)}$ are (ϵ, δ)-DP for $j = 1, \dots, m$, we say that the testing protocol (Y, T) is (ϵ, δ)-FDP. If $m = 1$, the setting reduces to the central DP setting, where all data is available on a single server. For $n = 1$, the setting corresponds to the local DP setting, where each server holds a single observation. For $n, m > 1$, the setting encompasses scenarios where multiple parties hold sensitive data and each publishes a differentially private summary without sharing raw data —- for example, separate hospitals with samples from the same population, where privacy concerns prevent direct data pooling.

We also note that in Definition 1, a shared source of randomness U does not compromise privacy guarantees, even if it is publicly available. This is because U is not used in the privacy mechanism itself; rather, it serves to 'synchronize' the transcripts. Such coordination enables the transcripts to be more informative about the underlying signal while each individual transcript reveals less about its local data.

1.2 Nonparametric goodness-of-fit testing

The white-noise-with-drift model serves as a benchmark model for nonparametric testing and has been extensively studied in the non-private setting Ermakov (1990); Ingster (1993); Lepskii (1992); Spokoiny (1996). Furthermore, the problem bares a close relationship with "classical" nonparametric goodness-of-fit testing developed by Kolmogorov (1933); Smirnov (1948) —- of which it can be viewed as the asymptotic limit and it also connects to broader nonparametric testing problems through asymptotic equivalence (see Section 1.4 in Ingster and Suslina (2003a) and references therein).

In the distributed version of the white-noise-with-drift model, the $j = 1, \dots, m$ servers each observe $i = 1, \dots, n$ i.i.d. $X_i^{(j)}$ taking values in $\mathcal{X} \subset L_2[0, 1]$; subject to the stochastic differential equation

$$dX_{t;i}^{(j)} = f(t)dt + \sigma dW_{t;i}^{(j)} \tag{2}$$

under P_f, with $t \mapsto W_{t;i}^{(1)}, \dots, t \mapsto W_{t;i}^{(m)}$ i.i.d. Brownian motions, $f \in L_2[0, 1]$ and with $\sigma > 0$ the known noise level for each observation. When $m = 1$, we recover the classical white-noise-with-drift model. We denote the total number of observations as $N = mn$ throughout the paper and will consider asymptotic regimes where $N \to \infty$.

We consider the canonical signal detection problem, where the goal is to test for the presence or absence of the "signal component" f. More formally, we consider testing the null hypothesis $H_0 : f \equiv 0$ against the alternative hypothesis that

$$f \in H_\rho^{s,R} := \left\{ f \in \mathcal{B}_{p,q}^s : \|f\|_{L_2} \geqslant \rho \text{ and } \|f\|_{\mathcal{B}_{p,q}^s} \leqslant R \right\}. \tag{3}$$

Here, the alternative hypothesis consists of s-smooth functions in a Besov space, with $\|\cdot\|_{\mathcal{B}_{p,q}^s}$ denoting the Besov-(s, p, q)-norm and $\mathcal{B}_{p,q}^s \subset L_2[0, 1]$ corresponds to the Besov

1. Local datasets $x, x' \in \mathcal{X}^n$ differ in at most one individual datum are said to be *neighboring*; which formally means that they are at most one apart in *Hamming distance* (see Section 1.6 for a definition).

space, see Section D in the Appendix for the definitions. Besov spaces, where $2 \leqslant p \leqslant \infty$, $1 \leqslant q \leqslant \infty$, are a very rich class of function spaces. They offer a framework for functions in (2) with specific smoothness characteristics. They include many traditional smoothness spaces such as Hölder and Sobolev spaces as special cases. We refer the reader to Triebel (1992) for a detailed discussion on Besov spaces.

Using a wavelet transform, the above testing problem is equivalent the observations under the Gaussian sequence model, where each of the $j = 1, \dots, m$ machines observes $i = 1, \dots, n$ observations $X_i^{(j)} := (X_{lk;i}^{(j)})_{l \geqslant 1, k=1,\dots,2^l}$

$$X_{lk;i}^{(j)} = f_{lk} + \sigma Z_{lk;i}^{(j)}, \tag{4}$$

where the $Z_{lk;i}^{(j)}$'s are i.i.d. standard Gaussian. The equivalent hypotheses (3) in the sequence model simply follows by replacing the $L_2[0,1]$-norm with the $\ell_2(\mathbb{N})$-norm and the Besov ball $\mathcal{B}_{p,q}^{s,R}$ set to $\{f \in \ell_2(\mathbb{N}) : \|f\|_{\mathcal{B}_{p,q}^s} \leqslant R\}$, where the Besov norm on the sequence space $\ell_2(\mathbb{N})$ is defined as

$$\|f\|_{\mathcal{B}_{p,q}^s} := \begin{cases} \left(\sum_{l=1}^{\infty} \left(2^{l(s+1/2-1/p)} \left\| (f_{lk})_{k=1}^{2^l} \right\|_p \right)^q \right)^{1/q} & \text{for } 1 \leqslant q < \infty, \\ \sup_{l \geqslant 1} 2^{l(s+1/2-1/p)} \left\| (f_{lk})_{k=1}^{2^l} \right\|_p & \text{for } q = \infty. \end{cases} \tag{5}$$

In other words, the results for testing under FDP derived for the sequence model of (4) with hypothesis (3) apply to the model described by (2) also, with the same corresponding hypothesis.

Given a $\{0,1\}$ valued test T, where $T(Y) = 1$ corresponds to rejecting the null hypothesis, we define the testing risk sum of the type I and worst case type II error over the alternative class;

$$\mathcal{R}(H_\rho^{s,R}, T) = \mathbb{P}_0 T(Y) + \sup_{f \in H_\rho^{s,R}} \mathbb{P}_f T(Y).$$

For the range of values $2 \leqslant p \leqslant \infty$, $1 \leqslant q \leqslant \infty$, the *minimax separation rate* in the non-private case is known to be $\rho \asymp (\sigma^2/N)^{\frac{s}{2s+1/2}}$ (see e.g. Ingster (1993)). This means that, for $\rho \gg (\sigma^2/N)^{\frac{s}{2s+1/2}}$, there exists a sequence of consistent tests $T \equiv T_N$ such that $\mathcal{R}(H_{\rho'}^{s,R}, T) \to 0$, whilst no such sequence of tests exists whenever $\rho \ll (\sigma^2/N)^{\frac{s}{2s+1/2}}$.

The minimax separation rate captures how the testing problem becomes easier, or more difficult, for different model characteristics. For (ϵ, δ)-FDP testing protocols $\mathscr{T}^{(\epsilon,\delta)}$, the minimax separation rate depends on the stringency of the privacy requirement, given by $\epsilon, \delta > 0$, as well as the model characteristics m, n, s and σ. That is, we aim to find ρ as a function of $m, n, s, \sigma, \epsilon, \delta$, such that $\inf_{T \in \mathscr{T}^{(\epsilon,\delta)}} \mathcal{R}(H_{\rho',R}^s, T)$ converges to either 0 or 1 depending on whether $\rho' \ll \rho$ or $\rho' \gg \rho$.

1.3 Main results and our contribution

For a quick overview, the main contributions of the paper are as follows:

- We derive the minimax separation rates for nonparametric goodness-of-fit testing under FDP constraints (Theorems 2 and 3), tight up to logarithmic factors.

- We exhibit adaptive FDP testing methods that achieve the oracle rates derived in Section 2 (Theorem 6).
- The theoretical lower bound derived (Theorem 12) involves several technical innovations necessitated by the FDP setting, which we summarize in Section 5.

Our findings offer entirely new insights when applied to both the purely local or purely central DP settings: By formulating our results in the more general FDP framework, we are able to address both settings at once, on top of bridge the gap between the local and central regimes.

Our analysis uncovers several intriguing findings, which we briefly highlight here. The performance guarantees for the methods demonstrated in Section 3, along with the lower bounds established in Section 5, indicate that the (ϵ, δ)-FDP testing problem for the hypotheses given in (3) is governed by the minimax separation rate (up to logarithmic factors)

$$\rho^2 \asymp \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^{3/2}\epsilon\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{m}n\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1/2}} + \frac{\sigma^2}{mn^2\epsilon^2}\right). \quad (6)$$

The precise statement is deferred to Theorem 2.

The derived rate reveals multiple phase transitions in the distributed testing problem under privacy constraints. These transitions create distinct regimes where privacy constraints affect the detection boundary differently. A smaller ϵ (stronger privacy[2]) increases the detection threshold, but its impact varies: while in some regimes a small ϵ substantially affects the rate, in others it has no effect on the detection boundary.

Which regime dominates depends on the degree of data distribution across servers. In the central DP setting ($m = 1$), privacy only affects the rate polynomially when $\epsilon \lesssim 1/\sqrt{n}$; otherwise, we recover the non-private minimax rate $\rho \asymp (\sigma^2/n)^{\frac{s}{2s+1/2}}$. In contrast, in the local DP setting ($n = 1$), privacy constraints substantially affect the rate whenever $\epsilon \lesssim 1$. In the general federated setting ($m \gg 1$), we observe similar effects: m and n enter the minimax rate with different powers whenever $\epsilon^2 \lesssim \sigma^{\frac{1}{2s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$. This demonstrates that distributing $N = mn$ observations across more servers makes the task more challenging than concentrating them on fewer servers – aligning with the intuition that privacy is easier to maintain with larger samples.

While the impact of data distribution across servers mirrors findings in nonparametric regression estimation Cai et al. (2023a), the phase transitions we observe are unique to testing. Notably, there are regimes where privacy constraints have a much smaller effect on the testing rate compared to analogous estimation settings. This difference stems from the nature of the problems: testing can leverage aggregated private test statistics, while nonparametric estimation inherently concerns a high-dimensional inference goal, which requires more extensive information sharing. We provide a detailed interpretation of these phase transitions in Section 2.

2. For a δ that decreases polynomially with N, its impact on the detection boundary is limited to a logarithmic factor, making its effect on the error rate minor compared to that of ϵ.

Our analysis also reveals that the minimax rate obtained in Theorem 2 becomes worse without access to shared randomness. This is revealed by Theorem 2 in Section 2. For certain values of ϵ, we show that the performance is strictly worse for methods that use only local randomness, and we exhibit optimal local and shared randomness methods for these regimes, respectively, in Sections 3.2 and 3.3.

1.4 Related Work

In the context of hypothesis testing, several works have studied testing under local differential privacy (DP) for discrete distributions Gaboardi et al. (2016); Sheffet (2018); Acharya et al. (2018b, 2019); Berrett and Butucea (2020); Acharya et al. (2020b, 2021). In this setting, each server holds a single multinomial observation. This framework has been extended to a nonparametric setting where servers hold a single draw a from continuous distribution on the unit interval. For such continuous settings, goodness-of-fit testing has been studied in Dubois et al. (2023); Lam-Weil et al. (2022), with the latter work also addressing adaptation. Related work by Butucea et al. (2023) examines the estimation of quadratic functionals under local DP constraints in the same setting, which has connections to goodness-of-fit testing. While there is some overlap in terms of rates, they do not capture the full spectrum of goodness-of-fit testing. For a detailed comparison with our results, see Remark 4.

For hypothesis testing under central differential privacy, where the complete dataset resides on a single server, Canonne et al. (2019) investigates simple hypothesis testing, while Acharya et al. (2018a) addresses uniformity and independence testing in the multinomial model. Alabi and Vadhan (2022) explores hypothesis testing in linear regression. Perhaps most relevant to our work in the central DP setting are Canonne et al. (2020) and Narayanan (2022b), which study signal detection in the many-normal-means model. Our results recover their rate as a special case, discussed in more detail in Remark .

Local DP estimation has been extensively studied across various settings: the many-normal-means model, discrete distributions, and parametric models Duchi et al. (2013, 2018); Acharya et al. (2020a); Ye and Barg (2018). Density estimation under local DP constraints has been examined by several authors Duchi et al. (2018); Sart (2023); Kroll (2021); Butucea et al. (2020), with the latter three works addressing adaptation. Central DP estimation has been investigated for various high-dimensional and nonparametric problems Smith (2011); Dwork et al. (2014b); Bassily et al. (2014); Kamath et al. (2019, 2020); Cai et al. (2021); Narayanan (2022a); Brown et al. (2023); Cai et al. (2024), while Lalanne et al. (2023) focuses on nonparametric density estimation with known smoothness.

Research in the broader federated setting has been more limited. For estimation, works include studies of discrete distributions Liu et al. (2020); Acharya et al. (2023), mean estimation Levy et al. (2021); Narayanan et al. (2022), nonparametric regression Cai et al. (2023a), and sparse linear regression Li et al. (2024). In the context of testing, Canonne and Sun (2023) examines discrete distribution testing in a two-server setting ($m = 2$) with varying DP constraints.

1.5 Organization of the paper

The rest of the paper is organized as follows. In Section 2, we present the main results of the paper, for the known smoothness case and the adaptive setting. Next, Section 3 presents the methods that achieve the optimal rates derived in Section 2. In Section 4, we extend these methods to be adaptive in the case that the smoothness is unknown. In Section 5, we present the lower bound theorems for the testing problem and give a sketch of its proof. Further proofs are deferred to the Appendix of the article.

1.6 Notation, and notions

Throughout the paper, we shall write $N := mn$. For two positive sequences a_k, b_k we write $a_k \lesssim b_k$ if the inequality $a_k \leqslant Cb_k$ 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 the notations $a \vee b$ and $a \wedge b$ for the maximum and minimum, respectively, between a and b. For $k \in \mathbb{N}$, $[k]$ shall denote the set $\{1, \dots, k\}$. We use c and C to 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 $\mathrm{Tr}(M)$ is its trace. Furthermore, we let I_d denote the $d \times d$ identity matrix. The Hamming distance on $\mathcal{X}^n$ is defined as $\mathrm{d}_H(x, \breve{x}) := \sum_{i=1}^n \mathbb{1}\{x_i \neq \breve{x}_i\}$ for $x = (x_i)_{i=1}^n, \breve{x} = (\breve{x}_i)_{i=1}^n \in \mathcal{X}^n$. Furthermore, for a vector space $\mathcal{X}$ and $x = (x_i)_{i\in[n]} \in \mathcal{X}^n$, we shall write $\overline{x}$ for the average $n^{-1}\sum_{i=1}^n x_i$.

Formally, shared randomness independent of the data means U is defined on a separate probability space $(\mathcal{U}, \mathscr{U}, \mathbb{P}^U)$ and for the joint law with the data we shall take the product space. No shared randomness corresponds to U being degenerate, i.e., $\mathscr{U} = \{\varnothing, \mathcal{U}\}$. We denote by $\mathscr{T}^{(\epsilon,\delta)}_{\mathtt{SHR}}$ the class of (ϵ, δ)-FDP shared randomness testing protocols (T, Y, U), where the law of Y satisfies Definition 1. The local randomness subclass, is denoted $\mathscr{T}^{(\epsilon,\delta)}_{\mathtt{LR}}$.

2 Minimax optimal testing rates under privacy constraints

In this section, we discuss the main results in detail. We start the discussion with results for the oracle case where the regularity parameter is known in Section 2.1. Section 2.2 describes the main results for when the regularity is not known.

2.1 Description of the minimax separation rate

We first give a precise statement concerning the minimax separation rate shown in (6).

Theorem 2. *Let $s, R > 0$ be given and consider any sequences of natural numbers $m \equiv m_N$ and $n := N/m$ such that $N = mn \to \infty$, $1/N \ll \sigma \equiv \sigma_N = O(1)$, $\epsilon \equiv \epsilon_N$ in $(N^{-1}, 1]$ and $\delta \equiv \delta_N \lesssim N^{-(1+\omega)}$ for any constant $\omega > 0$. Let $\rho > 0$ a sequence of numbers satisfying (6).*

Then,

$$\inf_{T\in\mathscr{T}^{(\epsilon,\delta)}_{\mathtt{SHR}}} \mathcal{R}(H^{s,R}_{\rho M_N}, T) \to \begin{cases} 0 & \text{for any } M_N^2 \gg \log\log(N)\log^{3/2}(N)\log(1/\delta), \\ 1 & \text{for any } M_N \to 0. \end{cases}$$

The proof of the theorem is given in Section B.4 of the Appendix. It is based on a combination of upper and lower bounds, where the lower bound is established in Section 5.

The upper bound is given in Section 3, where we present an (ϵ, δ)-DP testing protocol that attains the rate in Theorem 2. These upper and lower bounds are in fact non-asymptotic, meaning that they do not require the assumption that $N \to \infty$.

Theorem 2 shows multiple regime changes, where the distributed testing problem under privacy constraints undergoes a change in the minimax separation rate. Later on in this section, we highlight the different regimes and give an interpretation to each of them.

Theorem 2 considers the minimax rate for the class of distributed protocols with access to shared randomness, $\mathscr{T}^{(\epsilon,\delta)}_{\texttt{SHR}}$. Theorem 3 below considers the minimax rate for the (strictly smaller) class of distributed protocols without access to shared randomness, $\mathscr{T}^{(\epsilon,\delta)}_{\texttt{LR}}$. Here, transcripts depend *only* on their local data and possibly a local source of randomness.

Theorem 3. *Let $s, R > 0$ be given and consider any sequences of natural numbers $m \equiv m_N$ and $n := N/m$ such that $N = mn \to \infty$, $1/N \ll \sigma \equiv \sigma_N = O(1)$ and $\epsilon \equiv \epsilon_N$ in $(N^{-1}, 1]$ and $\delta \equiv \delta_N \lesssim N^{-(1+\omega)}$ for any constant $\omega > 0$. Let $\rho \equiv \rho_N$ a sequence of positive numbers satisfying*

$$\rho^2 \asymp \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{mn}\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)\right). \tag{7}$$

Then,

$$\inf_{T \in \mathscr{T}^{(\epsilon,\delta)}_{LR}} \mathcal{R}(H^{s,R}_{\rho M_N}, T) \to \begin{cases} 0 & \text{for any } M_N^2 \gg \log\log(N)\log^{3/2}(N)\log(1/\delta), \\ 1 & \text{for any } M_N \to 0. \end{cases}$$

The proof of Theorem 3 is given in Section B.4 of the Appendix. The theorem shows that, depending on the value of ϵ, the minimax rate for protocols that do not have access to shared randomness is strictly worse than those for protocols that do have access to shared randomness.

To facilitate comparison between these theorems, Table 1 organizes our results into six distinct regimes, each representing a fundamentally different behavior of the minimax separation rate. Each of the regimes correspond to the dominating term in the minimax separation rates of Theorems 2 and 3. The dominant term in each regime is determined by complex interplay between ϵ, in comparison to n, m, σ, s and the availability of shared randomness.

The six regimes naturally partition into two groups based on their privacy budget requirements and the role of shared randomness. Regimes 4, 5, and 6 – which we term the "low privacy-budget" regimes – occur when ϵ is relatively small. In these regimes, both types of protocols achieve identical rates through the same testing procedure, indicating that shared randomness offers no advantage when privacy constraints are stringent.

In contrast, Regimes 1, 2, and 3 – the "high privacy-budget" regimes – emerge when ϵ is relatively large. Table 1 reveals that shared randomness protocols achieve strictly better rates than local randomness protocols in Regimes 2 and 3, demonstrating how server coordination through shared randomization can meaningfully improve performance under moderate privacy constraints. This advantage disappears in the low privacy-budget case (Regimes 4, 5, and 6), where both protocols achieve identical rates, though the privacy thresholds ϵ at which regime transitions occur still differ between the two types of protocols.

Regime	Rate ρ^2	Range for ϵ
Using Shared Randomness		
1	$\left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}}$	$\epsilon \geqslant \sigma^{-\frac{2}{4s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$
2	$\left(\frac{\sigma^2}{mn^{3/2}\epsilon}\right)^{\frac{2s}{2s+1}}$	$\sigma^{-\frac{2}{4s+1}} m^{-\frac{2s}{4s+1}} n^{\frac{1/2-2s}{4s+1}} \leqslant \epsilon < \sigma^{-\frac{2}{4s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}, \epsilon \geqslant n^{-1/2}$
3	$\left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+1}}$	$\sigma^{-\frac{1}{2s}} m^{-\frac{1}{2}} n^{\frac{1-2s}{4s}} \leqslant \epsilon < n^{-1/2}$
4	$\left(\frac{\sigma^2}{\sqrt{mn}}\right)^{\frac{2s}{2s+1/2}}$	$n^{-1/2} \leqslant \epsilon < \sigma^{-\frac{2}{4s+1}} m^{-\frac{2s}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$
5	$\left(\frac{\sigma^2}{\sqrt{mn}^{3/2}\epsilon}\right)^{\frac{2s}{2s+1/2}}$	$\sigma^{\frac{1}{2s+1}} m^{-\frac{1}{2}} n^{-\frac{1+s}{2s+1}} \leqslant \epsilon < \sigma^{-\frac{1}{2s}} m^{-\frac{1}{2}} n^{\frac{1-2s}{4s}}, \epsilon < n^{-1/2}$
6	$\frac{\sigma^2}{mn^2\epsilon^2}$	$\epsilon < \sigma^{\frac{1}{2s+1}} m^{-\frac{1}{2}} n^{-\frac{1+s}{2s+1}}$
Using Only Local Randomness		
1	$\left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}}$	$\epsilon \geqslant \sigma^{-\frac{2}{4s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$
2	$\left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}}$	$\sigma^{-\frac{2}{4s+1}} m^{\frac{1/4-s}{4s+1}} n^{\frac{1/2-2s}{4s+1}} \leqslant \epsilon < \sigma^{-\frac{2}{4s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}, \epsilon \geqslant n^{-1/2}$
3	$\left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}}$	$\sigma^{-\frac{4}{4s-1}} m^{-\frac{1}{2}} n^{\frac{5/2-2s}{4s-1}} \leqslant \epsilon < n^{-1/2}$
4	$\left(\frac{\sigma^2}{\sqrt{mn}}\right)^{\frac{2s}{2s+1/2}}$	$n^{-1/2} \leqslant \epsilon < \sigma^{-\frac{2}{4s+1}} m^{\frac{1/4-s}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$
5	$\left(\frac{\sigma^2}{\sqrt{mn}^{3/2}\epsilon}\right)^{\frac{2s}{2s+1/2}}$	$\sigma^{\frac{1}{2s+1}} m^{-\frac{1}{2}} n^{-\frac{1+s}{2s+1}} \leqslant \epsilon < \sigma^{-\frac{4}{4s-1}} m^{-\frac{1}{2}} n^{\frac{5/2-2s}{4s-1}}, \epsilon < n^{-1/2}$
6	$\frac{\sigma^2}{mn^2\epsilon^2}$	$\epsilon < \sigma^{\frac{1}{2s+1}} m^{-\frac{1}{2}} n^{-\frac{1+s}{2s+1}}$

Table 1: Minimax separation rates and their corresponding ϵ ranges for both shared and local randomness settings. The rates are given up to logarithmic factors.

A particular distinction between shared and local randomness protocols emerges when examining the smoothness parameter regime $s \leqslant 1/4$. While shared randomness protocols can still achieve the high privacy-budget rates for sufficiently large ϵ, local randomness protocols remain trapped in the low privacy-budget regime across our entire parameter range ($N^{-1} \lesssim \sigma \lesssim 1$ and $N^{-1} < \epsilon \lesssim 1$). This limitation arises from the condition

$$\left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}} \asymp \left(\frac{\sigma^2}{\sqrt{mn}^{3/2}\epsilon}\right)^{\frac{2s}{2s+1/2}} \iff \epsilon^{\frac{2s-1/2}{2s+3/2}} \asymp \sigma^{-\frac{1}{2s+3/2}} m^{\frac{1}{2s+3/2}} n^{\frac{5/4-s}{2s+3/2}},$$

which cannot be satisfied when $s \leqslant 1/4$, $\epsilon \gtrsim (mn)^{-1}$ and $\sigma^2 = O(1)$. As a result, the minimax rate without shared randomness and $s \leqslant 1/4$ consolidates to Regimes 4, 5, and 6. This means in particular that the minimax rates corresponding to the high privacy-budget regimes cannot be attained for any value of ϵ when $s \leqslant 1/4$ and shared randomness is not available.

Within the low privacy-budget range, we observe three distinct regimes with remarkable properties. In Regime 4, the rate takes the form $(\frac{\sigma^2}{\sqrt{mn}})^{\frac{2s}{2s+1/2}}$. While this rate is polynomially worse in m than the unconstrained rate, it exhibits a surprising independence from privacy comes 'for free' through the locally optimal test statistic.

We achieve this rate through an (ϵ, δ)-DP testing protocol detailed in Section 3.1. The protocol operates in two steps: first computing locally optimal private test statistics from each server's data, then averaging these statistics to combine their power. Roughly speaking, the strategy's effectiveness under stringent privacy constraints stems from the idea that sharing a single private real-valued test statistic can preserve privacy more effectively than sharing private approximations of the full data.

The behavior changes dramatically when $\epsilon \lesssim n^{-1/2}$, where ϵ begins to affect the rate polynomially. Most striking is Regime 6, where the rate equals $\frac{\sigma^2}{mn^2\epsilon^2}$, which is independent of the regularity parameter s. This independence arises because signals of size $\frac{\sigma^2}{mn^2\epsilon^2}$ dominate the local estimation rate of $(\frac{\sigma^2}{n})^{\frac{2s}{2s+1/2}}$. In this regime, the privacy constraint – not the high-dimensional nature of the problem – becomes the sole bottleneck, as signals can be estimated locally with high accuracy.

The mathematical techniques required to establish optimality differ between the local and shared randomness protocol types. We construct explicit protocols achieving these optimal rates in Sections 3.2 and 3.3, while the corresponding lower bounds – which also require distinct proof techniques for local versus shared randomness – are developed in Section 5.

Remark 4. *Our results in Theorems 2 and 3 invite comparison with recent work on local differential privacy: Dubois et al. (2023); Lam-Weil et al. (2022), who study nonparametric goodness-of-fit testing, and Butucea et al. (2023), who examines the related question quadratic functional estimation. These papers consider the setting where each server holds a single observation ($n = 1$) from an s-smooth density f on $[0, 1]$. In their framework, the minimax rates take the form:*

$$\rho^2 \asymp \begin{cases} \left(\frac{1}{m\epsilon^2}\right)^{\frac{2s}{2s+1}} & \text{for interactive protocols} \\ \left(\frac{1}{m\epsilon^2}\right)^{\frac{2s}{2s+3/2}} & \text{for non-interactive protocols} \end{cases}$$

These rates correspond to Regime 2 in our Table 1. We observe a richer set of phenomena, even when $\sigma = 1$ and $n = 1$. This is surprising, as the model (2) is asymptotically equivalent to the density setting as observations tend to infinity (see e.g. Nussbaum (1996)). We conjecture that additional complexity would emerge if their framework is extended to each server observing $X^{(j)} = (X_k^{(j)})_{k=1}^K$ with i.i.d. $X_k^{(j)} \sim f$ and $X^{(j)}$ as unit of privacy, or if extended to the full FDP framework.

The distinction they observe between interactive and non-interactive protocols parallels our findings about shared versus local randomness. Indeed, when sequential- or interactive protocols are allowed, shared randomness can be employed in particular, suggesting that in non-interactive applications, shared randomness should be employed whenever possible.

Remark 5. *Within the central DP setting, another relevant comparison is with Canonne et al. (2020) and Narayanan (2022b), who investigate the many-normal-means model. Although their finite-dimensional framework does not require adaptation to unknown smoothness, we can compare our oracle rates by setting $m = 1$ and truncating our model (4) at a known threshold d. In this specialized scenario, our intermediate results recover the rates from those works (see Section B.4), with a logarithmic-factor improvement.*

2.2 Adaptation

In the previous section, we derived the minimax separation rate for the nonparametric distributed testing problem. However, the proposed tests constructed in Section 3 require knowledge of the regularity parameter s of the underlying f. Typically, the regularity of the function is unknown in practice, necessitating the use of data-driven methods to find the best adaptive testing strategies.

Given that the regularity of the underlying signal class is unknown, it makes sense to consider the minimax testing risk

$$\sup_{s\in[s_{\min},s_{\max}]} \mathcal{R}\left(H^{s,R}_{M_{N,s}\rho_s}, T\right),$$

for certain predetermined values $0 < s_{\min} < s_{\max} < \infty$. Here, we consider separation rates ρ_s depending on the underlying smoothness. In the case that the true underlying smoothness is $s = s_{\min}$, the separation rate is relatively larger than when (for example) $s = s_{\max}$. In the case that the true smoothness s is larger than $s_{\min}$, we would like to attain the smaller of the two rates ρ_s.

In the non-privacy constraint setting, adaptation for the above risk can be achieved with only a minor additional cost in the separation rate (a $\log\log N$ factor). See for example Theorem 2.3 in Spokoiny (1996) or Section 7 in Ingster and Suslina (2003a). Theorem 6 below shows that also under privacy constraint, the optimal private rate can be attained by a protocol that is adaptive to the regularity parameter s, with minimal additional cost; at most a logarithmic factor.

Theorem 6. *Let $0 < s_{\min} < s_{\max} < \infty$, $R > 0$ be given and consider any sequences of natural numbers $m \equiv m_N$ and $n := N/m$ such that $N = mn \to \infty$, $1/N \ll \sigma \equiv \sigma_N = O(1)$, $\epsilon \equiv \epsilon_N$ in $(N^{-1}, 1]$ and $\delta \equiv \delta_N \lesssim N^{-(1+\omega)}$ for any constant $\omega > 0$.*

If ρ a sequence of positive numbers satisfies (6), there exists a sequence of (ϵ,δ)-FDP testing protocols T_N such that

$$\sup_{s\in[s_{\min},s_{\max}]} \mathcal{R}(H^{s,R}_{\rho M_N}, T_N) \to \begin{cases} 0 & \text{for any } M_N^2 \gg \log\log(N)\log^{5/2}(N)\log(1/\delta) \\ 1 & \text{for any } M_N \to 0. \end{cases}$$

Furthermore, whenever ρ satisfies (7), there exists a sequence of (ϵ,δ)-FDP testing protocols T_N using only local randomness such that the above display holds as well.

We construct such adaptive (ϵ,δ)-FDP testing protocols in Section 4 and their resulting performance proofs the above theorem. The adaptive methods can be seen as extensions of the methods exhibited in Section 3 for when the smoothness is known. The adaptive methods can essentially be seen as a multiple testing extension of the known smoothness methods, testing along a grid of smoothness levels between $s_{\min}$ and $s_{\max}$. The strain on the privacy budget stemming from conducting multiple testing procedures is limited, due to the fact that the cardinality of this grid is order $\log(N)$. The Type I error control is assured by a Bonferroni correction, which leverages the exponential bounds on the Type I error of the individual "known smoothness tests".

3 Optimal differentially private testing procedures

In this section, we construct (ϵ, δ)-FDP testing procedures that attain the minimax separation rates derived in Section 2.

The testing procedures are constructed in three steps. First, in Section 3.1, we construct a (ϵ, δ)-FDP testing procedure that uses only local randomness and that is optimal in the low privacy-budget regime described in the previous section. We refer to this procedure as T_{I}. Second, we construct two (ϵ, δ)-FDP testing procedures that use local randomness and shared randomness, respectively, and that are optimal in their respective high privacy-budget regimes. We refer to these procedures as T_{II} and T_{III} and describe them in Sections 3.2 and 3.3, respectively.

The testing procedures differ in terms of the testing strategy. In the low privacy-budget case where T_{I} is optimal, the testing strategy can be seen to consist of first computing a locally optimal private test statistic in each machine; that is, a test statistic that would result in the optimal private test using just the local data. The locally optimal test statistic is based on the squared Euclidean norm of the truncated observation. To deal with the nonlinearity of the Euclidean norm, the strategy appropriately restricts the domain of the clipped locally optimal test statistic, after which we employ a Lipschitz-extension to obtain a test statistic that is well-defined on the sample space and more robust to outliers than the Euclidean norm itself. The noisy version of this test statistic is locally optimal under privacy constraints, in the sense that a corresponding (strict) p-value test attains the lower bound rate (up to a logarithmic factor) as established by Theorem 3 for the case where $m = 1$. When $m > 1$, the final test statistic is obtained by averaging the locally optimal private test statistics.

In the large ϵ regime, instead of computing a locally optimal test statistic, both T_{II} and T_{III} are based on truncated, clipped and noisy versions of the local observations. The key difference between the two is that the latter uses the same random rotation of the local observations, which is made possible by the availability of shared randomness.

Together, the methods prove Theorem 7 below, which forms the "upper bound" part of the minimax separation rate described by Theorems 3 and 2. Unlike the formulation of the latter theorems, we note that the result is not asymptotic.

Theorem 7. *Let $s, R > 0$ be given. For all $\alpha \in (0,1)$, there exists a constant $C_\alpha > 0$ such that if*

$$\rho^2 \geqslant C_\alpha \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{mn}n\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)\right), \quad (8)$$

there exists a (ϵ, δ)-FDP testing protocol $T \equiv T_{m,n,s,\sigma}$ such that

$$\mathcal{R}(H^{s,R}_{\rho M_N}, T) \leqslant \alpha, \quad (9)$$

for all natural numbers m, N and $n = N/m$, $\sigma \in [1/N, \sigma_{\max}]$, $\epsilon \in (N^{-1}, 1]$, $\delta \leqslant N^{-(1+\omega)}$ for any constant $\omega > 0$, $\sigma_{\max} > 0$ and a nonnegative sequence $M_N^2 \gtrsim \log\log(N)\log^{3/2}(N)\log(1/\delta)$.

