Instance-Optimal Differentially Private Estimation

Audra McMillan[*]   Adam Smith[†]   Jonathan Ullman[‡]

Abstract

In this work, we study local minimax convergence estimation rates subject to ε-differential privacy. Unlike worst-case rates, which may be conservative, algorithms that are locally minimax optimal must adapt to *easy* instances of the problem. We construct locally minimax differentially private estimators for one-parameter exponential families and estimating the tail rate of a distribution. In these cases, we show that optimal algorithms for simple hypothesis testing, namely the recent optimal private testers of Canonne et al. (2019), directly inform the design of locally minimax estimation algorithms.

1 Introduction

While the primary goal of statistical inference is to reveal properties of a population, many statistical estimators also reveal a significant amount of information about their sample, and this becomes a serious problem when the sample contains sensitive private information about individuals. As a response, *differential privacy* (Dwork et al., 2006) has emerged as a strong formal criterion for a statistical procedure to protect individual privacy. Differentially private algorithms are deployed in a variety of settings, from the public data products for the 2020 US decennial census to Google's keyboard prediction models (McMahan and Thakurta, 2022) and Apple device analytics (Apple Differential Privacy Team, 2017).

Differential privacy is a constraint on an estimator that requires the distribution of the estimator's outputs to be insensitive to changing a single individual's data, and it offers a strong semantic guarantee that no attacker can infer much more about any individual than they could have inferred had that individual's data never been collected (Kasiviswanathan and Smith, 2008). This semantic guarantee does not rely on any assumptions about the adversary's background knowledge and capabilities. In contrast, alternative approaches to protecting privacy have often been undermined by underestimating the abilities of the attacker. Although differential privacy is a constraint that significantly limits inference with small sample sizes, most statistical tasks are compatible with differential privacy given a large enough sample.

There is now a large body of work on differentially private estimation, which includes *minimax optimal* differentially private estimators for many estimation tasks (e.g. Duchi et al. (2013); Bun et al. (2014); Dwork et al. (2015)). A minimax optimal estimator is one that minimizes the maximum loss over all distributions in some family. However, even a minimax optimal estimator can be undesirable in practice because it might achieve the same error on all distributions, even if some distributions are *easier* than the worst-case distributions.

A more refined guarantee is called *local minimax optimality*. While the actual definition is necessarily subtle, intuitively a local minimax optimal estimator simultaneously has the best possible

---

[*]audra_mcmillan@apple.com. Apple, Inc.

[†]ads22@cs.bu.edu. Department of Computer Science, Boston University.

[‡]jullman@ccs.neu.edu Khoury College of Computer Sciences, Northeastern University.

error on every distribution, which means the error must automatically adapt to distributions that are easier. To illustrate this with a simple example from the non-private setting, suppose we are given a sample of size n from a Bernoulli distribution $P = \text{Ber}(\theta)$ and want to estimate the parameter $\theta \in (0, 1)$. The empirical mean has mean-squared error $\theta(1 - \theta)/n \leq 1/4n$. No estimator can have better error than $1/4n$ on all Bernoulli distributions (roughly because samples of size n from $\text{Ber}(\frac{1}{2} - \frac{1}{2\sqrt{n}})$ and $\text{Ber}(\frac{1}{2} + \frac{1}{2\sqrt{n}})$ are hard to reliably distinguish), so the empirical mean is (globally) minimax optimal. But it is also *locally* minimax optimal because it adapts automatically to the "easy" values of θ close to 0 or 1. In contrast, a hypothetical estimator that had mean-squared error exactly $1/4n$ for all values of θ would be minimax optimal but *not* locally minimax optimal.

We study the design of locally minimax differentially private estimators. We provide:

- A connection between locally minimax differentially private estimators and differentially private simple hypothesis testing: namely, the local estimation rate for the class of differentially private estimators is given by inverting the sample complexity of the optimal differentially private hypothesis test. Such a connection was previously shown in the non-private setting (Donoho and Liu, 1991) and in the more restrictive locally differentially private setting (Duchi and Ruan, 2018; Rohde and Steinberger, 2018).

- Locally minimax differentially private estimators for one-parameter exponential families. In the small ε (that is, $\varepsilon = O(1/\sqrt{n})$) regime, our estimator is directly informed by the locally[1] DP estimator introduced in (Duchi and Ruan, 2018), who show the locally differentially private version of this estimator is locally minimax optimal. For larger ε, our estimators are directly informed by the structure of the approximately optimal differentially private simple hypothesis tests of Canonne et al. (2019). In particular, our estimator critically relies on a refined version of their optimal test, introduced in this work, with additional properties.

- A general approach to nonparametric estimation of one-dimensional functionals. We illustrate its application to estimating tail decay rates.

Simple Hypothesis Testing and Local Estimation Rates (Section 3). As shown by Donoho and Liu (1991), local minimax estimation is closely related to simple hypothesis testing. The connection was originally developed in the non-private setting, but applies more generally to any restricted estimation setting. Suppose we have a population P from some family $\mathcal{P}$ and want to estimate a statistic $\theta(P)$. We have a sample $X \sim P^n$ and an estimator $\hat{\theta}(X)$. Given two distribution $P, Q \in \mathcal{P}$ we can use $\hat{\theta}$ as the basis for a simple hypothesis test that distinguishes P and Q by looking at $\hat{\theta}(X)$ and checking if it's closer to $\theta(P)$ or $\theta(Q)$, and this approach will be a successful hypothesis test if and only if $\hat{\theta}$ has sufficiently small error for both populations P and Q. See Figure 1 for a pictorial representation of how $\hat{\theta}$ can distinguish P and Q. Some pairs $P, Q \in \mathcal{P}$ cannot be reliably distinguished with a sample of size n and some can. Informally, we say that $\hat{\theta}$ is *locally minimax optimal* if it can be used in this fashion to obtain a hypothesis test for *any* pair of distributions in $\mathcal{P}$ that can be distinguished using n samples. This formulation makes it clear that lower bounds for simple hypothesis testing automatically give lower bounds on the local estimation rate. Although hypothesis tests for specific pairs of distributions do not inherently yield optimal estimators, the structure of optimal tests can guide the construction of locally minimax estimators. We show that this process of converting hypothesis testing results into estimation rates can be carried out in the private setting, and instantiate it for several univariate estimation problems.

---

[1] *Local differential privacy* refers to the model of differential privacy where data subjects randomize their own data points before sending it to the server. It is a more restricted model than central differential privacy, the main privacy model of interest in this paper. See Section 1.1
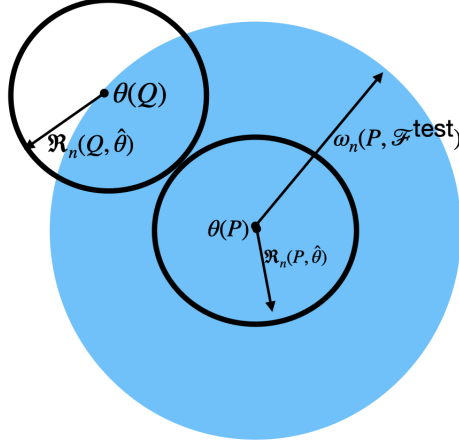
Figure 1: Graphical representation of the connection between simple hypothesis testing and local estimation rates.

In the non-private setting, the sample complexity of distinguishing between two distributions P and Q is $\Theta(1/H^2(P,Q))$, where $H(P,Q)$ is the Hellinger distance, and hence the Hellinger distance is the relevant distance when characterising local estimation rates in the non-private setting. Duchi and Ruan (2018) showed that in the *local DP* setting, the sample complexity is $\Theta(1/\varepsilon^2\mathrm{TV}(P,Q)^2)$, where $\mathrm{TV}(P,Q)$ is the total variation distance. Canonne et al. (2019) showed that the sample complexity in the central DP setting is more nuanced. However, in this work, we show that it has a simple form in the high privacy regime. When $\varepsilon = O(1/\sqrt{n})$, the sample complexity is $\Theta(1/\varepsilon\mathrm{TV}(P,Q))$, the square root of the sample complexity in the local DP setting. This move from the Hellinger distance to the total variation distance has implications for how well one can expect estimation algorithms to adapt to problem-specific difficulty. For example, the fact that non-private algorithms for Bernoulli parameter estimation can adapt to problem-specific difficulty, while local DP algorithms and central DP algorithms in the high privacy regime cannot, is a direct consequence of the fact $H(\mathrm{Bernoulli}(\theta),\mathrm{Bernoulli}(\theta+\alpha))$ is a function of θ, while $\mathrm{TV}(\mathrm{Bernoulli}(\theta),\mathrm{Bernoulli}(\theta+\alpha))$ is independent of θ. We discuss this further in Section 3.3.

While we show that this framework is suitable for univariate estimation problems, it is not generally suitable for estimating multivariate statistics, as this simple-hypothesis-testing formulation does not fully capture private estimation for multivariate statistics. In particular, one provably cannot achieve the local estimation rate even for simple tasks like estimating the mean of a multivariate Gaussian with identity covariance (Bun et al., 2014; Dwork et al., 2015) since the lower bounds on hypothesis testing and estimation depend on the dimension in different ways[2]. We leave it to future work to develop a suitable notion of local minimax estimation for higher-dimensional problems.