Similarly, for any $\alpha \in (0,1)$, there exists a constant $C_\alpha > 0$ such that if

$$\rho^2 \geqslant C_\alpha \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^{3/2}\epsilon\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{mn}}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)\right), \quad (10)$$

we have that there exists a (ϵ, δ)*-FDP shared randomness testing protocol* $T \equiv T_{m,n,s,\sigma}$ *such that*

$$\mathcal{R}(H^{s,R}_{\rho M_N}, T) \leqslant \alpha, \tag{11}$$

for all natural numbers m, N *and* $n = N/m$, $\sigma \in [1/N, \sigma_{\max}]$, $\epsilon \in (N^{-1}, 1]$, $\delta \leqslant N^{-(1+\omega)}$ *for any constant* $\omega > 0$ *and a nonnegative sequence* $M_N^2 \gtrsim \log\log(N)\log^{3/2}(N)\log(1/\delta)$.

The proof of the theorem follows directly from the guarantees proven for each of the three testing protocols; we defer it to Section B in the supplement. Before giving the detailed construction of the three tests, we introduce some common notation. Let Π_L denote the projection of elements $\mathbb{R}^{\mathbb{N}}$ onto the first $d_L := \sum_{l=1}^{L} 2^l$ coordinates, where the elements as ordered and indexed as $\Pi_L x = (x_{11}, \dots, x_{12}, x_{21}, \dots, x_{14}, \dots, x_{L1}, \dots, x_{L2^L}, 0, 0, 0, \dots)$.

We shall also use the notation $d_L := \sum_{l=1}^{L} 2^l$ and let $X^{(j)}_{L;i}$ denote vector in $\mathbb{R}^{d_L}$ formed by the first d_L coordinates of $\Pi_L X^{(j)}_i$ and let $X^{(j)}_L = (X^{(j)}_{L;i})_{i\in[n]}$. Furthermore, we recall that for $v = (v_1, \dots, v_n) \in \mathcal{X}^n$ for a vector space $\mathcal{X}$, $\bar{v}$ denotes the vector space average $n^{-1}\sum_{i=1}^{n} v_i$.

In order to obtain statistics with (uniformly) bounded sensitivity it is useful to bound quantities between certain thresholds. Formally, for $a, b, x \in \mathbb{R}$ with $a < b$, let $[x]^b_a$ denote x *clipped between* a *and* b, that is $[x]^b_a = \max(a, \min(b, x))$.

The distributed privacy protocols under consideration in this paper can be seen as noisy versions of statistics of the data. Roughly put, the "amount" of noise added depends on the *sensitivity* of the statistics. This brings us to the concept of sensitivity. Formally, consider a metric d on a set $\mathcal{Y}$. Given n elements $x = (x_1, \dots, x_n)$ in a sample space $\mathcal{X}$, the d-*sensitivity at* x of a map $S : \mathcal{X}^n \to \mathcal{Y}$ is $\Delta_S(x) := \sup_{\check{x}\in\mathcal{X}^n : \mathrm{d}_H(x,\check{x})\leqslant 1} \mathrm{d}\left(S(x), S(\check{x})\right)$, where d_H is the Hamming distance on $\mathcal{X}^n$ (see Section 1.6 for a definition). The d-*sensitivity* of S is defined as $\Delta_S := \sup_x \Delta_S(x)$. In this paper, the main noise mechanism is the *Gaussian mechanism*. The Gaussian mechanism yields (ϵ, δ)-differentially private transcripts for statistics that have bounded L_2-sensitivity, with the noise variance scaling with the L_2-sensitivity. See Dwork et al. (2014a) for a thorough treatment. We remark that for the rates in Regime 3 up until 6 in Table 1, $(\epsilon, 0)$-DP can be attained by employing a Laplace mechanism instead. That is, for the values of ϵ for which Regime 3 up until 6 in Table 1 are optimal, the test statistics in the sections have matching L_1- and L_2-sensitivity, so the Gaussian mechanism can be replaced by the Laplace mechanism instead in these regimes.

3.1 Private testing procedure I: low privacy-budget strategy

In the classical setting without privacy constraints (and $m = 1$), a rate optimal test for the hypotheses of (3) is given by

$$\mathbb{1}\left\{S^{(j)}_L > \kappa_\alpha\right\}, \text{ where } S^{(j)}_L := \frac{1}{\sqrt{d_L}}\left(\left\|\sigma^{-1}\sqrt{n}\overline{X^{(j)}_L}\right\|_2^2 - d_L\right), \tag{12}$$

where $d_L := \sum_{l=1}^{L} 2^l$ and the rate optimal choice of L is $L_* = \left\lceil \frac{1}{2s+1/2}\log_2(N/\sigma^2)\right\rceil$. Under the null hypothesis, $S^{(j)}_{L_*}$ is Chi-square distributed degrees of freedom. Under the alternative hypothesis, the test statistic picks up a positive "bias" as $\|\sigma^{-1}\sqrt{n}\overline{X^{(j)}_{L_*}}\|^2 \sim \chi^2_{L_*}(\|\Pi_{L_*} f\|_2^2)$

under $\mathbb{P}_f$, which could surpass the critical value κ_α if $\sigma^{-2}n\|\Pi_{L_*}f\|_2^2$ is large enough. Consequently, the level of the test is controlled by setting κ_α appropriately large. For a proof of its rate optimality, see e.g. Gine and Nickl (2016).

As is commonly the case for superlinear functions, the test statistic $S_{L;\tau}^{(j)}$ has poor sensitivity uniformly over the sample space, meaning that a change in just one datum can result in a large change in the test statistic. This means that it forms a poor candidate to base a privacy preserving transcript on. For example, one would need to add a substantial amount of noise guarantee DP for the statistic. To remedy this, we follow a similar strategy as proposed in Canonne et al. (2020) and improved upon by Narayanan (2022b). We construct a clipped and symmetrized version of the test statistic above, which has small sensitivity on a set $\mathcal{C}_{L;\tau}$, in which $X^{(j)}$ takes values with high probability. We define the test statistic explicitly on $\mathcal{C}_{L;\tau}$ only. By a version of the McShane–Whitney–Extension Theorem, we obtain a test statistic with the same sensitivity that is defined on the entire sample space.

Consider for $\tau > 0$, $L \in \mathbb{N}$, $d_L := \sum_{l=1}^{L} 2^l$ and $V_{L;\tau}^{(j)} \sim \chi^2_{d_L}$ independent of $X^{(j)}$ the random map from $(\mathbb{R}^{d_L})^n$ to $\mathbb{R}$ defined by

$$\tilde{S}_{L;\tau}^{(j)}(x) = \left[\frac{1}{\sqrt{d_L}}\left(\|\sigma^{-1}\sqrt{n}\overline{x}\|_2^2 - V_{L;\tau}^{(j)}\right)\right]_{-\tau}^{\tau}. \tag{13}$$

For any τ, this test statistic $\tilde{S}_{L;\tau}^{(j)}(X_L^{(j)})$ can be seen to have mean zero and bounded variance under the null hypothesis, by similar reasoning as for the test statistic in (12) (see the proof of Lemma 8 for details).

Loosely speaking, the test statistic $\tilde{S}_{L;\tau}^{(j)}(X_L^{(j)})$ retains the signal as long as $\tau > 0$ is chosen appropriately in comparison to the signal size (i.e. $\|\Pi_L f\|_2^2$) and has good sensitivity for "likely" values of $X^{(j)}$ under $\mathbb{P}_f$, but not uniformly over the sample space. We make the latter statement precise as follows.

Let $K_\tau = \lceil 2\tau D_\tau^{-1} \rceil$ and consider the set $\mathcal{C}_{L;\tau} = \mathcal{A}_{L;\tau} \cap \mathcal{B}_{L;\tau}$, where

$$\mathcal{A}_{L;\tau} = \left\{(x_i) \in (\mathbb{R}^\infty)^n : \left|\|\sigma^{-1}\textstyle\sum_{i\in\mathcal{J}} \Pi_L x_i\|_2^2 - kd_L\right| \leqslant \tfrac{1}{8}kD_\tau n\sqrt{d_L} \ \ \forall \mathcal{J} \subset [n], |\mathcal{J}| = k \leqslant K_\tau\right\}, \tag{14}$$

$$\mathcal{B}_{L;\tau} = \left\{(x_i) \in (\mathbb{R}^\infty)^n : \left|\langle \sigma^{-1}\Pi_L x_i, \sigma^{-1}\textstyle\sum_{k\neq i} \Pi_L x_k\rangle\right| \leqslant \frac{1}{8}kD_\tau n\sqrt{d_L}, \ \ \forall i = 1, \ldots, n\right\}.$$

Lemma 26 in the supplement shows that $X^{(j)}$ concentrates on $\mathcal{C}_{L;\tau}$ when the underlying signal is, roughly speaking, not too large compared to τ (in particular under the null hypothesis).

It can be shown that, on the set $\mathcal{C}_{L;\tau}$, $x \mapsto S^{(j)}(x)$ is D_τ-Lipschitz with respect to the Hamming distance, see Lemma 27 in the supplement. Lemma 28 in the supplement shows that there exists a measurable function $S_{L;\tau}^{(j)} : (\mathbb{R}^{d_L})^n \to \mathbb{R}$, D_τ-Lipschitz with respect to the Hamming distance, such that $S_{L;\tau}^{(j)}(X_L^{(j)}) = \tilde{S}_{L;\tau}^{(j)}(X_L^{(j)})$ whenever $X^{(j)} \in \mathcal{C}_{L;\tau}$. Lemma 28 is essentially the construction of McShane McShane (1934) for obtaining a Lipschitz extension with respect to the Hamming distance, but our lemma verifies in addition the Borel measurability of the resulting map.

The Lipschitz constant upper bounds the sensitivity of a test statistic that is Lipschitz continuous with respect to the Hamming distance. Specifically, we have that

$$\Delta_{S^{(j)}} = \sup_{x,\breve{x}\in\ell_2(\mathbb{N})^n:\mathrm{d}_H(x,\breve{x})\leqslant 1} \left|S^{(j)}(x) - S^{(j)}(\breve{x})\right| \leqslant D_\tau.$$

Using the Gaussian mechanism, the transcripts

$$Y^{(j)}_{L;\tau} = \gamma_\tau \breve{S}^{(j)}_{L;\tau}(X^{(j)}_L) + W^{(j)}_\tau, \quad \text{where } W^{(j)}_\tau \sim N(0,1) \text{ independent for } j \in [m], \tag{15}$$

$\gamma_\tau = \epsilon/(D_\tau\sqrt{2\mathfrak{c}\log(2/\delta)})$ and $\tau > 0$, are $(\epsilon/\sqrt{\mathfrak{c}}, \delta)$-differentially private for any $\epsilon > 0$ (see e.g. Dwork et al. (2014a)). These transcripts are mean zero and have bounded variance under the null hypothesis, so a test of the form

$$\varphi_\tau := \mathbb{1}\left\{\frac{1}{\sqrt{m}}\sum_{j=1}^m Y^{(j)}_{L;\tau} \geqslant \kappa(\gamma_\tau \vee 1)\right\} \tag{16}$$

has an arbitrarily small level for large enough $\kappa > 0$ (see Lemma 29 in the supplement). Furthermore, the lemma below shows that, if the signal size is large enough in the $\sum_{l=1}^L 2^l$ first coordinates, the above test enjoys a small Type II error probability as well.

Lemma 8. *Consider the test φ_τ as defined by* (16)*. If*

$$\tau/4 \leqslant \frac{n\|f_L\|_2^2}{\log(N)\sqrt{2\mathfrak{c}\log(2/\delta)}\sigma^2\sqrt{d}} \leqslant \tau/2 \tag{17}$$

and

$$\|\Pi_L f\|_2^2 \geqslant C_\alpha \kappa \log(N)\sqrt{\mathfrak{c}\log(1/\delta)}\left(\frac{\sqrt{2^L}}{\sigma^2\sqrt{N}\sqrt{n}(\sqrt{n}\epsilon \wedge 1)}\right) \bigvee \left(\frac{1}{\sigma^2 N n \epsilon^2}\right) \tag{18}$$

for $C_\alpha > 0$ large enough, it holds that $\mathbb{P}_f(1-\varphi_\tau) \leqslant \alpha$.

A proof of the above lemma is given in Section B.1 of the supplement. The above test is calibrated for the detection of signals size between $\tau/4$ and $\tau/2$. In order to detect signals of any size larger than the right-hand side of (18), we follow what is essentially a multiple testing procedure. For large signals, we need a larger clipping to detect them, as well as a larger set $\mathcal{C}_{L;\tau}$ to assure that the data is in $\mathcal{C}_{L;\tau}$ with high probability, as larger signals increase the probability of "outliers" from the perspective of the sensitivity of the L_2-norm.

It turns out that a sufficient range of clipping thresholds to consider (for detecting the signals $f \in \mathcal{B}^{s,R}_{p,q}$ under consideration in Lemma 9) is given by

$$\tau \in \mathrm{T}_L := \left\{2^{-k+2}\frac{n(1-2^{-s})^{2-2/q}R^2}{\sigma^2\sqrt{2^L}} : k = 1, \dots, \lceil 1 + 2\log_2(NR/\sigma)\rceil\right\}. \tag{19}$$

The (ϵ,δ)-differentially private testing procedure T_{I} is now constructed as follows. For each $\tau \in \mathrm{T}_L$, the machine transfers (15) with $\mathfrak{c} = |\mathrm{T}_L|$. By the independence of the Gaussian noise added in (15) for each $\tau \in \mathrm{T}_L$, the transcript $Y^{(j)} = \{Y^{(j)}_{L;\tau} : \tau \in \mathrm{T}_L\}$ is (ϵ,δ)-differentially private (see e.g. Theorem A.1 in Dwork et al. (2014a)).

The test

$$T_{\mathrm{I}} := \mathbb{1}\left\{\max_{\tau\in \mathrm{T}_L} \frac{1}{\sqrt{m}}\sum_{j=1}^{m} Y_{L;\tau}^{(j)} \geqslant \kappa_\alpha \left(\frac{\epsilon}{D_\tau\sqrt{2|\mathrm{T}_L|\log(2/\delta)}} \vee 1\right)\sqrt{\log|\mathrm{T}_L|}\right\} \tag{20}$$

then satisfies $\mathbb{P}_0 T_{\mathrm{I}} \leqslant \alpha$ via a union bound and sub-exponential tail bound, we defer the reader to the proof of Lemma 9 for details. Furthermore, for $f \in \mathcal{B}_{p,q}^{s,R}$, we have $\|\Pi_L f\|_2 \leqslant \|f\|_2 \lesssim R$. If f in addition satisfies (18), there exists $\tau^* \in \mathrm{T}_L$ such that (17) is satisfied and consequently $\mathbb{P}_f(1-T_{\mathrm{I}}) \leqslant \mathbb{P}_f(1-\varphi_{\tau^*}) \leqslant \alpha/2$.

The optimal choice of L depends on the regularity level of the signal f, balancing the approximation error $\|f-\Pi_L f\|_2^2$ and the right-hand side of (18), for which we defer the details to Section B.1 in the supplement. To summarize, we have obtained the following lemma.

Lemma 9. *For all $R>0$, $\alpha\in(0,1)$ there exists $\kappa_\alpha>0$ and $C_\alpha>0$ such that the test T_I defined in (20) satisfies $\mathbb{P}_0 T_I \leqslant \alpha$. Furthermore, if $f\in\mathcal{B}_{p,q}^{s,R}$ is such that for some L and $M_{N,\delta,\tau} = \log(N)\sqrt{\log\log(NR/\sigma)\log(NR/\sigma)\log(1/\delta)}$,*

$$\|\Pi_L f\|_2^2 \geqslant C_\alpha M_{N,\delta,\tau}\left(\frac{\sqrt{2^L}}{\sigma^2\sqrt{N}\sqrt{n}(\sqrt{n}\epsilon\wedge 1)}\right)\bigvee\left(\frac{1}{\sigma^2 N n\epsilon^2}\right),$$

we have that $\mathbb{P}_f(1-T_I)\leqslant\alpha$.

3.2 Private testing procedure II: high privacy-budget strategy

In the high-privacy budget regime, we construct a testing procedure that consists essentially of two steps. In the first step, the data is truncated, clipped and averaged over the coordinates, after which Gaussian noise is added to obtain a private summary of the original data. Then, as a second step, the transcripts are averaged, and based on this average, a test statistic that is reminiscent of a chi-square test is computed in the central server. This is in contrast to the strategy of the previous section, where each server computes a (private version of) a chi-square test statistic.

The approach taken here is to divide the servers equally over the first d_L coordinates (i.e. as uniformly as possible), where we recall the notation $d_L := \sum_{l=1}^{L} 2^l$. That is to say, for $L, K_L \in \mathbb{N}$, we partition the coordinates $\{1,\dots,d_L\}$ into approximately d_L/K_L sets of size K_L. The servers are then equally divided over each of these partitions and communicate the sum of the clipped $X_{L;i}^{(j)}$'s coefficients corresponding to their partition, were we also recall that the notation $X_{L;i}^{(j)}$ denotes the vector in $\mathbb{R}^{d_L}$ formed by the first d_L coordinates of $\Pi_L X_i^{(j)}$.

More formally, take $K_L = \lceil n\epsilon^2 \wedge d_L\rceil$ and consider sets $\mathcal{J}_{lk;L}\subset[m]$ for indexes $(l,k)\in\{l=1,\dots,L, k=1,\dots,2^l\} =: I_L$, such that $|\mathcal{J}_{lk;L}| = \lceil\frac{mK_L}{d_L}\rceil$ and each $j\in\{1,\dots,m\}$ is in $\mathcal{J}_{lk;L}$ for at least K_L different indexes $k\in\{1,\dots,d_L\}$. For $(l,k)\in I_L$, $j\in\mathcal{J}_{lk;L}$, generate the transcripts according to

$$Y_{lk;L}^{(j)}|X^{(j)} \equiv Y_{lk;L}^{(j)}(X^{(j)}) = \gamma_L\sum_{i=1}^{n}[\sigma^{-1}(X_i^{(j)})_{lk}]_{-\tau}^{\tau} + W_{lk}^{(j)} \tag{21}$$

with $\gamma_L = \epsilon/(2\sqrt{2K_L \log(2/\delta)}\tau)$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N/\sigma)}$ and $(W_{lk}^{(j)})_{j\in[m],(l,k)\in I_L}$ i.i.d. standard Gaussian noise.

Since $x \mapsto \sum_{i=1}^n [\sigma(x_i^{(j)})_{lk}]_{-\tau}^{\tau}$ has sensitivity bounded by 2τ, for $k = 1, \dots, K$, releasing

$$Y_L^{(j)}(X^{(j)}) = (Y_{L,l_1k_1}^{(j)}(X^{(j)}), \dots, Y_{L,l_{K_L}k_{K_L}}^{(j)}(X^{(j)}))$$

satisfies (ϵ, δ)-DP, see Lemma 31 in the supplement for details.

If the privacy budget were of no concern, submitting the above transcripts with $2^L \asymp N^{1/(2s+1/2)}$ would be sufficient to construct a test statistic that attains the unconstrained rate of $\rho^2 \asymp N^{-2s/(2s+1/2)}$. Under (more stringent) privacy constraints, however, the optimal number of coordinates to be transmitted should depend on the privacy budget. Whenever $\epsilon \lesssim 1/\sqrt{n}$, it turns out that submitting just one coordinate is in fact rate optimal. Sending more than one coordinate leads to worse rates as the noise overpowers the benefit of having a higher dimensional transcript. As ϵ increases, the optimal number of coordinates to be transmitted increases as well. Whenever $\epsilon \gtrsim \sigma^{-\frac{2}{4s+1}} m^{\frac{1}{4s+1}} n^{\frac{1/2-2s}{4s+1}}$, the optimal number of coordinates to be transmitted is $2^L \asymp N^{1/(2s+1/2)}$.

The test

$$T_{\mathrm{II}} = \mathbb{1}\left\{ \frac{1}{\sqrt{d_L}} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j\in\mathcal{J}_{lk;L}} Y_{lk;L}^{(j)} \right)^2 - \frac{n\epsilon^2}{4K_L\tau^2} - 1 \right] \geqslant \kappa_\alpha \left(\frac{n\epsilon^2}{4K_L\tau^2} \vee 1 \right) \right\} \tag{22}$$

satisfies $\mathbb{P}_0 T_{\mathrm{II}} \leqslant \alpha$ by Lemma 32 in the supplement whenever $\tilde{\kappa}_\alpha > 0$ and $\kappa_\alpha > 0$ are chosen large enough.

The power that the test attains depends on the signal size up until resolution level L, i.e. $\|\Pi_L f\|_2$. Specifically, the test Type II error $\mathbb{P}_f(1 - T_{\mathrm{II}}) \leqslant \alpha$ whenever

$$\|\Pi_L f\|_2^2 \geqslant C_\alpha \frac{\log\log(N)\log(N)\log(1/\delta)2^{(3/2)L}}{mn^2\epsilon^2}. \tag{23}$$

The optimal choice of L for is determined by the trade-off between the approximation error $\|f - \Pi_L f\|_2^2$ and the right-hand side of (23). The proof of the following lemma is given in Section B.2 of the supplement.

Lemma 10. *Take $\alpha \in (0,1)$. Suppose f satisfies (23) and that $\epsilon \geqslant \frac{2^{L+1}}{\sqrt{mn}}$ for some $L \in \mathbb{N}$. Then, the (ϵ, δ)-FDP testing protocol T_{II} of level α has Type II error $\mathbb{P}_f(1 - T) \leqslant \alpha$ for a large enough constant $C_\alpha > 0$ and $\tilde{\kappa}_\alpha > 0$, depending only on α.*

3.3 Private testing procedure III: high privacy-budget shared randomness strategy

In this section, we construct a testing procedure that is based on the same principles as the one in the previous section, but with the difference that the servers a source of randomness. The transcripts are still based on the clipped and averaged coordinates of the truncated data, but instead of dividing the servers across the coordinates, we apply the same random rotation across the servers.

Next, we describe the testing procedure in detail. Consider for $L \in \mathbb{N}$ the quantities $d_L = \sum_{l=1}^{L} 2^l$ and $K_L = \lceil n\epsilon^2 \wedge d_L \rceil$ and let U_L denote a random rotation uniformly drawn (i.e. from the Haar measure) on the group of random orthonormal $d_L \times d_L$-matrices.

For $I_L := \{(l,k) : l = 1, \dots, \lceil \log_2(K_L) \rceil,\ k = 1, \dots, 2^l\}$, $(l,k) \in I_L$ and $j = 1, \dots, m$, generate the transcripts according to

$$Y^{(j)}_{lk;L} | (X^{(j)}, U) = \gamma_L \sum_{i=1}^{n} [(UX^{(j)}_{L;i})_{lk}]^{\tau}_{-\tau} + W^{(j)}_{lk}, \tag{24}$$

with $\gamma_L = \frac{\epsilon}{2\sqrt{2K_L \log(2/\delta) \log(N)}\tau}$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N/\sigma)}$, $\tilde{\kappa}_\alpha > 0$ and $(W^{(j)}_l)_{j,l}$ i.i.d. centered standard Gaussian noise. By an application of Lemma 33, the transcript $Y^{(j)}_L := (Y^{(j)}_{lk;L})_{(l,k)\in I_L}$ is (ϵ, δ)-differentially private.

In the shared randomness strategy above, we essentially only send the first $\sum_{l=1}^{\lceil \log_2(K_L) \rceil} 2^l$ coordinates. The random rotation U_L ensures that, roughly speaking, a sufficient amount of the signal is present in these first coordinates, with high probability.

We then construct the test

$$T_{\mathrm{III}} = \mathbb{1}\left\{ \frac{1}{\sqrt{K_L}} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y^{(j)}_{lk;L} \right)^2 - n\gamma_L^2 - 1 \right] \geqslant \kappa_\alpha \left(n\gamma_L^2 \vee 1 \right) \right\}, \tag{25}$$

which satisfies $\mathbb{P}_0 \varphi \leqslant \alpha/2$ by Lemma 34 in the supplement, for $\kappa_\alpha > 0$ large enough. The lemma below is proven in Section B of the supplement, and yields that the Type II error of the test satisfies $\mathbb{P}_f(1 - T_{\mathrm{III}}) \leqslant \alpha$ whenever the coordinates up to resolution level L are of sufficient size. The optimal value for L depends on the truncation level s, and is chosen by balancing the approximation error $\|f - \Pi_L f\|_2^2$ and the right-hand side of (26), we defer the reader to Section B.3 of the supplement for details.

Lemma 11. *The testing protocol T_{III}, with level α and has corresponding Type II error probability $\mathbb{P}_f(1 - T_{III}) \leqslant \alpha$ whenever*

$$\|\Pi_L f\|_2^2 \geqslant C_\alpha \frac{2^L \log(1/\delta) \log(N)}{mn\sqrt{n\epsilon^2 \wedge 2^L}\sqrt{n\epsilon^2 \wedge 1}} \tag{26}$$

for constant $C_\alpha > 0$ and $\tilde{\kappa}_\alpha > 0$ depending only on α.

4 Adaptive tests under DP constraints

In the previous section we have derived methods that match (up to logarithmic factors) the theoretical lower bound established in Section 2. The proposed tests, however, depend on the regularity parameter s of the functional parameter of interest f.

In this section we derive an (ϵ, δ)-FDP testing protocol that adapts to the regularity when it is unknown. This method attains the optimal rate of Theorem 6, and consequently proves the aforementioned theorem.

The adaptive procedure builds on the tests constructed in Section 3 and combines them using essentially a multiple testing strategy. Roughly speaking, the method consists of

taking approximately a $1/\log N$-mesh-size grid in the regularity interval $[s_{\min}, s_{\max}]$, constructing optimal tests for each of the grid points and combining them using a type of Bonferroni's correction. By design, the tests constructed in Section 3 are based on sub-exponential private test statistics, which allows a combination of the test statistics with a Bonferroni correction of the order of $\log\log N$.

Combining $\log N$ many (ϵ', δ)-differentially private transcripts using Gaussian mechanisms, results in a (ϵ, δ)-differentially private protocol, with $\epsilon = \epsilon'\sqrt{\log N}$. This means that the erosion of the privacy budget by conducting a test for each grid-point is limited to a logarithmic factor, means the method greatly improves over the potentially polynomially worse rate of a non-adaptive method.

The detailed adaptive testing procedures are given as follows. Let ρ_s equal the right-hand side of (7) in case there is access to local randomness only, or the right-hand side of (6) in case shared randomness is available. Let $L_s = \lfloor s^{-1}\log_2(1/\rho_s)\rfloor \vee 1$ and define furthermore $\mathcal{S} := \{L_{s_{\min}}, \dots, L_{s_{\max}}\}$ such that $L_s \in \mathcal{S}$ for all $s \in [s_{\min}, s_{\max}]$. Furthermore, we note that the resulting "collection of resolution levels" satisfies $|\mathcal{S}| \leqslant C_{s_{\max}}\log N$ for some constant $C_{s_{\max}} > 0$ depending only on $s_{\max}$.

Consider the first the case without access to shared randomness. We partition the collection of resolution levels $\mathcal{S}$, depending on the model characteristics, as follows.

$$\mathcal{S}_{\mathtt{LR}}^{\mathtt{LOW}} = \left\{L \in \mathcal{S} : 2^L \leqslant \epsilon\sqrt{mn}(1 + \sqrt{n}\mathbb{1}_{\{\sqrt{n}\epsilon > 1\}})\right\}, \quad \mathcal{S}_{\mathtt{LR}}^{\mathtt{HIGH}} = \mathcal{S}\backslash\mathcal{S}_{\mathtt{LR}}^{\mathtt{LOW}}. \tag{27}$$

If the true regularity s_0 is such that $L_{s_0} \in \mathcal{S}^{\mathtt{LOW}}$, the low privacy-budget test of Section 3.1 (with $L = L_{s_0}$) is a rate optimal strategy. If $L_{s_0} \in \mathcal{S}^{\mathtt{HIGH}}$, the high privacy-budget test of Section 3.2 is rate optimal.

For the case of shared randomness, the phase transitions occur for different values of $s \in [s_{\min}, s_{\max}]$, or their respective resolution levels L_s. So in this case, we partition the collection of resolution levels as

$$\mathcal{S}_{\mathtt{SHR}}^{\mathtt{LOW}} = \left\{L \in \mathcal{S} : 2^L \leqslant \epsilon^2 mn\right\}, \quad \mathcal{S}_{\mathtt{SHR}}^{\mathtt{HIGH}} = \mathcal{S}\backslash\mathcal{S}_{\mathtt{SHR}}^{\mathtt{LOW}}. \tag{28}$$

Consider some $\mathcal{S}' \subset \mathcal{S}$ and let $\mathrm{T} \cup_{L\in\mathcal{S}'} \mathrm{T}_L$, where T_L is as defined in (19). The "adaptive version" of the low privacy-budget test defined in (20) takes the form

$$T_{\mathrm{I}}^S := \mathbb{1}\left\{\max_{L\in\mathcal{S}', \tau\in\mathrm{T}} \frac{1}{\sqrt{m}\,(\gamma_L \vee 1)\sqrt{\log|\mathrm{T}||\mathcal{S}'|}} \sum_{j=1}^m Y_{L;\tau}^{(j)} \geqslant \kappa_\alpha\right\}, \tag{29}$$

where $Y_L^{(j)} = \{Y_{L;\tau}^{(j)} : \tau \in \mathrm{T}_L\}$ is generated according to (15) for $L \in S$ with

$$\gamma_\tau = \frac{\epsilon}{2D_\tau\sqrt{|\mathrm{T}||\mathcal{S}'|\log(4/\delta)}}.$$

The above choice of ϵ yields that $(Y_L^{(j)})_{L\in S}$ is $(\epsilon/2, \delta/2)$-DP due to the Gaussian mechanism. The enlargement of the critical region, which is now effectively rescaled by $\sqrt{\log|\mathrm{T}||\mathcal{S}'|}$ instead of $\sqrt{\log|\mathrm{T}|}$, accounts for the potentially larger set of test statistics over which the maximum is taken. In the case of having access only to local sources of randomness, we

set $\mathcal{S}' = \mathcal{S}_{\texttt{LR}}^{\texttt{LOW}}$. If $\mathcal{S}_{\texttt{LR}}^{\texttt{LOW}}$ is empty, we set $T_{\mathrm{I}} = 0$ instead, which forms an $(0, 0)$-differentially private protocol.

In the case of having access to local sources of randomness only; if $\mathcal{S}_{\texttt{LR}}^{\texttt{HIGH}}$ is non-empty, the adaptive version of the high privacy-budget test defined in (B.3) is given by

$$T_{\mathrm{II}} = \mathbb{1}\left\{ \max_{L \in \mathcal{S}_{\texttt{LR}}^{\texttt{HIGH}}} \frac{1}{\sqrt{d_L}\,(\eta_L \vee 1)} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} Y_{lk;L}^{(j)} \right)^2 - \eta_L - 1 \right] \geqslant \kappa_\alpha \sqrt{\log |\mathcal{S}^{\texttt{HIGH}}|} \right\}, \tag{30}$$

where the transcripts are generated according to (95) for $L \in \mathcal{S}^{\texttt{HIGH}}$, with $\gamma_L = \epsilon/(4\sqrt{|\mathcal{S}^{\texttt{HIGH}}| K_L \log(4/\delta)}\tau)$, $\eta_L = \frac{n\epsilon^2}{4K_L\tau^2}$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N/\sigma)}$. Due to the Gaussian mechanism, the transcripts satisfy an $(\epsilon/2, \delta/2)$-DP constraint. As before, if $\mathcal{S}^{\texttt{HIGH}}$ is empty, set $T_{\mathrm{II}} = 0$ instead.

In the case of having access to local randomness only, the adaptive testing procedure then consists of computing the tests $T_{\mathrm{I}}^{\mathcal{S}^{\texttt{LOW}}}$ and T_{II}, for which the released transcripts satisfy (ϵ, δ)-DP. The final test is then given by

$$T = T_{\mathrm{I}}^{\mathcal{S}^{\texttt{LOW}}} \vee T_{\mathrm{II}}. \tag{31}$$

In Section B in the supplement, it is shown that this test is adaptive and rate optimal (up to logarithmic factors), proving the first part of Theorem 6.

In case of shared randomness, the adaptive version of the high privacy-budget test defined in (103) is given by

$$T_{\mathrm{III}} = \mathbb{1}\left\{ \max_{L \in S_{\texttt{SHR}}^{\texttt{HIGH}}} \frac{1}{\sqrt{K_L}\left(n\gamma_L^2 \vee 1\right)} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{lk;L}^{(j)} \right)^2 - n\gamma_L^2 - 1 \right] \geqslant \kappa_\alpha \sqrt{\log |\mathcal{S}|} \right\}, \tag{32}$$

where the transcripts are generated according to (102) for $L \in S_{\texttt{SHR}}^{\texttt{HIGH}}$, $\gamma_L = \frac{\epsilon}{4\sqrt{K_L |S_{\texttt{SHR}}^{\texttt{HIGH}}| \log(4/\delta)\log(N)}\tau}$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N/\sigma)}$. By similar reasoning as earlier, the transcripts $\{Y_L^{(j)} : L \in \S_{\texttt{SHR}}^{\texttt{HIGH}}\}$ are $(\epsilon/2, \delta/2)$-DP. If $S_{\texttt{SHR}}^{\texttt{HIGH}}$ is empty, we set $T_{\mathrm{III}} = 0$ instead.

The adaptive testing procedure in the case of shared randomness then consists of computing the tests $T_{\mathrm{I}}^{S^{\texttt{LOW}}}$ and T_{III}, for which the released transcripts satisfy (ϵ, δ)-DP. The final test is then given by

$$T = T_{\mathrm{I}}^{S_{\texttt{SHR}}^{\texttt{LOW}}} \vee T_{\mathrm{III}}. \tag{33}$$

In the supplement's Section B, we prove that this test is adaptive, attaining the optimal rate for shared randomness protocols (up to logarithmic factors), giving us the second statement Theorem 6.

5 The minimax private testing lower bound

In this section, we present a single theorem outlining the lower bound for the detection threshold for testing protocols that adhere to DP constraints, with and without the use of shared randomness. The theorem directly yields the "lower bound part" of Theorems 3 and 2 presented in Section 2. In conjunction with Theorem 7, the theorem shows that the tests constructed in Section 3 are rate optimal up to logarithmic factors.

Theorem 12. *Let $s, R > 0$ be given. For all $\alpha \in (0,1)$, there exists a constant $c_\alpha > 0$ such that if*

$$\rho^2 \leqslant c_\alpha \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)^{\frac{2s}{2s+3/2}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{m}n\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)\right), \quad (34)$$

it holds that

$$\inf_{T \in \mathscr{T}^{(\epsilon,\delta)}} \mathcal{R}(H_\rho^{s,R}, T) > 1 - \alpha, \quad (35)$$

for all natural numbers m, N and $n = N/m$, $\sigma > 0$, $\epsilon \in (N^{-1}, 1]$ and $\delta \leqslant N^{-(1+\omega)}$ for any constant $\omega > 0$.

Similarly, for any $\alpha \in (0,1)$, there exists a constant $c_\alpha > 0$ such that if

$$\rho^2 \leqslant c_\alpha \left(\frac{\sigma^2}{mn}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^{3/2}\epsilon\sqrt{1 \wedge n\epsilon^2}}\right)^{\frac{2s}{2s+1}} \wedge \left(\left(\frac{\sigma^2}{\sqrt{m}n}\right)^{\frac{2s}{2s+1/2}} + \left(\frac{\sigma^2}{mn^2\epsilon^2}\right)\right), \quad (36)$$

we have that there exists a (ϵ,δ)-FDP shared randomness testing protocol $T \equiv T_{m,n,s,\sigma}$ such that

$$\inf_{T \in \mathscr{T}_{\mathit{SHR}}^{(\epsilon,\delta)}} \mathcal{R}(H_\rho^{s,R}, T) > 1 - \alpha, \quad (37)$$

for all natural numbers m, N and $n = N/m$, $\sigma > 0$, $\epsilon \in (N^{-1}, 1]$ and $\delta \leqslant N^{-(1+\omega)}$ for any constant $\omega > 0$.

The theorem states that, whenever the signal-to-noise ratio ρ is below a certain threshold times the minimax separation rate, no testing protocol can achieve a combined Type I and Type II error rate below α. Its proof is lengthy and involves a combination of various techniques. We defer the full details of the proof to Section A of the supplement, but provide an overview of the main steps below.

For Steps 1, 2 and 3, there is no distinction between local and shared randomness. We use the same notation for distributed protocols in these steps, but simply assume U is degenerate in the case of local randomness.

Step 1: The first step is standard in minimax testing analysis: we lower bound the testing risk by a Bayes risk,

$$\inf_{T \in \mathscr{T}} \mathcal{R}(H_\rho, T) \geqslant \inf_{T \in \mathscr{T}} \sup_{\pi} \left(\mathbb{P}_0(T(Y) = 1) + \int \mathbb{P}_f(T(Y) = 0) d\pi(f) - \pi(H_\rho^c)\right), \quad (38)$$

where $\mathscr{T}$ denotes either the class of local randomness or shared randomness (ϵ,δ)-DP protocols. This inequality allows the prior π to be chosen adversarially to the distribution of the transcripts. This turns out to be crucial in the context of local randomness protocols, as is further highlighted in Step 4. The specific prior distribution is chosen to be a centered Gaussian distribution, with a finite rank covariance, where the rank is of the order 2^L, for some $L \in \mathbb{N}$. This covariance is constructed in a way that it puts most of its mass in the dimensions in which the privacy protocol is the least informative, whilst at the same time it assures that the probability mass outside of

the alternative hypothesis $\pi(H_\rho^c)$ is small. The particular choice for a Gaussian prior (instead of e.g. the two point prior in Ingster and Suslina (2003b)) is motivated by Step 3.

Step 2: In this step, we approximate the distribution of the transcripts with another distribution that results in approximately the same testing risk, but has two particular favorable properties for our purposes.

- Whenever the distribution of a transcript satisfies an (ϵ, δ)-DP constraint with $\delta > 0$, the transcript's density can be unbounded on a set of small probability mass (proportional to δ). Consequently, the local likelihood of the transcript can have erratic behavior in the tails. To remedy this, we consider approximations to the transcript with a bounded likelihood. These approximating transcripts satisfy a $(\epsilon, 2\delta)$-DP privacy constraint. These bounded likelihoods enable the argument of Step 5.
- Similarly, whenever $\delta > 0$, the distribution of the data conditionally on the transcript no longer has a bounded density. For the argument employed in Step 3, we require a uniform abound on the density of the distribution of $X|Y$. We mitigate this by approximating the transcript with another one such that the data has a bounded density conditionally on the approximating transcript. The approximating transcript satisfies a $(\epsilon, 3\delta)$-DP constraint.

Furthermore, we show that both approximations can be done in a way that the approximating transcript distribution is $(\epsilon, 6\delta)$-DP.

Step 3: By standard arguments, on can further lower bound the testing risk in (38) for a particular transcript distribution $\mathbb{P}^{Y|X,U} = \bigotimes_{j=1}^m \mathbb{P}^{Y^{(j)}|X^{(j)},U}$ and prior distribution π by a quantity depending on the chi-square divergence between $\mathbb{P}_\pi^{Y|U=u}$ and $\mathbb{P}_0^{Y|U=u}$;

$$1 - \left(\sqrt{(1/2) \int \mathbb{E}_0^{Y|U=u} \left(\left(\frac{d\mathbb{P}_\pi^{Y|U=u}}{d\mathbb{P}_0^{Y|U=u}} \right)^2 - 1 \right)^2 d\mathbb{P}^U(u)} + \pi(H_\rho^c) \right). \tag{39}$$

The likelihood ratio of the transcripts depends on the privacy protocol, and is difficult to analyze directly. We employ the technique developed in Szabó et al. (2023). Specifically, Lemma 10.1 in Szabó et al. (2023), which states, roughly speaking, that the inequality

$$\mathbb{E}_0^{Y|U=u} \left(\frac{d\mathbb{P}_\pi^{Y|U=u}}{d\mathbb{P}_0^{Y|U=u}} \right)^2 \leqslant G \prod_{j=1}^m \mathbb{E}_0^{Y^{(j)}|U=u} \left(\frac{d\mathbb{P}_\pi^{Y^{(j)}|U=u}}{d\mathbb{P}_0^{Y^{(j)}|U=u}} \right)^2 \tag{40}$$

holds for a finite constant $0 < G < \infty$ and equality with the smallest possible G is attained whenever the conditional distribution of the data given the transcripts is Gaussian in an appropriate sense (we defer the details here to Section A.3 in the supplement). This result is a type of Brascamp-Lieb inequality of Lieb (1990). That

(40) has a "Gaussian maximizer" allows tractable analysis of the chi-square divergence in (39), yielding that the latter display is further lower bounded by

$$1 - \sqrt{(1/2)\int (\mathrm{A}_u^\pi \mathrm{B}_u^\pi - 1)\, d\mathbb{P}^U(u)} + \pi(H_\rho^c), \tag{41}$$

where

$$\mathrm{A}_u^\pi := \int e^{f^\top \sum_{j=1}^m \Xi_u^j g} d(\pi \times \pi)(f,g), \quad \mathrm{B}_u^\pi := \prod_{j=1}^m \mathbb{E}_0^{Y^{(j)}|U=u} \left(\frac{d\mathbb{P}_\pi^{Y^{(j)}|U=u}}{d\mathbb{P}_0^{Y^{(j)}|U=u}} \right), \tag{42}$$

where Ξ_u^j denotes the covariance of (a subset of) the data $X_L^{(j)}$ (defined as in (43)) conditionally on the transcript $Y^{(j)}$ and $U = u$;

$$\Xi_u^j := \mathbb{E}_0^{Y^{(j)}|U=u} \mathbb{E}_0 \left[\sum_{i=1}^n \sigma^{-1} X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right] \mathbb{E}_0 \left[\sum_{i=1}^n \sigma^{-1} X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right]^\top. \tag{43}$$

Whilst the quantities A_u and B_u^π are still not fully tractable, sharp bounds for both are possible and form the content of Steps 4 and 5, respectively.

Step 4: As remarked earlier, class of local randomness protocols is a strictly smaller class. To attain the sharper (i.e. larger) lower bound for local randomness protocols, we exploit the fact that Step 2 allows us to choose the prior adversarially to the distribution of the transcripts. In particular, since U is degenerate in the case of local randomness only, this means that the covariance of π to be more diffuse in the directions in which Ξ_u^j is the smallest. When considering shared randomness protocols, U is not degenerate, and the lower bound follows by taking the covariance of π to be an order 2^L-rank approximation of the identity map on $\ell_2(\mathbb{N})$. The bounds for A_u^π are

$$\mathrm{A}_u^\pi = \exp\left(C \frac{\rho^4}{c_\alpha 2^{3L}} \mathrm{Tr}\left(\Xi_u\right)^2 \right) \quad \text{and} \quad \mathrm{A}_u^\pi \leqslant \exp\left(C \frac{\rho^4}{c_\alpha 2^{2L}} \|\Xi_u\| \mathrm{Tr}\left(\Xi_u\right) \right) \tag{44}$$

for local randomness protocols and shared randomness protocols, respectively.

Step 5: So far, Steps 1-4 have not used the fact that the transcripts are necessarily less informative than the original data, as a consequence of the transcripts being (ϵ, δ)-DP. In this step, we exploit the privacy constraint to argue that A_u^π and B_u^π are small at the detection boundary for ρ.

In order to capture the information loss due to privacy in A_u^π, it suffices to bound the trace and operator norm of Ξ_u. The quantity Ξ_u can be seen as the Fisher information of the finite dimensional submodel spanned by the covariance of π. This quantity, loosely speaking, captures how much information the transcript contains on the original data. In order to analyze Ξ_u, we rely on a "score attack" type of technique, as employed in Cai et al. (2023b,a).

The quantity B_u^π corresponds to the (product of) the local likelihoods of the transcripts. Whenever $\epsilon \geqslant 1/\sqrt{n}$, it suffices to consider the trivial bound

$$\mathbb{E}_0^{Y^{(j)}|U=u}\left(\frac{d\mathbb{P}_\pi^{Y^{(j)}|U=u}}{d\mathbb{P}_0^{Y^{(j)}|U=u}}\right) \leqslant \mathbb{E}_0^{X^{(j)}|U=u}\left(\frac{d\mathbb{P}_\pi^{X^{(j)}}}{d\mathbb{P}_0^{X^{(j)}}}\right), \tag{45}$$

and further bounding the right-hand side without privacy specific arguments. Whenever $\epsilon < 1/\sqrt{n}$, more sophisticated methods are need to capture the effect of privacy. Our argument uses a coupling method, which, combined with the fact that the likelihoods of the transcripts are bounded in our construction, allows us to obtain a sharp bound for B_u^π. After obtaining the bounds in terms of the rank 2^L and ρ, the proof is finished by choosing L such that the second and third term in (39) are balanced (minimizing their sum).

6 Discussion

The findings in this paper highlight the trade-off between statistical accuracy and privacy in federated goodness-of-fit testing under federated differential privacy (FDP) constraints. We characterize the problem in terms of the minimax separation rate, which quantifies the difficulty of the testing problem based on the regularity of the underlying function, the sample size, the degree of data distribution, and the stringency of the DP constraint. The minimax separation rate varies depending on whether the testing protocol has access to local or shared randomness. Furthermore, we construct data-driven adaptive testing procedures that achieve the same optimal performance, up to logarithmic factors, even when the regularity of the functional parameter is unknown.

One possible extension of this work is to consider a more general distribution of the privacy budget across the servers. Our current analysis supports differing budgets to the extent that $\epsilon_j \asymp \epsilon_k$, $\delta_j \asymp \delta_k$, and $n_j \asymp n_k$. However, one could explore more heterogeneous settings where severs differ significantly in their differential privacy constraints and number of observations. Although this would complicate the presentation of results, the techniques developed in this paper could, in principle, be extended to such settings.

Another interesting direction is to consider multiple testing problems, where the goal is to test multiple hypotheses simultaneously. We anticipate that the framework, insights, and theoretical results provided in the current paper will serve as valuable resources for future studies in this domain.

Regarding adaptation, not much is known about the cost of fundamental privacy. Interestingly, the cost of adaptation is minimal in the privacy setting considered in this paper. It remains an open question whether this minimal cost is a general phenomenon, whether it can be characterized exactly, or whether the cost of adaptation is more severe in other settings. We leave these questions for future research.

Acknowledgments and Disclosure of Funding

The research was supported in part by NSF grant NSF DMS-2413106 and NIH grants R01-GM123056 and R01-GM129781.

References

Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. In *Advances in Neural Information Processing Systems*, volume 31, 2018a.

Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. *Advances in Neural Information Processing Systems*, 31, 2018b.

Jayadev Acharya, Clement Canonne, Cody Freitag, and Himanshu Tyagi. Test without trust: Optimal locally private distribution testing. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89, pages 2067–2076. PMLR, 2019.

Jayadev Acharya, Kallista Bonawitz, Peter Kairouz, Daniel Ramage, and Ziteng Sun. Context aware local differential privacy. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 52–62. PMLR, 2020a.

Jayadev Acharya, Clément L. Canonne, and Himanshu Tyagi. Inference under information constraints i: Lower bounds from chi-square contraction. *IEEE Transactions on Information Theory*, 66(12):7835–7855, 2020b.

Jayadev Acharya, Clément L. Canonne, Cody Freitag, Ziteng Sun, and Himanshu Tyagi. Inference under information constraints iii: Local privacy constraints. *IEEE Journal on Selected Areas in Information Theory*, 2(1):253–267, 2021.

Jayadev Acharya, Yuhan 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 of *Proceedings of Machine Learning Research*, pages 8561–8585. PMLR, 25–27 Apr 2023.

Daniel Alabi and Salil Vadhan. Hypothesis testing for differentially private linear regression. In *Advances in Neural Information Processing Systems*, volume 35, pages 14196–14209, 2022.

Pathum Chamikara Mahawaga Arachchige, Peter Bertok, Ibrahim Khalil, Dongxi Liu, Seyit Camtepe, and Mohammed Atiquzzaman. Local differential privacy for deep learning. *IEEE Internet of Things Journal*, 7(7):5827–5842, 2019.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th annual symposium on foundations of computer science*, pages 464–473. IEEE, 2014.

Francoise Beaufays, Kanishka Rao, Rajiv Mathews, and Swaroop Ramaswamy. Federated learning for emoji prediction in a mobile keyboard. 2019.

Thomas Berrett and Cristina Butucea. Locally private non-asymptotic testing of discrete distributions is faster using interactive mechanisms. *Advances in Neural Information Processing Systems*, 33:3164–3173, 2020.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: a nonasymptotic theory of independence*. Oxford University Press, Oxford, 1st ed edition, 2013. OCLC: ocn818449985.

Gavin Brown, Samuel Hopkins, and Adam Smith. Fast, sample-efficient, affine-invariant private mean and covariance estimation for subgaussian distributions. In *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pages 5578–5579. PMLR, 12–15 Jul 2023.

Cristina Butucea, Amandine Dubois, Martin Kroll, and Adrien Saumard. Local differential privacy: Elbow effect in optimal density estimation and adaptation over Besov ellipsoids. *Bernoulli*, 26(3):1727 – 1764, 2020.

Cristina Butucea, Angelika Rohde, and Lukas Steinberger. Interactive versus noninteractive locally differentially private estimation: Two elbows for the quadratic functional. *Annals of Statistics*, 51(2), April 2023.

T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825–2850, 2021.

T. Tony Cai, Abhinav Chakraborty, and Lasse Vuursteen. Optimal federated learning for nonparametric regression with heterogenous distributed differential privacy constraints. *preprint*, 2023a.

T Tony Cai, Yichen Wang, and Linjun Zhang. Score attack: A lower bound technique for optimal differentially private learning. *arXiv preprint arXiv:2303.07152*, 2023b.

T Tony Cai, Dong Xia, and Mengyue Zha. Optimal differentially private pca and estimation for spiked covariance matrices. *arXiv preprint arXiv:2401.03820*, 2024.

Clément L Canonne and Yucheng Sun. Private distribution testing with heterogeneous constraints: Your epsilon might not be mine. *arXiv preprint arXiv:2309.06068*, 2023.

Clément L Canonne, Gautam Kamath, Audra McMillan, Adam Smith, and Jonathan Ullman. The structure of optimal private tests for simple hypotheses. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 310–321, 2019.

Clément L Canonne, Gautam Kamath, Audra McMillan, Jonathan Ullman, and Lydia Zakynthinou. Private identity testing for high-dimensional distributions. In *Advances in Neural Information Processing Systems*, volume 33, pages 10099–10111, 2020.

Albert Cohen, Ingrid Daubechies, and Pierre Vial. Wavelets on the interval and fast wavelet transforms. *Applied and computational harmonic analysis*, 1993.

Ingrid Daubechies. *Ten lectures on wavelets*. SIAM, 1992.

Amandine Dubois, Thomas Berrett, and Cristina Butucea. Goodness-of-Fit Testing for Hölder Continuous Densities Under Local Differential Privacy. In *Foundations of Modern Statistics*, volume PROMS-425, pages 53–119. Springer International Publishing, 2023.

John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, 2013.

John C Duchi, Michael I Jordan, and Martin J Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521): 182–201, 2018.

Cynthia Dwork and Adam Smith. Differential privacy for statistics: What we know and what we want to learn. *Journal of Privacy and Confidentiality*, 1(2), 2010.

Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014a.

Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 11–20, 2014b.

Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4:61–84, 2017.

MS Ermakov. Asymptotically minimax tests for nonparametric hypotheses concerning the distribution density. *Journal of Soviet Mathematics*, 52:2891–2898, 1990.

Marco Gaboardi, Hyun Lim, Ryan Rogers, and Salil Vadhan. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 2111–2120. PMLR, 2016.

Evarist Gine and Richard Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press, Cambridge, 2016.

Yu. I. Ingster and Irina A. Suslina. *Nonparametric Goodness-of-Fit Testing Under Gaussian Models*, volume 169 of *Lecture Notes in Statistics*. Springer New York, New York, NY, 2003a.

Yuri Ingster and Irina A Suslina. *Nonparametric goodness-of-fit testing under Gaussian models*, volume 169. Springer Science & Business Media, 2003b.

Yuri I Ingster. Asymptotically minimax hypothesis testing for nonparametric alternatives. i, ii, iii. *Math. Methods Statist*, 2(2):85–114, 1993.

Iain M Johnstone. *Function Estimation and Gaussian Sequence Models*. Unpublished manuscript, 2019.

Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Conference on Learning Theory*, pages 1853–1902. PMLR, 2019.

Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pages 2204–2235. PMLR, 2020.

Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. *arXiv preprint arXiv:1711.03908*, 2017.

A. N. Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell'Istituto Italiano degli Attuari*, 4:83–91, 1933.

Martin Kroll. On density estimation at a fixed point under local differential privacy. *Electronic Journal of Statistics*, 15(1):1783 – 1813, 2021.

Clément Lalanne, Aurélien Garivier, and Rémi Gribonval. About the cost of central privacy in density estimation. *Transactions on Machine Learning Research*, 2023.

Joseph Lam-Weil, Béatrice Laurent, and Jean-Michel Loubes. Minimax optimal goodness-of-fit testing for densities and multinomials under a local differential privacy constraint. *Bernoulli*, 28(1):579–600, 2022.

OV Lepskii. Asymptotically minimax adaptive estimation. i: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications*, 36(4):682–697, 1992.

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.

Mengchu Li, Ye Tian, Yang Feng, and Yi Yu. Federated transfer learning with differential privacy. 2024.

Elliott H. Lieb. Gaussian kernels have only Gaussian maximizers. *Inventiones Mathematicae*, 102(1):179–208, December 1990. Publisher: Springer New York.

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.

Edward James McShane. Extension of range of functions. 1934.

Shyam Narayanan. Private high-dimensional hypothesis testing. In *Conference on Learning Theory*, pages 3979–4027. PMLR, 2022a.

Shyam Narayanan. Private high-dimensional hypothesis testing. In *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178, pages 3979–4027. PMLR, 2022b.

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.

Michael Nussbaum. Asymptotic equivalence of density estimation and Gaussian white noise. *The Annals of Statistics*, 24(6):2399 – 2430, 1996.

Liudas Panavas, Amit Sarker, Sara Di Bartolomeo, Ali Sarvghad, Cody Dunne, and Narges Mahyar. Illuminating the landscape of differential privacy: An interview study on the use of visualization in real-world deployments. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–16, 2024.

Mathieu Sart. Density estimation under local differential privacy and Hellinger loss. *Bernoulli*, 29(3):2318 – 2341, 2023.

Or Sheffet. Locally private hypothesis testing. In *International Conference on Machine Learning*, pages 4605–4614. PMLR, 2018.

N. Smirnov. Table for Estimating the Goodness of Fit of Empirical Distributions. *The Annals of Mathematical Statistics*, 19(2):279 – 281, 1948.

Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822, 2011.

V. G. Spokoiny. Adaptive hypothesis testing using wavelets. *The Annals of Statistics*, 24 (6), December 1996.

Botond Szabó, Lasse Vuursteen, and Harry Van Zanten. Optimal distributed composite testing in high-dimensional gaussian models with 1-bit communication. *IEEE Transactions on Information Theory*, 68(6):4070–4084, 2022.

Botond Szabó, Lasse Vuursteen, and Harry van Zanten. Optimal high-dimensional and nonparametric distributed testing under communication constraints. *The Annals of Statistics*, 51(3):909 – 934, 2023.

H. Thorisson. *Coupling, Stationarity, and Regeneration*. Probability and Its Applications. Springer New York, 2000.

H. Triebel. *Theory of Function Spaces II*. Monographs in mathematics. Springer, 1992.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, 2009.

Aad W Van Der Vaart and Jon A Wellner. *Weak convergence*. Springer, 1996.

Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 1 edition, September 2018.

Lasse Vuursteen. Optimal private and communication constraint distributed goodness-of-fit testing for discrete distributions in the large sample regime. *Advances in Neural Information Processing Systems*, 37, 2024.

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

Appendix A. Proof of the lower bound Theorem 12

In this section, we provide a proof for Theorem 12. The proof is divided into several steps, following the outline provided in Section 5.

It is convenient to introduce the following notation. Given an (ϵ, δ)-DP protocol triplet

$$(T, \{\left(\mathbb{P}^{Y^{(j)}|X^{(j)}=x,U=u}\right)_{x\in\mathcal{X}^n,u\in\mathcal{U}}\}_{j=1}^m, (\mathcal{U}, \mathscr{U}, \mathbb{P}^U)),$$

we shall use the notation $P_f = \mathbb{P}_f^{X^{(j)}}$. For the Markov kernel $(A, (x, u)) \mapsto \mathbb{P}^{Y^{(j)}|(X^{(j)},U)=(x,u)}(A)$, we shall use the shorthand $(A, (x, u)) \mapsto K^j(A|(x, u))$. If the transcript $Y^{(j)}$ is (ϵ, δ)-DP, the Markov kernel satisfies

$$K^j(A|x_1, \ldots, x_i, \ldots, x_n, u) \leqslant e^\epsilon K^j(A|x_1, \ldots, x_i', \ldots, x_n, u) + \delta$$
$$\text{for all } A \in \mathscr{Y},\ x_i', x_1, \ldots, x_i, \ldots, x_n \in \mathcal{X},\ i \in \{1, \ldots, n\}.$$

For local randomness protocols, the probability space $(\mathcal{U}, \mathscr{U}, \mathbb{P}^U)$ has the trivial sigma-algebra; $\mathscr{U} = \{\varnothing, \mathcal{U}\}$. This allows streamlining the argument for both lower bounds of Theorem 12.

When the parameter underlying the true distribution of the data F is drawn from a prior distribution on $\ell_2(\mathbb{N})$, we obtain that the distributed testing protocol satisfies the Markov chain $F \to (X, U) \to Y \to T$, with $\mathbb{P}_f^{Y^{(j)}|U=u} = P_f K^j(\cdot|X^{(j)}, u)$ for all $j \in [m]$. Let $K_u = \bigotimes_{j=1}^m K^j(\cdot|\cdot, u)$ denote the product conditional distribution with $U = u$, such that distribution of the collection of transcripts conditionally on $U = u$ then satisfies $\mathbb{P}_f^{Y|U=u} = P_f^m K_u()$.

A.1 Step 1: Lower bounding the testing risk by the Bayes risk

As a first step, we lower bound the testing risk by the Bayes risk. Following the notation introduced above, let $T \equiv (T, \{K^j\}_{j=1}^m, (\mathcal{U}, \mathscr{U}, \mathbb{P}^U))$ be an (ϵ, δ)-DP distributed testing protocol and let $K_u = \bigotimes_{j=1}^m K^j(\cdot|\cdot, u)$. The testing risk for T can be written as

$$\mathcal{R}(H_\rho^{s,R}, T) = \int P_0^m K_u T d\mathbb{P}^U(u) + \sup_{f\in H_\rho^{s,R}} \int P_f^m K_u (1 - T) d\mathbb{P}^U(u). \tag{46}$$

Consider $L \in \mathbb{N}$, $d_L = \sum_{l=1}^L 2^l$ and consider $\pi = N(0, c_\alpha^{-1/2} d_L^{-1} \rho^2 \bar{\Gamma})$ for a symmetric idempotent matrix $\bar{\Gamma} \in \mathbb{R}^{d_L \times d_L}$. Consider also the linear operator $\Psi_L : \mathbb{R}^{d_L} \to \ell_2(\mathbb{N})$ defined by $\Psi_L \tilde{f} = f$ for $f_{lk} = \tilde{f} lk \cdot \mathbb{1}\{l \leqslant L\}$, $\tilde{f} = (\tilde{f}_{11}, \ldots, \tilde{f}_{L2^L}) \in \mathbb{R}^{d_L}$. Since Ψ_L is measurable, any probability distribution π_L on $\mathbb{R}^{2^L}$, $\pi_L \circ \Psi_L^{-1}$ defines a probability measure on the Borel sigma algebra of $\ell_2(\mathbb{N})$.

Using that $0 \leqslant T \leqslant 1$, the above testing risk is lower bounded by the Bayes risk

$$\int \left(P_0^m K_u T + \int P_f^m K_u (1 - T) d\pi \circ \Psi_L^{-1}(f) \right) d\mathbb{P}^U(u) - \pi \circ \Psi_L^{-1} \left((H_\rho^{s,R})^c \right). \tag{47}$$

We highlight here that $\bar{\Gamma}$ can depend on $\{K^j\}_{j=1}^m$.

To lower bound the testing risk further, it suffices to show that the prior π has little mass outside of $H_\rho^{s,R}$. This follows by a standard Gaussian concentration argument, provided in Lemma 13 below, from which it follows that, for $c_\alpha > 0$ small enough, it holds that $\pi \circ \Psi^{-1}((H_\rho^{s,R})^c) \leqslant \alpha/4$. This yields that (47) is lower bounded by

$$\int \left(P_0^m K_u T + \int P_f^m K_u (1-T) d\pi \circ \Psi_L^{-1}(f) \right) d\mathbb{P}^U(u) - \alpha/4. \tag{48}$$

Lemma 13. *Suppose $\rho \lesssim c_\alpha 2^{-Ls}$. Then, for any $c_\alpha > 0$ small enough, it holds that $\pi \circ \Psi_L^{-1}((H_\rho^{s,R})) \leqslant \alpha/4$.*

Proof Let $f \in \ell_2(\mathbb{N})$. It holds that $f \in H_\rho^{s,R}$ if and only if $\|f\|_2^2 \geqslant \rho^2$ and $\|f\|_{\mathcal{B}_{p,q}^s} \leqslant R$. For the first of these events, we have that

$$\pi \circ \Psi_L^{-1}(f : \|f\|_2^2 \geqslant \rho^2) = \Pr\left(Z^\top \bar{\Gamma} Z \geqslant \sqrt{c_\alpha} d_L \right), \tag{49}$$

where $Z \sim N(0, I_{d_L})$. Using that $\bar{\Gamma}$ is idempotent, the right-hand side of the above display can be made arbitrarily small for small enough choice of $c_\alpha > 0$, by a concentration argument for Chi-square random variables, see e.g. Lemma A.13 in Szabó et al. (2023).

To assure that $f \sim \pi \circ \Psi_L^{-1}$ concentrates on a Besov ball, we recall the definition of the Besov norm as given in (5). We have that

$$\begin{aligned}
\pi \circ \Psi_L^{-1}(f : \|f\|_{\mathcal{B}_{p,q}^s} \leqslant R) &= \pi \circ \Psi_L^{-1}\left(2^{L(s+1/2-1/p)} \left\| (f_{Lk})_{k=1}^{2^l} \right\|_p \leqslant R \right) \\
&= \Pr\left(2^{-L/p} \|Z\|_p \leqslant CR/c_\alpha^{3/4} \right)
\end{aligned}$$

where $Z \sim N(0, I_{2^L})$ and $C > 0$ a constant. The right-hand side of the above display can be made arbitrarily small for small enough choice of $c_\alpha > 0$, following from the fact that Z is a standard normal vector $\mathbb{E}\|Z\|_p^p \lesssim 2^L$, where the constant depends on p (see e.g. Proposition 2.5.2 in Vershynin (2018)) and Markov's inequality. ■

The larger L, the larger the effective dimension of the signal under the alternative hypothesis. Setting L such that $2^L \asymp \rho^{-1/s}$ means that the requirement of the above lemma are satisfied and consequently the Gaussian prior most of it its mass in the alternative hypothesis. Recalling that $d_L \asymp 2^L$, the condition (34) in the case of local randomness protocols can be written as

$$\rho^2 \lesssim c_\alpha \sigma^2 \left(\frac{d_L^{3/2}}{mn(n\epsilon^2 \wedge d_L)} \bigwedge \left(\frac{\sqrt{d_L}}{\sqrt{m} n \sqrt{n\epsilon^2 \wedge 1}} \bigvee \frac{1}{mn^2\epsilon^2} \right) \right), \tag{50}$$

and in the case of shared randomness protocols, (36) can be written as

$$\rho^2 \lesssim c_\alpha \sigma^2 \left(\frac{d_L}{mn\sqrt{n\epsilon^2 \wedge 1}\sqrt{n\epsilon^2 \wedge d_L}} \bigwedge \left(\frac{\sqrt{d_L}}{\sqrt{m} n \sqrt{n\epsilon^2 \wedge 1}} \bigvee \frac{1}{mn^2\epsilon^2} \right) \right). \tag{51}$$

A.2 Step 2: Approximating the distribution of the transcripts

In Step 3, we aim to use the Brascamp-Lieb type inequality Lemma 10.1 of Szabó et al. (2023), which we restate as Lemma 20 below. However, the *forward-backward channel* corresponding to K_u, $(x_1, x_2) \mapsto q_u(x_1, x_2)$, defined as

$$q_u(x_1, x_2) := \int \frac{dK(\cdot|x_1, u)}{d\mathbb{P}_0^{Y|U=u}}(y) \frac{dK(\cdot|x_2, u)}{d\mathbb{P}_0^{Y|U=u}}(y) d\mathbb{P}_0^{Y|U=u}(y), \tag{52}$$

is possibly unbounded when $\delta > 0$. To overcome this, we use Lemma 14 below to construct $(\epsilon, 3\delta)$-DP Markov kernels $\{\tilde{K}^j\}_{j=1}^m$ such that the corresponding forward-backward channel $\tilde{q}_u(x_1, x_2)$ is bounded. The construction of $\{\tilde{K}^j\}_{j=1}^m$ is such that the total variation distance between $\tilde{K}^j$ and $\tilde{K}^j$ is small.