Exponential Families (Section 4) We give a DP estimator for one-parameter exponential families that uniformly achieves the private, locally minimax-optimal error under suitable regularity conditions. The estimator works (and is optimal) for any setting of $\varepsilon = O(1)$. We identify two qualitatively different regimes: the "low privacy" regime, $\varepsilon = \Omega(1/\sqrt{n})$, and the "high privacy" regime, $\varepsilon = O(1/\sqrt{n})$. In the low-privacy regime, privacy can be achieved without increasing the asymptotic error of the estimator, while in the high-privacy regime, the error due to privacy dominates the

---

[2]For example, the sample complexity for privately distinguishing between two Gaussian distributions with identity covariance at total variation distance α is $O(\sqrt{d})$ (for constant α and ε) (see, e.g., Narayanan (2022)), while the sample complexity required for privately estimating a Gaussian with identity covariance to within total variation distance α is $\Omega(d/\log(d))$ (Kamath et al., 2019).

sampling error. A weaker version of the low-privacy result appears in Smith (2011); however, that result matches the best nonprivate error only for $\varepsilon = \omega(1/\sqrt[4]{n})$, instead of $\varepsilon = \Omega(1/\sqrt{n})$.

In both regimes, our algorithm first uses a subroutine of Karwa and Vadhan (2018) to identify a rough, initial approximation $\hat{\theta}_0$ to the true parameter. The next step is to compute and release a (noisy) test statistic $\hat{f} = f_{n,\varepsilon,\hat{\theta}_0}(X)$. In the low privacy regime, this statistic is, very roughly, the same one that arises in the private simple hypothesis test of Canonne et al. (2019) for distinguishing $\hat{\theta}_0$ from $\hat{\theta}_0 + \alpha$, where α is roughly the local minimax error at $\hat{\theta}_0$. The exact form of the statistic is more subtle, and relies on a linearization of the model in a neighborhood of $\hat{\theta}_0$. The statistic takes a simpler form in the high privacy regime. Finally, we take the estimate $\hat{\theta}$ to be the unique solution to $\hat{f} = \mathbb{E}_{X \sim P_{\hat{\theta}}^n}\left(f_{n,\varepsilon,\hat{\theta}_0}(X)\right)$, which finds the value $\hat{\theta}$ for which the expected value of the test statistic matches the observation $\hat{f}$. The key in both regimes is to prove that $f_{n,\varepsilon,\hat{\theta}_0}$ is a good test statistic not only for distinguishing $\hat{\theta}_0$ from $\hat{\theta}_0 + \alpha$, but for distinguishing all pairs of the form $(\theta, \theta + \alpha)$ for θ in a neighborhood of $\hat{\theta}_0$.

Our approach parallels that of Duchi and Ruan (2018), who developed a similar result for the more restricted setting of *locally differentially private* algorithms. Indeed, in the high privacy regime, the structure of the optimal estimator is very similar to theirs, and the asymptotic sample complexity of the optimal (central-model) private estimator is exactly the square root of that of the optimal locally-private estimator. In the low-privacy regime, however, the estimators' structure differs. In all cases, the lower bound techniques are quite different.

Estimation of More General Functionals (Section 5). In addition to parametric estimation problems, our framework applies to the estimation of one-dimensional functionals $T(P)$ of distributions, even when the functional of interest does not completely describes the underlying data distribution P. We discuss general approaches to such problems and explore the estimation of tail decay rates in real-valued distributions, an example also studied in depth by Donoho and Liu (1991).

There are several natural meanings to local optimality in such a setting. Following Donoho and Liu (1991), we seek estimation algorithms that, for each θ, achieves error rate $\mathfrak{R}(\theta)$ for all distributions P in the subfamily $\{P \in \mathcal{P} : T(P) = \theta\}$, where $\mathcal{P}$ is the family of distributions to which the true population is assumed to belong and $\mathfrak{R}(\theta)$ is the optimal local estimation rate for at least one distribution in this set. Fairly generically, one can devise near-optimal differentially private algorithms whenever testing the compound hypothesis $(\{P \in \mathcal{P} : T(P) \le \theta_0\}, \{P \in \mathcal{P} : T(P) \ge \theta_1\})$ is equivalent to a simple hypothesis testing problem of distinguishing two specific distributions (with parameters θ_0 and θ_1, respectively). We illustrate this with the design of near-optimal estimators for tail decay rates.

1.1 Related Work

While the literature on differentially private statistical inference is too vast to survey, we give an overview of the most closely related work. For additional discussion of the literature, we direct the reader to the survey of Kamath and Ullman (2020).

Minimax Optimality Under Privacy Constraints. There is now an extensive body of literature on differentially private estimation, which is too large to fully survey here. The most technically relevant prior work to our work are the results of Canonne et al. (2019) characterizing optimal differentially private simple hypothesis testing. The first global minimax lower bounds for multivariate differentially private estimation were given by Bun et al. (2014); Dwork et al. (2015); Steinke and Ullman (2017), based on a technique called *fingerprinting* or *tracing*. Work by Barber and Duchi (2014); Kamath et al. (2020) also gave minimax lower bounds for private mean estimation of univariate heavy-tailed

statistics, and Alon et al. (2019) give minimax lower bounds for privately estimating a univariate distribution in CDF distance.

There are also numerous constructions of minimax optimal differentially private estimators for specific tasks. Perhaps most closely related to our work are the estimators of Karwa and Vadhan (2018) who construct locally minimax optimal estimators for the parameters of a univariate Gaussian, which is a special case of our constructions.

Beyond Global Sensitivity. Several works in the differential privacy literature give general purpose techniques for privately estimating *empirical quantities* in a way that adapts to easy datasets (datasets on which the empirical quantity is stable). These techniques include smooth sensitivity (Nissim et al., 2007), propose-test-release (Dwork and Lei, 2009) and the use of Lipschitz extensions to extend regions of low variablility in the quantity of interest (Chen and Zhou, 2013; Blocki et al., 2013; Kasiviswanathan et al., 2013). The most closely related work to ours is that of Asi and Duchi (2020), who give a general class of differentially private estimators for computing empirical quantities that are locally optimal (under some regularity assumptions). However, in this work we study estimators for *population quantities*. While estimating empirical and population quantities are very related, they are fundamentally distinct. To see why, consider the example of computing the mean of a Gaussian random variable $N(\mu, \sigma^2)$. In the non-private setting, the empirical mean gives a locally minimax optimal estimator for μ. However, applying the locally minimax optimal estimator of Asi and Duchi for the empirical mean will have mean-squared error ∞ for any sample size. In contrast, there is a differentially private estimator for the quantity μ that has mean-squared error roughly $\sigma^2/n + \sigma^2/\varepsilon^2 n^2$ for $\varepsilon \leq 1$ (e.g., Karwa and Vadhan (2018)). Thus, we have to reason directly about population statistics when we try to construct locally minimax private estimators, and cannot simply apply the transformation of Asi and Duchi to an arbitrary locally minimax non-private estimator.

Local Differential Privacy. Our work studies the standard *centralized model* of differential privacy, where we assume that the estimator M receives the samples $X_1, \ldots, X_n$ as input. There is also a large body of research on so-called *local differential privacy* (Kasiviswanathan et al., 2008), where we assume that differential privacy is applied to each sample before it is collected. In its most basic non-interactive form, this means that the mechanism can be written in the form $A(M(X_1), \ldots, M(X_n))$ where M is differentially private and A is arbitrary.

Locally differentially private estimators are known to have significantly worse rates than general differentially private estimators (Kasiviswanathan et al., 2008; Beimel et al., 2008; Chan et al., 2011; Duchi et al., 2013; Edmonds et al., 2020). Recent work gives locally minimax optimal estimators subject to local differential privacy (Duchi and Ruan, 2018; Rohde and Steinberger, 2018). In addition to different minimax rates, there are key conceptual differences between the local and centralized settings that make the centralized setting more complex to reason about. In particular: (1) The complexity of simple hypothesis testing under local differential privacy is characterized by the total variation distance between the two distributions, whereas a much more subtle notion is required for centralized differential privacy, and (2) The local minimax rate subject to local differential privacy is always larger than that of non-private estimation, whereas our results show that the local minimax rate subject to centralized differential privacy can be either the same or larger than non-private estimation in different ranges of the privacy parameter.

2 Local Estimate Rates and Simple Hypothesis Testing

2.1 Local Estimation Rates and Uniform Achievability

Let $\Delta(\chi)$ be the set of all distributions on a space χ and $\mathcal{P} \subset \Delta(\chi)$ be a set of distributions on χ. Let $\theta : \mathcal{P} \to \mathbb{R}$ be a functional on $\mathcal{P}$, so for any distribution $P \in \mathcal{P}$, $\theta(P)$ is the parameter that we want to estimate. Let $\mathcal{F}$ be a class of (potentially randomised) functions $\hat{\theta} : \chi^n \to \mathbb{R}$. For any estimator $\hat{\theta}$ in $\mathcal{F}$, $\hat{\theta}$ has local error rate $\mathfrak{R}_n(P, \hat{\theta})$ if for all $P \in \mathcal{P}$ and $n \in \mathbb{N}$, if $X_1, \cdots, X_n \sim P$ then with probability 0.75:

$$|\hat{\theta}(X_1, \cdots, X_n) - \theta(P)| \leq \mathfrak{R}_n(P, \hat{\theta}).$$

Notice that this error rate is *instance specific* in the sense that the error rate is a function of the distribution being sampled from. Since worst-case analysis can be too pessimistic in practice, and the local rate allows the error rate to adapt to *easy* instances of the problem. Defining a notion of *instance optimality* is nuanced since no algorithm can be optimal for all P; that is, one can not define an algorithm $\hat{\theta}$ such that $\mathfrak{R}_n(P, \hat{\theta}) \leq \min_{\theta' \in \mathcal{F}} \mathfrak{R}_n(P, \hat{\theta}')$ for all P. This is easy to see since for any algorithm $\hat{\theta}$ and distribution P, the algorithm $\hat{\theta}'(X_1, \cdots, X_n) = P$ satisfies $\mathfrak{R}_n(P, \hat{\theta}') = 0 \leq \mathfrak{R}_n(P, \hat{\theta})$. Of course, this algorithm is not a good point of comparison because it does poorly on distributions that are not P. Thus, we want to compare to algorithms that perform well on at least two distributions. This leads us to the following definition of the *optimal local estimation rate* at P by:

$$\mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta) = \sup_{Q \in \mathcal{P}} \inf_{\hat{\theta} \in \mathcal{F}} \max\{\mathfrak{R}_n(P, \hat{\theta}), \mathfrak{R}_n(Q, \hat{\theta})\}. \tag{1}$$

We call this definition the local estimation rate based on the intuition that the hardest distributions to distinguish from P are those that are "close" or "local" to P (Fig. 1). The local estimation rate is also sometimes to referred to as the rate of the *hardest one dimensional sub-problem*. We say an algorithm $\hat{\theta}$ is *instance optimal* if $\mathfrak{R}_n(P, \hat{\theta}) = \mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta)$ for all $P \in \mathcal{P}$. Intuitively, if $\hat{\theta}$ is instance optimal then for every distribution P, if $\hat{\theta}$ performs poorly on P, then there exists another distribution Q such that no algorithm $\hat{\theta}'$ performs well on both P and Q. In contrast, the trivial algorithm $\hat{\theta}'(X_1, \cdots, X_n) = P$ performs well on P, but unnecessarily sacrifices performance on distributions Q far from P. Hence the optimal local estimation rate gives a specific kind of lower bound on the performance of any algorithm.