Lemma 14. *For any $\alpha \in (0, 1)$ and (ϵ, δ)-DP collection of Markov kernels $\{K^j\}_{j=1}^m$, there exists a collection of $(\epsilon, 3\delta)$-DP kernels $\{\tilde{K}^j\}_{j=1}^m$ such that for some fixed constant $C > 0$,*

$$\sup_{x \in \mathcal{X}^n} \frac{d\tilde{K}^j(\cdot|x)}{dP_0\tilde{K}^j(\cdot|X^{(j)})}(y) < C, \; P_0\tilde{K}^j(\cdot|X^{(j)})\text{-almost surely}, \tag{53}$$

whilst

$$\|P_f(K^j(\cdot|X^{(j)}) - \tilde{K}^j(\cdot, X^{(j)}))\|_{\mathrm{TV}} \leqslant \frac{\alpha}{2m}.$$

Proof For any $x \in \mathcal{X}^n$ and set $A \in \mathscr{Y}^{(j)}$, we have that

$$K^j(A|x) = \int_A \frac{dK^j(\cdot|x)}{dP_0K^j(\cdot|X^{(j)})}(y) dP_0K^j(y|X^{(j)}) \leqslant 1,$$

where we note that the density of the integrand exists. So, by Markov's inequality, there exists a set $A_x^M \in \mathscr{Y}^{(j)}$ such that

$$\frac{dK^j(\cdot|x)}{dP_0K^j(\cdot|X^{(j)})}(y) \leqslant M \;\text{ on } A_x^M,$$

whilst

$$K^j\left((A_x^M)^c|x\right) \leqslant 1/M. \tag{54}$$

Define for all $x \in \mathcal{X}$,

$$\tilde{K}^j(B|x) := K^j\left(B \cap A_x^M|x\right) + K^j\left((A_x^M)^c|x\right) \frac{K^j(B \cap A_x^M|x)}{K^j(A_x^M|x)}. \tag{55}$$

Then, $\tilde{K}^j$ is $(\epsilon, 3\delta)$-DP whenever $M > 4\delta^{-1}$; for any $x, x' \in \mathcal{X}^n$ that are Hamming distance 1-apart and $B \in \mathscr{Y}^{(j)}$,

$$\begin{aligned}
\tilde{K}^j(B|x) &\leqslant K^j\left(B|x\right) + K^j\left((A_x^M)^c|x\right) \frac{K^j\left(B \cap A_x^M|x\right)}{K^j\left(A_x^M|x\right)} \\
&= K^j\left(B \cap A_{x'}^M|x\right) + K^j\left(B \cap (A_{x'}^M)^c|x\right) + K^j\left((A_x^M)^c|x\right) \frac{K^j\left(B \cap A_x^M|x\right)}{K^j\left(A_x^M|x\right)} \\
&\leqslant e^\epsilon K^j\left(B \cap A_{x'}^M|x'\right) + e^\epsilon K^j\left(B \cap (A_{x'}^M)^c|x'\right) + 2\delta + \frac{1}{M} \\
&\leqslant e^\epsilon \tilde{K}^j\left(B|x'\right) + (1 + e^\epsilon)M^{-1} + 2\delta,
\end{aligned}$$

where the second to last inequality follows by (54) and the last inequality follows by simply adding the nonnegative second term in (55). Furthermore, its Radon-Nikodym derivative satisfies

$$\frac{d\tilde{K}^j(\cdot|x)}{dP_0\tilde{K}^j(\cdot|X^{(j)})}(y) \leqslant 2\mathbb{1}_{A_x^M}\frac{dK^j(\cdot|x)}{dP_0K^j(\cdot|X^{(j)})}(y) \leqslant 2M$$

$P_0\tilde{K}^j(\cdot|X^{(j)})$-almost surely. Moreover, it holds for any $f \in \ell_2(\mathbb{N})$ that

$$\begin{aligned}\|P_f^n(K^j(\cdot|X^{(j)}) - \tilde{K}^j(\cdot, X^{(j)}))\|_{\mathrm{TV}} &\leqslant \int \|K^j(\cdot|x) - \tilde{K}^j(\cdot|x)\|_{\mathrm{TV}} dP_f^n(x) \\ &\leqslant 2\int |K^j\left((A_x^M)^c|x\right)| dP_f^n(x) \leqslant \frac{2}{M}.\end{aligned}$$

Since a choice $M > \delta^{-1} \vee 2m/\alpha$ yields the bound uniformly in $x \in \mathcal{X}^n$, the result follows. ■

By standard arguments (see also Lemma 17), the first term in (48) can be lower bounded as follows;

$$\begin{aligned}P_0^m K_u T &= P_0^m \tilde{K}_u T + P_0^m K_u T - P_0^m \tilde{K}_u T \\ &\geqslant P_0^m \tilde{K}_u T - \sum_{j=1}^m \|P_0\tilde{K}^j(\cdot|X^{(j)}) - P_0K^j(\cdot|X^{(j)})\|_{\mathrm{TV}}.\end{aligned}$$

The same argument on the second term yields that, we obtain that (48) is lower bounded by

$$\int \left(P_0^m \tilde{K}_u T + \int P_f^m \tilde{K}_u (1-T) d\pi \circ \Psi_L^{-1}(f)\right) d\mathbb{P}^U(u) - 3\alpha/8, \tag{56}$$

for an $(\epsilon, 3\delta)$-DP distributed testing protocol $\tilde{T} = (T, \{\tilde{K}^j\}_{j=1}^m, (\mathcal{U}, \mathscr{U}, \mathbb{P}^U))$.

Another issue suffered by (ϵ, δ)-DP Markov kernels with $\delta > 0$, is that one has very poor control over the higher moments of the local likelihoods

$$\mathcal{L}_{\pi,u}^j(y) := \frac{dP_\pi\tilde{K}^j(\cdot|X^{(j)}, u)}{dP_0\tilde{K}^j(\cdot|X^{(j)}, u)}(y),$$

where $P_\pi(A) := \int P_f(A) d\pi \circ \Psi_L^{-1}(f)$, which are required to sufficiently bound the corresponding quantity B_u defined in (70) in Step 5. Using similar ideas as in the proof of Lemma 14, we can construct approximating kernels $\{\breve{K}^j\}_{j=1}^m$ such that the likelihoods $\mathcal{L}_{\pi,u}^j$ are bounded. This is the content of the following lemma.

Lemma 15. *Let $\alpha \in (0,1)$, $\pi = N(0, c_\alpha^{-1/2} d_L^{-1/2} \rho^2 \bar{\Gamma})$ for an arbitrary positive semidefinite $\bar{\Gamma}$ and let $\{K^j\}_{j=1}^m$ correspond to a (ϵ, δ)-DP distributed protocol T (i.e. K^j satisfies (**??**)). Furthermore, assume that $\epsilon \leqslant 1/\sqrt{n}$ and define for $j = 1, \ldots, m$ the events*

$$A_{j,u} := \left\{y : |\mathcal{L}_{\pi,u}^j(y) - 1| \leqslant 4m^{1/2}\alpha^{-1}\right\}$$

and

$$\tilde{K}^j(B|x,u) := K^j(B \cap A_{j,u}|x,u) + K^j(A_{j,u}^c|x,u)\frac{P_0K^j(B \cap A_{j,u}|X^{(j)}, u)}{P_0K^j(A_{j,u}|X^{(j)}, u)}.$$

Suppose in addition that $\delta \leqslant c_\alpha/m$. Then,

(a) The collection $\{\tilde{K}^j\}_{j=1}^m$ *are* $(\epsilon, 2\delta)$*-DP Markov kernels.*

(b) It holds $P_0\tilde{K}^j(\cdot|X^{(j)}, u)$*-a.s. that* $\tilde{\mathcal{L}}^j_{\pi,u}(y) := \int \frac{d\tilde{K}^j(y|x,u)}{dP_0\tilde{K}^j(y|X^{(j)},u)} dP^n_\pi(x)$ *satisfies*

$$|\tilde{\mathcal{L}}^j_{\pi,u}(y) - 1| \leqslant \frac{5m^{1/2}}{\alpha}. \tag{57}$$

(c) If ρ *satisfies* (34) *or* (36) *with* $c_\alpha > 0$ *small enough, it holds that*

$$P_0^m KT + \int P_f^m K(1-T)d\pi \circ \Psi_L^{-1}(f) \geqslant P_0^m \tilde{K}T + \int P_f^m \tilde{K}(1-T)d\pi \circ \Psi_L^{-1}(f) - \alpha,$$

where $\tilde{K}$ *is the product kernel corresponding to* $\{\tilde{K}^j\}_{j=1}^m$*.*

(d) If K^j *satisfies* (53)*, then* $\tilde{K}^j$ *satisfies the same bound for some constant* $C > 0$*.*

We prove this lemma at the end of this section. The lemma above, in combination with the same argument as before (e.g. Lemma 17) allows us to replace the $(\epsilon, 3\delta)$-DP Markov kernel $\tilde{K}^j$ of (56) with an $(\epsilon, 6\delta)$-DP Markov kernel $\breve{K}^j$, whose transcripts have bounded likelihoods in the sense of (57). Furthermore, by part *(c)* of the lemma, the Bayes risk of (56) is further lower bounded by

$$\int \left(P_0^m \breve{K}_u T + \int P_f^m \breve{K}_u (1-T) d\pi \circ \Psi_L^{-1}(f) \right) d\mathbb{P}^U(u) - \alpha/2, \tag{58}$$

where the $\breve{K}^j$'s satisfy (53).

A.2.1 Proof of Lemma 15

Proof The first statement follows by Lemma 16 below. For the second statement, we first note that by Lemma 25 proven in Section A.5, it holds that

$$\begin{aligned}
\frac{c_\alpha^{1/4}}{m^{1/2}} + \delta + \frac{c_\alpha}{m^{3/2}} &\geqslant (P_\pi - P_0)K^j \left(\{|\mathcal{L}^j_{\pi,u} - 1| \geqslant 4m^{1/2}/\alpha\} | X^{(j)}, u \right) \\
&= P_0 K^j \left((\mathcal{L}^j_{\pi,u} - 1) \mathbb{1}\{|\mathcal{L}^j_{\pi,u} - 1| \geqslant 4m^{1/2}/\alpha\} | X^{(j)}, u \right) \\
&\geqslant 4\frac{m^{1/2}}{\alpha} P_0 K^j \left(|\mathcal{L}^j_{\pi,u} - 1| \geqslant 4m^{1/2}/\alpha | X^{(j)}, u \right),
\end{aligned}$$

where the second inequality follows from the fact that $m^{1/2}/\alpha \geqslant 1$ and $\mathcal{L}^j_{\pi,u} \geqslant 0$. Using that $\delta \leqslant c_\alpha/m$, we obtain that

$$P_0 K^j(A^c_{j,u}|X^{(j)}, u) \leqslant (4m)^{-1}\alpha(c_\alpha^{1/4} + c_\alpha(1 + m^{-1})) := \eta_\alpha. \tag{59}$$

Since $K^j(B|x,u) \leqslant \tilde{K}^j(B|x,u)$ for all measurable $B \subset A_{j,u}$ and $P_0\tilde{K}^j(\cdot|X^{(j)}, u)$ has no support outside of $A_{j,u}$, it holds that

$$\frac{dK^j(\cdot|x,u)}{dP_0\tilde{K}^j(\cdot|X^{(j)},u)}(y) \leqslant \frac{dK^j(\cdot|x,u)}{dP_0K^j(\cdot|X^{(j)},u)}(y),$$

for all $y \in A_{j,u}$ (and hence $P_0\tilde{K}^j(\cdot|X^{(j)},u)$-a.s.). Similarly, we have for P_π-a.s. all x's that

$$\frac{K^j(A_{j,u}^c|x,u)}{P_0K^j(A_{j,u}|X^{(j)},u)}\frac{dP_0K^j(\cdot\cap A_{j,u}|X^{(j)},u)}{dP_0\tilde{K}^j(\cdot|X^{(j)},u)}(y) \leqslant \frac{K^j(A_{j,u}^c|x,u)}{P_0K^j(A_{j,u}|X^{(j)},u)} \leqslant \frac{1}{1-\eta_\alpha},$$

using that $K^j \leqslant 1$ and $P_0K^j(A_{j,u}|X^{(j)},u) \geqslant 1-\eta_\alpha$. By standard arguments and the above two statements, it follows that

$$\begin{aligned}\int \frac{d\tilde{K}^j(y|x,u)}{dP_0\tilde{K}^j(y|X^{(j)},u)}dP_\pi^n(x) &\leqslant \mathbb{1}_{A_{j,u}}(y)\int \frac{dK^j(y|x,u)}{dP_0K^j(y|X^{(j)},u)}dP_\pi^n(x) + \frac{1}{1-\eta_\alpha}\\ &= \mathbb{1}_{A_{j,u}}(y)\mathcal{L}_{\pi,u}^j(y) + (1-\eta_\alpha)^{-1}.\end{aligned}$$

Applying the definition of the event $A_{j,u}$ and using that $\alpha \leqslant 1$, we obtain that for $c_\alpha > 0$ small enough $\tilde{\mathcal{L}}_{\pi,u}^j - 1 \leqslant \alpha^{-1}4m^{1/2} + (1-\eta_\alpha)^{-1} - 1 \leqslant 5m^{1/2}\alpha^{-1}$. Using that $\tilde{\mathcal{L}}_{\pi,u}^j - 1 \geqslant -1$, we obtain (57), proving statement *(b)*.

For the third statement, we will aim to apply Lemma 17. By the construction of $\tilde{K}^j$ and the triangle inequality,

$$\|P_0K^j(\cdot|X^{(j)},u) - P_0\tilde{K}^j(\cdot|X^{(j)},u)\|_{\mathrm{TV}} \leqslant 2\left\|P_0K^j(\cdot\cap A_{j,u}^c|X^{(j)},u)\right\|_{\mathrm{TV}}.$$

The latter is bounded by $\alpha/(2m)$ (see (59)). Similarly,

$$\|P_\pi K^j(\cdot|X^{(j)},u) - P_\pi\tilde{K}^j(\cdot|X^{(j)})\|_{\mathrm{TV}} \leqslant 2P_\pi K^j(A_{j,u}^c|X^{(j)},u).$$

By Lemma 25,

$$P_\pi K^j(A_{j,u}^c|X^{(j)},u) \leqslant \left(1 + c_\alpha^{1/4}m^{-1/2}\right)P_0K^j(A_j^c|X^{(j)},u) + \delta + m^{-3/2}c_\alpha.$$

Again using (59) and the fact that $\delta \leqslant c_\alpha/m$ yield that the latter is also bounded by $\alpha/4m$ for $c_\alpha > 0$ small enough. The condition and small enough choice of $c_\alpha > 0$ yields that the conditions of Lemma 17 and the conclusion of (c) follows. Finally, if K^j satisfies (53), the last statement follows directly by the construction of $\tilde{K}^j$. ■

We finish the section by providing the two technical lemmas mentioned in the earlier proofs above. We omit the presence of the shared randomness U in the statement of the lemmas, as it is of no consequence to the arguments below.

Lemma 16. *Let K be a Markov kernel from $(\mathcal{X},\mathscr{X})^n$ to $(\mathcal{Y},\mathscr{Y})$ satisfying an (ϵ,δ)-DP constraint (i.e. (**??**)) and define for a $A \in \mathscr{Y}$ and a probability measure μ on $\mathscr{Y}$*

$$\tilde{K}(B|x) := K(B\cap A|x) + K(A^c|x)\mu(B), \text{ for } x \in \mathcal{X},\ B \in \mathscr{Y}.$$

Then, $\tilde{K}$ is a Markov kernel $(\mathcal{X},\mathscr{X})$ to $(\mathcal{Y},\mathscr{Y})$ satisfying an $(\epsilon,2\delta)$-DP constraint.

Proof First, $\tilde{K}$ can be seen to be a Markov kernel, as the necessary measurability assumptions hold by construction and $\tilde{K}(\mathcal{Y}|x) = K(\mathcal{Y}\cap A|x) + K(A^c|x) = 1$, where it is used

that μ is a probability measure. Furthermore, for arbitrary B and $x, x' \in \mathcal{X}^n$ such that $d_H(x, x') \leqslant 1$, it holds that

$$\tilde{K}(B|x) \leqslant e^{\epsilon} K(B \cap A|x') + \delta + e^{\epsilon} K(A^c|x')\mu(B) + \mu(B)\delta \leqslant\leqslant e^{\epsilon}\tilde{K}(B|x') + 2\delta.$$

■

Lemma 17. *Let $\alpha \in (0,1)$ be given. Let $(T, \{K^j\}_{j=1}^m)$ be a distributed testing protocol and suppose that there exist kernels $\{\tilde{K}^j\}_{j=1}^m$ such that for $j = 1, \ldots, m$,*

$$\|P_I(K^j(\cdot|X^{(j)}, u) - \tilde{K}^j(\cdot|X^{(j)}, u))\|_{\mathrm{TV}} \leqslant \frac{\alpha}{2m} \ \mathbb{P}^U\text{-a.s},$$

where $\mu \in \{\delta_0, \pi\}$. Then,

$$\mathbb{P}^U P_0^m K(T(Y)|X, U) + \mathbb{P}^U \int P_f^m K(1 - T(Y)|X, U))d\pi(f) \geqslant$$
$$\mathbb{P}^U P_0^m \tilde{K}(T(Y)|X, U) + \mathbb{P}^U \int P_f^m \tilde{K}(1 - T(Y)|X, U))d\pi(f) - \alpha,$$

for the same collection of distributions.

Proof We omit the dependence of u in the proof, as it is of no consequence to the arguments below. We have that

$$P_0^m K(T(Y) = 1|X) + \int P_f^m K(T(Y) = 0|X)d\pi(f) \geqslant$$
$$P_0^m \tilde{K}(T(Y) = 1|X) + \int P_f^m \tilde{K}(T(Y) = 0|X)d\pi(f) - \|P_0^m(K(\cdot|X) - \tilde{K}(\cdot|X))\|_{\mathrm{TV}}$$
$$- \|P_\pi^m(K(\cdot|X) - \tilde{K}(\cdot|X))\|_{\mathrm{TV}}.$$

Standard arguments (see e.g. Lemma 13 in Vuursteen (2024)) yield

$$\|P_\pi^m(K(\cdot|X) - \tilde{K}(\cdot|X))\|_{\mathrm{TV}} \leqslant \sum_{j=1}^m \|P_\pi(K^j(\cdot|X^{(j)}) - \tilde{K}^j(\cdot|X^{(j)}))\|_{\mathrm{TV}}.$$

By applying the same lemma to $\|P_0^m(K(\cdot|X) - \tilde{K}(\cdot|X))\|_{\mathrm{TV}}$, combined with what is assumed in this lemma, we obtain the result. ■

A.3 Step 3: Bounding the Chi-square divergence using the Brascamp-Lieb inequality

We proceed with lower bounding (58), where in a slight abuse of notation, we shall denote $\check{K}^j$ by K^j for the remainder of the section.

To lower bound the Bayes risk, in light of the Neyman-Pearson lemma, it should suffice to show that $\mathcal{L}_\pi^{Y|U=u}(Y)$ is close to 1 with high probability. This is made precise below by showing that the right-hand side of (58) is further bounded from below by

$$1 - \left(\sqrt{(1/2) \int \mathbb{E}_0^{Y|U=u} \left(\mathcal{L}_\pi^{Y|U=u}(Y) - 1 \right)^2 d\mathbb{P}^U(u)} - \alpha/2 \right). \tag{60}$$

To see this, note that any $T : \mathcal{Y}^m \to \{0,1\}$, we can write $A_T = T^{-1}(\{0\})$ and note that

$$P_0^m \tilde{K}_u T(Y) + P_\pi^m \tilde{K}_u (1 - T(Y)) = 1 - \left(P_0^m \tilde{K}_u \left(Y \in A_T \right) - P_\pi^m \tilde{K}_u \left(Y \in A_T \right) \right).$$

We obtain that

$$\int \left(P_0^m \tilde{K}_u T + \int P_f^m \tilde{K}_u (1 - T) d\pi \circ \Psi^{-1}(f) \right) d\mathbb{P}^U(u) \geqslant 1 - \sup_A \left| \int P_0^m \tilde{K}_u(A) - P_\pi^m \tilde{K}_u(A) d\mathbb{P}^U(u) \right|.$$

We find by Jensen's inequality that $\| \mathbb{P}_0^Y - \mathbb{P}_\pi^Y \|_{\mathrm{TV}} \leqslant \int \| \mathbb{P}_0^{Y|U=u} - \mathbb{P}_\pi^{Y|U=u} \|_{\mathrm{TV}} d\mathbb{P}^U(u)$. Combining the above with Pinsker's second inequality a standard bound for the Kullback-Leibler divergence (see Lemma 2.7 of Tsybakov (2009)), we obtain (60).

This brings us to a crucial part of the proof; the application of Lemma 10.1 in Szabó et al. (2023), which we restate as Lemma 18 below. We first introduce some notation.

We note that for $f \in \mathbb{R}^{d_L}$, it holds that

$$\frac{dP_{\Psi f}}{dP_0}(X_i^{(j)}) \overset{d}{=} \frac{dN\left(f, \sigma^2 I_{d_L}\right)}{dN\left(0, \sigma^2 I_{d_L}\right)} =: \mathscr{L}_f^{ji}(X_i^{(j)}),$$

where the equality in distribution is true under $\mathbb{P}_0^X$.

Denote the "local" and "global" likelihoods of the data as $\mathscr{L}_f^j(X^{(j)}) = \prod_{i=1}^n \mathscr{L}_f^{ji}(X_i^{(j)})$, $\mathscr{L}_f(X) := \prod_{j=1}^m \mathscr{L}_f^j(X^{(j)})$, and the mixture likelihoods as $\mathscr{L}_\pi^j(X) = \int \mathscr{L}_f(X^{(j)}) d\pi(f)$ and $\mathscr{L}_\pi(X) = \int \mathscr{L}_f(X) d\pi(f)$.

In view of the Markov chain structure, the probability measure $d\mathbb{P}_\pi(x,u,y)$ disintegrates as $d\mathbb{P}^{Y|(X,U)=(x,u)} d\mathbb{P}_f^X(x) d\mathbb{P}^U(u) d\pi(f)$. Using this, $\mathbb{E}_0^{Y|U=u} \left(L_\pi^{Y|U=u}(Y) \right)^2$ can be seen to equal

$$\mathbb{E}_0^{Y|U=u} \mathbb{E}_0 \left[\mathscr{L}_\pi(X) \middle| Y, U = u \right]^2 = \int \left(\int \mathscr{L}_\pi(x) \frac{dK(\cdot|x,u)}{d\mathbb{P}_0^{Y|U=u}}(y) d\mathbb{P}_0^X(x) \right)^2 d\mathbb{P}_0^{Y|U=u}(y), \tag{61}$$

where it is used that $K(\cdot|x,u) \ll \mathbb{P}_0^{Y|U=u}(\cdot)$, $\mathbb{P}_f^{(X,U)}$-almost surely. Using Fubini's theorem ("decoupling" in X), we can write the above display as

$$\int \mathscr{L}_\pi(x_1) \mathscr{L}_\pi(x_2) q_u(x_1, x_2) d(\mathbb{P}_0^X \times \mathbb{P}_0^X)(x_1, x_2), \tag{62}$$

where

$$q_u(x_1, x_2) := \int \frac{dK(\cdot|x_1,u)}{d\mathbb{P}_0^{Y|U=u}}(y) \frac{dK(\cdot|x_2,u)}{d\mathbb{P}_0^{Y|U=u}}(y) d\mathbb{P}_0^{Y|U=u}(y). \tag{63}$$

Since $K(\cdot|x,u)$ and $\mathbb{P}_0^{Y|U=u}$ are product measures on $\mathcal{Y} = \mathcal{Y}^m$, we can write $q_u(x_1,x_2) = \Pi_{j=1}^m q_u^j(x_1,x_2)$ where

$$q_u^j(x_1,x_2) = \int \frac{K^j(y^j|x_1^j,u)K^j(y^j|x_2^j,u)}{\mathbb{P}_0^{Y^j|U=u}(y^j)} d\mathbb{P}_0^{Y^{(j)}|U=u}(y). \tag{64}$$

The map $(x_1,x_2) \mapsto q_u(x_1,x_2)$ can be seen as capturing the dependence between the original data X and a random variable X' with conditional distribution

$$X'|X = x \sim \int d\mathbb{P}_0^{X|(Y,U)=(y,u)} d\mathbb{P}^{Y|(X,U)=(x,u)}, \tag{65}$$

which is sometimes referred to as the "forward-backward channel", stemming from the fact that $X \to Y \to X'$ forms a Markov chain. An easy computation using the law of total expectation shows that the covariance of $q_u(x_1,x_2)d(P_0 \times P_0)(x_1,x_2)$,

$$\int \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \begin{pmatrix} x_1^\top & x_2^\top \end{pmatrix} q_u(x_1,x_2) d(P_0 \times P_0)(x_1,x_2) \in \mathbb{R}^{2mnd_L \times 2mnd_L}, \tag{66}$$

is equal to $\Sigma_u := \mathrm{Diag}\left(\Sigma_u^{11}, \ldots, \Sigma_u^{1n}, \ldots, \Sigma_u^{m1}, \ldots, \Sigma_u^{mn}\right) \in \mathbb{R}^{2mnd_L \times 2mnd_L}$ for

$$\Sigma^{ji} := \sigma^2 \begin{pmatrix} I_{d_L} & \Xi_u^{ji} \\ \Xi_u^{ji} & I_{d_L} \end{pmatrix},$$

with $\Xi_u^{ji} := \mathbb{E}_0^{Y^{(j)}|U=u} \mathbb{E}_0 \left[\sigma^{-1} X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right] \mathbb{E}_0 \left[\sigma^{-1} X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right]^\top$. Define also

$$\Xi_u^j := \mathbb{E}_0^{Y^{(j)}|U=u} \mathbb{E}_0 \left[\sigma^{-1} \sum_{i=1}^n X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right] \mathbb{E}_0 \left[\sigma^{-1} \sum_{i=1}^n X_{L;i}^{(j)} \middle| Y^{(j)}, U = u \right]^\top. \tag{67}$$

We are now ready to state the lemma that forms the crux of our distributed testing lower bound proof.

Lemma 18. *Suppose that $(x_1,x_2) \mapsto q_u(x_1,x_2)$ is bounded and that π is a centered Gaussian distribution on $\mathbb{R}^d$. Then,*

$$\frac{\int \mathscr{L}_\pi(x_1)\mathscr{L}_\pi(x_2) q_u(x_1,x_2) d(\mathbb{P}_0^X \times \mathbb{P}_0^X)(x_1,x_2)}{\prod_{j=1}^m \int \mathscr{L}_\pi^j(x_1^j)\mathscr{L}_\pi^j(x_2^j) q_u^j(x_1^j,x_2^j) d(\mathbb{P}_0^{X^{(j)}} \times \mathbb{P}_0^{X^{(j)}})(x_1^j,x_2^j)} \tag{68}$$

is bounded above by

$$\frac{\int \mathscr{L}_\pi(x_1)\mathscr{L}_\pi(x_2) dN(0,\Sigma)(x_1,x_2)}{\prod_{j=1}^m \int \mathscr{L}_\pi^j(x_1^j)\mathscr{L}_\pi^j(x_2^j) dN(0,\Sigma^j)(x_1^j,x_2^j)}.$$

The lemma has the following interpretation: the ratio of the second moment of the Bayes factor of the "global Bayesian hypothesis test" that of the product of second moments of the "local Bayes factors", is maximized over the class of forward-backward channel with

covariance Σ when the forward-backward channel is Gaussian. For a proof, we refer to Lemma 10.1 in Szabó et al. (2023), which uses that the prior π is Gaussian, exploiting the conjugacy between the prior and the model which enables the use of techniques from Lieb (1990).

The main consequence of the above lemma is analytic expressions, which can be used to upper bound the chi-square divergence in (60) which is the content of Lemma 19 below.

Lemma 19. *Define*

$$\mathrm{A}_u^\pi := \int e^{f^\top \sum_{j=1}^m \Xi_u^j g} d(\pi \times \pi)(f,g) \tag{69}$$

and

$$\mathrm{B}_u^\pi := \prod_{j=1}^m \mathbb{E}_0^{Y^{(j)}|U=u} \mathbb{E}_0 \left[\mathscr{L}_\pi \left(X^{(j)} \right) \middle| Y^{(j)}, U = u \right]^2. \tag{70}$$

If $(x_1, x_2) \mapsto q_u(x_1, x_2)$ is bounded and if π is a centered Gaussian distribution on $\mathbb{R}^{d_L}$, it holds that

$$\mathbb{E}_0^{Y|U=u} \left(L_\pi^{Y|U=u}(Y) \right)^2 \leqslant \mathrm{A}_u^\pi \cdot \mathrm{B}_u^\pi.$$

The above lemma describes how the variance of the Bayes factor given U is bounded by two factors. One factor depends on the Fisher information of the transcript's likelihood at $f = 0$ given $U = u$; $\Xi_u := \sum_{j=1}^m \Xi_u^j$. In this sense, A_u^π captures how well the transcript allows for "estimation" of f. The second factor can be seen as the m-fold product of the local Bayes factors, capturing essentially the power of combining the locally most powerful test statistics; the likelihood ratios.

Proof [Proof of Lemma 19] We start by noting that B_u^π is equal to the denominator of (68). By Lemma 18,

$$\mathbb{E}_0^{Y|U=u} \left(L_\pi^{Y|U=u}(Y) \right)^2 \leqslant \frac{\int \mathscr{L}_\pi(x_1)\mathscr{L}_\pi(x_2) dN(0,\Sigma)(x_1,x_2)}{\prod_{j=1}^m \int \mathscr{L}_\pi^j(x_1^j)\mathscr{L}_\pi^j(x_2^j) dN(0,\Sigma^j)(x_1^j,x_2^j)} \cdot \mathrm{B}_u^\pi.$$

By the block diagonal matrix structure of Σ, the denominator in the first factor of the right-hand side equals

$$\begin{aligned} \prod_{j=1}^m \int e^{\frac{1}{2}\left(\|\sqrt{\Sigma^j}(f,g)\|_2^2 - \|(f,g)\|_2^2\right)} d(\pi \times \pi)\,(f,g) &= \prod_{j=1}^m \int e^{f^\top \Xi_u^j g} d(\pi \times \pi)(f,g) \\ &\geqslant \prod_{j=1}^m e^{\int f^\top \Xi_u^j g \, d(\pi\times\pi)(f,g)} = 1. \end{aligned}$$

Through the expression for the moment generating function of the Gaussian, the numerator of A_u^π is equal to $\int e^{f^\top \sum_{j=1}^m \Xi_u^j g} d(\pi \times \pi)(f,g)$. ∎

A.4 Step 4: Adversarially choosing the prior based on shared or local randomness

Suppose that for some constant $c > 0$,

$$\varrho^2 \|\sqrt{\bar{\Gamma}}^\top \Xi_u \sqrt{\bar{\Gamma}}\| \leqslant c. \tag{71}$$

If $\bar{\Gamma} \in \mathbb{R}^{d_L \times d_L}$ is symmetric, idempotent with rank proportional to d_L and $\pi = N(0, \varrho^2 \bar{\Gamma})$, standard results for the Gaussian chaos, e.g. Lemma 6.2.2 in Vershynin (2018) combined with (71) and the fact that $\|\sqrt{\bar{\Gamma}}\| \leqslant 1$, yield that

$$\mathrm{A}_u^\pi \leqslant \exp\left(C\sigma^{-2}\varrho^4 \mathrm{Tr}\left((\sqrt{\bar{\Gamma}}^\top \Xi_u \sqrt{\bar{\Gamma}})^2\right)\right),$$

for a constant $C > 0$ depending only on c. As a final step of the testing risk lower bound technique, we use essentially a geometric argument to sharpen this bound in case the distributed protocol does not enjoy shared randomness. The $d_L \times d_L$ matrix $\Xi_u := \sum_{j=1}^m \Xi_u^j$ geometrically captures how well Y allows to "reconstruct" the compressed sample X. When U is degenerate, Ξ_u is "known" to the prior, and $\bar{\Gamma}$ can be chosen to exploit "direction" in which Ξ_u contains the least information. The lemma below makes this notion precise. For a proof, we refer to Szabó et al. (2023), Section 9.

Lemma 20. *Let $\alpha \in (0,1)$ and suppose that the map $(x_1, x_2) \mapsto q_u(x_1, x_2)$ defined in (63) is bounded for all distributed testing protocols in $\mathscr{T}$. Let $\pi = N(0, \varrho^2\bar{\Gamma})$, with $\varrho := \frac{\rho}{c_\alpha^{1/4} d_L^{1/2}}$ and $\bar{\Gamma} \in \mathbb{R}^{d_L \times d_L}$ is symmetric, idempotent with rank proportional to d_L. Assume that ρ is such that $\varrho^2\|\Xi_u\| \leqslant c$ $\mathbb{P}^U$-a.s. for some constant $c > 0$. Then, (44) holds for shared and local randomness protocols respectively.*

A.5 Step 5: Capturing the cost of privacy in trace of Ξ_u and the local Bayes factors

The cost of privacy is captured through bounds on A_u^π and B_u^π. These bounds specifically use the fact that the Markov kernels that underlie these quantities are $(\epsilon, 6\delta)$-differentially private.

We start with the bound on A_u^π, for which we proceed by a data processing argument for the matrix Ξ_u under the $(\epsilon, 6\delta)$-DP constraint. This comes in the guise of Lemma 21 below. Its proof is deferred to the end of the section.