The estimator $\hat{\theta}$ in eqn (1) has the advantage of being told the two distributions P and Q. Hence, unlike worst-case optimality, which is always achieved by some algorithm, an instance optimal algorithm does not necessarily exist for every estimation problem. In fact, a main question in this area is *when* do instance optimal algorithms exist? When an instance optimal algorithm exists we will say the estimation problem satisifies *uniform achievability*. This question of uniform achievability, under the constraint of differential privacy, is the main question of interest in this work. This question has been studied previously in the non-private setting (Donoho and Liu, 1991) and under the constraint of local differential privacy (Rohde and Steinberger, 2018; Duchi and Ruan, 2018). We will refer to the subset of $\mathcal{F}$ that contains all ε-DP estimators (defined in Section 3.1) as $\mathcal{F}_\varepsilon$.

2.2 Simple Hypothesis Testing

The crucial insight for understanding the optimal local estimation rate is the connection to simple hypothesis testing. In simple hypothesis testing, we are given two distributions P and Q and the goal is to design an algorithm that given n samples drawn from either P or Q, will, with high probability,

correctly guess which distribution the samples were drawn from. We say a test $T : \chi^n \to \{0,1\}$ distinguishes between P and Q with n samples if $\mathbb{P}(T(P^n) = 0) \geq 0.75$ and $\mathbb{P}(T(Q^n) = 1) \geq 0.75$, where the probability is taken over both the randomness in the sample, and the randomness in T. Let $\chi^* = \cup_{n \in \mathbb{N}} \chi^n$. We will use $\mathrm{SC}_T(P,Q)$ to denote the sample complexity of a test T, i.e.,

$$\mathrm{SC}_T(P,Q) = \inf\{n \in \mathbb{N} \mid \text{for all } N \geq n, \mathbb{P}(T(P^N) = 0) \geq 0.75 \text{ and } \mathbb{P}(T(Q^N) = 1) \geq 0.75\}.$$

For every estimator class $\mathcal{F}$, we can define an associated class of binary testing algorithms, $\mathcal{F}^{\text{test}}$, to be the class of binary (potentially randomised) functions $T : \chi^n \to \{0,1\}$ obtained from $\mathcal{F}$ by thresholding:

$$\mathcal{F}^{\text{test}} = \left\{ T_{f,\tau}(X) = \begin{cases} 0 & \text{if } f(X) < \tau \\ 1 & \text{otherwise} \end{cases} \middle| f \in \mathcal{F}, \tau \in \mathbb{R} \right\}. \tag{2}$$

We will use this translation throughout this work. Given a class of tests $\mathcal{F}^{\text{test}}$, define $\mathrm{SC}_{\mathcal{F}^{\text{test}}}(P,Q) = \inf_{T \in \mathcal{F}^{\text{test}}} \mathrm{SC}_T(P,Q)$. That is, $\mathrm{SC}_{\mathcal{F}^{\text{test}}}(P,Q)$ is the smallest n such that there exists a test $T \in \mathcal{F}^{\text{test}}$ that distinguishes P and Q.

2.3 Connecting Local Estimation Rates and Simple Hypothesis Testing

Consider the definition of the optimal local estimation rate given in Eqn 1. Given two distributions P and Q, if $\theta(P)$ and $\theta(Q)$ are close then it is easy to find an estimator that performs well on both P and Q (e.g. the estimator that outputs $\frac{1}{2}|\theta(P) - \theta(Q)|$). Similarly, if there exists a test that distinguishes P and Q, then it is easy to define an estimator that performs well on both P and Q (e.g. by outputting the test result). Thus, the supremum in the definition is achieved at a distribution Q that is as far as possible from P, while still being indistinguishable from P. This intuition gives rise to the definition of the *modulus of continuity* at $P \in \mathcal{P}$:

$$\omega_n(P, \mathcal{F}, \mathcal{P}, \theta) = \sup\{ |\theta(P) - \theta(Q)| \mid \mathrm{SC}_{\mathcal{F}^{\text{test}}}(P,Q) > n \text{ and } Q \in \mathcal{P} \}.$$

The following theorem formalises the intuition above and allows us to translate the question of characterizing $\mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta)$ into characterizing $\mathrm{SC}_{\mathcal{F}^{\text{test}}}$. This is useful since characterizations of $\mathrm{SC}_{\mathcal{F}^{\text{test}}}$ in a variety of settings already exist, in particular a characterization of $\mathrm{SC}_{\mathcal{F}^{\text{test}}}$ when $\mathcal{F}$ is the class of all differentially private estimators was given in Canonne et al. (2019). We say $\mathcal{F}$ is closed under post-processing if for any $\hat{\theta} \in \mathcal{F}$ and $f : \mathbb{R} \to \mathbb{R}$, $f \circ \hat{\theta} \in \mathcal{F}$. Donoho and Liu (1987) studied the characterisation of $\mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta)$ where $\mathcal{F}$ is the class of all possible estimators; their work can be extended to work for any class of estimators closed under post-processing.

Proposition 2.1. *For any $\mathcal{P} \subset \Delta(\chi)$, statistic θ, and class of estimators $\mathcal{F}$, if $\mathcal{F}$ is closed under post-processing and contains all constant functions then for all $P \in \mathcal{P}$ and $n \in \mathbb{N}$,*

$$\mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta) = \tfrac{1}{2}\omega_n(P, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta),$$

where $\mathcal{F}^{\text{test}}$ is as defined in eqn (2).

When $\mathcal{P}$ and θ are clear from context, we write $\omega_n(P, \mathcal{F}^{\text{test}})$ for $\omega_n(P, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta)$, and similarly $\mathcal{E}_n^{\text{loc}}(P, \mathcal{F})$ for $\mathcal{E}_n^{\text{loc}}(P, \mathcal{F}, \mathcal{P}, \theta)$. We will primarily be concerned with the class of differentially private estimators in this paper, which is closed under post-processing and contains all constant functions. We include the proof below to build intuition for this connection.

Proof. Let us first prove that $\mathcal{E}_n^{\mathrm{loc}}(P, \mathcal{P}, \mathcal{F}, \theta) \geq \frac{1}{2}\omega_n(P, \mathcal{F}^{\mathrm{test}}, \mathcal{P}, \theta)$. Suppose for sake of contradiction that $\mathcal{E}_n^{\mathrm{loc}}(P, \mathcal{P}, \mathcal{F}, \theta) < \frac{1}{2}\omega_n(P, \mathcal{F}^{\mathrm{test}}, \mathcal{P}, \theta)$. Then there exists $\hat{\theta} \in \mathcal{F}$ and $Q \in \mathcal{P}$ such that $\mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}(P, Q) > n$, $\frac{1}{2}\omega_n(P, \mathcal{F}^{\mathrm{test}}, \mathcal{P}, \theta) \leq \frac{1}{2}|\theta(P) - \theta(Q)|$ and,

$$\mathfrak{R}_n(P, \hat{\theta}) < \tfrac{1}{2}\omega_n(P, \mathcal{F}^{\mathrm{test}}, \mathcal{P}, \theta) \quad \text{and} \quad \mathfrak{R}_n(Q, \hat{\theta}) < \tfrac{1}{2}\omega_n(P, \mathcal{F}^{\mathrm{test}}, \mathcal{P}, \theta)$$

Therefore, $T_{\hat{\theta}, \frac{1}{2}(\theta(P)+\theta(Q))}$ (as defined in eqn (2)) distinguishes P and Q with n samples, which is a contradiction since $\mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}(P, Q) > n$. Figure 1 gives a graphical representation of this, if the balls do not overlap then we have a test that distinguishes P and Q.