Lemma 21. *Let $0 < \epsilon \leqslant 1$ and let $Y^{(j)}$ be a transcript generated by an (ϵ, δ)-DP constraint distributed protocol, with $0 < \epsilon \leqslant 1$ and $0 \leqslant \delta \leqslant \left(\left(nd_L{}^{-1} \wedge n^{1/2} d_L{}^{-1/2}\right)\epsilon^2\right)^{1+\omega}$ for some $\omega > 0$. The matrix Ξ_u^j as defined in (67) satisfies $Tr\left(\Xi_u^j\right) \leqslant (Cn^2\epsilon^2) \wedge (nd_L)$ for a fixed constant $C > 0$. Furthermore, it holds that $\Xi_u^j \leqslant nI_{d_L}$.*

The lemma implies in particular that $\|\Xi_u^j\| \leqslant (Cn^2\epsilon^2) \wedge n$, as Ξ_u^j is symmetric and positive definite. Combining this with (51) and the triangle inequality, we obtain

$$\varrho^2\|\Xi_u\| \leqslant \varrho^2 \sum_{j=1}^m \|\Xi_u^j\| \leqslant \frac{m\left((Cn^2\epsilon^2) \wedge n\right)\rho^2}{\sigma^2\sqrt{c_\alpha} d_L} \leqslant C\sqrt{c_\alpha}. \tag{72}$$

Similarly, (50) yields

$$\frac{2\rho^2}{\sqrt{c_\alpha} d_L^2}\mathrm{Tr}\left(\Xi_u^j\right) \leqslant \frac{m\left((Cn^2\epsilon^2) \wedge (nd_L)\right)\rho^2}{\sigma^2\sqrt{c_\alpha} d_L^2} \leqslant C\sqrt{c_\alpha}/\sqrt{d_L}. \tag{73}$$

The last two displays together finish the verification of the conditions of Lemma 20. The above data processing inequalities for Ξ_u^j and bounds on ρ^2 also yield a bound on A_u^π as defined in Lemma 20. In case of shared randomness protocols, using (**??**), (72), Lemma 21 and (51), we obtain $\mathrm{A}_u^\pi \leqslant \exp\left(C^2 c_\alpha\right)$.

In case of local randomness protocols, combining (**??**) with (73) and (50) yields the above bound on A_u^π.

Next, we turn to B_u^π. The first bound, given in Lemma 22, does not use privacy at all. This bound is only tight whenever $\epsilon \gtrsim 1/\sqrt{n}$ (as a bound on B_u^π), in which case it corresponds to the regime where majority voting bares no privacy cost. The proof can be found in Section 9 Step 2 of Szabó et al. (2023).

Lemma 22. *Consider* B_u^π *as in* (70) *with* $\pi = N(0, \rho^2 c_\alpha^{-1/2} d_L^{-1} \bar{\Gamma})$ *with* $\bar{\Gamma}$ *idempotent. It holds that*

$$\mathrm{B}_u^\pi \leqslant \exp\left(C \frac{mn^2\rho^4}{c_\alpha \sigma^2 d_L}\right).$$

Whenever $\epsilon \leqslant n^{-1/2}$, a much more involved data processing argument is needed than the one above. In the argument that follows, we will use the bound on $\mathcal{L}_{\pi,u}^j$ obtained through Lemma 15 in Step 2. The data processing argument leads to the bound of Lemma 23 below. Its proof is based on coupling arguments, where the two different couplings constructed result in the different rates observed in the condition of the theorem.

Lemma 23. *Let* $\pi = N(0, d_L^{-1}\rho^2\bar{\Gamma})$*, with* $\bar{\Gamma} \in \mathbb{R}^{d_L \times d_L}$ *a symmetric idempotent matrix,*

$$\rho^2 \leqslant \sigma^2 c_\alpha (d_L^{1/2}/(\sqrt{m} n^{\frac{3}{2}} \epsilon) \vee 1/(mn^2\epsilon^2))$$

and $\{K^j\}_{j=1}^m$ *correspond to a* (ϵ, δ)*-DP distributed protocol with transcripts* $Y^{(j)}$ *such that* $0 < \epsilon \leqslant 1$*,* $\delta \lesssim c_\alpha(m^{-1} \wedge \epsilon)$ *and* $|\mathcal{L}_{\pi,u}^j(y) - 1| \leqslant \frac{5m^{1/2}}{\alpha} \quad P_0 K^j(\cdot|X^{(j)}, u)$*-a.s..*

Then, there exists a universal constant $C > 0$ *such that* $\mathrm{B}_u^\pi \leqslant e^{C\sqrt{c_\alpha}}$*.*

Combining the lemma above with the bound $\mathrm{B}_u^\pi \leqslant \exp(C\frac{mn^2\rho^4}{\sigma^2 d_L})$ (which follows from Lemma 22), we obtain that $\mathrm{B}_u^\pi \leqslant e^{C\sqrt{c_\alpha}}$ whenever ρ satisfies (50) or (51). Combining this with the bounds on A_u^π derived earlier and considering (60) lower bounds the testing risk, we obtain that $\mathcal{R}(H_\rho^{s,R}, T) > 1 - \alpha$ for $c_\alpha > 0$ small enough, from which the result of Theorem 12 follows. To complete the proof, we provide the proofs of Lemmas 21 and 23 in the following subsection.

A.5.1 Proof of Lemma 23

Before providing the proof of Lemma 23, we first develop additional tools.

First off is the following general coupling lemma that will be used in the proof of Lemma 23. The lemma is in essence Lemma 6.1 in Karwa and Vadhan (2017), but its proof might be easier to verify and is provided for completeness in Section C.3.

Lemma 24. *Consider random variables* $X_1, \ldots, X_n \overset{i.i.d.}{\sim} P_1$ *and* $\tilde{X}_1, \ldots, \tilde{X}_n \overset{i.i.d.}{\sim} P_2$ *defined on the same space. Write* $X = (X_1, \ldots, X_n)$*,* $\tilde{X} = (\tilde{X}_1, \ldots, \tilde{X}_n)$ *and let* K *be a Markov kernel between the sample space of* X *(equivalently* $\tilde{X}$*) and an arbitrary target space, satisfying an* (ϵ, δ)*-DP constraint with* $\epsilon \leqslant 1$*.*

Suppose that there exists a coupling $\mathbb{P}$ *of* $(\tilde{X}, X)$ *such that* $\mathbb{P}^{\tilde{X}} = P_1^n$, $\mathbb{P}^X = P_2^n$ *and*

$$D_i := \mathbb{1}\left\{\tilde{X}_i \neq X_i\right\} \sim Ber(p)\,, \ \textit{i.i.d. for } i = 1, \ldots, n, \ p \in [0,1]$$

under $\mathbb{P}$*. Then, it holds that*

$$P_1^n K\left(A|\tilde{X}\right) \leqslant e^{4\epsilon np} P_2^n K\left(A|X\right) + 2\delta e^{4np\epsilon}. \tag{74}$$

Next, we construct the coupling that will be used in the proof of Lemma 23, in conjunction with Lemma 24 above. The couplings use similar ideas to the one constructed in Narayanan (2022a). A proof is provided in Section C.2.

Lemma 25. *Let* K^j *satisfy an* (ϵ, δ)*-DP constraint for* $0 < \epsilon \leqslant 1$*.*

Consider $\pi = N(0, c_\alpha^{1/2} d_L^{-1} \rho^2 \bar{\Gamma})$*, with* ρ^2 *satisfying* (50) *or* (51)*, with* $\epsilon \leqslant 1/\sqrt{n}$ *and* $\delta \leqslant c_\alpha(m^{-1} \wedge \epsilon)$*.*

For all measurable sets A *it holds that*

$$P_\pi K^j\left(A|X^{(j)}, u\right) \leqslant \left(1 + c_\alpha^{1/4} m^{-1/2}\right) P_0 K^j\left(A|X^{(j)}, u\right) + 2\delta + \frac{c_\alpha}{m^{3/2}} \tag{75}$$

and

$$P_\pi K^j\left(A|X^{(j)}, u\right) \geqslant \left(1 - c_\alpha^{1/4} m^{-1/2}\right) P_0 K^j\left(A|X^{(j)}, u\right) - 2\delta - \frac{c_\alpha}{m^{3/2}} \tag{76}$$

for all $c_\alpha > 0$ *small enough.*

With the above two lemmas in hand, we are ready to prove Lemma 23.

Proof of Lemma 23: Write $\mathcal{L}_{\pi,u}^j(Y^{(j)}) \equiv \mathcal{L}_\pi^j$ and let $V_\pi \equiv V_\pi^j := \mathcal{L}_\pi^j - 1$. Using that $\mathbb{E}_0 \mathscr{L}_\pi(\tilde{X}^{(j)}) = 1$ and that by the law of total probability

$$\mathbb{E}_0^{Y^{(j)}|U=u} \mathbb{E}_0[\mathscr{L}_\pi(\tilde{X}^{(j)}) \mid |Y^{(j)}, U = u] = 1,$$

it follows that $\mathbb{E}_0^{Y^{(j)}|U=u} V_\pi = 0$ and

$$\mathbb{E}_0^{Y^{(j)}|U=u}\left(\mathcal{L}_\pi^j\right)^2 = 1 + \mathbb{E}_0^{Y^{(j)}|U=u} \mathcal{L}_\pi^j(\mathcal{L}_\pi^j - 1) = 1 + \mathbb{E}_\pi^{Y^{(j)}|U=u} V_\pi.$$

Define $V_\pi^+ := 0 \vee V_\pi$ and let $V_\pi^- = -(0 \wedge V_\pi)$, which are both nonnegative random variables, with $V_\pi = V_\pi^+ - V_\pi^-$. We have

$$\mathbb{E}_\pi^{Y^{(j)}|U=u} V_\pi^+ = \int_0^T \mathbb{P}_\pi^{Y^{(j)}|U=u}\left(V_\pi^+ \geqslant t\right) dt + \int_T^\infty \mathbb{P}_\pi^{Y^{(j)}|U=u}\left(V_\pi^+ \geqslant t\right) dt. \tag{77}$$

Taking $T = \frac{5m^{1/2}}{\alpha}$, the second term is equal to zero as $V_\pi^+ \leqslant |\mathcal{L}_{\pi,u}^j(y) - 1| \leqslant 5m^{1/2}\alpha^{-1}$, $P_0 K^j(\cdot|X^{(j)}, u)$-a.s. and $P_\pi \sim P_0$ (which in turn implies $\mathbb{P}_\pi^{Y^{(j)}|U=u} \sim \mathbb{P}_0^{Y^{(j)}|U=u}$). The integrand of the first term equals $P_\pi^n K^j(\{V_\pi \geqslant t\}|X^{(j)}, u)$. By Lemma 25, it holds that

$$P_\pi K^j(\{V_\pi \geqslant t\}|X^{(j)}, u) \leqslant (1 + c_\alpha^{1/4} m^{-1/2}) P_0 K^j(\{V_\pi \geqslant t\}|X^{(j)}, u) + \delta + \frac{c_\alpha}{m^{3/2}}.$$

It follows that (77) is bounded from above by

$$\left(1+c_\alpha^{1/4}m^{-1/2}\right)\int_0^T \mathbb{P}_0^{Y^{(j)}|U=u}\left(V_\pi^+ \geqslant t\right)dt + T\delta + T\frac{c_\alpha}{m^{3/2}} \leqslant \left(1+c_\alpha^{1/4}m^{-1/2}\right)\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^+ + T\delta + T\frac{c_\alpha}{m^{3/2}}.$$

Similarly, we have

$$\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi^- = \int_0^T \mathbb{P}_\pi^{Y^{(j)}|U=u}\left(V_\pi^- \geqslant t\right)dt + \int_T^\infty \mathbb{P}_\pi^{Y^{(j)}|U=u}\left(V_\pi^- \geqslant t\right)dt. \tag{78}$$

Choosing $T \geqslant 1$ here results in the second term being zero, as $L_\pi^j \geqslant 0$. Applying Lemma 25, the right-hand side of the above display is further bounded from below by

$$\left(1-c_\alpha^{1/4}m^{-1/2}\right)\int_0^T \mathbb{P}_0^{Y^{(j)}|U=u}\left(V_\pi^- \geqslant t\right)dt - T\delta - T\frac{c_\alpha}{m^{3/2}} \geqslant \left(1-c_\alpha^{1/4}m^{-1/2}\right)\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^- - T\delta - T\frac{c_\alpha}{m^{3/2}},$$

where the inequality uses $V_\pi^- \leqslant 1$.

Combining the above bounds with the fact that $V_\pi^+ + V_\pi^- = |V_\pi|$ and $\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi = 0$ yields that

$$\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi = \mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi^+ - \mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi^- \leqslant \quad \frac{c_\alpha^{1/4}\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi|}{\sqrt{m}} + 2T\delta + 2T\frac{c_\alpha}{m^{3/2}}.$$

Plugging in the choice of $T = 5m^{1/2}/\alpha$ and using that $\delta \leqslant c_\alpha m^{-3/2}$, we obtain

$$\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi \leqslant m^{-1/2}c_\alpha^{1/4}\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi| + 20c_\alpha(m\alpha)^{-1}.$$

If $\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi| \lesssim m^{-1/2}$, we obtain $\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi \lesssim m^{-1}(c_\alpha^{1/4} + c_\alpha/\alpha)$. Assume next that $\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi| \gtrsim m^{-1/2}$. Then, $\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi \lesssim m^{-1/2}c_\alpha^{1/4}\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi|$.

By the fact that $\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi = \mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^2$ and using that by Cauchy-Schwarz, $\mathbb{E}_0^{Y^{(j)}|U=u}|V_\pi|$ is bounded above by $\sqrt{\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^2}$. Hence, $\sqrt{\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^2} \lesssim C'c_\alpha^{1/4}m^{-1/2}$, for a universal constant $C' > 0$ depending only on α. In both cases, we obtain that

$$\mathrm{B}_u^\pi = \prod_{j=1}^m \left(1+\mathbb{E}_\pi^{Y^{(j)}|U=u}V_\pi\right) = \prod_{j=1}^m \left(1+\mathbb{E}_0^{Y^{(j)}|U=u}V_\pi^2\right) \leqslant e^{C\sqrt{c_\alpha}}$$

for universal constant $C > 0$, finishing the proof of the lemma.

Appendix B. Proofs related to the optimal testing strategies

The proofs concerning the three Sections 3.1, 3.2 and 3.3 are given in this section, divided across the subsections B.1, B.2 and B.3 respectively.

We recall the notations $d_L = \sum_{l=1}^L 2^l$ for $L \in \mathbb{N}$, $f_L = \Pi_L f$ for $f \in \ell_2(\mathbb{N})$ and $\overline{X_L^{(j)}} = n^{-1}\sum_{i=1}^n X_{L;i}^{(j)}$ for $j = 1, \ldots, m$.

B.1 Procedure I

We briefly recall the testing procedure outlined in Section 3.1. Let $\tau > 0$, $L \in \mathbb{N}$, $d_L := \sum_{l=1}^{L} 2^l$ and $V_{L;\tau}^{(j)} \sim \chi^2_{d_L}$ independent of $X^{(j)}$ the random map from $(\mathbb{R}^{d_L})^n$ to $\mathbb{R}$ defined by

$$\tilde{S}_{L;\tau}^{(j)}(x) = \left[\frac{1}{\sqrt{d_L}} \left(\left\| \sigma^{-1} \sqrt{n} \overline{x} \right\|_2^2 - V_{L;\tau}^{(j)} \right) \right]_{-\tau}^{\tau}. \tag{79}$$

Define furthermore

$$D_\tau = (n\sqrt{d_L})^{-1} \tilde{\kappa}_\alpha \log(N) \left(\sqrt{n\sqrt{d_L}\tau} \vee \sqrt{n d_L} \right). \tag{80}$$

Set $K_\tau = \lceil 2\tau D_\tau^{-1} \rceil$ and consider the set $\mathcal{C}_{L;\tau} = \mathcal{A}_{L;\tau} \cap \mathcal{B}_{L;\tau}$, where

$$\mathcal{A}_{L;\tau} = \left\{ (x_i) \in (\mathbb{R}^\infty)^n : \left| \left\| \sigma^{-1} \textstyle\sum_{i \in \mathcal{J}} \Pi_L x_i \right\|_2^2 - k d_L \right| \leqslant \tfrac{1}{8} k D_\tau n \sqrt{d_L} \ \ \forall \mathcal{J} \subset [n], |\mathcal{J}| = k \leqslant K_\tau \right\}, \tag{81}$$

$$\mathcal{B}_{L;\tau} = \left\{ (x_i) \in (\mathbb{R}^\infty)^n : \left| \left\langle \sigma^{-1} \Pi_L x_i, \sigma^{-1} \textstyle\sum_{k \neq i} \Pi_L x_k \right\rangle \right| \leqslant \frac{1}{8} k D_\tau n \sqrt{d_L}, \ \ \forall i = 1, \dots, n \right\}.$$

Lemma 26 below shows that $X^{(j)}$ concentrates on $\mathcal{C}_{L;\tau}$ under the null hypothesis. We defer its proof to Section C.1.

Lemma 26. *Whenever $\sigma^{-2} n \|\Pi_L f\|_2^2 d_L^{-1/2} \leqslant \tau/2$, $\tau \leqslant nR^2/\sqrt{d_L}$ and $\tilde{\kappa}_\alpha$ is taken large enough, it holds that*

$$\mathbb{P}_f \left(X^{(j)} \notin \mathcal{C}_{L;\tau} \right) \leqslant \frac{\alpha}{2N}.$$

In Lemma 27, it is shown that $x \mapsto S^{(j)}(x)$ is D_τ-Lipschitz with respect to the Hamming distance on $\mathcal{C}_{L;\tau}$, with D_τ as defined in (80).

Lemma 27. *The map $x \mapsto S_\tau^{(j)}(x)$ defined in* (13) *is D_τ-Lipschitz with respect to $(\mathbb{R}^d)^n$-Hamming distance on $\mathcal{C}_{L;\tau}$.*

Proof Consider $x = (x_i)_{i \in [n]}, \breve{x} = (\breve{x})_{i \in [n]} \in \mathcal{C}_{L;\tau}$ with $k := \mathrm{d}_H(x, \breve{x})$. If $k > \lceil 2\tau D_\tau^{-1} \rceil$, we have $|S_\tau^{(j)}(x) - S_\tau^{(j)}(\breve{x})| \leqslant 2\tau \leqslant D_\tau k$. If $k \leqslant \lceil 2\tau D_\tau^{-1} \rceil$, let $\mathcal{J} \subset [n]$ denote the indexes of columns in which x and $\breve{x}$ differ. Define the sum of the elements that x and $\breve{x}$ have in common as $v = \sigma^{-1} \sum_{i \in [n] \setminus \mathcal{J}} x_i$, such that $\sigma^{-1} \sum_{i=1}^{n} x_i = v + w$ and $\sigma^{-1} \sum_{i=1}^{n} \breve{x}_i = v + \breve{w}$. We have

$$\begin{aligned} S_\tau^{(j)}(x) - S_\tau^{(j)}(\breve{x}) &= \frac{n}{\sqrt{d_L}} \left(\left\| n^{-1}(v + w) \right\|^2 - n^{-1} V^{(j)} \right) - \frac{n}{\sqrt{d_L}} \left(\left\| n^{-1}(v + \breve{w}) \right\|^2 - n^{-1} V^{(j)} \right) \\ &= \frac{1}{n\sqrt{d_L}} \left(2\langle w, v \rangle - 2\langle \breve{w}, v \rangle + \|w\|_2^2 - \|\breve{w}\|_2^2 \right). \end{aligned}$$

The last two terms are bounded by $kD_\tau/4$ since $x, \breve{x} \in \mathcal{A}_{L;\tau}$. The first two terms equal

$$\frac{2}{n\sqrt{d_L}} \left(\langle w, v + w \rangle - \langle \breve{w}, v + \breve{w} \rangle + \|\breve{w}\|_2^2 - \|w\|_2^2 \right),$$

where the last two terms are bounded by $kD_\tau/2$. It holds that

$$\langle w, v+w\rangle - \langle \breve{w}, v+\breve{w}\rangle = \sigma^{-2}\sum_{i\in\mathcal{J}}(\langle x_i, \sum_{i\in[n]\setminus\mathcal{J}} x_i\rangle - \langle \breve{x}_i, \sum_{i\in[n]\setminus\mathcal{J}} \breve{x}_i\rangle + \|x_i\|_2^2 - \|\breve{x}_i\|_2^2),$$

which is bounded by $kD_\tau/4$ for $x \in \mathcal{A}_{L;\tau} \cap \mathcal{B}_{L;\tau}$. Putting it all together and by symmetry of the argument, we obtain that $\left|S_\tau^{(j)}(x) - S_\tau^{(j)}(\breve{x})\right| \leqslant D_\tau k$. ∎

Lemma 28 below shows that there exists a measurable function $S_{L;\tau}^{(j)} : \mathbb{R}^{d_L} \to \mathbb{R}$, D_τ-Lipschitz with respect to the Hamming distance, such that $S_{L;\tau}^{(j)}(X_L^{(j)}) = \tilde{S}_{L;\tau}^{(j)}(X_L^{(j)})$ whenever $X^{(j)} \in \mathcal{C}_{L;\tau}$. That is, letting $\Psi_L : \mathbb{R}^\infty \to \mathbb{R}^{d_L}$ be the coordinate projection for the first d_L coordinates, the lemma allows a Lipschitz extension of the test statistic defined in (79) on $\Psi_L\mathcal{C}_{L;\tau}$ to all of $\mathbb{R}^{d_L}$. The proof follows essentially the construction of McShane McShane (1934) for obtaining a Lipschitz extension with respect to the Hamming distance, but our lemma verifies in addition the Borel-measurability of the resulting map.

We follow a construction that is in essence that of McShane McShane (1934), whilst also verifying that such an extension is Borel measurable. A proof is given in Section C.5.

Lemma 28. *Let $\mathcal{C} \subset (\mathbb{R}^{d_L})^n$ and $S : \mathcal{C} \to \mathbb{R}$ be a (Borel) measurable D-Lipschitz map with respect to the Hamming distance on $(\mathbb{R}^{d_L})^n$. Then, there exists a map $\tilde{S} : (\mathbb{R}^{d_L})^n \to \mathbb{R}$ measurable with respect to the Borel sigma algebra such that it is D-Lipschitz with respect to the Hamming distance on $(\mathbb{R}^{d_L})^n$ and $\tilde{S} = S$ on $\mathcal{C}$.*

B.1.1 Proof of Lemma 8

Consider the transcript of Section 3.1, which we recall here as

$$Y_{L;\tau}^{(j)} = \gamma_\tau \breve{S}_{L;\tau}^{(j)}(X^{(j)}) + W_\tau^{(j)}, \qquad \text{with } \gamma_\tau := \frac{\epsilon}{D_\tau\sqrt{2\mathfrak{c}\log(2/\delta)}}, \tag{82}$$

$\mathfrak{c} > 0$, $W_\tau^{(j)} \sim N(0,1)$ independent for $j = 1, \dots, m$ and $\tau > 0$. These transcripts are $(\epsilon/\mathfrak{c}, \delta)$-differentially private for any $\epsilon > 0$ (see e.g. Dwork et al. (2014a)). Define the test

$$\varphi_\tau := \mathbb{1}\left\{\frac{1}{\sqrt{m}}\sum_{j=1}^m Y_{L;\tau}^{(j)} \geqslant (\gamma_\tau \vee 1)\,\kappa\right\}. \tag{83}$$

We will prove the following more general version of Lemma 8.

Lemma 29. *Consider the test φ_τ as defined by (16). Whenever $\tau/4 \leqslant \frac{n\|f_L\|_2^2}{\sigma^2\sqrt{d_L}} \leqslant \tau/2$ and*

$$\|f_L\|_2^2 \geqslant C_\alpha\kappa\sigma^2\sqrt{\mathfrak{c}\log(1/\delta)}\log(N)\left(\frac{\sqrt{d_L}}{\sqrt{N}\sqrt{n}(\sqrt{n}\epsilon \wedge 1)}\right)\bigvee\left(\frac{1}{Nn\epsilon^2}\right), \tag{84}$$

for $C_\alpha > 0$ large enough, it holds that $\mathbb{P}_f(1-\varphi_\tau) \leqslant \alpha$.

Proof On the event that $X^{(j)} \in \mathcal{C}_\tau$ for all $j \in [m]$, we have that

$$\sum_{j=1}^{m} Y_\tau^{(j)} = \sum_{j=1}^{m} \left(\gamma_\tau \check{S}_\tau(X^{(j)}) + W_\tau^{(j)} \right) = \sum_{j=1}^{m} \left(\gamma_\tau S_\tau(X^{(j)}) + W_\tau^{(j)} \right). \tag{85}$$

Consequently, $\mathbb{P}_f(1 - \varphi_\tau)$ is bounded above by

$$\mathbb{P}_f \left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} \left(\gamma_\tau S_\tau(X^{(j)}) + W_\tau^{(j)} \right) \geqslant (\gamma_\tau \vee 1)\, \kappa \right) + \mathbb{P}_f \left(\exists j : X^{(j)} \notin \mathcal{C}_\tau \right). \tag{86}$$

By Lemma 26 and a union bound, the second term is bounded above by $\alpha/2$.

Under $\mathbb{P}_f$, it holds that

$$\frac{n}{\sqrt{d_L}} \left(\left\| \sigma^{-1} \overline{X_L^{(j)}} \right\|_2^2 - \frac{V^{(j)}}{n} \right) \stackrel{d}{=} \frac{n\sigma^{-2}\|f_L\|_2^2}{\sqrt{d_L}} + 2\frac{\sqrt{n}}{\sqrt{d_L}} \langle Z, \sigma^{-1} f \rangle + \frac{\|Z\|_2^2 - V^{(j)}}{\sqrt{d_L}}, \tag{87}$$

where $Z \sim N(0, I_{d_L})$. By assumption, $\frac{n\|f_L\|_2^2}{\sqrt{d_L}\sigma^2} \leqslant \tau/2$, $\mathrm{Var}(\frac{\sqrt{n}}{\sqrt{d_L}} \langle Z, \sigma^{-1} f_L \rangle) = n\sigma^{-2}\|f_L\|_2^2/d_L \leqslant \tau/2$ and $(\|Z\|_2^2 - V^{(j)})/\sqrt{d_L}$ tends to a Gaussian with variance 4 for large d_L. The second and third term in (87) are symmetric in distribution about 0, have uniformly bounded densities (since the Chi-square and normal densities are bounded, and the third term tends weakly to a Gaussian in d_L) and $\sigma^{-2} d_L^{-1/2} n \|f_L\|_2^2 \leqslant \tau/2$, which means that the conditions of Lemma 36 are satisfied. Applying said lemma (with $\mu = \sigma^{-2} d_L^{-1/2} n \|f_L\|_2^2$), we get that there exists a uniform constant $c > 0$ such that

$$\mathbb{E}_f \frac{1}{\sqrt{m}} \sum_{j=1}^{m} \left(\gamma_\tau S_\tau(X^{(j)}) + W_\tau^{(j)} \right) \geqslant c \frac{\sqrt{m} n \|f_L\|_2^2 \gamma_\tau}{\sigma^2 \sqrt{d_L}}.$$

Under $\mathbb{P}_f$, by independence of the data and the Gaussian noise,

$$\mathrm{Var}_f \left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} \gamma_\tau S_\tau(X^{(j)}) + W_\tau^{(j)} \right) = 1 + \mathrm{Var}_f \left(\gamma_\tau S_\tau(X^{(1)}) \right).$$

Since

$$\mathbb{E}_f \frac{n}{\sqrt{d_L}} \left(\left\| \sigma^{-1} \overline{X^{(j)}}_L \right\|_2^2 - \frac{V^{(j)}}{n} \right) = \frac{n\sigma^{-2}\|f_L\|_2^2}{\sqrt{d_L}} \leqslant \tau/2,$$

the fact that clipping reduces the variance yields

$$\mathrm{Var}_f \left(\gamma_\tau S_\tau(X^{(1)}) \right) \leqslant \gamma_\tau^2 \mathrm{Var}_f \left(\frac{n}{\sqrt{d_L}} \left(\left\| \sigma^{-1} \overline{X^{(j)}} \right\|_2^2 - \frac{V^{(j)}}{n} \right) \right) \leqslant \gamma_\tau^2 \left(\frac{4n\|f_L\|_2^2}{\sigma^2 d_L} + 4 \right).$$

Assume now that for all $C_\alpha > 0$ large enough,

$$(\gamma_\tau \vee 1)\, \kappa \leqslant c \frac{1}{2} \sqrt{m} n \sigma^{-2} \|f_L\|_2^2 \gamma_\tau / \sqrt{d_L}, \tag{88}$$

which is a claim we shall prove later on. Then, the first term in (86) is bounded above by

$$\mathbb{P}_f\left(\frac{1}{\sqrt{m}}\sum_{j=1}^{m}\left(\gamma_\tau S_\tau(X^{(j)}) + W_\tau^{(j)} - \mathbb{E}_f(S_\tau(X^{(j)}) + W_\tau^{(j)})\right) < -c\frac{\sqrt{m}n\sigma^{-2}\|f_L\|_2^2\gamma_\tau}{2\sqrt{d_L}}\right), \quad (89)$$

which, by Chebyshev's inequality is bounded by

$$\left(\frac{\sqrt{m}n\sigma^{-2}\|f_L\|_2^2\gamma_\tau}{\sqrt{d_L}}\right)^{-2} + \left(mn\sigma^{-2}\|f_L\|_2^2\right)^{-1} + \left(\frac{\sqrt{m}n\sigma^{-2}\|f_L\|_2^2}{\sqrt{d_L}}\right)^{-2}.$$

For f satisfying (18), the last two terms are easily seen to be smaller than $\alpha/6$ for a large enough to choice for C_α. To see that this is also true for the first term, recall that $\gamma_\tau := \frac{\epsilon}{D_\tau\sqrt{2\mathfrak{c}\log(2/\delta)}}$, with $D_\tau = \frac{\tilde{\kappa}_\alpha \log(N)(\sqrt{n\sqrt{d_L}\tau}\vee\sqrt{nd_L})}{n\sqrt{d_L}}$, which yields that the square root of the first term equals

$$\frac{\sqrt{m}n\sigma^{-2}\|f_L\|_2^2\epsilon}{\sqrt{d_L}D_\tau\sqrt{2\mathfrak{c}\log(2/\delta)}} = \frac{\sqrt{m}n^2\sigma^{-2}\|f_L\|_2^2\epsilon}{\tilde{\kappa}_\alpha\log(N)(\sqrt{n\sqrt{d_L}\tau}\vee\sqrt{nd_L})\sqrt{2\mathfrak{c}\log(2/\delta)}}.$$

When the maximum is taken in $\sqrt{nd_L}$, (84) leads to the latter being larger than C_α. When the maximum is taken in $\sqrt{n\sqrt{d_L}\tau}$, using that $4\sigma^{-2}n\|f\|_2^2/\sqrt{d_L} \geqslant \tau$ yields that the above display is bounded from below by

$$\frac{\sqrt{m}\sigma^{-1}n\|f_L\|_2\epsilon}{\tilde{\kappa}_\alpha\sqrt{\log(N)\sqrt{2\mathfrak{c}\log(2/\delta)}}} \geqslant C_\alpha.$$

In either case, it follows that the Type II error (i.e. (89)) can be made arbitrarily small per large enough choice of $C_\alpha > 0$.

We return to the claim of (88). When $\gamma_\tau \geqslant 1$, the claim is satisfied whenever $\kappa \lesssim \log(N)\sqrt{\mathfrak{c}\log(1/\delta)}$ for $C_\alpha > 0$ large enough. When $\gamma_\tau < 1$, it is required that

$$\frac{\sqrt{m}\sigma^{-2}n\gamma_\tau\|f_L\|_2^2\epsilon}{\sqrt{d_L}} = \frac{\sqrt{m}\sigma^{-2}n^2\|f_L\|_2^2\epsilon}{\tilde{\kappa}_\alpha\log(N)\sqrt{2\mathfrak{c}\log(2/\delta)}(\sqrt{n\sqrt{d_L}\tau}\vee\sqrt{nd_L})} \gtrsim \kappa,$$

which is satisfied whenever (84) holds. ∎

B.1.2 Proof of Lemma 9

We start by recalling the notation, before proving a slightly more general result. Consider for a set $\mathcal{S} \subset \mathbb{N}$ the test

$$T_{\mathrm{I}} := \max_{L\in\mathcal{S},\,\tau\in\mathrm{T}_L} \mathbb{1}\left\{\frac{1}{\sqrt{m}}\sum_{j=1}^{m} Y_{L;\tau}^{(j)} \geqslant \kappa_\alpha\left(\frac{\epsilon}{D_\tau\sqrt{2|\mathrm{T}_L||\mathcal{S}|\log(2/\delta)}}\vee 1\right)\sqrt{\log|\mathrm{T}_L||\mathcal{S}|}\right\}, \quad (90)$$

where

$$\mathrm{T}_L := \left\{2^{-k+2}\frac{nR^2}{\sqrt{2^{L+1}}} : k = 1,\dots,\lceil 1 + 2\log_2(NR/\sigma)\rceil\right\}, \quad (91)$$

and $Y_L^{(j)} = \{Y_{L;\tau}^{(j)} : \tau \in \mathrm{T}_L\}$ is generated according to (82) with $\gamma_\tau = \frac{\epsilon}{D_\tau \sqrt{2|\mathrm{T}_L||\mathcal{S}|\log(2/\delta)}}$. By the same reasoning as in the proof of Lemma 29, the test T_{I} is (ϵ, δ)-differentially private. We will prove the following more general version of Lemma 9.

Lemma 30. *For all $M > 0$, $\alpha \in (0,1)$ there exists $\kappa_\alpha > 0$ and $C_\alpha > 0$ such that the test defined by* (90) *satisfies $\mathbb{P}_0 T_I \leqslant \alpha$. Furthermore, for any $f \in \mathcal{B}_{p,q}^{s,R}$ such that for some $L^* \in \mathcal{C}$ it holds that*

$$\|f_{L^*}\|_2^2 \geqslant C_\alpha M_N \left(\frac{\sqrt{2^{L^*}}}{\sqrt{m} n (\sqrt{n}\epsilon \wedge 1)} \right) \bigvee \left(\frac{1}{mn^2\epsilon^2} \right), \tag{92}$$

where $M_N \gtrsim \log(N)\sqrt{2\log\log(NR/\sigma)\log(NR/\sigma)|\mathcal{S}|\log(2/\delta)}$, we also have that $\mathbb{P}_f(1-T_I) \leqslant \alpha$,where $C_\alpha > 0$ depends only on α.

Proof We start by proving that the level of the test is controlled. Using a union bound and writing $\mathfrak{r} = (\gamma_\tau \vee 1)\sqrt{\log|\mathrm{T}_L||\mathcal{S}|}$, we have that

$$\begin{aligned} \mathbb{P}_0 T_{\mathrm{I}} &= \mathbb{P}_0 \left(\max_{L \in \mathcal{S},\, \tau \in \mathrm{T}_L} \frac{1}{\mathfrak{r}\sqrt{m}} \sum_{j=1}^m [\gamma_L \breve{S}_{L;\tau}^{(j)}(X^{(j)}) + W_\tau^{(j)}] \geqslant \kappa_\alpha \right) \\ &\leqslant \mathbb{P}_0 \left(\max_{L \in \mathcal{S},\, \tau \in \mathrm{T}_L} \frac{1}{\mathfrak{r}\sqrt{m}} \sum_{j=1}^m [\gamma_L \tilde{S}_{L;\tau}^{(j)}(X^{(j)}) + W_\tau^{(j)}] \geqslant \kappa_\alpha \right) + \mathbb{P}_0 \left(\exists (j, L, \tau) \,:\, X^{(j)} \notin \mathcal{C}_{L;\tau} \right) \end{aligned} \tag{93}$$

where it is used that the Lipschitz extension $\breve{S}_{L;\tau}^{(j)}$ of $\tilde{S}_{L;\tau}^{(j)}$ equals the latter function on $\mathcal{C}_{L;\tau}$. By Lemma 26 and a union bound, the second probability on the right-hand side is bounded above by $\alpha m|\mathcal{S}||\mathrm{T}_L|/(2N)$, for a large enough choice of $\tilde{\kappa}_\alpha > 0$.

Considering the first probability on the right-hand side of the inequality displayed above, we first note that by Lemma 35, $\tilde{S}_{L;\tau}^{(j)}(X^{(j)}) = \left[\frac{1}{\sqrt{d_L}} \left(\left\| \sqrt{n}\overline{X_L^{(j)}} \right\|_2^2 - V_{L;\tau}^{(j)} \right) \right]_{-\tau}^{\tau}$ is $1/(4\sqrt{d_L})$-sub-exponential, where the sub-exponentiality parameter follows from a straightforward calculation (see e.g. Lemma Szabó et al. (2022)). Since $W_{L;\tau}^{(j)}$ is independent $N(0,1)$, it follows by Bernstein's inequality (see e.g. Theorem 2.8.1 in Vershynin (2018)) and a union bound (e.g. Lemma 37) that the first probability in (93) is bounded as follows

$$\mathbb{P}_0 \left(\max_{L \in \mathcal{S},\, \tau \in \mathrm{T}_L} \frac{1}{\mathfrak{r}\sqrt{m}} \sum_{j=1}^m [\gamma_L \tilde{S}_{L;\tau}^{(j)}(X^{(j)}) + W_\tau^{(j)}] \geqslant \kappa_\alpha \right) \leqslant \sum_{L \in \mathcal{S},\, \tau \in \mathrm{T}_L} \frac{1}{(|\mathrm{T}_L||\mathcal{S}|)^{\kappa_\alpha^2/2}},$$

which is less than $\alpha/2$ for a large enough choice of $\kappa_\alpha > 0$. For $f \in \mathcal{B}_{p,q}^{s,R}$, Lemma 42 yields that

$$\|f_L\|_2 \leqslant \|f\|_2 \leqslant (1 - 2^{-s})^{1/q-1} \|f\|_{s,p,q} \leqslant R(1 - 2^{-s})^{1/q-1}. \tag{94}$$

Consequently, when f satisfies (18), it holds in particular that $\|f_L\|_2^2 \geqslant N^{-1}$. Thus, there exists $\tau^* \in \mathrm{T}_{L^*}$ such that the condition of Lemma 29 is satisfied, which now yields that $\mathbb{P}_f(1 - T_{\mathrm{I}}) \leqslant \mathbb{P}_f(1 - \varphi_{\tau^*}) \leqslant \alpha/2$. ■

B.2 Procedure II

Let $\epsilon, \delta > 0$ and $\mathcal{S} \subset \mathbb{N}$ be given. Consider for $L \in \mathcal{S}$, $K_L = \lceil n\epsilon^2 \wedge d_L \rceil$ and take sets $\mathcal{J}_{lk;L} \subset [m]$ such that $|\mathcal{J}_{lk;L}| = \lceil \frac{mK_L}{d_L} \rceil$ and each $j \in \{1, \dots, m\}$ is in $\mathcal{J}_{lk;L}$ for K_L different indexes $k \in \{1, \dots, d_L\}$. For $(l,k) \in \{l = 1, \dots, L, k = 1, \dots, 2^l\} =: I_L$, $j \in \mathcal{J}_{lk;L}$, generate the transcripts according to

$$Y^{(j)}_{L;lk} | X^{(j)} = \gamma_L \sum_{i=1}^{n} \sigma^{-1} \left[(X^{(j)}_{L;i})_{lk} \right]^{\tau}_{-\tau} + W^{(j)}_{L;lk} \tag{95}$$

with $\gamma_L = \frac{\epsilon}{2\sqrt{|\mathcal{S}| K_L \log(2/\delta)\tau}}$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N)}$ and $(W^{(j)}_{L;lk})_{j,k}$ i.i.d. standard Gaussian noise.

Define $Y^{(j)}_L := (Y^{(j)}_{L;lk})_{(l,k) \in I_L : j \in \mathcal{J}_{lk;L}}$ and consider the transcripts $Y^{(j)} := (Y^{(j)}_L)_{L \in \mathcal{S}}$, for $j = 1, \dots, m$. The lemma below shows that $Y^{(j)}$ satisfies DP.

Lemma 31. *The transcript defined $Y^{(j)}$ is (ϵ, δ)-differentially private.*

Proof The rescaled and clipped sums have at most L_2-sensitivity less than or equal to one:

$$\sup_{\breve{x} \in (\mathbb{R}^{\mathbb{N}})^n : \mathrm{d}_H(x, \breve{x}) \leqslant 1} \gamma_L \| (\sum_{i=1}^{n} \sigma^{-1} [(x_i)_{lk}]^{\tau}_{-\tau} - \sum_{i=1}^{n} \sigma^{-1} [(\breve{x}_i)_{lk}]^{\tau}_{-\tau})_{(l,k) \in I_L, L \in \mathcal{S}} \|_2 \leqslant \tag{96}$$
$$\gamma_L \sqrt{\sum_{L \in \mathcal{S}} \sum_{(l,k) \in I_L} (\sup_{\breve{x} \in (\mathbb{R}^{\mathbb{N}})^n : \mathrm{d}_H(x, \breve{x}) \leqslant 1} \sigma^{-1} [(x_i)_{lk}]^{\tau}_{-\tau} - \sigma^{-1} [(\breve{x}_i)_{lk}]^{\tau}_{-\tau})^2} \leqslant 1.$$

Consequently, the addition of the Gaussian noise assures that the transcript $Y^{(j)}$ is (ϵ, δ)-differentially private (see e.g. Appendix A in Dwork et al. (2014a)). ■

Consider the test given by

$$T_{\mathrm{II}} := \mathbb{1} \left\{ \max_{L \in \mathcal{S}} \frac{1}{\sqrt{d_L} \left(n\gamma_L^2 \vee 1 \right)} \sum_{(l,k) \in I_L} \left[\left(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} Y^{(j)}_{L;lk} \right)^2 - \nu_{\epsilon,L} \right] \geqslant \kappa_\alpha \log\log(e \vee |\mathcal{S}|) \right\}, \tag{97}$$

with $\nu_{\epsilon,L} := \mathbb{E}_0 (|\mathcal{J}_{lk;L}|^{-1/2} \sum_{j \in \mathcal{J}_{lk;L}} Y^{(j)}_{L;lk})^2$.

Lemma 32. *For all $M > 0$, $\alpha \in (0,1)$ there exists $\kappa_\alpha > 0$ and $C_\alpha > 0$ such that the test defined by (97) satisfies $\mathbb{P}_0 T_{II} \leqslant \alpha$. Furthermore, for any $f \in \mathcal{B}^{s,R}_{p,q}$ such that*

$$\|f_{L^*}\|_2^2 \geqslant C_\alpha \log\log(e \vee |\mathcal{S}|) \log(N) |\mathcal{S}| \log(1/\delta) \left(\frac{\sqrt{2^{(3/2)L^*}}}{\sqrt{m} n (n\epsilon^2 \wedge d_{L^*})} \right) \bigvee \left(\frac{1}{mn^2\epsilon^2} \right), \tag{98}$$

for some $L^* \in \mathcal{S}$ and $C_\alpha > 0$ large enough depending only on α, it holds that $\mathbb{P}_f (1 - T_{II}) \leqslant \alpha$.

Proof Under the null hypothesis, $(X^{(j)}_{L;i})_{lk}$ are independent standard Gaussian for $(l,k) \in I_L$, $j \in [m]$. Hence, by Lemma 35, the random variables $Y^{(j)}_{L;lk}$ are sub-gaussian, mean zero and i.i.d. for $k \in [d_L]$, $j \in [m]$ under the null hypothesis. More specifically, we have

that $\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} Y_{L;lk}^{(j)}$ is $(\frac{\sqrt{n}\epsilon}{2\sqrt{|\mathcal{S}|K_L \log(2/\delta)\tau}} + 1)$-sub-gaussian, which in turn implies that the random variables

$$\frac{1}{\sqrt{d_L}\left(n\gamma_L^2 \vee 1\right)} \sum_{(l,k) \in I_L} [(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} Y_{L;lk}^{(j)})^2 - \nu_{\epsilon,L}]$$

are C-sub-exponential for $L \in \mathcal{S}$, for some constant $C > 0$ (see e.g. Section 2.7 in Vershynin (2018)). It now follows from Lemma 37 that, for $\kappa_\alpha > 0$ large enough, $\mathbb{P}_0 T_{\mathrm{II}} \leqslant \alpha/2$.

We now turn our attention to the Type II error. Assume now that, for $L \in \mathcal{S}$, f satisfies (98). We have that

$$1 - T_{\mathrm{II}} \leqslant \mathbb{1}\left\{ \frac{1}{\sqrt{d_L}\left(n\gamma_L^2 \vee 1\right)} \sum_{(l,k) \in I_L} \left[\left(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} Y_{L;lk}^{(j)} \right)^2 - \nu_{\epsilon,L} \right] < \kappa_\alpha \log\log(e \vee |\mathcal{S}|) \right\},$$

so it suffices to bound the $\mathbb{E}_f$-expectation of the right-hand side.

Under $\mathbb{P}_f$, $(X_i^{(j)})_{lk} \stackrel{d}{=} f_{lk} + \sigma^{-1} Z_{lk;i}^{(j)}$ with i.i.d. $Z_{lk;i}^{(j)} \sim N(0,1)$ and is independent of the centered i.i.d. $W_{lk}^{(j)}$. Hence, the quantity

$$V_{lk} := \left(\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} \left[\gamma_L \sum_{i=1}^{n} \sigma^{-1} (X_i^{(j)})_{lk} + W_{lk}^{(j)} \right] \right)^2$$

is in distribution equal to

$$\left(\gamma_L \sqrt{|\mathcal{J}_{lk;L}|} n \sigma^{-1} f_{lk} + \gamma_L \sqrt{n} \eta + \frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} W_{lk}^{(j)} \right)^2$$

under $\mathbb{P}_f$, with $\eta \sim N(0,1)$ independent. Therefore, a straightforward calculation shows that V_{lk} has mean $\sigma^{-1}\gamma_L^2 n^2 |\mathcal{J}_{lk;L}| f_{lk}^2 + n\gamma_L^2 + \mathbb{E}(W_{lk}^{(j)})^2$ under $\mathbb{P}_f$. Since $\mathbb{E}(Z_{lk;i}^{(j)})^4 = 3$ and $\mathbb{E}(W_{lk}^{(j)})^4 \asymp 1$, its variance equals

$$\begin{aligned} n^2\gamma_L^4 \mathrm{Var}\left(\eta^2\right) + \mathrm{Var}((\frac{1}{\sqrt{|\mathcal{J}_{lk;L}|}} \sum_{j \in \mathcal{J}_{lk;L}} W_{lk}^{(j)})^2) + \gamma_L^4 |\mathcal{J}_{lk;L}| n^3 \sigma^{-2} f_{lk}^2 \mathbb{E}\eta^2 \\ + \gamma_L^2 |\mathcal{J}_{lk;L}| n^2 \sigma^{-2} f_{lk}^2 \mathbb{E}(W_l^{(j)})^2 + n\gamma_L^2 \mathbb{E}(W_{lk}^{(j)})^2 \mathbb{E}\eta^2, \end{aligned} \tag{99}$$

which is of the order

$$(\gamma_L^4 |\mathcal{J}_{lk;L}| n^3 \sigma^{-2} f_{lk}^2) \vee (\gamma_L^2 |\mathcal{J}_{lk;L}| n^2 \sigma^{-2} f_{lk}^2) \vee \gamma_L^4 n^2 \vee 1. \tag{100}$$

Since (94) holds, we have that for $\tilde{\kappa}_\alpha$ large enough, $\max_{((l,k) \in I_L)} |f_{lk}| \leqslant \tau/2/$Consequently, an application of the triangle inequality and a standard result for the maximum of Gaussian vectors (see e.g. Lemma 16 in Vuursteen (2024)) yield that, for $\tilde{\kappa}_\alpha > 0$ large enough, we have with probability at least $1 - 2Nd_l e^{-\tau^2/4} \geqslant 1 - (1+N)^{2-\tilde{\kappa}_\alpha^2/4} \geqslant 1 - \alpha/4$ that

$\max_{i\in[n],j\in[m],(l,k)\in I_L} |(X_i^{(j)})_{lk}| \leqslant \tau$. Consequently, under the null hypothesis ($f=0$), using that $|\mathcal{J}_{lk;L}| = |\mathcal{J}_1|$, $\mathbb{P}_0\varphi$ is bounded above by

$$\mathbb{P}_0\left(\frac{1}{\sqrt{d_L}|\mathcal{J}_{1d_L;L}|}\sum_{(l,k)\in I_L}[(\sum_{j\in\mathcal{J}_{lk;L}}(\gamma_L\sigma^{-1}(n\bar{X^{(j)}})_{lk}+W_{lk}^{(j)}))^2-\mathbb{E}_0V_l]\geqslant\kappa(\gamma_L^2n\vee 1)\right)+\frac{\alpha}{4},$$

where we write $\kappa=\kappa_\alpha\log\log(e\vee|\mathcal{S}|)$. Chebyshev's inequality yields that the first term on the left-hand side is bounded $\alpha/4$ for $\kappa_\alpha>0$ large enough. In the case that $\max_{((l,k)\in I_L)}|f_{lk}|\leqslant \tau/2$, we also have that $\mathbb{P}_f(1-T_{\mathrm{II}})$ is bounded above by

$$\Pr\left(\frac{1}{\sqrt{d_L}|\mathcal{J}_{1d_L;L}|}\sum_{(l,k)\in I_L}[(\sum_{j\in\mathcal{J}_{lk;L}}(\gamma_L(nf_{lk}+\sqrt{n}Z)+W_{lk}^{(j)}))^2-\mathbb{E}_0V_l]<\kappa(\gamma_L^2n\vee 1)\right)+\frac{\alpha}{4}.$$

Subtracting $d_L^{-1/2}\sum_{(l,k)\in I_L}\gamma_L^2n^2|\mathcal{J}_{lk;L}|\sigma^{-1}f_{lk}^2$ on both sides, the first term is bounded above by

$$\Pr\left(\frac{1}{\sqrt{d_L}|\mathcal{J}_{1d_L;L}|}\sum_{(l,k)\in I_L}[(\sum_{j\in\mathcal{J}_{lk;L}}(\gamma_L(n\sigma^{-1}f_{lk}+\sqrt{n}Z)+W_l^{(j)}))^2-\mathbb{E}_fV_l]<-\frac{\mathfrak{c}}{2}\right)$$

with $\mathfrak{c}=\sigma^{-2}\frac{\gamma_L^2n^2|\mathcal{J}_{lk;L}|}{\sqrt{d_L}}\|f_L\|_2^2$, whenever

$$\mathfrak{c}\geqslant 2\kappa(\gamma_L^2n\vee 1)\iff\frac{2\sigma^2\kappa_\alpha\log\log(e\vee|\mathcal{S}|)(\gamma_L^2n\vee 1)d_L\sqrt{d_L}}{\gamma_L^2n^2mK_L\|f_L\|_2^2}\leqslant 1.$$

This follows from the assumed inequality (98) in the lemma's statement.

An application of Chebyshev's inequality, the variance bound computed in (100) and plugging in $|\mathcal{J}_{lk;L}|\asymp mK_L/d_L$, $\gamma_L=\frac{\epsilon}{2\sqrt{|\mathcal{S}|K_L\log(2/\delta)\tau}}$ and $\tau=\tilde{\kappa}_\alpha\sqrt{\log(N)}$ now yields that the probability in the second last display is of the order

$$\frac{(\frac{m\epsilon^4}{16K_Ld_L\Lambda_{N,\delta,|\mathcal{S}|}^2}n^3\sigma^{-2}\|f_L\|_2^2)\vee(\frac{m\epsilon^2}{4K_Ld_L\Lambda_{N,\delta,|\mathcal{S}|}}n^2\sigma^{-2}\|f_L\|_2^2)\vee(\frac{\epsilon^4}{16K_L\Lambda_{N,\delta,|\mathcal{S}|}^2})n^2\vee 1}{acm\epsilon^2n^24K_Ld_L\Lambda_{N,\delta,|\mathcal{S}|}\sqrt{d_L}\sigma^{-2}\|f_L\|_2^2)^2}\tag{101}$$

with $\Lambda_{N,\delta,|\mathcal{S}|}=|\mathcal{S}|\log(2/\delta)\tilde{\kappa}_\alpha^2\log(N)$. For $1/\sqrt{n}\lesssim\epsilon$ and $n\epsilon^2\lesssim d_L$, this expression is of the order $\frac{\Lambda_{N,\delta,|\mathcal{S}|}^2d_L^3}{m^2n^4\epsilon^4\|f_L\|_2^4}$, which is bounded by $1/C_\alpha^2$ when f satisfies (98). This can be seen to be arbitrarily small whenever f is such that (98) by taking $C_\alpha>0$ large enough. Lastly, when $n\epsilon^2\gtrsim d_L$, (107) is of the order $\frac{\kappa_\alpha^2\Lambda_{N,\delta,|\mathcal{S}|}d_L}{m^2n^2\|f_L\|_2^4}$, which also holds for $C_\alpha>0$ large enough for f satisfying (98). Consequently, we obtain that $\mathbb{P}_f(1-T_{\mathrm{II}}^{\epsilon,\delta})\leqslant\alpha/2$ for C_α large enough, as desired.

■

B.3 Procedure III

Let $\mathcal{S} \subset \mathbb{N}$ and consider for $L \in \mathbb{N}$, $d_L := \sum_{l=1}^{L} 2^l$, $K_L := \lceil n\epsilon^2 \wedge 2^L \rceil$, $I_L := \{(l,k) : l = 1, \dots, \lceil \log_2(K_L) \rceil, k = 1, \dots, 2^l\}$ and $j = 1, \dots, m$ the transcripts

$$Y_{lk}^{(j)} | (X^{(j)}, U_L) = \gamma_L \sum_{i=1}^{n} [(U_L X_{L;i}^{(j)})_{lk}]_{-\tau}^{\tau} + W_{lk}^{(j)}, \tag{102}$$

with $\gamma_L = \frac{\epsilon}{2\sqrt{K_L |\mathcal{S}| \log(2/\delta)\tau}}$, $\tau = \tilde{\kappa}_\alpha \sqrt{\log(N)}$, U_L a d_L-by-d_L random rotation drawn uniformly (i.e. from the Haar measure on the special orthogonal group in $\mathbb{R}^{d_L \times d_L}$) and $(W_{lk}^{(j)})_{j,l,k}$ i.i.d. centered standard Gaussian noise. By the same proof as Lemma 31, the transcript is (ϵ, δ)-differentially private.

Lemma 33. *The transcript defined in* (102) *is* (ϵ, δ)*-differentially private.*

Based on these transcripts, one can compute

$$T_{\mathrm{III}} = \mathbb{1}\left\{ \max_{L \in \mathcal{L}} \frac{1}{\sqrt{K_L}(n\gamma_L^2 \vee 1)} \sum_{(l,k) \in I_L} [(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{lk}^{(j)})^2 - \nu_L] \geqslant \kappa_\alpha \log\log(e \vee |\mathcal{S}|) \right\}, \tag{103}$$

where $\nu_L = n\gamma_L^2 + 1$. The following lemma establishes the Type I and Type II error probability guarantees for the above test.

Lemma 34. *For all* $s, R > 0$, $\alpha \in (0,1)$ *there exists* $\kappa_\alpha, \tilde{\kappa}_\alpha > 0$ *and* $C_\alpha > 0$ *such that the test* T_{III} *defined by* (103) *satisfies* $\mathbb{P}_0 T_{III} \leqslant \alpha$*. Furthermore, for any* $f \in \mathcal{B}_{p,q}^{s,R}$ *such that*

$$\|f_{L^*}\|_2^2 \geqslant C_\alpha \log\log(e \vee |\mathcal{S}|) \log(N) |\mathcal{S}| \log(1/\delta) \left(\frac{2^{L^*}}{\sqrt{m} n (\sqrt{n}\epsilon \wedge 1)} \right) \bigvee \left(\frac{1}{mn^2\epsilon^2} \right), \tag{104}$$

for some $L^* \in \mathcal{S}$ *and* $C_\alpha > 0$ *large enough depending only on* α*, it holds that* $\mathbb{P}_f(1 - T_{II}^{\epsilon,\delta}) \leqslant \alpha$*.*

Proof Under $\mathbb{P}_f$, $(U_L X_{L;i}^{(j)})_{lk} \overset{d}{=} (U_L f)_{lk} + (U_L Z_{L;i}^{(j)})_{lk}$, where $Z_{L;i}^{(j)} \sim N(0, I_{d_L})$ independent of U_L and the centered i.i.d. $W_{lk}^{(j)}$. Furthermore, $U_L Z_{L;i}^{(j)} \overset{d}{=} Z_{L;i}^{(j)}$, $(U_L Z_{L;i}^{(j)})_{lk} \sim N(0,1)$, still independent of $(W_{lk}^{(j)})_{l,k,j}$. We obtain that $\gamma_L \sum_{i=1}^{n} (U_L X_{L;i}^{(j)})_{lk} + W_{lk}^{(j)}$ is in law equal to $\gamma_L n (U_L f_L)_{lk} + \gamma_L \sqrt{n}\eta + W_{lk}^{(j)}$, with $\eta \sim N(0,1)$ and all three terms independent.

A first consequence of the relationship above is that under the null hypothesis, the random variables $Y_{L;k}^{(j)}$ are sub-gaussian, mean zero and i.i.d. for $(l,k) \in I_L$, $j \in [m]$ under the null hypothesis (see Lemma 35). More specifically, we have that $\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{lk}^{(j)}$ is $(\frac{\sqrt{n}\epsilon}{2\sqrt{|\mathcal{S}| K_L \log(2/\delta)\tau}} + 1)$-sub-gaussian, which in turn implies that the random variables $\frac{1}{\sqrt{K_L}} \sum_{(l,k) \in I_L} [(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{L;k}^{(j)})^2 - \nu_L]$ are C-sub-exponential for $L \in \mathcal{S}$, for some constant $C > 0$ (see e.g. Section 2.7 in Vershynin (2018)). It now follows from Lemma 37 that, for $\kappa_\alpha > 0$ large enough, $\mathbb{P}_0 T_{\mathrm{III}} \leqslant \alpha/2$.

Next, we turn to the Type II error probability. Assume now that, for $L \in \mathcal{S}$, f satisfies (98). We have that

$$1 - T_{\mathrm{III}} \leqslant \mathbb{1}\left\{ \frac{1}{\sqrt{K_L}} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{lk}^{(j)} \right)^2 - \nu_L \right] < \kappa_\alpha \left(n\gamma_L^2 \vee 1 \right) \log\log(e \vee |\mathcal{S}|) \right\},$$

so it suffices to bound the $\mathbb{E}_f$-expectation of the right-hand side. Since (94) holds, we have that for $\tilde{\kappa}_\alpha$ large enough, $\max_{((l,k)\in I_L)} |f_{lk}| \leqslant \tau/2$. By a similar computation as earlier, it suffices to show that

$$\mathbb{P}_f \left(\frac{1}{\sqrt{K_L}} \sum_{(l,k)\in I_L} \left[\left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} Y_{lk}^{(j)} \right)^2 - \nu_L \right] < \kappa_\alpha \left(n\gamma_L^2 \vee 1 \right) \log\log(e \vee |\mathcal{S}|) \right) \leqslant \alpha/4, \quad (105)$$

for $C_\alpha > 0$ large enough. Define $V_{lk} := \left(\frac{1}{\sqrt{m}} \sum_{j=1}^{m} \left[\gamma_L \sum_{i=1}^{n} (U_L X_i^{(j)})_{lk} + W_{lk}^{(j)} \right] \right)^2$. Under $\mathbb{P}^U$, we have that $(U_L f_L)_{lk} \overset{d}{=} \|f_L\|_2 \frac{Z_{lk}}{\|Z\|_2}$, for $Z = (Z_{11}, \ldots, Z_{1d_L}) \sim N(0, I_{d_L})$. As $\sum_{(l,k)\in I_L} \mathbb{E} Z_{lk}^2/\|Z\|_2^2 = 1$, $\mathbb{E} Z_{lk}^2/\|Z\|_2^2 = 1/d_L$ by symmetry. Hence,

$$\mathbb{E}_f V_{lk} = \gamma_L^2 mn^2 \mathbb{E}^U (U_L f_L)_{lk}^2 + n\gamma_L^2 + \mathbb{E}(W_{lk}^{(j)})^2 = \frac{\gamma_L^2 mn^2 \sqrt{K_L} \|f_L\|_2^2}{d_L} + n\gamma_L^2 + \mathbb{E}(W_{lk}^{(j)})^2.$$

Subtracting $d^{-1}\gamma_L^2 mn^2 \sqrt{K_L} \|f_L\|_2^2$ on both sides, the first term in (105) is bounded above by

$$\Pr\left(\frac{1}{\sqrt{K_L}} \sum_{(l,k)\in I_L} [V_{lk} - \mathbb{E}_f V_{lk}] < -\frac{\gamma_L^2 n^2 m \sqrt{K_L} \|f_L\|_2^2}{2d_L} \right)$$

whenever

$$\gamma_L^2 mn^2 \sqrt{K_L} \|f_L\|_2^2 d_L^{-1} \geqslant 2\kappa_\alpha (\gamma_L^2 n \vee 1). \quad (106)$$

The latter is indeed satisfied whenever (104). Under the alternative hypothesis, a straightforward calculation shows that V_l has expectation conditionally on U_L equal to $\gamma_L^2 mn^2 (U_L f)_{lk}^2 + n\gamma_L^2 + \mathbb{E}(W_{lk}^{(j)})^2$. Since $\mathbb{E}(Z_{L;i})_{lk}^4 = 3$ and $\mathbb{E}(W_{lk}^{(j)})^4 \asymp 1$, its variance conditionally on U_L equals

$$\begin{aligned} &n^2\gamma_L^4 \mathrm{Var}\left(\eta^2 \mid |U_L\right) + \mathrm{Var}((\frac{1}{\sqrt{m}} \sum_{j=1}^{m} W_{lk}^{(j)})^2 \mid |U_L) + 2\gamma_L^4 n^3 m (U_L f)_{lk}^2 \mathbb{E}\eta^2 \\ &\quad + 2\gamma_L^2 mn^2 (U_L f)_{lk}^2 \mathbb{E}(W_{lk}^{(j)})^2 + 2\gamma_L^2 n \mathbb{E}(W_{lk}^{(j)})^2 \mathbb{E}\eta^2, \end{aligned}$$

which is of the order $(\gamma_L^4 mn^3 (U_L f)_{lk}^2) \vee (\gamma_L^2 mn^2 (U_L f)_{lk}^2) \vee \gamma_L^4 n^2 \vee 1$. Consequently, by applying Chebyshev's inequality, the probability on the left-hand side of (105) is of the order

$$\left(\frac{d_L}{mn\sqrt{K_L}\|f_L\|_2^2} \right) \vee \left(\frac{d_L \Lambda_{N,\mathcal{S},\delta} \tilde{\kappa}_\alpha^2}{mn^2\epsilon^2 \|f_L\|_2^2} \right) \vee \left(\frac{\kappa_\alpha^2 d_L^2}{m^2 n^2 K_L \|f_L\|_2^4} \right) \vee \left(\frac{\kappa_\alpha^2 \tilde{\kappa}_\alpha^4 \Lambda_{N,\mathcal{S},\delta}^2 d_L^2}{m^2 n^4 \epsilon^4 K_L \|f_L\|_2^4} \right), \quad (107)$$

where $\Lambda_{N,\mathcal{S},\delta} := \log(N)|\mathcal{S}|\log(1/\delta)\log$, and we have plugged in $\gamma_L = \frac{\epsilon}{2\sqrt{K_L|\mathcal{S}|\log(2/\delta)\tau}}$, $\tau = \tilde{\kappa}_\alpha\sqrt{\log(N)}$. The latter can be made arbitrarily small per choice of $C_\alpha > 0$ in (104).

To see this, first consider the case where $\epsilon \leqslant n^{-1/2}$, such that $K_L \asymp 1$ and the above display is of the order of the second and fourth term in the above display. Whenever f satisfies (104), the first term is of the order $1/C_\alpha$. For $1/\sqrt{n} \leqslant \epsilon \leqslant d_L$, $K_L \asymp n\epsilon^2 \wedge d_L \gtrsim 1$ and (104), (107) is bounded by $1/C_\alpha^2$. Lastly, when $n\epsilon^2 \gtrsim d_L$, it holds that $K_L \asymp d_L$, so (107) is of the order $\kappa_\alpha^2 d_L/(m^2n^2\|f_L\|_2^4)$. When (104) holds, the above display is bounded by $1/C_\alpha^2$. We conclude that $\mathbb{P}_f(1-T_{\mathrm{III}}) \leqslant \alpha/4$ for C_α large enough, as desired. ■

B.4 Proofs of the theorems in Section 2 and Theorem 7

The proofs of Theorems 3 and 2 in Section 2 are direct consequences of Theorem 7 and 12. To see this, note that for any sequences $m_N \equiv m$, $n_N \equiv N/m$, $\sigma_N \equiv \sigma$, $\epsilon_N \equiv \epsilon$ and $\delta_N \equiv \delta$, Theorem 12 and Theorem 7 can be applied with an arbitrarily slow decreasing sequence $\alpha_N \to 0$ or $\alpha_N \to 1$ in order to obtain the desired convergence of the minimax risk.

Theorem 7 and Theorem 6 are consequences the lemmas proven earlier in the section. Below, we tie together the results of these lemmas to obtain both theorems.