For the opposite inequality, we need to show that for all $Q \in \mathcal{P}$, there exists an estimator $\hat{\theta} \in \mathcal{F}$ such that $\max\{\mathfrak{R}_n(P, \hat{\theta}), \mathfrak{R}_n(Q, \hat{\theta})\} \leq \frac{1}{2}\omega_{n, \mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}}(P)$. First suppose that $\mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}(P, Q) > n$, that Q lies inside the blue ball around P in Figure 1, and $\frac{1}{2}|\theta(P) - \theta(Q)| \leq \frac{1}{2}\omega_{n, \mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}}(P)$. Let $\hat{\theta}$ be the constant function that always outputs $\frac{1}{2}|\theta(P) - \theta(Q)|$ so $\mathfrak{R}_n(P, \hat{\theta}) = \frac{1}{2}|\theta(P) - \theta(Q)|$ and $\mathfrak{R}_n(Q, \hat{\theta}) = \frac{1}{2}|\theta(P) - \theta(Q)|$, so we are done. Finally, suppose that $\mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}(P, Q) \leq n$ so there exists $\hat{\theta} \in \mathcal{F}$ and $\tau \in \mathbb{R}$ such that

$$\mathbb{P}[\hat{\theta}(P^n) \leq \tau] \geq 0.75 \quad \text{and} \quad \mathbb{P}[\hat{\theta}(Q^n) \leq \tau] \geq 0.75.$$

Let $f : \mathbb{R} \to \mathbb{R}$ be defined by $f(x) = \theta(P)$ if $\hat{\theta}(x) \leq \tau$ and $f(x) = \theta(Q)$ is $\hat{\theta}(x) \geq \tau$ so $f \circ \hat{\theta} \in \mathcal{F}$ and $\mathfrak{R}_n(P, f \circ \hat{\theta}) = 0$ and $\mathfrak{R}_n(Q, f \circ \hat{\theta}) = 0$, so we are done. $\square$

For distributions P and Q, let

$$\mathrm{H}(P, Q) = \sqrt{\int_\chi (P(x) - Q(x))^2 \, \mathrm{d}x}$$

be the Hellinger distance between P and Q. It is well known that for the class of all estimators,

$$\mathrm{SC}_{\mathcal{F}^{\mathrm{test}}}(P, Q) = \Theta\left(\frac{1}{\mathrm{H}^2(P, Q)}\right),$$

so the following corollary is an immediate consequence of Theorem 2.1. Define the H-information by:

$$J_{\mathrm{H}, P}^{-1}(\beta) = \sup\left\{|\theta(P) - \theta(Q)| \,\Big|\, \mathrm{H}(P, Q) \leq \beta, Q \in \mathcal{P}\right\}.$$

Corollary 2.2 (Non-private optimal local estimation rate (Donoho and Liu, 1987)). *Let $\mathcal{F}$ be the set of all functions, then there exists constants C_1 and C_2 such that for any family $\mathcal{P}$, any statistic θ, any distribution P, and $n \in \mathbb{N}$,*

$$\mathcal{E}_n^{\mathrm{loc}}(P, \mathcal{P}, \mathcal{F}, \theta) \in \left[J_{\mathrm{H}, P}^{-1}\left(\frac{C_1}{\sqrt{n}}\right), J_{\mathrm{H}, P}^{-1}\left(\frac{C_2}{\sqrt{n}}\right)\right].$$

2.4 Super efficiency

The optimal local estimation rate $\mathcal{E}_n^{\mathrm{loc}}(\cdot, \mathcal{F}, \mathcal{P}, \theta)$ has the property that for any estimator $\hat{\theta}$, if $\hat{\theta}$ achieves better accuracy than $\mathcal{E}_n^{\mathrm{loc}}(P, \mathcal{F}, \mathcal{P}, \theta)$ at some distribution P, then there exists a distribution Q such that the accuracy of $\hat{\theta}$ at Q is *at least* as bad as $\mathcal{E}_n^{\mathrm{loc}}(Q, \mathcal{F}, \mathcal{P}, \theta)$. One can also ask if an estimation rate satisfies the stronger condition of having a *super-efficiency* result. Roughly, an estimation rate R has a super-efficiency result if for any estimator $\hat{\theta}$ that achieves better accuracy than $R(P)$ at a particular value P, there exists another value Q where the accuracy of $\hat{\theta}$ is *strictly*

8

worse than $R(Q)$. A super-efficiency result for a given rate R shows, in a sense, that that R is a meaningful target rate. The optimal local estimation rate does not necessarily satisfy a super-efficiency result for general families. Super-efficiency may hold for *specific* families but a general result seems to require further assumptions. We leave the question of super-efficiency of the optimal local estimation rate to future work, since our focus is the general regime.

3 Differentially Private Simple Hypothesis Testing and the Optimal Local Estimation Rate in the High Privacy Setting

In this section, we will discuss the optimal test statistic for differentially private simple hypothesis testing and characterise the optimal local estimation rate in the high privacy setting. The test statistic we give is a slight variant on that presented in Canonne et al. (2019), who first characterised the sample complexity of differentially private simple hypothesis testing. The test statistic given here is more efficient and more amenable to the estimation problem. The characterisation of the local estimation rate in the high privacy regime is simpler than in other regimes, and offers a direct comparison to the local estimation rates in the non-private and local differential privacy regimes.

3.1 Differential Privacy

In this work we are concerned with estimators that satisfy *differential privacy*, which we will formally define in this section. Let $\mathcal{X}$ be a data universe and $\mathcal{X}^n$ be the space of datasets of size n. Two datasets $d, d' \in \mathcal{X}^n$ are neighboring, denoted $d \sim d'$, if they differ on a single record. Let $\mathcal{Y}$ be an output space.

Definition 1 (ε-Differential Privacy (Dwork et al., 2006))**.** Given privacy parameters $\varepsilon \geq 0$ and $\delta \in [0,1]$, a randomized mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ is ε-*differentially private* if for all datasets $d \sim d' \in \mathcal{X}^n$, and events $E \subseteq \mathcal{Y}$,

$$\Pr[M(d, \text{hyperparams}) \in E] \leq e^\varepsilon \cdot \Pr[M(d', \text{hyperparams}) \in E] + \delta,$$

where the probabilities are taken over the randomness induced by M.

The key intuition for this definition is that the distribution of outputs on input dataset d is almost indistinguishable from the distribution of outputs on input dataset d'. Therefore, given the output of a differentially private mechanism, it is impossible to confidently determine whether the input dataset was d or d'. For strong privacy guarantees, the privacy-loss parameter is typically taken to be a small constant less than 1 (note that $e^\varepsilon \approx 1 + \varepsilon$ as $\varepsilon \to 0$) and δ is taken to be very small (say 10^{-6}). In fact, for simple hypothesis testing, we can show that if $\varepsilon > 1$, then for any $\delta \in [0,1]$, the private sample complexity within a constant factor of the non-private sample complexity, i.e., $\text{SC}_{\varepsilon,\delta}(P,Q) = \Theta(\text{SC}(P,Q))$. Hence, for the remainder of this work, we will assume that $\varepsilon \leq 1$. Note if $\mathcal{F}_\varepsilon$ is the set of all ε-DP estimators, then $\mathcal{F}_\varepsilon^{\text{test}}$ is the set of all ε-DP tests.

3.2 An Optimal Differentially Private Simple Hypothesis Test

A characterisation of the sample complexity of differentially private simple hypothesis testing was given in Canonne et al. (2019). They showed that a simple *noisy and clamped* version of the log likelihood ratio test gave an optimal sample complexity differentially private simple hypothesis test.

Given distributions P and Q, let cLLR_a^b be the clamped log-likelihood statistic with thresholds a and b, and ncLLR_a^b be a noisy version:

$$\mathrm{cLLR}_a^b(X) = \sum_{i=1}^n \left[\ln \frac{P(x_i)}{Q(x_i)} \right]_a^b \quad \text{and} \quad \mathrm{ncLLR}_a^b(X) = \mathrm{cLLR}_a^b(X) + \mathrm{Lap}\left(\frac{|b-a|}{\varepsilon} \right). \tag{3}$$

In the original version of this test, the authors' proved that this test statistic gave rise to an optimal test when one set $b = \varepsilon$ and $a = -\varepsilon'$, where ε' is some function of ε, P and Q. In Appendix A, we improve on their results to show that setting $b = \Theta(\varepsilon)$ and $-a = \Theta(\varepsilon)$ is sufficient. This extension is crucial to us in our estimation algorithm where ε' can not be computed. This is also of independent interest as an improvement of the testing result: unlike the original test presented in Canonne et al. (2019), setting $a = -\varepsilon$ and $b = \varepsilon$, results in an efficient test which only requires oracle access to P and Q. The original result in Canonne et al. (2019) required full knowledge of the distributions P and Q in order to compute ε'. In order to simplify notation we use $\mathrm{SC}_\varepsilon(P,Q) := \mathrm{SC}_{\mathcal{F}_\varepsilon}(P,Q)$ to denote the optimal sample complexity for distinguishing P and Q using an ε-DP algorithm. The proof of the following proposition is found in Section A.1

Proposition 3.1. *If $\varepsilon = O(1)$ then for all $a = \Theta(\varepsilon)$ and $b = \Theta(\varepsilon)$, there exists constants C_1 and C_2 such that for all distributions P and Q,*

$$\mathrm{SC}_{\mathrm{ncLLR}_a^b}(P,Q) \in [C_1 \cdot \mathrm{SC}_\varepsilon(P,Q), C_2 \cdot \mathrm{SC}_\varepsilon(P,Q)].$$

The sample complexity of $\mathrm{ncLLR}_\varepsilon^\varepsilon$, characterised in Canonne et al. (2019), has a nuanced dependence on ε, P and Q. If ε is large enough, privacy comes for free, and $\mathrm{SC}_\varepsilon(P,Q) = \Theta(\mathrm{SC}(P,Q))$. As ε decreases the dependence becomes more complicated. However, in Lemma 3.2 we will show that once ε is small enough, $\varepsilon_n \leq \frac{1}{\sqrt{n}}$, the dependence is once again simple.

For hypothesis tests with constant error probabilities, the sample complexity bounds are equivalent, up to constant factors, for pure ε-differential privacy, and the less strict notions of approximate (ε, δ)-differential privacy and concentrated differential privacy (Dwork and Rothblum, 2016; Bun and Steinke, 2016) (see Acharya et al. (2018, Lemma 5)). Consequently, the test $\mathrm{ncLLR}_\varepsilon^\varepsilon$ is optimal (up to constants) for each of these weaker notions. The class of estimators defined by each of these notions is closed under post-processing and thus, by Theorem 2.1, the optimal local estimation rate is, up to constants, the same for each of these notions. This may seem like a contradiction since there are many well-known cases of asymptotic gaps in the estimation rate of pure differential privacy and approximate differential privacy. However, the optimal local estimation rate need not be *uniformly achievable* under all (or any) of these notions of privacy, leaving room for a gap in the achievable estimation rate under pure, concentrated and approximate DP.