Proof of Theorem 7: Consider $\mathcal{S} = \{L\}$, for a given $L \in \mathbb{N}$ which is to be determined. In the case of local randomness only protocols, let $T = T_{\mathrm{I}} \vee T_{\mathrm{II}}$, where T_{I} and T_{II} are the tests defined in (90) and (97), respectively, with their critical regions such that $\mathbb{P}_0 T_{\mathrm{I}} \leqslant \alpha/4$ and $\mathbb{P}_0 T_{\mathrm{II}} \leqslant \alpha/4$ (see the first statements of Lemma 30 and Lemma 32). If, for some $L \in \mathbb{N}$, f satisfies

$$\|f_L\|_2^2 \geqslant C_\alpha M_N \sigma^2 \left(\frac{2^{(3/2)L}}{mn(n\epsilon^2 \wedge 2^L)} \bigwedge \left(\frac{\sqrt{2^L}}{\sqrt{m}n\sqrt{n\epsilon^2 \wedge 1}} \bigvee \frac{1}{mn^2\epsilon^2} \right) \right), \tag{108}$$

where $M_N \gtrsim \log(N)\sqrt{2\log\log(N)\log(N)|\mathcal{S}|\log(2/\delta)}$, Lemma 30 and Lemma 32 yield that

$$\mathbb{P}_f(1-T) \leqslant \mathbb{P}_f(1-T_{\mathrm{I}}) + \mathbb{P}_f(1-T_{\mathrm{II}}) \leqslant \alpha/2. \tag{109}$$

In view of $(a+b)^2/2 - b^2 \leqslant a^2$, $\|f_L\|_2^2 \geqslant \frac{\|f\|_2^2}{2} - \|f - f_L\|_2^2$. Since $f \in \mathcal{B}_{p,q}^{s,R}$, we have that $\|f - f_L\|_2^2 \leqslant 2^{-2Ls}R^2$ (see e.g. Lemma 45). Consequently, taking $L = 1 \vee \lceil -\frac{1}{s}\log_2(\rho) \rceil$,

$$\|f_L\|_2^2 \geqslant \rho^2 M_N C_\alpha^2/2 - R^2 2^{-2Ls} \geqslant \rho M_N^2 (C_\alpha^2/2 - R^2).$$

We note here that the latter bound could be sharper by choosing $L = 1 \vee \lceil -\frac{1}{s}\log_2(\rho) \rceil$, but we choose the simpler bound for the sake of clarity. Given the above display, this choice of L means that (108) is satisfied for

$$\rho^2 \gtrsim \sigma^2 \left(\frac{(1 \vee \rho^{-1/s})^{3/2}}{mn(n\epsilon^2 \wedge (1 \vee \rho^{-1/s}))} \bigwedge \left(\frac{\sqrt{1 \vee \rho^{-1/s}}}{\sqrt{m}n\sqrt{n\epsilon^2 \wedge 1}} \bigvee \frac{1}{mn^2\epsilon^2} \right) \right). \tag{110}$$

Solving this for ρ, we obtain the rate in Theorem 7, (8).

In case of shared randomness, the same argument applies, with the only difference that the test $T = T_{\mathrm{I}} \vee T_{\mathrm{III}}$ is used. Lemmas 30 and 34 yield that this test satisfies $\mathbb{P}_0 T \leqslant \alpha/2$ and also $\mathbb{P}_f(1-T) \leqslant \alpha/2$ whenever

$$\rho^2 \gtrsim \sigma^2 \left(\frac{1 \vee \rho^{-1/s}}{mn\sqrt{n\epsilon^2 \wedge 1}\sqrt{n\epsilon^2 \wedge (1 \vee \rho^{-1/s})}} \bigwedge \left(\frac{\sqrt{1 \vee \rho^{-1/s}}}{\sqrt{m}n\sqrt{n\epsilon^2 \wedge 1}} \bigvee \frac{1}{mn^2\epsilon^2} \right) \right)$$

and M_N satisfying the condition of the theorem. Solving the above display for ρ, we obtain the rate in Theorem 7, (10).

Proof of Theorem 6: The proof is similar to that of Theorem 7 in the sense that it follows from the lemmas proven earlier in the section.

In the case of local randomness, consider T as defined in (31), with T_{I} and T_{II} as defined in (29) and (30), respectively. Applying Lemma 30 with $\mathcal{S} = S^{\mathtt{LOW}}$ and Lemma 32 with $\mathcal{S} = S^{\mathtt{HIGH}}$ yields that $\mathbb{P}_0 T_{\mathrm{I}} \leqslant \alpha/4$ and $\mathbb{P}_0 T_{\mathrm{II}} \leqslant \alpha/4$, so T has the correct level.

If $f \in \mathcal{B}^{s,R}_{p,q}$, with $s \in [s_{\min}, s_{\max}]$, we have that L_s is either in $S^{\mathtt{LOW}}$ or $S^{\mathtt{HIGH}}$. By applying Lemma 30 and Lemma 32, a similar argument as given in the proof of Theorem 7 above yields that $\mathbb{P}_f(1-T) \leqslant \alpha/2$. The shared randomness case is analogous, with T defined as in (33).

Appendix C. Auxiliary lemmas and proofs

C.1 Proof of Lemma 26

Since $\mathcal{C}_{L;\tau} = \mathcal{A}_{L;\tau} \cap \mathcal{B}_{L;\tau}$, it suffices to show that $\mathcal{A}^c_{L;\tau}$ and $\mathcal{B}^c_{L;\tau}$ as defined in (81) are small in $\mathbb{P}_f$-probability for a large enough choice of $\tilde{\kappa}_\alpha > 0$.

Define $\eta_\tau := D_\tau n\sqrt{d_L}/8$. For both sets, we proceed via a union bound:

$$\begin{aligned}\mathbb{P}_f\left(X^{(j)} \notin \mathcal{A}_{L;\tau}\right) &= \mathbb{P}_f\left(\exists \mathcal{J} \subset [n],\, |\mathcal{J}| \leqslant K_\tau : \left| \| \sum_{i \in \mathcal{J}} \sigma^{-1} X^{(j)}_{L;i} \|_2^2 - |\mathcal{J}| d_L \right| > |\mathcal{J}| \eta_\tau \right) \\ &\leqslant \sum_{k=1}^{K_\tau} \binom{n}{k} \Pr\left(\left| \|\sigma^{-1}\sqrt{k} f_L - Z\|_2^2 - d_L \right| > \eta_\tau \right) \end{aligned} \tag{111}$$

where we recall that f_L is the projection of f onto its first d_L coordinates and $Z \sim N(0, I_{d_L})$. We have

$$\|\sigma^{-1}\sqrt{k} f_L - Z\|_2^2 = \sigma^{-2} k \|f_L\|_2^2 - 2\sigma^{-1}\sqrt{k} f_L^\top Z + \|Z\|_2^2.$$

Recalling that $K_\tau = \lceil 2\tau D_\tau^{-1} \rceil$, we obtain that

$$K_\tau \leqslant \frac{2\tau n\sqrt{d_L}}{\tilde{\kappa}_\alpha \log(N)(\sqrt{n\sqrt{d_L}\tau} \vee \sqrt{n d_L})} \leqslant \frac{\sqrt{n\sqrt{d_L}\tau}}{\tilde{\kappa}_\alpha \log(N)} \lesssim \frac{\eta_\tau}{\tilde{\kappa}_\alpha \log(N)}. \tag{112}$$

By the assumptions of the lemma ($\sigma^{-2} n \|f_L\|_2^2 \leqslant \tau\sqrt{d_L}/2$ and $\tau \leqslant nR^2/\sqrt{d_L}$), we obtain that $\sigma^{-2} k \|f_L\|_2^2 \leqslant K_\tau R^2$. Consequently, we have that for $\tilde{\kappa}_\alpha > 0$ large enough $\sigma^{-2} K_\tau \|f_L\|_2^2 < \eta_\tau/2$, so it holds that

$$\Pr\left(\|\sigma^{-1}\sqrt{k} f - Z\|_2^2 - d_L > \eta_\tau \right) \leqslant \Pr\left(\|Z\|_2^2 - d_L - 2\sigma^{-1}\sqrt{k} f_L^\top Z > \eta_\tau/2 \right).$$

Using that $\Pr(A \cap B) + \Pr(A \cap B^c) \leqslant \Pr(A') + \Pr(A \cap B^c)$ for $A \cap B \subset A'$, it follows that the latter display is bounded above by

$$\Pr\left(\|Z\|_2^2 - d_L > \eta_\tau/4\right) + \Pr\left(-2\sigma^{-1}\sqrt{k} f_L^\top Z > \eta_\tau/4\right).$$

By a standard concentration argument for Chi-square random variables, e.g. Lemma A.13 in Szabó et al. (2023), the first probability is bounded by $e^{-d\eta_\tau/8}$. Again using $K_\tau \|f_L\|_2^2 < \eta_\tau/2$, the second term is bounded by $e^{-\eta_\tau/32}$, where we note that the second term equals zero in the case that $f = 0$. The bound

$$\Pr\left(\|\sigma^{-1}\sqrt{k} f_L - Z\|_2^2 - d_L < -\eta_\tau\right) \leqslant e^{-d_L\eta_\tau/4} + e^{-\eta_\tau/8}$$

follows by similar reasoning. Combining the above with the elementary bound $\sum_{k=1}^{K_\tau} \binom{n}{k} \leqslant e^{K_\tau \log(n)}$ and (112) means that

$$\mathbb{P}_f\left(X^{(j)} \notin \mathcal{A}_{L;\tau}\right) \leqslant 2\exp\left(K_\tau \log(N) - \frac{(1 + d_L/2)\eta_\tau}{4}\right) \leqslant \alpha/(4mn).$$

Turning our attention to $\mathcal{B}_{L;\tau}$, we find that $\mathbb{P}_f\left(X^{(j)} \notin \mathcal{B}_{L;\tau}\right)$ is equal to

$$\mathbb{P}_f\left(\max_{i\in[n]} \sigma^{-2}\left|\sum_{k\in[n]\setminus\{i\}} \langle X_i^{(j)}, X_k^{(j)}\rangle\right| > \eta_\tau\right) \leqslant n\Pr\Big(\left|\langle \sigma^{-1}f + Z, (n-1)\sigma^{-1}f + \sqrt{n-1}Z'\rangle\right| > \eta_\tau\Big),$$

where Z and Z' are independent $N(0, I_d)$ random vectors. Using another union bound, the above is further bounded by

$$\Pr(\sqrt{n-1}\langle Z, Z'\rangle > \eta_\tau/2 - (n-1)\|\sigma^{-1}f_L\|_2^2) + \Pr((n-1)\langle \sigma^{-1}f, Z'\rangle + \sqrt{n-1}\langle \sigma^{-1}f, Z\rangle > \eta_\tau/2). \tag{113}$$

Using that $n\|\sigma^{-1}f_L\|_2^2 \leqslant \tau\sqrt{d_L}/2$ and $\tau \leqslant nR^2/\sqrt{d_L}$ by assumption of the lemma, we see that

$$n\|\sigma^{-1}f_L\|_2^2 \leqslant \tau\sqrt{d_L}/2 \leqslant \sqrt{n\sqrt{d}\tau}\frac{\sqrt{\sqrt{d}\tau}}{2\sqrt{n}} \leqslant \frac{R}{\log(N)\tilde{\kappa}_\alpha}\eta_\tau. \tag{114}$$

For $\tilde{\kappa}_\alpha > 0$, the latter can be seen to be larger than $\eta_\tau/4$. Consequently, the first term in (113) can be seen to be bounded by

$$\Pr(\sqrt{n-1}\langle Z, Z'\rangle > \eta_\tau/4) \leqslant e^{-\frac{\eta_\tau}{4\sqrt{nd_L}}} \leqslant e^{-\kappa_\alpha \log(N)},$$

where the inequality follows from Proposition 2.7.1 in Vershynin (2018). Since Z and Z' are independent standard Gaussian vectors, the second term in (113) is bounded above by

$$\Pr(\sqrt{2}(n-1)\langle \sigma^{-1}f, Z\rangle > \eta_\tau/2) \leqslant e^{-\frac{\eta_\tau^2}{8n^2\|\sigma^{-2}f_L\|_2^2}}.$$

Since $n\|\sigma^{-1}f_L\|_2^2 \leqslant \tau\sqrt{d_L}/2$, $n^2\|\sigma^{-1}f_L\|_2^2 \leqslant n\tau\sqrt{d_L}/2 \leqslant \frac{\eta_\tau^2}{\log^2(N)\tilde{\kappa}_\alpha^2}$. Hence, we have obtained that $\mathbb{P}_f\left(X^{(j)} \notin \mathcal{B}_{L;\tau}\right) \leqslant \frac{\alpha}{4mn}$, for $\tilde{\kappa}_\alpha > 0$ large enough. This concludes the proof of the lemma.

C.2 Proof of Lemma 25

We suppress the dependence of the Markov kernels on the draw of the shared randomness u, as it is of no relevance to the proof. Consider $\tilde{X}^{(j)} \sim P_\pi$ and $X^{(j)} \sim P_0$. Since $X^{(j)}_{L'k} \overset{d}{=} \tilde{X}^{(j)}_{L'k}$ for all $L' > L$, it suffices to construct couplings for $(\tilde{X}^{(j)}_L, X^{(j)}_L)$. We construct two couplings, one for each of the different regimes of ϵ. That is, for each of the regimes, we derive a joint distribution of $(\tilde{X}^{(j)}_L, X^{(j)}_L)$ called $\mathbb{P}_{\pi,0}$ such that $\tilde{X}^{(j)}_L \sim \mathbb{P}_{\pi,0}^{\tilde{X}^{(j)}_L} = P_\pi$ and $X^{(j)}_L \sim \mathbb{P}_{\pi,0}^{X^{(j)}_L} = P_0$. The specific couplings that we construct aim at assuring $\mathrm{d}_H(\tilde{X}^{(j)}_L, X^{(j)}_L)$ is small with high probability. When constructing the coupling below, it suffices to consider $\sigma = 1$ without loss of generality (i.e. by rescaling $X^{(j)}$ and $\tilde{X}^{(j)}$). After the construction of both the couplings, the lemma follows by an application of Lemma 24.

Case 1: Consider $1/\sqrt{n} \geqslant \epsilon \geqslant 1/\sqrt{mnd_L}$. The construction we follow is similar to that used in Theorem D.6 of Narayanan (2022b), whose dependencies are favorable for our purpose.

If $n = 1$, Pinsker's inequality followed by bounding the KL-divergence by the Chi-square-divergence (e.g. Lemma 2.7 of Tsybakov (2009)) and Lemma 22 applied with $m = 1$ yields that

$$\|P_0 - P_\pi\|_{\mathrm{TV}} \leqslant \sqrt{\frac{1}{2} D_{\chi^2}(P_0; P_\pi)} \leqslant C \frac{\sqrt{c_\alpha}\rho^2}{\sqrt{d_L}}$$

for a universal constant $C > 0$ contingent only on $\|\bar{\Gamma}\|$. By Lemma 39, there exists a coupling $\mathbb{P}_{\pi,0}$ such that $\tilde{X}^{(j)} \sim \mathbb{P}_{\pi,0}^{\tilde{X}^{(j)}_L} = P_\pi$ and $X^{(j)} \sim \mathbb{P}_{\pi,0}^{X^{(j)}_L} = P_0$ and

$$p := \mathbb{P}\left(\tilde{X}^{(j)}_L \neq X^{(j)}_L\right) \leqslant \left(C \frac{\sqrt{c_\alpha}\rho^2}{\sqrt{d_L}}\right) \wedge 1.$$

Applying Lemma 24, it follows that

$$P_\pi^n K^j(A|\tilde{X}^{(j)}_L) = \mathbb{E}_{\pi,0} K^j(A|\tilde{X}^{(j)}_L) \leqslant e^{4\epsilon np} P_0 K^j(A|X^{(j)}_L) + 2\delta np e^{\epsilon np}.$$

By applying condition (50) or (51) and the bound on p, we obtain that $\exp(C\frac{\sqrt{c_\alpha}\epsilon\rho^2}{d_L^{1/2}}) \leqslant 1 + C\sqrt{c_\alpha}/\sqrt{m}$. Similarly, using that $\delta \leqslant \epsilon/\sqrt{m}$, we get $\delta p e^{\epsilon p} \leqslant \delta + C' c_\alpha / m^{3/2}$. The first identity we wish to show, i.e. (75), now follows for $n = 1$ and a sufficiently small enough choice of $c_\alpha > 0$.

For what follows, take $n > 1$. Consider V a uniform draw from the unit sphere in $\mathbb{R}^{d_L}$ and $Z \sim N(0, I_{d_L})$, both independent of the other random variables considered. We have

$$\overline{X^{(j)}_L} := \frac{1}{n}\sum_{i=1}^n X^{(j)}_{L;i} \overset{d}{=} \frac{\|Z\|_2}{\sqrt{n}} V$$

for $X^{(j)}_L \sim \mathbb{P}_0^{X^{(j)}_L}$ (see e.g. Vershynin (2018) Exercise 3.3.7). Similarly,

$$\overline{\tilde{X}^{(j)}_L} \overset{d}{=} V^{-1} \left\| (I_{d_L} + n^{-1/2} n c_\alpha^{1/2} \rho^2 d_L \bar{\Gamma})^{1/2} Z \right\|_2.$$

Next, we note that for $\eta_1, \ldots, \eta_n \sim N(0, I_{d_L})$ independent of $X^{(j)} = (X_1^{(j)}, \ldots, X_n^{(j)})$, we have

$$X_L^{(j)} \stackrel{d}{=} \left(\overline{X^{(j)}} + \eta_i - \frac{1}{n}\sum_{i=1}^n \eta_i \right)_{1 \leqslant i \leqslant n}. \tag{115}$$

To see this, note that both the left- and right-hand side are mean zero Gaussian and

$$\mathbb{E}\left(\overline{X_L^{(j)}} + \eta_i - \frac{1}{n}\sum_{i=1}^n \eta_i \right)\left(\overline{X_L^{(j)}} + \eta_k - \frac{1}{n}\sum_{i=1}^n \eta_i \right)^\top = \mathbb{1}_{i=k} I_{d_L}, \tag{116}$$

which means that the covariances of the left-hand side and right-hand side of (115) are equal too. Noting that $\tilde{X}_L^{(j)} \stackrel{d}{=} (F + X_{L;i}^{(j)})_{i\in[n]}$ and $\overline{\tilde{X}_L^{(j)}} \stackrel{d}{=} F + \overline{X_L^{(j)}}$, where $F \sim N(0, \sqrt{c_\alpha} d_L^{-1} \rho^2 \bar{\Gamma})$ is independent of $\overline{X_L^{(j)}}$, it follows that $\tilde{X}_L^{(j)} \stackrel{d}{=} (\overline{\tilde{X}_L^{(j)}} + \eta_i - \frac{1}{n}\sum_{i=1}^n \eta_i)_{1\leqslant i\leqslant n}$ by similar reasoning. Since the matrix $(I - VV^\top)$ is idempotent, we have that $\eta_i = VV^\top \eta_i + (I - VV^\top)\eta_i$; where $VV^\top \eta_i$ is independent from $(I - VV^\top)\eta_i$ and $V^\top \eta_i$ is standard normally distributed, both conditionally and unconditionally on V. We can write

$$\eta_i - \frac{1}{n}\sum_{i=1}^n \eta_i = VV^\top \eta_i - \frac{1}{n}\sum_{i=1}^n VV^\top \eta_i + G_i,$$

where $G_i \equiv G_i(\eta_i) := (I - VV^\top)\eta_i - \frac{1}{n}\sum_{i=1}^n (I - VV^\top)\eta_i$ and G_i independent of $VV^\top \eta_i - \frac{1}{n}\sum_{i=1}^n VV^\top \eta_i$. Let $\tilde{\eta}_i$ be identically distributed to η_i, for $i = 1, \ldots, n$. Combining the above, we have that

$$X_L^{(j)} \stackrel{d}{=} \left\{ V\left(\frac{\|Z\|_2}{\sqrt{n}} + V^\top \eta_i - \frac{1}{n}\sum_{i=1}^n V^\top \eta_i \right) + G_i \right\}_{i\in[n]} =: (C_i)_{i\in[n]}, \tag{117}$$

$$\tilde{X}_L^{(j)} \stackrel{d}{=} \left\{ V\left(n^{-1/2} \|(I_d + n c_\alpha^{1/2} \rho^2 d_L^{-1} \bar{\Gamma})^{1/2} Z\|_2 + V^\top \tilde{\eta}_i - \frac{1}{n}\sum_{i=1}^n V^\top \tilde{\eta}_i \right) + G_i \right\}_{i\in[n]} =: (\tilde{C}_i)_{i\in[n]}. \tag{118}$$

As further notations, we introduce $\zeta_i := \|Z\|_2/\sqrt{n} + V^\top \eta_i - \frac{1}{n}\sum_{i=1}^n V^\top \eta_i$ and

$$\tilde{\zeta}_i := \|(I_{d_L} + n c_\alpha^{1/2} \rho^2 d_L^{-1} \bar{\Gamma})^{1/2} Z\|_2/\sqrt{n} + V^\top \tilde{\eta}_i - \frac{1}{n}\sum_{i=1}^n V^\top \tilde{\eta}_i.$$

We have that $\zeta_i | Z \sim N\left(\frac{\|Z\|_2}{\sqrt{n}}, \left(1 - \frac{1}{n}\right) \right)$ and

$$\tilde{\zeta}_i | Z \sim N\left(n^{-1/2} \|(I_{d_L} + d^{-1} n c_\alpha^{1/2} \rho^2 \bar{\Gamma})^{1/2} Z\|_2, \left(1 - n^{-1}\right) \right).$$

By e.g. Lemma 40, we find that their respective push forward measures $\mathbb{P}^{\zeta_i|Z}$ and $\mathbb{P}^{\tilde{\zeta}_i|Z}$ satisfy

$$\begin{aligned}\|\mathbb{P}^{\zeta_i|Z} - \mathbb{P}^{\tilde{\zeta}_i|Z}\|_{\mathrm{TV}} &\leqslant \frac{1}{2(1-1/n)\sqrt{n}}|\|Z\|_2 - \|(I_{d_L} + nc_\alpha^{1/2}\rho^2 d_L^{-1}\bar{\Gamma})^{1/2}Z\|_2| \\ &\leqslant \frac{\sqrt{n}c_\alpha^{1/2}\rho^2}{d_L}\frac{Z^\top\bar{\Gamma}Z}{\sqrt{Z^\top I_{d_L}Z} + \sqrt{Z^\top(I_{d_L} + \frac{nc_\alpha^{1/2}\rho^2}{d_L}\bar{\Gamma})Z}} \leqslant \|\bar{\Gamma}\|\frac{\sqrt{n}c_\alpha^{1/2}\rho^2}{d_L}\|Z\|_2,\end{aligned}$$

where the second inequality follows from $n > 1$ in addition to the identity $(\sqrt{a} - \sqrt{b})(\sqrt{a} + \sqrt{b}) = a - b$ and the final inequality follows from the fact that $\bar{\Gamma}$ is positive semidefinite and $\bar{\Gamma} \leqslant \|\bar{\Gamma}\|I_{d_L}$. By Lemma 39, there exists a coupling of $\zeta_i|Z$ and $\tilde{\zeta}_i|Z$ such that

$$\mathbb{P}\left(\zeta_i \neq \tilde{\zeta}_i|Z\right) \leqslant (\sqrt{n}c_\alpha^{1/2}\rho^2\|\bar{\Gamma}\|\|Z\|_2(2d_L)^{-1}) \wedge 1. \tag{119}$$

Take $\mathbb{P}^{\zeta_i,\tilde{\zeta}_i|Z}$ satisfying (119), $G_i = \tilde{G}_i$ and set $(X_L^{(j)}, \tilde{X}_L^{(j)}) = (C, \tilde{C})$ under $\mathbb{P}_{\pi,0}$. We obtain that $(\tilde{C}_i, C_i)$ is independent of $(\tilde{C}_k, C_k)$ for $k \neq i$ and $\tilde{\zeta}_i = \zeta_i$ implies $\tilde{C}_i = C_i$. Consequently, $\sum_{i=1}^n \mathbb{1}\{\tilde{C}_i \neq C_i\} \mid |\nu \sim \mathrm{Bin}(n, p_\nu)$, with $\nu = (Z, n^{-1}\sum_{i=1}^n \eta_i, n^{-1}\sum_{i=1}^n \tilde{\eta}_i)$. By (119), $\|\bar{\Gamma}\| \asymp 1$ and the fact that $\|Z\|_2$ is $\sqrt{d}$-sub-exponential (using e.g. Proposition 2.7.1 in Vershynin (2018)), we obtain that

$$p := \mathbb{P}\left(\tilde{C}_i \neq C_i\right) = \mathbb{E}^Z\mathbb{P}\left(\zeta_i \neq \tilde{\zeta}_i|Z\right) \leqslant (\tilde{C}n^{1/2}(c_\alpha^{1/4}\rho)^2 d_L^{-1/2}) \wedge 1,$$

for a universal constant $\tilde{C} > 0$. Since $[\zeta_i = \tilde{\zeta}_i]$ implies $[C_i = \tilde{C}_i]$, we have that

$$p = \mathbb{P}(C_i \neq \tilde{C}_i) \leqslant \mathbb{E}^Z\mathbb{P}\left(\zeta_i \neq \tilde{\zeta}_i|Z\right) \leqslant \tilde{C}n^{1/2}c_\alpha^{1/4}\rho^2 d_L^{-1/2}.$$

To summarize, we have now obtained that there exists a joint distribution $\mathbb{P}_{\pi,0}$ of $(X^{(j)}, \tilde{X}^{(j)})$ such that $\left(X^{(j)}, \tilde{X}^{(j)}\right)$ satisfy

$$S_\nu := \sum_{i=1}^n \mathbb{1}\{\tilde{X}_i^{(j)} \neq X_i^{(j)}\} \mid |\nu \sim \mathrm{Bin}(n, p_\nu), \text{ with } p = \mathbb{P}\left(X_i^{(j)} \neq \tilde{X}_i^{(j)}\right) \leqslant \frac{\tilde{C}n^{1/2}(c_\alpha^{1/4}\rho)^2}{d_L^{1/2}} \wedge 1.$$

Let $\mathbb{E}_{\pi,0}$ denote its corresponding expectation. Consequently, by applying Lemma 24, we have for any measurable A that

$$\begin{aligned}P_\pi^n K^j(A|\tilde{X}^{(j)}) &= \mathbb{E}_{\pi,0}K^j(A|\tilde{X}^{(j)}) = \mathbb{E}_{\pi,0}^{\tilde{X}^{(j)},X^{(j)}}K^j(A|\tilde{X}^{(j)}) \\ &\leqslant \mathbb{E}^\nu e^{4\epsilon np_\nu}P_0K^j(A|X^{(j)}) + 2\delta n\mathbb{E}^\nu p_\nu e^{2\epsilon np_\nu}.\end{aligned}$$

It follows that $\mathbb{E}^\nu e^{4\epsilon np_\nu} \leqslant 1 + C'\frac{\epsilon n^{3/2}\rho^2}{\sqrt{c_\alpha}d_L^{1/2}}$, for a universal constant $C' > 0$, where the inequality follows from the fact that under the assumptions on ρ^2 (i.e. condition (50) or (51)) that $d_L^{-1/2}\epsilon n^{3/2}c_\alpha^{1/2}\rho^2 \leqslant \sqrt{c_\alpha}/\sqrt{m}$ and a sufficiently small enough choice of $c_\alpha > 0$. Similarly, using that $\delta \leqslant \epsilon/\sqrt{m}$, we have $\mathbb{E}^\nu\delta np_\nu e^{2\epsilon np_\nu} \leqslant \delta + C'c_\alpha/m^{3/2}$.

Case 2: Consider $\epsilon \leqslant 1/\sqrt{mnd_L}$. We will make use of the total variation coupling between $\tilde{X}_i^{(j)} \sim N(f, I_{d_L})$ and $X_i^{(j)} \sim N(0, I_{d_L})$, as given by Lemma 39. Since

$$\|N(0, I_{d_L}) - N(f, I_{d_L})\|_{\mathrm{TV}} \leqslant \frac{1}{2}\|f\|_2 \wedge 2$$

(see e.g. Lemma 40), we can couple the two data sets observation wise independently (simply taking the product space) such that $\sum_{i=1}^{n} \mathbb{1}\{\tilde{X}_i^{(j)} \neq X_i^{(j)}\}|f \sim \mathrm{Bin}(n, p_f)$, where $p_f = (\|f\|_2/4) \wedge 1$. Given $k \in \mathbb{N}$, $\|f\|_2 \overset{d}{=} d_L^{-1/2} c_\alpha^{1/4} \rho \|N(0, I_{d_L})\|_2$ and $\|N(0, I_{d_L})\|_2$ is $\sqrt{d_L}$-sub-exponential we obtain (using e.g. Proposition 2.7.1 in Vershynin (2018))

$$\int p_f^k d\pi(f) \leqslant \int (\|f\|_2/4)^k d\pi(f) \leqslant \tilde{C}^k k^k (c_\alpha^{1/4}\rho)^k,$$

for a universal constant $\tilde{C} > 0$. The assumed condition on ρ yields $\epsilon n c_\alpha^{1/4} \rho \leqslant c_\alpha^{1/4}/\sqrt{m}$, which by similar arguments as before implies

$$e^{4\epsilon np} \leqslant 1 + C' c_\alpha^{1/4}/\sqrt{m} \quad \text{and} \quad \delta np e^{2\epsilon np} \leqslant \delta + C' c_\alpha^{1/2}/m^{3/2},$$

for a universal constant $C' > 0$. By applying the claim at the start of the lemma and using the assumptions on ρ, we obtain that

$$\begin{aligned} P_\pi^n K^j(A|\tilde{X}^{(j)}) = \mathbb{E}_{\pi,0} K^j(A|\tilde{X}^{(j)}) &= \int \mathbb{E}_{f,0} K^j(A|\tilde{X}^{(j)}) d\pi(f) \\ &\leqslant (1 + C c_\alpha^{1/4}/\sqrt{m}) P_0 K^j(A|X^{(j)}) + \delta + C c_\alpha^{1/2}/m^{3/2} \end{aligned}$$

as desired. Again, (76) follows by similar steps.

C.3 Proof of Lemma 24

Let $\mathbb{E}$ denote expectation with respect to $\mathbb{P}$ and write $D = (D_i)_{i\in[n]}$, $S := \sum_{i=1}^n D_i$. We start by noting that

$$\mathbb{E}\left[K(A|\tilde{X})|S = 0\right] = \mathbb{E}\left[K(A|X)|S = 0\right]. \tag{120}$$

Next, we show that for all $k \in [n]$,

$$e^{-\epsilon}\mathbb{E}\left[K(A|\tilde{X})|S = k-1\right] - \delta \leqslant \mathbb{E}\left[K(A|\tilde{X})|S = k\right] \leqslant e^{\epsilon}\mathbb{E}\left[K(A|\tilde{X})|S = k-1\right] + \delta. \tag{121}$$

Write $v_{(-i)} = (v_i)_{[n]\setminus\{i\}}$ for a vector $v \in \mathbb{R}^n$. Let $k \in [n]$ be given and let $\mathcal{V}_k$ denote the set of $v \in \{0,1\}^n$'s such that $\sum_{i=1}^n v_i = k$. Using the definition of DP, the integrand in the conditional expectation satisfies

$$e^{-\epsilon}K(A|\tilde{X}_1, \ldots, \breve{X}_i, \ldots, \tilde{X}_n) - \delta \leqslant K(A|\tilde{X}) \leqslant e^{\epsilon}K(A|\tilde{X}_1, \ldots, \breve{X}_i, \ldots, \tilde{X}_n) + \delta, \tag{122}$$

for any random variable $\breve{X}_i$ taking values in the sample space of $\tilde{X}_i$. In particular, if $v_i = 1$ it holds that

$$\begin{aligned} \mathbb{E}\left[K(A|x_1, \ldots, X_i, \ldots, x_n)\middle|D_i = v_i, \tilde{X}_{(-i)} = x_{(-i)}, D_{(-i)} = v_{(-i)}\right] \leqslant \\ e^{\epsilon}\mathbb{E}\left[K(A|x_1, \ldots, X_i, \ldots, x_n)\middle|D_i = 0, \tilde{X}_{(-i)} = x_{(-i)}, D_{(-i)} = v_{(-i)}\right] + \delta, \end{aligned}$$

for all x in the sample space of $\tilde{X}$. It follows by the law of total probability that

$$\mathbb{E}\left[K(A|\tilde{X})|D=v\right] \leqslant e^{\epsilon}\mathbb{E}\left[K(A|\tilde{X})|D_i=0,\, D_k=v_k \text{ for } k\in[n]\backslash\{i\}\right]+\delta,$$

for all $i\in[n]$. The event $\{D=v\}$ equals the event $\{D=v,\, S=k\}$, and similarly it holds that

$$\{D_k=v_k \text{ for } k\in[n]\backslash\{i\},\, D_i=0\} = \{D_k=v_k \text{ for } k\in[n]\backslash\{i\},\, D_i=0,\, S=k-1\}\,.$$

Consider now the sets

$$\begin{aligned}\mathcal{V}_{k-1}(v) &:= \left\{v'\in\mathcal{V}_{k-1} \,:\, v_k=v'_k \text{ except for one } i\in[n]\right\} \text{ for } v\in\mathcal{V}_k,\\ \mathcal{V}_k(v') &:= \left\{v\in\mathcal{V}_k \,:\, v_k=v'_k \text{ except for one } i\in[n]\right\} \text{ for } v'\in\mathcal{V}_{k-1}.\end{aligned}$$

By what we have derived so far, it holds that any $v\in\mathcal{V}_k$ and $v'\in\mathcal{V}_{k-1}(v)$,

$$\mathbb{E}\left[K(A|\tilde{X})|D=v,\, S=k\right] \leqslant e^{\epsilon}\mathbb{E}\left[K(A|\tilde{X})|D=v',\, S=k-1\right]+\delta.$$

Consider $\{I_k(v) : v\in\mathcal{V}_k\}$ independent random variables (on a possibly enlarged probability space) taking values in $[n]$ such that $\mathbb{P}(I_k(v)=i)=1/k$ whenever $v_i=1$. Combining the above with the total law of probability we find that

$$\begin{aligned}&\mathbb{E}\left[K(A|\tilde{X})|S=k\right] =\\ &\frac{1}{\binom{n}{k}}\sum_{v\in\mathcal{V}_k}\mathbb{E}\left[K(A|\tilde{X})|D=v,\, S=k\right] \leqslant\\ &e^{\epsilon}\frac{1}{\binom{n}{k}}\sum_{v\in\mathcal{V}_k}\mathbb{E}\left[K(A|\tilde{X})|D_{I(v)}=0,\, D_{-I(v)}=v_{-I(v)},\, S=k-1\right]+\delta =\\ &e^{\epsilon}\frac{1}{\binom{n}{k}}\frac{1}{k}\sum_{v\in\mathcal{V}_k}\sum_{v'\in\mathcal{V}_{k-1}(v)}\mathbb{E}\left[K(A|\tilde{X})|D=v',\, S=k-1\right]+\delta =\\ &e^{\epsilon}\frac{1}{\binom{n}{k}}\frac{1}{k}\sum_{v'\in\mathcal{V}_{k-1}}\sum_{v\in\mathcal{V}_k(v')}\mathbb{E}\left[K(A|\tilde{X})|D=v',\, S=k-1\right]+\delta =\\ &e^{\epsilon}\frac{1}{\binom{n}{k-1}}\sum_{v'\in\mathcal{V}_{k-1}}\mathbb{E}\left[K(A|\tilde{X})|D=v',\, S=k-1\right]+\delta =\\ &e^{\epsilon}\mathbb{E}\left[K(A|\tilde{X})|S=k-1\right]+\delta,\end{aligned}$$

where it is used that $|\mathcal{V}_k|=\binom{n}{k}$,

$$\mathbb{P}(D_1=v_1,\dots,D_n=v_n|S=k) = \mathbb{P}(D_1=\tilde{v}_1,\dots,D_n=\tilde{v}_n|S=k)$$

for all $v=(v_i)_{i\in[n]}, \tilde{v}=(\tilde{v}_i)_{i\in[n]}\in\mathcal{V}_k$ and for any $v'\in\mathcal{V}_{k-1}$ there are $n-k+1$ ways to obtain $v\in\mathcal{V}_k$ such that $v_k=v'_k$ except for one $i\in[n]$.