3.3 A Lower Bound for Instance Optimal Estimation in the High Privacy Regime

The characterization of the local estimation rate is significantly more complex in the central DP regime than in the local DP or non-private regimes. This is a direct consequence of the characterisation of the optimal sample complexity of simple hypothesis testing being more nuanced in the central DP regime than the local DP or non-private regimes. However, the existence of a simple characterisation of the sample complexity in the high privacy regime allows us to give a simple characterisation of the local estimation rate in that regime.

For distributions P and Q, let $\mathrm{TV}(P,Q) = (1/2) \int |P(x) - Q(x)| dx$ be the total variation distance. For $\beta \in [0,1]$, we define the L_1-information at a distribution $P \in \mathcal{P}$ by

$$J_{\mathrm{TV},P}^{-1}(\beta) = (1/2) \cdot \sup\left\{ |\theta(P) - \theta(Q)| \,\Big|\, \mathrm{TV}(P,Q) \leq \beta, Q \in \mathcal{P} \right\}. \tag{4}$$

Note that the L_1-information is the analogue of the H-information, which characterizes the sample complexity in the non-private setting, using the total variation distance (also known as the L_1-norm) instead of the Hellinger distance. Our estimation rate in the high privacy regime is characterized by the L_1-information. This follows immediately from Theorem 2.1 and Lemma 3.2 which we'll state below.

Theorem 2. *Let $\mathcal{F}_\varepsilon$ be the set of all ε-differentially private estimators. For any constant k there exists a constants C_1 and C_2 such that if $\varepsilon_n \leq \frac{k}{\sqrt{n}}$, then for all families $\mathcal{P}$, $P \in \mathcal{P}$, and $n \in \mathbb{N}$,*

$$J_{\mathrm{TV},P}^{-1}\left(\frac{C_1}{n\varepsilon_n}\right) \leq \mathcal{E}_n^{\mathrm{loc}}(P,\mathcal{P},\mathcal{F}_\varepsilon,\theta) \leq J_{\mathrm{TV},P}^{-1}\left(\frac{C_2}{n\varepsilon_n}\right).$$

Theorem 2 is interesting to contrast with Corollary 2.2, which gives the estimation rate in the non-private regime. Note first that if $\varepsilon_n = O(\frac{1}{\sqrt{n}})$ then $n\varepsilon_n < \sqrt{n}$, so the estimation rate is indeed slower under the constraint of privacy. Further, the metric characterizing the problem changes from the Hellinger distance to the total variance distance. A similar phenomenon is observed under local differential privacy constraints in Duchi and Ruan (2018).

Theorem 3 (Local DP (Duchi and Ruan, 2018))**.** *Let $\mathcal{F}_{local,\varepsilon}$ be the set of all ε-locally differentially private functions, there exists constants C_1 and C_2 such that for all families $\mathcal{P}$, estimators θ, any $P \in \mathcal{P}$, and $n \in \mathbb{N}$,*

$$J_{\mathrm{TV},P}^{-1}\left(\frac{C_1}{\varepsilon\sqrt{n}}\right) \leq \mathcal{E}_n^{\mathrm{loc}}(P,\mathcal{F}_{local,\varepsilon}) \leq J_{\mathrm{TV},P}^{-1}\left(\frac{C_2}{\varepsilon\sqrt{n}}\right).$$

The corresponding class of testing functions $\mathcal{F}^{\mathrm{test}}$ contains the set of all ε-local DP binary functions. Duchi et al. (2013) showed that the sample complexity for distinguishing between two distributions P and Q under local differential privacy is $\Theta\left(\frac{1}{\varepsilon^2 TV^2(P,Q)}\right)$.

As discussed in Duchi and Ruan (2018), the change from the Hellinger modulus of continuity to the total variation modulus of continuity has implications for how well one can expect estimation algorithms in the high privacy setting to adapt to problem-specific difficulty. For example in the case of Bernoulli estimation, the non-private local estimation rate for a Bernoulli with parameter $p \in [0,1]$ is $\Theta(\sqrt{p(1-p)/n})$, which shows that estimation algorithms in the non-private (and low central privacy setting) are able to adapt to "easy" instances of the problem. In contrast, in the high privacy setting, the local estimation rate is $\Theta\left(\frac{1}{\varepsilon n}\right)$, which is the same for all p, showing that private algorithms in this regime are not able to adapt to "easy" instances. As mentioned earlier, this is a direct consequence of the fact that the Hellinger distance between Bernoulli(p) and Bernoulli($p + \alpha$) is a function of p, while the total variation distance between these two distributions is independent of p.

Theorem 2 is a direct consequence of the following characterisation of the sample complexity of private hypothesis testing in the high privacy regime. The proof follows from the fact that in the high privacy regime, $\varepsilon \leq k/\sqrt{n}$, a noisy Scheffé test performs as well as the optimal test ncLLR$_\varepsilon^\varepsilon$.

Lemma 3.2 (High Privacy Sample Complexity Characterisation)**.** *For any constant k, there exists constants C_1 and C_2 such that for any distributions P and Q, if $\varepsilon_n \leq \frac{k}{\sqrt{n}}$ then*

$$SC_{\varepsilon_n}(P,Q) \in \left[\frac{C_1}{\varepsilon_n \cdot \mathrm{TV}(P,Q)}, \frac{C_2}{\varepsilon_n \cdot \mathrm{TV}(P,Q)}\right].$$

Before we prove Lemma 3.2, a quick note on the privacy parameters. We will allow our privacy parameter, ε, to vary with the size of the database, n, so let ε_n be a sequence and $n : [0,\infty) \to \mathbb{N}$ be

such that $\varepsilon_{n(\varepsilon)} = \varepsilon$ and $n(\varepsilon_n)$. We will often abuse notation and drop the argument of the function, e.g., referring to ε_n as simply ε. We will assume that ε is decreasing, so the larger the dataset, the more private we require our algorithm to be. We will say a simple hypothesis testing problem has sample complexity $n = SC_{\varepsilon_n}(P, Q)$ if $n = n(\varepsilon)$ is the smallest value such that $SC_\varepsilon(P, Q)$ and $n(\varepsilon)$ intersect.

Proof. The lower bound portion of this lemma is not specific to the high privacy setting; there exists C_1 such that for all ε, $SC_\varepsilon(P, Q) \geq \frac{C_1}{\varepsilon \mathrm{TV}(P,Q)}$. One way to prove this is as a direct consequence of (Acharya et al., 2018, Theorem 11). This theorem argues that one can lower bound the sample complexity of an ε-DP test by upper bounding the Hamming distance between two datasets of size n drawn from either P and Q, i.e., the Hamming distance between X and Y where $X \sim P^n$ and $Y \sim Q^n$.

For the upper bound, we will show that a noisy version of the simple Scheffé test has sample complexity $O(1/\varepsilon \mathrm{TV}(P, Q))$ in the high privacy regime. Let $E = \{x \parallel P(x) > Q(x)\}$ be the Scheffé set and define the test statistic f_E by, for any database $X = \{x_1, \cdots, x_n\}$,

$$f_E(X) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{x_i \in E} + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right).$$

Then by definition of the total variation distance,

$$\mathbb{E}_{X \sim P^n}[f_E(X)] - \mathbb{E}_{X \sim Q^n}[f_E(X)] = \Pr_{x \sim P}(x \in E) - \Pr_{x \sim Q}(x \in E) = \mathrm{TV}(P, Q).$$

Further,

$$\max\{\mathrm{var}_{X \sim P^n}(f_E(X)), \mathrm{var}_{X \sim Q^n}(f_E(X))\} \leq \frac{1}{n} + \frac{1}{\varepsilon^2 n^2} \leq \frac{1+k^2}{\varepsilon^2 n^2},$$

where the last inequality follows since $\varepsilon \leq \frac{k}{\sqrt{n}}$. Therefore, if $n \geq \sqrt{\frac{1+k^2}{12}} \frac{1}{\varepsilon_n \mathrm{TV}(P,Q)}$, we have that

$$\mathbb{E}_{X \sim P^n}[f_E] - \mathbb{E}_{X \sim Q^n}[f_E] \geq \frac{1}{12} \max\{\sqrt{\mathrm{var}_{X \sim P^n}(f_E(X))}, \sqrt{\mathrm{var}_{X \sim Q^n}(f_E(X))}\}.$$

A simple application of Chebyshev's inequality (for details see (Canonne et al., 2019, Lemma 2.6)) implies that there exists a threshold τ such that the test that outputs P if $f_E(X) \geq \tau$ and Q otherwise, distinguishes between P and Q with sample complexity $\sqrt{\frac{1+k^2}{12}} \frac{1}{\varepsilon_n \mathrm{TV}(P,Q)}$. $\qquad\square$

4 One-Parameter Exponential Families: Characterising the Optimal Local Estimation Rate and Uniform Achievability

We now turn our attention to an example where uniform achievability is possible under differential privacy: one-parameter exponential families. In this section we will characterize the optimal local estimation rate of estimating the parameter in a one-parameter exponential family, then show that this optimal local estimation rate is uniformly achievable under differential privacy. In particular, we will see how the results of Section 3.2 on the form of the optimal DP simple hypothesis test, and it's sample complexity, inform the design of the locally minimax estimator.

One parameter exponential families are a broad class of families of distributions that encompasses many natural distributions. Examples of exponential families include Poisson distributions, Binomial distributions, normal distributions with known variance and normal distributions with known mean.