By applying the privacy lower bound of (122) and repeating the same steps, we also find that

$$e^{-\epsilon}\mathbb{E}\left[K(A|\tilde{X})|S=k-1\right]-\delta \leqslant \mathbb{E}\left[K(A|\tilde{X})|S=k\right].$$

This proves (121), which, applying iteratively, results in the bound

$$e^{-\epsilon k}\mathbb{E}\left[K(A|\tilde{X})|S=0\right]-\delta k\leqslant\mathbb{E}\left[K(A|\tilde{X})|S=k\right]\leqslant e^{\epsilon k}\mathbb{E}\left[K(A|\tilde{X})|S=0\right]+\delta k e^{\epsilon k}, \tag{123}$$

for $k=0,1,\ldots,n$. By symmetry of the argument, the same inequalities hold for X in place of $\tilde{X}$. Using the above inequalities, we can bound

$$P_1K(A|\tilde{X})=\mathbb{E}K(A|\tilde{X})=\mathbb{E}^S\mathbb{E}\left[K(A|\tilde{X})|S\right],$$

by $\mathbb{E}^S e^{S\epsilon}\mathbb{E}\left[K(A|\tilde{X})|S=0\right]+\delta\mathbb{E}Se^{S\epsilon}$.. Similarly, applying (121) with X in place of $\tilde{X}$, we find

$$P_2K(A|X)=\mathbb{E}^S\mathbb{E}\left[K(A|X)|S\right]\geqslant\mathbb{E}^S e^{-S\epsilon}\mathbb{E}\left[K(A|X)|S=0\right]-\mathbb{E}\delta S.$$

Combining the two inequalities with (120), we obtain that

$$P_1K(A|\tilde{X})\leqslant\frac{\mathbb{E}^S e^{S\epsilon}}{\mathbb{E}^S e^{-S\epsilon}}\left(P_2K(A|X)+\mathbb{E}^S\delta S\right)+\delta\mathbb{E}Se^{S\epsilon}. \tag{124}$$

In view of the moment generating function of the binomial distribution,

$$\frac{\mathbb{E}^S e^{S\epsilon}}{\mathbb{E}^S e^{-S\epsilon}}=\left(\frac{1+p(e^{\epsilon}-1)}{1+p(e^{-\epsilon}-1)}\right)^n\leqslant e^{4np\epsilon},$$

where the inequality follows from $0\leqslant\epsilon,p\leqslant 1$, the inequality $e^x-e^{-x}\leqslant 3x$ for $0\leqslant x\leqslant 1$ and Taylor expanding $\log(1+x)=x-x^2/2+\ldots$. By Chebyshev's association inequality (e.g. Theorem 2.14 in Boucheron et al. (2013)), $\mathbb{E}^S S\mathbb{E}^S e^{S\epsilon}\leqslant\mathbb{E}^S Se^{S\epsilon}$. Consequently, using the nonnegativity of S,

$$\delta\left(\frac{\mathbb{E}^S e^{S\epsilon}}{\mathbb{E}^S e^{-S\epsilon}}\mathbb{E}^S S+\mathbb{E}Se^{S\epsilon}\right)\leqslant 2\delta\mathbb{E}Se^{S\epsilon}.$$

Lemma 41 (a straightforward calculation) now finishes the proof.

C.4 Proof of Lemma 21

Proof Consider without loss of generality $\sigma=1$ (the general result follows by the σ^{-1} rescaling). The bound $\mathrm{Tr}(\Xi_u^j)\leqslant nd_L$ follows by the fact that conditional expectation contracts the L_2-norm; let $v\in\mathbb{R}^{d_L}$, then

$$v^\top\Xi_u^j v=\mathbb{E}_0^{Y^{(j)}}\mathbb{E}_0^{Y|U=u}[v^\top(\sum_{i=1}^n X_i^{(j)})\mid|Y^{(j)},U=u]^2.$$

Since the conditional expectation contracts the L_2-norm, we obtain that the latter is bounded by

$$\mathbb{E}_0 v^\top(\sum_{i=1}^n X_i^{(j)})(\sum_{i=1}^n X_i^{(j)})^\top v=n\|v\|_2^2,$$

which completes the proof of the statement "$\Xi_u^j \leqslant nI_{d_L}$" and "$\mathrm{Tr}(\Xi_u^j) \leqslant n\overline{d_L}$".

For second other bound on the trace, we start introducing the notations $\overline{X_L^{(j)}} = n^{-1}\sum_{i=1}^n X_{L;i}^{(j)}$ and

$$G_i = \langle \mathbb{E}_0[n\overline{X^{(j)}}_L \mid |Y^{(j)}, U=u], X_{L;i}^{(j)} \rangle.$$

For the remainder of the proof, consider versions of $X_L^{(j)}$ and $Y^{(j)}$ defined on the same probability given $U=u$, and we shall write as a shorthand

$$\mathbb{P}^j \equiv \mathbb{P}_0^{(X^{(j)},Y^{(j)})|U=u} \text{ and } \mathbb{E}^j \equiv \mathbb{E}_0^{(X^{(j)},Y^{(j)})|U=u}.$$

Consider random variables V, W defined on the same probability space. It holds that $\mathbb{E}W\mathbb{E}[W|V] = \mathbb{E}\mathbb{E}[W|V]\mathbb{E}[W|V]$, since $W - \mathbb{E}[W|V]$ is orthogonal to $\mathbb{E}[W|V]$. Combining this fact with the linearity of the inner product and conditional expectation, we see that

$$\mathrm{Tr}(\Xi_u^j) = \mathbb{E}_0^{Y^{(j)}|U=u} \left\| \mathbb{E}_0[n\overline{X^{(j)}}|Y^{(j)}, U=u] \right\|_2^2 = \sum_{i=1}^n \mathbb{E}^j G_i. \tag{125}$$

Define also $\breve{G}_i = \left\langle \mathbb{E}_0[n\overline{X^{(j)}}|Y^{(j)}, U=u], \breve{X}_i^{(j)} \right\rangle$, where $\breve{X}_i^{(j)}$ is an independent copy of $X_i^{(j)}$ (defined on the same, possibly enlarged probability space) and note that $\mathbb{E}^j \breve{G}_i = 0$. Write $G_i^+ := 0 \vee G_i$ and $G_i^- = -(0 \wedge G_i)$. We have

$$\begin{aligned}
\mathbb{E}^j G_i^+ &= \int_0^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) dt = \int_0^T \mathbb{P}^j\left(G_i^+ \geqslant t\right) + \int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) \\
&\leqslant e^\epsilon \int_0^T \mathbb{P}^j\left(\breve{G}_i^+ \geqslant t\right) dt + T\delta + \int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) \\
&\leqslant \int_0^T \mathbb{P}^j\left(\breve{G}_i^+ \geqslant t\right) dt + 2\epsilon \int_0^T \mathbb{P}^j\left(\breve{G}_i^+ \geqslant t\right) dt + T\delta + \int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) \\
&\leqslant \int_0^\infty \mathbb{P}_0^j\left(\breve{G}_i^+ \geqslant t\right) dt + 2\epsilon \int_0^\infty \mathbb{P}^j\left(\breve{G}_i^+ \geqslant t\right) dt + T\delta + \int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right),
\end{aligned}$$

where in the second to last inequality follows by Taylor expansion and the fact that $\epsilon \leqslant 1$. Similarly, we obtain

$$\begin{aligned}
\mathbb{E}^j G_i^- &\geqslant \int_0^T \mathbb{P}^j\left(G_i^- \geqslant t\right) \geqslant e^{-\epsilon} \int_0^T \mathbb{P}^j\left(\breve{G}_i^- \geqslant t\right) dt - T\delta \\
&\geqslant \int_0^\infty \mathbb{P}^j\left(\breve{G}_i^- \geqslant t\right) dt - 2\epsilon \int_0^\infty \mathbb{P}^j\left(\breve{G}_i^- \geqslant t\right) dt - T\delta - \int_T^\infty \mathbb{P}^j\left(\breve{G}_i^- \geqslant t\right) dt.
\end{aligned}$$

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

$$\mathbb{E}^j G_i \leqslant \mathbb{E}^j \breve{G}_i + 2\epsilon \mathbb{E}^j|\breve{G}_i| + 2T\delta + \int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) dt + \int_T^\infty \mathbb{P}^j\left(\breve{G}_i^- \geqslant t\right) dt. \tag{126}$$

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

$$\breve{G}_i \Big| \left[Y^{(j)}, X^{(j)}, U=u\right] \sim N(0, \|\mathbb{E}_0[n\overline{X^{(j)}}|Y^{(j)}, U=u]\|_2^2),$$

so by Cauchy–Schwarz, $\mathbb{E}^j|\breve{G}_i| = \mathbb{E}\|\mathbb{E}[n\overline{X^{(j)}}|Y^{(j)}, U = u]\|_2 \leqslant \sqrt{\mathrm{Tr}(\Xi_u^j)}$. To bound the last two terms in (126) we shall employ tail bounds, which follow after showing that G_i is $\sqrt{d_L n}$-sub-exponential. To see this, note that by Jensen's inequality followed by the law of total probability,

$$\begin{aligned}\mathbb{E}^j e^{t|G_i|} &= \int\int e^{t|\langle \mathbb{E}_0[n\overline{X_L^{(j)}}|Y^{(j)}=y,U=u],x\rangle|} d\mathbb{P}_0^{X^{(j)}|Y^{(j)}=y,U=u}(x) d\mathbb{P}_0^{Y^{(j)}|U=u}(y) \\ &\leqslant \int\int \mathbb{E}_0\left[e^{t|\langle n\overline{X_L^{(j)}},x\rangle|}|Y^{(j)}=y,U=u\right] d\mathbb{P}_0^{X^{(j)}|Y^{(j)}=y,U=u}(x) d\mathbb{P}_0^{Y^{(j)}|U=u}(y).\end{aligned}$$

Since conditional expectation contracts the L_1-norm and using that U is independent of $X^{(j)}$, latter is further bounded by $\mathbb{E}_0^{X^{(j)}} e^{t|\langle n\overline{X^{(j)}}_L, X_{L;i}^{(j)}\rangle|}$. Next, we bound $\mathbb{E}_0^{X^{(j)}} e^{t|\langle n\overline{X_L^{(j)}}, X_{L;i}^{(j)}\rangle|}$.

By the triangle inequality and independence,

$$\mathbb{E}_0^{X^{(j)}} e^{t|\langle n\overline{X_L^{(j)}}, X_{L;i}^{(j)}\rangle|} \leqslant \mathbb{E}_0^{X^{(j)}} e^{t|\langle \sum_{k\neq i}^n X_{L;k}^{(j)}, X_{L;i}^{(j)}\rangle|} \mathbb{E}_0^{X_i^{(j)}} e^{t|\langle X_{L;i}^{(j)}, X_{L;i}^{(j)}\rangle|}.$$

The random variable $\langle X_{L;i}^{(j)}, X_{L;i}^{(j)}\rangle$ is $\chi^2_{d_L}$-distributed, so by standard computations (see e.g. Lemma 12 in Szabó et al. (2022)) we obtain that

$$\mathbb{E}_0^{X_i^{(j)}} e^{t|\langle X_{L;i}^{(j)}, X_{L;i}^{(j)}\rangle|} = \left(\mathbb{E}e^{tN(0,1)^2\rangle|}\right)^{d_L} \leqslant e^{td_L + 2t^2 d_L},$$

whenever $t \leqslant 1/4$. A straightforward computation shows that

$$\mathbb{E}_0^{X^{(j)}} e^{t|\langle \sum_{k\neq i}^n X_{L;k}^{(j)}, X_{L;i}^{(j)}\rangle|} = \left(\mathbb{E}e^{\frac{t^2(n-1)}{2}N(0,1)^2}\right)^{d_L} \leqslant e^{\frac{1}{2}(t^2(n-1)d_L + t^4(n-1)^2 d_L)},$$

where the inequality follows again by e.g. Lemma 12 in Szabó et al. (2022) if $t^2(n-1)^2 \leqslant 1/2$. By the fact that $G_i^+ \leqslant |G_i|$ and Markov's inequality,

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

Combined with the above bounds for the moment generating function means that for $\delta = 0$, the result follows from letting $T \to \infty$. If $\delta > 0$, take $T = 8(d + L \vee \sqrt{nd_L})\log(1/\delta)$ to obtain that $\int_T^\infty \mathbb{P}^j\left(G_i^+ \geqslant t\right) dt \leqslant e^{-\log(1/\delta)}$.

It is easy to see that the same bound applies to $\int_T^\infty \mathbb{P}_0^j\left(\breve{G}_i^- \geqslant t\right) dt$. We obtain that

$$\sum_{i=1}^n \mathbb{E}^j G_i \leqslant 2n\epsilon\sqrt{\mathrm{Tr}(\Xi_u^j)} + 16\delta(d_L \vee \sqrt{nd_L})\log(1/\delta) + 2n\delta.$$

If $\sqrt{\mathrm{Tr}(\Xi_u^j)} \leqslant n\epsilon$, the lemma holds (there is nothing to prove). So assume instead that $\sqrt{\mathrm{Tr}(\Xi_u^j)} \geqslant n\epsilon$. Combining the above display with (125), we get

$$\sqrt{\mathrm{Tr}(\Xi_u^j)} \leqslant 2n\epsilon + 16\delta\frac{d_L \vee \sqrt{nd_L}}{n\epsilon}\log(1/\delta) + \frac{2}{\epsilon}\delta.$$

Since $x^p \log(1/x)$ tends to 0 as $x \to 0$ for any $p > 0$, the result follows for $\delta \leqslant \left(\left(\frac{n}{d_L} \wedge \frac{n^{1/2}}{\sqrt{d_L}}\right)\epsilon^2\right)^{1+p}$ for some $p > 0$ as this implies that the last two terms are $O(n\epsilon)$. ∎

C.5 Proof of Lemma 28

The case where $\mathcal{C}$ is empty is trivial, so we shall assume $\mathcal{C}$ to be nonempty. Consider the map $\breve{S} : (\mathbb{R}^{d_L})^n \to [-\infty, \infty)$ given by

$$\breve{S}(x) = \begin{cases} S(x) \ \text{ if } x \in \mathcal{C}, \\ \inf\{S(c) + D\mathrm{d}_H(c, x) \,:\, c \in \mathcal{C}\} \ \text{ otherwise.} \end{cases}$$

Fix any $c' \in \mathcal{C}$. Since S is D-Lipschitz with respect to the Hamming distance, we have for all $c \in \mathcal{C}$ that

$$S(c) + D\mathrm{d}_H(c, x) \geqslant S(c') - D\mathrm{d}_H(c', c) + D\mathrm{d}_H(c, x) \geqslant S(c') - D\mathrm{d}_H(c', x) > -\infty$$

where the last step follows from the triangle inequality. So, $\breve{S}$ is real valued. For all $x \in (\mathbb{R}^{d_L})^n$ and $\gamma > 0$, there exists $c_\gamma \in \mathcal{C}$ such that $\breve{S}(x) \geqslant S(c_\gamma) + D\mathrm{d}_H(c_\gamma, x) - \gamma$. So for $x, x' \in \ell_2(\mathbb{N})^n$,

$$\breve{S}(x') - \breve{S}(x) \leqslant \breve{S}(c_\gamma) + D\mathrm{d}_H(c_\gamma, x') - \breve{S}(c_\gamma) - D\mathrm{d}_H(c_\gamma, x) + \gamma \leqslant D\mathrm{d}_H(x, x') + \gamma.$$

By symmetry of the argument and since $\gamma > 0$ is given arbitrarily, we conclude that $\breve{S}$ is D-Lipschitz with respect to the Hamming distance. Note however, that this construction does not guarantee that $\breve{S}$ is measurable.

For any map $H : (\mathbb{R}^{d_L})^n \to [-\infty, \infty]$, let H^* denote its minimal Borel-measurable majorant. That is, a measurable map $H^* : (\mathbb{R}^{d_L})^n \to [-\infty, \infty]$ such that

1. $H \leqslant H^*$ and
2. $H^* \leqslant T$ $\mathbb{P}_0$-a.s. for every measurable $T : (\mathbb{R}^{d_L})^n \to [-\infty, \infty]$ with $T \geqslant H$.

Such a map exists by e.g. Lemma 1.2.1 in Van Der Vaart and Wellner (1996). The map $\tilde{S} : (\mathbb{R}^{d_L})^n \to \mathbb{R}$ defined by

$$\tilde{S}(x) = \breve{S}^*(x)\mathbb{1}_{x \notin \mathcal{C}} + S(x)\mathbb{1}_{x \in \mathcal{C}}$$

is measurable and can be seen to be a Borel-measurable majorant of $\breve{S}$; following from the fact that sums and products of measurable functions are measurable, $\breve{S} \leqslant \breve{S}^*$ and $\breve{S}(x) = S(x)$ for $x \in \mathcal{C}$.

Furthermore, by combining the fact that $\tilde{S}$ is measurable with e.g. Lemma 1.2.2 in Van Der Vaart and Wellner (1996), we get

$$|\tilde{S}(x) - \tilde{S}(x')| = |(\tilde{S}(x) - \tilde{S}(x'))^*| \leqslant |\breve{S}(x) - \breve{S}(x')|^*, \tag{127}$$

where $(x, x') \mapsto |\breve{S}(x) - \breve{S}(x')|^*$ is minimal Borel-measurable majorant of $(x, x') \mapsto |\breve{S}(x) - \breve{S}(x')|$. Since $\breve{S}$ is D-Lipschitz with respect to the Hamming distance $(x, x') \mapsto \mathrm{d}_H(x, x')$, which is a measurable map, $|\breve{S}(x) - \breve{S}(x')|^* \leqslant D\mathrm{d}_H(x, x')$. From (127) it follows that for all $x, x' \in (\mathbb{R}^{d_L})^n$, $|\tilde{S}(x) - \tilde{S}(x')| \leqslant D\mathrm{d}_H(x, x')$. We have obtained a map $\tilde{S}$ that is D-Lipschitz with respect to the Hamming distance, measurable and $\tilde{S} = S$ on $\mathcal{C}$, concluding the proof.

C.6 General privacy related lemmas

A random variable V is called *ν-sub-gaussian* if $\mathbb{P}(|V| \geqslant t) \leqslant 2e^{-t^2/\nu^2}$, for all $t > 0$. It is well known (see e.g. Vershynin (2018)) that if $\mathbb{E}V = 0$, the above inequality holds if and only if $\mathbb{E}e^{tV} \leqslant e^{C\nu^2 t^2}$ for all $t \geqslant 0$ and a fixed constant $C > 0$. A random variable V is called *ν-sub-exponential* if $\mathbb{P}(|V| \geqslant t) \leqslant 2e^{-t/\nu}$, for all $t > 0$. If $\mathbb{E}V = 0$, the above inequality holds if and only if $\mathbb{E}e^{tV} \leqslant e^{C\nu^2 t^2}$ for all $0 \leqslant t \leqslant 1/(c\nu)$, with constants $c, C > 0$.

The following lemma shows that clipping symmetric, mean zero random variables preserves sub-gaussianity and sub-exponentiality.

Lemma 35. *Let $V_1, \ldots, V_m$ denote independent random variables, each symmetrically distributed around zero. If $\sum_{j=1}^m V_j$ are sub-gaussian (resp. sub-exponential) with parameter ν, then so are the random variables $\sum_{j=1}^m [V_j]_{-\tau}^{\tau}$, for any $\tau > 0$.*

Proof For any symmetric about 0 function $g : \mathbb{R} \to \mathbb{R}$ such that $x \mapsto g(x)$ is increasing on $[0, \infty)$, it holds that

$$g\left([x]_{-\tau}^{\tau}\right) \leqslant g\left(x\right), \quad \text{for all } \tau > 0. \tag{128}$$

Since V_j is symmetric about zero, so is $[V_j]_{-\tau}^{\tau}$. For an independent Rademacher random variable R_j, we have by the afformentioned symmetry about zero that $[V_j]_{-\tau}^{\tau} \overset{d}{=} R_j[V_j]_{-\tau}^{\tau}$ and consequently

$$\mathbb{E}e^{t[V_j]_{-\tau}^{\tau}} = \mathbb{E}e^{tR_j[V_j]_{-\tau}^{\tau}} = \mathbb{E}\cosh\left(t[V_j]_{-\tau}^{\tau}\right).$$

Using the fact that $V_1, \ldots, V_m$ are independent, we obtain that

$$\mathbb{E}e^{t\sum_{j=1}^m V_j} = \prod_{j=1}^m \mathbb{E}\cosh\left(t[V_j]_{-\tau}^{\tau}\right) \leqslant \prod_{j=1}^m \mathbb{E}\cosh\left(tV_j\right),$$

where the inequality follows from (128). The conclusion can now be drawn from the moment generating function characterization of sub-gaussianity (resp. sub-exponentiality). ■

The next lemma gives a lower bound on the expectation of a clipped random variable that is symmetric around a real number $\mu > 0$.

Lemma 36. *Let $\tau, \mu > 0$ satisfy $\tau/4 \leqslant \mu \leqslant \tau/2$, let V be a random variable symmetric about zero ($V \overset{d}{=} -V$) with Lebesgue density bounded by $M > 0$ and $Pr(|V| \leqslant \frac{1}{12M} \vee (\tau/2)) \geqslant c$ for some constant $c > 0$. It then holds that*

$$\mathbb{E}\left[\mu + V\right]_{-\tau}^{\tau} \geqslant (c \wedge 1/2)\mu. \tag{129}$$

Proof By definition of clipping, $\mathbb{E}\left[\mu + V\right]_{-\tau}^{\tau} = \mathbb{E}\left[V\right]_{-\tau-\mu}^{\tau-\mu} + \mu$. The first term equals

$$\begin{aligned}
&\mathbb{E}\left[V\right]_{-(\tau-\mu)}^{\tau-\mu} + \mathbb{E}\mathbb{1}_{\{V \in [-\tau-\mu, -\tau+\mu]\}}\left(\left[V\right]_{-\tau-\mu}^{-\tau+\mu} + (\tau - \mu)\right) \geqslant \\
&\mathbb{E}\left[V\right]_{-(\tau-\mu)}^{\tau-\mu} - (\tau+\mu)\mathrm{Pr}\left(-\tau-\mu \leqslant V \leqslant -\tau+\mu\right) = \\
&-(\tau+\mu)\mathrm{Pr}\left(-\tau-\mu \leqslant V \leqslant -\tau+\mu\right),
\end{aligned}$$

where the last equality follows from the symmetry of V. By the condition on the Lebesgue density of V, the second term in the above display can be further bounded from below by $-2M(\tau+\mu)\mu$. When $3M\tau < 1/2$, we obtain (129) with the constant $1/2$. Assume $3M\tau \geqslant 1/2$. Then, since $\mu > 0$ and V is symmetric about zero,

$$\begin{aligned}\mathbb{E}\left[\mu+V\right]_{-\tau}^{\tau} &= \mathbb{E}\left(\mu+V\right)\mathbb{1}\{|\mu+V| \leqslant \tau\} + \tau\left(\Pr\left(\mu+V>\tau\right) - \Pr\left(\mu+V<\tau\right)\right) \\ &\geqslant \mathbb{E}\left(\mu+V\right)\mathbb{1}\{|V| \leqslant \tau-\mu\} \geqslant \mathbb{E}\left(\mu+V\right)\mathbb{1}\{|V| \leqslant \tau/2\},\end{aligned}$$

where the last inequality follows from $\mu \leqslant \tau/2$. From the symmetry of V about zero, the right-hand side equals $\mu\Pr\left(|V| \leqslant (1/12M) \vee (\tau/2)\right)$. ∎

C.7 General auxiliary lemmas

The following lemmas are well known but included for completeness.

The following lemma is a standard tail bound for the maximum of sub-gaussian (resp. sub-exponential) random variables that is used to control the type I error of the adaptive tests. For a proof, see Vershynin (2018).

Lemma 37. *Consider a sequence of subsets $\mathcal{S}_n \subset \mathbb{N}$ and let $S_n(L)$ be ν-sub-Gaussian. Then,*

$$\Pr\Big(\max_{L\in\mathcal{S}_n}|S_n(L)| \geqslant t_n\Big) \leqslant 2|\mathcal{S}_n|\,\exp\Big(-\frac{t_n^2}{2\nu^2}\Big).$$

If, instead, $S_n(L)$ is ν-sub-exponential (with sub-exponentiality parameter ν), then for all $t_n > 0$,

$$\Pr\Big(\max_{L\in\mathcal{S}_n}|S_n(L)| \geqslant t_n\Big) \leqslant 2|\mathcal{S}_n| \cdot \exp\Big(-\min\Big\{\tfrac{t_n^2}{2\nu^2}, \tfrac{t_n}{2\nu}\Big\}\Big).$$

Consider the following formal definition of coupling.

Definition 38. *Consider probability measures P and Q on a measurable space $(\mathcal{X},\mathscr{X})$. A* coupling *of P and Q is any probability measure $\mathbb{P}$ on $(\mathcal{X}\times\mathcal{X}, \mathscr{X}\otimes\mathscr{X})$ such that $\mathbb{P}$ has marginals P and Q; $P = \mathbb{P}\circ\pi_1^{-1}$, $Q = \mathbb{P}\circ\pi_2^{-1}$, where $\pi_i : \mathcal{X}\times\mathcal{X} \to \mathcal{X}$ is the projection onto the i-th coordinate (i.e. $\pi_i(x_1,x_2) = x_i$ for $i = 1,2$).*

Lemma 39 below is a well known result showing that, for random variables X and $\tilde{X}$ defined on a Polish space, small total variation distance between their corresponding laws guarantees the existence of a coupling such that they are equal with high probability.

Lemma 39. *For any two probability measures P and Q on a measurable space $(\mathcal{X},\mathscr{X})$ with $\mathcal{X}$ a Polish space and $\mathscr{X}$ its Borel sigma-algebra. There exists a coupling $\mathbb{P}^{X,\tilde{X}}$ such that $\|P-Q\|_{TV} = 2\mathbb{P}^{X,\tilde{X}}(X \neq \tilde{X})$.*

For a proof, see e.g. Section 8.3 in Thorisson (2000).

The following lemma is well known and included for completeness, it follows from Pinsker's inequality and a straightforward calculation.

Lemma 40. *Let P_f denote the distribution of a $N(f, \sigma I_d)$ distributed random vector for $f \in \mathbb{R}^d$ and let P_f^n denote the distribution of n i.i.d. draws (i.e. $P_f^n = \bigotimes_{i=1}^n P_f$).*

It holds that

$$\left\|P_f^n - P_g^n\right\|_{\mathrm{TV}} \leqslant \frac{n}{2\sigma} \|f - g\|_2 .$$

The following lemma bounds the maximum of a possibly correlated Gaussian random vector.

Lemma 41. *Let $S \sim Bin(p, n)$ and $0 \leqslant \epsilon \leqslant 1$. It holds that $\mathbb{E}Se^{\epsilon S} \leqslant e^{4\epsilon np}$.*

Proof Write $S = \sum_{i=1}^n B_i$, with $B_1, \ldots, B_n \overset{\text{i.i.d.}}{\sim} \mathrm{Ber}(p)$. Then, $\mathbb{E}Se^{\epsilon S}$ can be written as

$$\sum_{i=1}^n \mathbb{E}B_i e^{\epsilon S} = \sum_{i=1}^n pe^{\epsilon}\mathbb{E}e^{\epsilon \sum_{k\neq i} B_k} = npe^{\epsilon}\left(\mathbb{E}e^{\epsilon B_1}\right)^{n-1} = npe^{\epsilon}\left(1 + p(e^{\epsilon} - 1)\right)^{n-1} \leqslant e^{4np\epsilon},$$

where the inequality follows from the fact that $e^x - 1 \leqslant 2x$ for $0 \leqslant x \leqslant 1$. ■

Appendix D. Appendix on Besov norms in sequence space

In this section we briefly introduce Besov spaces as subspaces of $\ell_2(\mathbb{N})$ and collect some properties used in the paper.

The definition of the Besov norm given in (5) is equivalent to the Besov norm as typically defined on a function space when considering the wavelet transform of a function in $L_2[0,1]$, $f = \sum_{j=j_0}^{\infty}\sum_{k=0}^{2^j-1} f_{jk}\psi_{jk}$, where $\left\{\phi_{j_0 m}, \psi_{jk} : m \in \{0, \ldots, 2^{j_0} - 1\}, \quad j > j_0, \quad k \in \{0, \ldots, 2^j - 1\}\right\}$, are the orthogonal wavelet basis functions, for father $\phi(.)$ and mother $\psi(.)$ wavelets with N vanishing moments and bounded support on $[0, 2N-1]$ and $[-N+1, N]$, respectively, following e.g. the construction of Cohen, Daubechies and Vial Cohen et al. (1993); Daubechies (1992), with $N > s$. The Besov norm on the function space is then equivalent to the one defined above for the wavelet coefficients $(f_{jk})_{j \geqslant j_0, k \in \{0,\ldots,2^j-1\}}$. See e.g. Chapter 4 in Gine and Nickl (2016) for details.

Let $\mathcal{B}_{p,q}^s(R) \subset \ell_2(\mathbb{N})$ denote the Besov ball of radius R. The following lemmas are standard results, see e.g. Chapter 9 in Johnstone (2019) for proofs.

Lemma 42. *There exists a constant $C_{s,q} > 0$ such that $\|f\|_2 \leqslant C_{s,q} R$ for all $f \in \mathcal{B}_{p,q}^s(R)$ with $1 \leqslant q \leqslant \infty$ and $2 \leqslant p \leqslant \infty$.*

Lemma 43. *Let f_{lk} are the wavelet coefficients of the function $f \in B_{p,q}^{\alpha}(R)$. For any $1 \leqslant q \leqslant \infty$, $2 \leqslant p \leqslant \infty$, $L > 0$, we have $\sum_{l>L}\sum_{k=0}^{2^l-1} f_{lk}^2 \leqslant c_\alpha 2^{-2L\alpha} R^2$, where $c_\alpha > 0$ is a universal constant depending only on α.*

Lemma 44. *There exists a constant $C_{\alpha,R} > 0$ such that $\|f\|_\infty \leqslant C_{\alpha,R}$ for all $f \in B_{p,q}^{\alpha}(R)$ with $\alpha - 1/2 - 1/p > 0$, $1 \leqslant q \leqslant \infty$ and $2 \leqslant p \leqslant \infty$.*

Lemma 45. *Let f_{lk} are the wavelet coefficients of the function $f \in \mathcal{B}_{p,q}^s(R)$. For any $1 \leqslant q \leqslant \infty$, $2 \leqslant p \leqslant \infty$, $L > 0$, we have $\sum_{l>L}\sum_{k=0}^{2^l-1} f_{lk}^2 \leqslant C_{s,q} 2^{-2Ls} R^2$, where $C_{s,q} > 0$ is a universal constant depending only on s and q.*