Formally, a *one parameter exponential family*, $\mathcal{P}_\mu = \{P_\theta\}$, is determined by a base measure μ such that for each θ, the distribution P_θ has density

$$p_\theta(x) := e^{\theta x - A(\theta)} \quad \text{(relative to } \mu\text{)},$$

where $A(\theta) = \ln \int e^{\theta x} d\mu(x)$ is the normalisation.[3] Note that the mean and the variance have the following simple formulations: $\mathbb{E}_\theta[x] = A'(\theta)$ and $\text{var}_\theta(x) = A''(\theta)$. The formula for p_θ does not give a well defined distribution for values of θ for which $A(\theta) = \infty$, so each measure μ has an associated range which we will denote $\Phi_\mu = \{\theta \mid A(\theta) < \infty\}$. When μ is clear from context, we will drop the dependence on μ and refer to Φ_μ simply as Φ.

Let us begin with the characterization of the optimal local estimation rate. The formal version of this theorem is a combination of Corollary 4.9 and Corollary 4.14, which characterize the optimal local estimation rate separately for the high and low privacy regimes.

Theorem 4 (Characterization of Optimal Local Estimation Rate—Simplified from Corollaries 4.9 and 4.14). *For all exponential families (i.e., measures μ), $\delta > 0$, all sequences of privacy parameters $\varepsilon_n \in [0, 1]$, $n \in \mathbb{N}$, and $\theta_0 \in \Phi_\mu$,*

$$\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}}) = \Theta_\mu\left(\frac{1}{\sqrt{A''(\theta_0)}\min\{n\varepsilon_n, \sqrt{n}\}}\right),$$

where the Θ_μ notation hides constants depending only on μ (but not θ_0).

This convergence result is uniform in a fairly strong sense. Given a family defined by a measure μ, there exists constants C_1, C_2 such that for all sequences $\varepsilon_n \in [0, 1]$, sufficiently large n, and $\theta_0 \in \Phi_\mu$,

$$\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}}) \in \left[\frac{C_1}{\sqrt{A''(\theta_0)}\min\{n\varepsilon_n, \sqrt{n}\}}, \frac{C_2}{\sqrt{A''(\theta_0)}\min\{n\varepsilon_n, \sqrt{n}\}}\right].$$

The formal statements of this theorem are slightly stronger than Theorem 4 in that we show that the constants C_1 and C_2 depend only on a few properties of μ. We will discuss these properties later in this section. The non-private local estimation rate for exponential families is $\omega_n(P_{\theta_0}, \mathcal{F}^{\text{test}}) = \Theta_\mu\left(\frac{1}{\sqrt{A''(\theta_0)n}}\right)$, so we can see that in the low privacy regime, privacy comes for free. In the high privacy regime, this characterisation matches the L1-information at P as expected from Theorem 2.

Under some mild conditions, this optimal local estimation rate is actually uniformly achievable. That is, there exists an algorithm that achieves the optimal local estimation rate. The following is an informal statement of Proposition 4.10 and Proposition 4.15, which contain the formal uniform achievability statements in the high and low privacy regimes separately.

Theorem 5 (Uniform Achievability——Simplified from Propositions 4.10 and 4.15). *For all exponential families (i.e., all measures μ), there is an algorithm $\mathcal{A}_\mu$ such that for all $\theta_0 \in \Phi$, and for all sequences $\varepsilon_n \in [0, 1]$, $\delta \in [0, 1]$ and $n \in \mathbb{N}$, $\mathcal{A}_\mu(\varepsilon, \delta, \cdot)$ is (ε, δ)-DP and*

$$\mathfrak{R}_n(\theta_0, \mathcal{A}_\mu) = O_\mu(\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})).$$

The O_μ notation hides constants depending on μ (but not θ_0).

---

[3]It is common to see a sufficient statistic, $T(x)$, included in the definition of an exponential family so that $p_\theta(x) := e^{\theta T(x) - A(\theta)}$. Defined in this way, an exponential family can be defined over any space, not simply $\mathbb{R}$. However, for the purpose of estimating θ, the two definitions are equivalent up to a change in the base measure, μ.

In the low privacy regime, the optimal test is based on the clamped log-likelihood ratio test from Canonne et al. (2019). Much of the work of both Theorem 4 and Theorem 5 goes into finding the right conditions for uniform convergence. There are several key quantities that determine the optimal local estimation rate, and when it is uniformly achievable.

- Define the radius of smoothness of A'' around θ as

$$\kappa(\theta) = \max\left\{ r \,\Big|\, \forall \theta' \in [\theta - r, \theta + r], \frac{A''(\theta)}{A''(\theta')} \in \left[\frac{1}{2}, 2\right] \right\}.$$

By the continuity of A'', $\kappa(\theta) > 0$ for all θ. Recall that $A''(\theta)$ is the variance of P_θ so $\kappa(\theta)$ is related to the smoothness of the variance. Our theorems will be strongest for families where $\kappa(\theta)$ is large for most θ of interest. Given the characterisation of $\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})$ in Theorem 4, if $\kappa(\theta)$ is large, then this means that the local rate varies slowly. The parameter $\kappa(\theta)$ affects the achievability in two main ways.

  - The form we give for the local estimation rate holds for sample sizes n above some threshold that depends on $\kappa(\theta)$. Specifically, one requirement is that $J_{\text{TV},P}^{-1}(\kappa(\theta)) \geq \frac{1}{\varepsilon n}$. That is, n must be large enough that if θ satisfies $TV(P_\theta, P_{\theta'}) \leq 1/\varepsilon n$, then $\theta' \in \Phi_\mu(\theta)$, where $\Phi_\mu(\theta) = [\theta - \kappa(\theta), \theta + \kappa(\theta)]$. This condition ensures that with high probability our private estimate lies within $\kappa(\theta)$ of θ.

  - In order for our procedure to succeed (that is, produce an accurate estimate) with probability at least $1 - \beta$, we require that there exists a constant $C > 0$ such that $\kappa(\theta) \geq \frac{1}{C} \cdot \frac{\sqrt{\log(2/\beta)}}{\sqrt{A''(\theta)}}$ for all θ. Under this condition, the distributions P_θ are sub-Gaussian, that is

$$\mathbb{P}_{x \sim P_\theta}\left(|x - A'(\theta)| \geq (2 + C)\sqrt{A''(\theta)}\sqrt{\ln(2/\beta)} \right) \leq \beta.$$

This light-tailed property ensures that with high probability a dataset sampled from P_θ lies mostly in an interval of width $O(\sqrt{A''(\theta)}\sqrt{\ln(2/\beta)})$. This allows us to limit the amount of noise added for privacy to also scale with the standard deviation $\sqrt{A''(\theta_0)}$. Without a light tailed assumption, additional noise needs to be added to maintain privacy, resulting in a worse estimation rate. We see this effect in estimating the parameter of a Bernoulli distribution, where the scale of the noise needed to maintain privacy scales with $\frac{1}{\varepsilon n}$, rather than $\frac{\sqrt{p(1-p)}}{\varepsilon n}$ which would be predicted by Theorem 4. The family of Bernoulli distributions fails to satisfy this assumption unless we constrain $\min(p, 1 - p)$ to be at least a constant.

- We will also require that the central standardised fourth moment is bounded. That is, there exists a constant ζ such that

$$\frac{\mathbb{E}_\theta\left(x - A'(\theta) \right)^4}{A''(\theta)^2} = \frac{\mathbb{E}_\theta\left(x - \mathbb{E}_\theta x \right)^4}{\text{var}(P_\theta)^2} \leq \zeta.$$

The central standardised fourth moment is also known as the *kurtosis*, this assumption allows us to give a *lower* bound on the tails of P_θ. That is, there exists a constant c such that $\text{Pr}_{P_\theta}[X \geq A'(\theta) + \frac{1}{2}\sqrt{A''(\theta)}] \geq c$. This assumption is required for our algorithm to properly estimate the standard deviation $\sqrt{A''(\theta)}$, which plays a crucial role in our estimator. It is possible that this assumption can be weakened with an improved private variance estimator.

4.1 Examples of Exponential Families

Before we move onto the proofs of Theorems 4 and 5, let us consider a few examples of simple exponential families and the implications of these theorems.

Example 6 (Gaussian mean with known variance). Note we can write

$$\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} = e^{x\frac{\mu}{\sigma}-\frac{1}{2}\left(\frac{\mu}{\sigma}\right)^2}e^{-\frac{1}{2}\frac{x^2}{\sigma^2}}.$$

Thus, if σ is known we can define an exponential family by $\theta = \mu/\sigma$, $A(\theta) = \frac{1}{2}\theta^2$ and $d\mu(x) = e^{-\frac{1}{2}\frac{x^2}{\sigma^2}}dx$. Notice that $A''(\theta) = 1$ so $\kappa(\theta) = \infty$ for all θ. Further, the central standardised fourth moment is 3. This is the ideal behavior for the conditions needed for Theorem 4 and Theorem 5 to hold. Therefore, the local minimax optimal rate for privately estimating θ is $\left[\frac{C_1}{\min\{n\varepsilon_n,\sqrt{n}\}}, \frac{C_2}{\min\{n\varepsilon_n,\sqrt{n}\}}\right]$, which implies that local minimax optimal rate for privately estimating the mean μ is

$$\left[\frac{C_1\sigma}{\min\{n\varepsilon_n,\sqrt{n}\}}, \frac{C_2\sigma}{\min\{n\varepsilon_n,\sqrt{n}\}}\right].$$

This recovers a result of Karwa and Vadhan (2018).

Example 7 (Poisson family). Recall that the Poission distribution, characterized by paramater $\lambda > 0$, assigns mass to nonnegative integers according to

$$\frac{\lambda^x e^{-\lambda}}{x!} = e^{x\ln\lambda-\lambda}\frac{1}{x!}.$$

We can define an exponential family by taking $\theta = \ln\lambda$, $A(\theta) = \lambda = e^\theta$ and the base measure μ that assigns mass $\frac{1}{x!}$ to all nonnegative integers x. Thus, $A''(\theta) = e^\theta = \lambda$, which implies that $\kappa(\theta) = \ln 2$ for all θ. Further, the central fourth moment is $3\lambda^2 + \lambda$; normalized by the square of the variance, we get $3 + \frac{1}{\lambda}$. Suppose there exists a constant $c > 0$ such that it is guaranteed that $\lambda > c$. Once n is sufficiently large, the local minimax optimal rate for privately estimating θ is $\left[\frac{C_1}{\sqrt{\lambda}\min\{n\varepsilon_n,\sqrt{n}\}}, \frac{C_2}{\sqrt{\lambda}\min\{n\varepsilon_n,\sqrt{n}\}}\right]$. Using the first-order Taylor approximation for $\lambda = e^\theta$, we see that the local minimax optimal rate for privately estimating λ is in

$$\left[\frac{C_1\sqrt{\lambda}}{\min\{n\varepsilon_n,\sqrt{n}\}}, \frac{C_2\sqrt{\lambda}}{\min\{n\varepsilon_n,\sqrt{n}\}}\right]$$

for constants C_1 and C_2 depending on c.

4.2 Basic Facts about Exponential Families

Let us begin by reviewing some basic properties of exponential families. A family $\{p_\theta(x)\}$ has monotone likelihood ratio if for all $\theta < \theta'$, $\frac{p_{\theta'}(x)}{p_\theta(x)}$ is a non-decreasing function of x. Exponential families have monotone likelihood ratio.

Lemma 4.1 (Lehmann & Romano, Lemma 3.4.2). *Let $\{p_\theta(x)\}$ be a family with monotone likelihood ratio, then*

- *If g is a nondecreasing function of x, then $\mathbb{E}_\theta g(x)$ is a nondecreasing function of θ.*

- *For any $\theta < \theta'$, and any t, $\mathbb{P}_\theta(x > t) < \mathbb{P}_{\theta'}(x > t)$.*

Corollary 4.2. *Assume that A'' is continuous. Then for all θ, the function $h \mapsto \mathrm{TV}(P_\theta, P_{\theta+h})$ is continuous and monotonically increasing on $h \geq 0$.*

Lemma 4.3. *For any θ, θ', if $|A'(\theta') - A'(\theta)| \leq \frac{1}{2} A''(\theta) \kappa(\theta)$ then $\theta' \in \Phi(\theta)$.*

Proof. If $|\theta - \theta'| \geq \kappa(\theta)$ then $|A'(\theta) - A'(\theta')| \geq \min_{\theta'' \in [\theta, \theta']} A''(\theta'') |\theta - \theta'| \geq \frac{1}{2} A''(\theta) \kappa(\theta)$. □

The following concentration inequality is proved in Section B.2.

Lemma 4.4. [Concentration Inequality for Exponential Families] *For all measures μ, $\theta \in \Phi$, and $\beta \in [0, 1]$,*

$$\mathbb{P}_{x \sim P_\theta} \left(|x - A'(\theta)| \geq 2\sqrt{A''(\theta)} \sqrt{\ln(2/\beta)} + \frac{\ln(2/\beta)}{\kappa(\theta)} \right) \leq \beta.$$

In particular, if $\kappa(\theta) \geq \frac{1}{C} \cdot \frac{\sqrt{\log(2/\beta)}}{\sqrt{A''(\theta)}}$,

$$\mathbb{P}_{x \sim P_\theta} \left(|x - A'(\theta)| \geq (2 + C)\sqrt{A''(\theta)} \sqrt{\ln(2/\beta)} \right) \leq \beta.$$

Lemma 4.4 shows that the tail of a distribution in an exponential family transitions from exponential to Gaussian as we move further out into the tail. How far into the tail one has to move for the tails to be Gaussian is a function of the standard deviation $\sqrt{A''(\theta)}$ and the stability of the standard deviation $\kappa(\theta)$.

4.3 Non-Private Estimation

Before we start designing our differentially private locally optimal estimator, let us first discuss the locally optimal estimator in the non-private setting. Given a sample $X \sim P_\theta^n$, let

$$\mathcal{A}_{opt}(X) = A'^{-1} \left(\frac{1}{n} \sum_{i=1}^n x_i \right).$$

Proposition 4.5 (Characterization of Optimal Local Estimation Rate in Non-Private Regime (Barndorff-Nielsen, 1978)). *$\mathcal{A}_{opt}$ is the optimal non-private estimation algorithm and, for all measures μ and $\theta_0 \in \Phi$, has rate*

$$\mathcal{E}_n^{\mathrm{loc}}(P_\theta, \mathcal{P}_\mu, \mathcal{Q}_{\mathrm{est}}, \theta_0) = \Theta\left(\frac{1}{\sqrt{A''(\theta_0)n}} \right).$$

Proposition 4.6. [Uniform Achievability in Non-Private Regime] *For all measures μ and $\theta_0 \in \Phi$ and $n \in \mathbb{N}$, if $n \geq \frac{36}{\kappa(\theta_0)^2 A''(\theta_0)}$, then*

$$\mathfrak{R}_n(\theta_0, \mathcal{A}_{opt}) \leq \frac{6}{\sqrt{nA''(\theta_0)}}.$$

Proof of Proposition 4.6. Note that $\mathbb{E}_{\theta_0}[\frac{1}{n} \sum_{i=1}^n x_i] = nA'(\theta_0)$ and $\mathrm{var}_{\theta_0}[\frac{1}{n} \sum_{i=1}^n x_i] = \frac{A''(\theta_0)}{n}$. Thus, with probability 8/9,

$$\left| \frac{1}{n} \sum_{i=1}^n x_i - A'(\theta_0) \right| \leq 3\sqrt{\frac{A''(\theta_0)}{n}}. \tag{5}$$

16

Now, $\frac{6}{\sqrt{nA''(\theta_0)}} \leq \kappa(\theta_0)$ implies that $3\sqrt{\frac{A''(\theta_0)}{n}} \leq \frac{1}{2}A''(\theta_0)\kappa(\theta_0)$. So, by Equation 4.3, if Equation (5) holds then $\mathcal{A}_{opt}(X) \in \Phi(\theta_0)$. Therefore, with probability 8/9-0.1,

$$
\begin{aligned}
\left| A'^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}x_i\right) - \theta_0 \right| &\leq \max_{t \in [\frac{1}{n}\sum_{i=1}^{n}x_i, A'(\theta_0)]} (A'^{-1})'(t) \left| \frac{1}{n}\sum_{i=1}^{n}x_i - A'(\theta_0) \right| \\
&\leq \max_{\theta' \in [A'^{-1}(\frac{1}{n}\sum_{i=1}^{n}x_i), \theta_0]} \frac{1}{A''(\theta')} 3\sqrt{\frac{A''(\theta_0)}{n}} \\
&\leq 6\frac{1}{\sqrt{nA''(\theta_0)}}. \qquad \square
\end{aligned}
$$

4.4 Initial Estimator

In both the high and low privacy settings, our first step will be to a get a crude estimate of $\mathbb{E}_{\theta_0}[x]$. This initial estimate will then be used to obtain a more refined estimate of θ_0. In both cases a sufficient initial estimate is given by a slight variation of the mean estimator given in Karwa and Vadhan (2018). Note that we could use this estimate of $A'(\theta_0) = \mathbb{E}_{\theta_0}[x]$ to get an estimate of θ_0 in the same way we did in $\mathcal{A}_{opt}$. However, the resulting estimator of θ is suboptimal by a factor of $\sqrt{\ln n}$. A full description of the initial estimator is given in Appendix B.3, we will denote it by $\mathcal{M}_{\zeta,C}$.

Theorem 8. *There exists constants $c > 0$ and $b > 0$ such that for all $\varepsilon > 0$, $\delta \in [0,1]$, $\zeta > 0$, and $C > 0$, there exists an (ε, δ)-DP algorithm, $\mathcal{M}_{\zeta,C}$, such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\frac{\mathbb{E}_{\theta_0}[|x-A'(\theta_0)|^3]}{\sqrt{A''(\theta_0)}^3} \leq \zeta$

- $\kappa(\theta_0) \geq \frac{1}{C}\frac{1}{\sqrt{A''(\theta_0)}}$

then for all $n \in \mathbb{N}$ such that $n \geq \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon}$, if $X \sim P_{\theta_0}^n$, then with probability 0.8,

$$
\left| \mathcal{M}_{\zeta,C}(X) - \frac{1}{n}\sum_{x \in X}x \right| \leq b(6+C)\left(\frac{\sqrt{A''(\theta_0)}}{n\varepsilon}\sqrt{\ln(n)}\right).
$$

Corollary 4.7. *There exists a constant $c > 0$ such that for all $\varepsilon_n = \Omega(\frac{\ln n}{n})$, $\delta \in [0,1]$, $\zeta > 0$, $C > 0$ there exists an (ε_n, δ)-DP algorithm, $\mathcal{M}_{\zeta,C}$ and constants $D > 0$ and $N \in \mathbb{N}$ such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\frac{\mathbb{E}_{\theta_0}[|x-A'(\theta_0)|^3]}{\sqrt{A''(\theta_0)}^3} \leq \zeta$

- $\kappa(\theta_0) \geq \frac{1}{C}\frac{1}{\sqrt{A''(\theta_0)}}$

then for all $n \in \mathbb{N}$ such that $n \geq \max\{N, \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon_n}\}$, with probability at least 0.8,

$$
|A'^{-1}(\mathcal{M}_{\zeta,C}(X)) - \theta_0| \leq D\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon_n\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right).
$$

17

Proof. By Theorem 4.5 and Theorem 8, there exists a constant $C > 0$ such that for all $\theta_0 \in \Phi$ satisfying the two conditions and $n \geq \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon_n}$ with probability 0.8 we have,

$$|\mathcal{M}_{\zeta,C}(X) - A'(\theta_0)| = \frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}}{n\varepsilon_n}\sqrt{\ln(n)}\right).$$

Now, since $\kappa(\theta_0)\sqrt{A''(\theta_0)} \geq B$ and $\varepsilon_n = \Omega(\frac{\ln n}{n})$, there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}\sqrt{\ln n}}{\varepsilon_n n}\right) \leq \frac{1}{2}A''(\theta_0)\kappa(\theta_0)$$

combined with Lemma 4.3 implies that $A'^{-1}(\mathcal{M}_{\zeta,C}(X)) \in \Phi(\theta_0)$. Therefore,

$$
\begin{aligned}
|(A')^{-1}(\mathcal{M}_{\zeta,C}(X)) - \theta_0| &\leq \max_{t \in [\mathcal{M}_{\zeta,C}(X), A'(\theta_0)]} (A')^{-1})'(t)|\mathcal{M}_{\zeta,C}(X) - \theta_0| \\
&= \max_{t \in [\mathcal{M}_{\zeta,C}(X), A'(\theta_0)]} \frac{1}{A''(A'^{-1}(t))}|\mathcal{M}_{\zeta,C}(X) - \theta_0| \\
&\leq 2\frac{1}{A''(\theta_0)}\frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}}{n\varepsilon_n}\sqrt{\ln(n)}\right) \\
&= C\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon_n\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right)
\end{aligned}
$$

$\square$

4.5 High Privacy Regime

We begin with the high privacy regime. While the noisy clamped log-likelihood ratio test is optimal in general for private simple hypothesis testing, a simpler test works in the high privacy regime. This test, which informs our design of the private estimator in this section, is a simple noisy counting test, and looks very similar to the optimal test in the local DP setting, presented in Duchi and Ruan (2018). The form of the estimation rate is also simpler in this section since, as we saw in Lemma 3.2, the sample complexity of the differentially private simple hypothesis testing takes on a simpler form in this regime.

4.5.1 Characterising the Optimal Local Estimation Rate in High Privacy Regime

Recall from Corollary 2 that the optimal local estimation rate in the high privacy regime is characterized by the L_1-information, defined in Equation (4): for $\beta \in [0, 1]$

$$J_{\mathrm{TV},\theta}^{-1}(\beta) = \sup\{|h| \mid \mathrm{TV}(P_\theta, P_{\theta+h}) \leq \beta\}.$$

The following lemma characterizes the L_1-information, and hence the optimal local estimation rate, in terms of properties of the one-parameter exponential family. The proof can be found in Appendix B.5.

Lemma 4.8. *For all $\zeta > 0$, there exists a constant C such that for all measures μ, $\theta_0 \in \Phi$, and $\beta \in [0, 1]$, if $\frac{\mathbb{E}_{P_\theta}(X - A'(\theta_0))^4}{A''(\theta_0)^2} \leq \zeta \leq \frac{9}{128\beta}$ and $\kappa(\theta) \geq J_{\mathrm{TV},\theta}^{-1}(\beta)$ then,*

$$J_{\mathrm{TV},\theta}^{-1}(\beta) \in \left[\frac{1}{\sqrt{2}}\frac{\beta}{\sqrt{A''(\theta)}}, C\frac{\beta}{\sqrt{A''(\theta)}}\right].$$

18

The following corollary follows immediately from Theorem 2 and Lemma 4.8.

Corollary 4.9 (Optimal Local Estimation Rate in the High Privacy Regime). *For all constants k and $\zeta > 0$, there exists a constants C_1, C_2 and C_3 such that for all measures μ and $\theta_0 \in \Phi$, if $\frac{\mathbb{E}_{P_\theta}(X-A'(\theta_0))^4}{A''(\theta_0)^2} \leq \zeta$, $\varepsilon_n \leq \frac{k}{\sqrt{n}}$, $\kappa(\theta) \geq J^{-1}_{\mathrm{TV},\theta}(\frac{C_3}{n\varepsilon_n})$ and $\frac{C_3}{n\varepsilon_n} \leq \frac{9}{128\zeta}$ then for all $n \in \mathbb{N}$,*

$$\mathcal{E}^{\mathrm{loc}}_n(P_{\theta_0}, \mathcal{P}_\mu, \mathcal{Q}_{\mathrm{est},\varepsilon}, \theta) \in \left[\frac{C_1}{n\varepsilon_n\sqrt{A''(\theta_0)}}, \frac{C_2}{n\varepsilon_n\sqrt{A''(\theta_0)}} \right].$$

4.5.2 Uniform Achievability in High Privacy Regime

In this section we show that in the high privacy regime, uniform achievability is achieved using a simple estimator based on estimating $\mathbb{P}_\theta(x > \mathbb{E}_\theta[x])$. Let $f_t(X) = \frac{1}{n}\sum_{i=1}^n \mathbb{1}_{x_i > t}$. Note that by Lemma 4.1,

$$g_t(\theta) := \mathbb{E}_\theta[f_t(x)] = \mathbb{P}_\theta(x > t)$$

is monotone and invertible in θ, so Algorithm 1 is well-defined. Our estimator $\mathcal{A}_{\mathrm{high}}$ requires as input a DP mean estimator, so $\hat{t}$ is an estimate of $\mathbb{E}_{\theta_0}[x]$. We refine the estimate $\hat{t}$ by using an estimator with lower sensitivity. Any sufficiently accurate mean estimator can be used for $\mathcal{M}$, but we note that the estimator described in Theorem 8 (derived from Karwa and Vadhan (2018)) is sufficient.

Algorithm 1 `Count`

Require: Sample $X \sim P^n_{\theta_0}$, $\hat{t}$, ε
  1: $\hat{\theta} = g^{-1}_{\hat{t}}(f_{\hat{t}}(X) + \mathrm{Lap}(\frac{1}{\varepsilon n}))$.

Algorithm 2 $\mathcal{A}_{\mathrm{high}}$

Require: Sample $X_1 \sim P^n_{\theta_0}$ and $X_2 \sim P^n_{\theta_0}$, an $(\varepsilon/2, \delta)$-DP mean estimator $\mathcal{M}$
  1: $\hat{t} = \mathcal{M}(X_1)$.
  2: $\hat{\theta} = \texttt{Count}(X_2, \hat{t}, \varepsilon/2)$.

Proposition 4.10 (Uniform Achievability in High Privacy Regime). *For any $\varepsilon_n > 0$ and $\delta_n \in [0, 1]$, $\mathcal{A}_{high}$ is $(\varepsilon_n, \delta_n)$-DP. Further, there exists a constant $c > 0$ such that for all constants $\zeta > 0$, $C > 0$, and $\delta \in [0, 1]$, there exists an estimator $\mathcal{M}_{\zeta,C}$ such that there exists constants $N \in \mathbb{N}$ and $C' > 0$ such that for all exponential families (i.e. any measure μ) and $\theta_0 \in \Phi$, if*

- $\Omega\left(\frac{\ln n}{n}\right) \leq \varepsilon_n \leq O\left(\frac{1}{\sqrt{n}}\right)$

- $\frac{\mathbb{E}_{P_{\theta_0}}(X-A'(\theta_0))^4}{A''(\theta_0)^2} \leq \zeta,$

- $\kappa(\theta_0) \geq \frac{1}{C}\frac{1}{\sqrt{A''(\theta_0)}}$

then for all $n \in \mathbb{N}$ such that $\kappa(\theta_0) \geq J^{-1}_{\mathrm{TV},\theta_0}(\frac{1}{\varepsilon_n n})$ and $n \geq \max\{N, \frac{c\zeta^2 \ln(1/\delta_n)}{\varepsilon_n}\}$,

$$\mathfrak{R}_n(\theta_0, \mathcal{A}_{\mathrm{high},C,\zeta}) \leq C' \cdot J^{-1}_{\mathrm{TV},\theta_0}\left(\frac{1}{\varepsilon_n n}\right) = O\left(\frac{1}{n\varepsilon_n\sqrt{A''(\theta_0)}}\right),$$

where $\mathcal{A}_{\mathrm{high},C,\zeta}$ is $\mathcal{A}_{\mathrm{high}}$ with initial mean estimator $\mathcal{M}_{\zeta,C}$.

Note that the upper bound in Proposition 4.10 matches the characterization of the optimal local estimation rate given in Lemma 4.9. Thus, Proposition 4.10 implies uniform achievability; there exists an algorithm that achieves the optimal local estimation rate for every $\theta_0 \in \Phi$. Note that many of the conditions required in Proposition 4.10 were already present in Corollary 4.7. Indeed, we will primarily use these conditions to ensure that our initial estimate $\mathcal{M}(X_1)$ is sufficiently accurate.

The proof proceeds by arguing that the test defined by the test statistic $f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)$ is good enough to distinguish P_{θ_0} from P_{θ_1} provided $|\theta_0 - \theta_1| \geq C J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{n\varepsilon_n}\right)$, and thus the estimator inherited by this tester will, with high probability, not output such a θ_1. The main technical challenge in this section will be showing that $f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)$ is a good test statistic for distinguishing between P_{θ_0} and P_{θ_1} when $|\theta_0 - \theta_1| \geq C J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{n\varepsilon_n}\right)$. We first show that if $\hat{t}$ is a good enough estimate for $A'(\theta_0)$, then $|\mathbb{P}_{\theta_0}(X > \hat{t}) - \mathbb{P}_{\theta_1}(X > \hat{t})| \approx \mathrm{TV}(P_{\theta_0}, P_{\theta_1})$. Note that this would be obviously true if $\hat{t} = \frac{A(\theta_0) - A(\theta_1)}{\theta_0 - \theta_1}$, so the majority of the work goes into proving that $\hat{t}$ is close enough to this ideal boundary point. Then, we show that the standard deviation of the statistic $f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right) \approx \frac{1}{\varepsilon n}$, so the test will distinguish P_{θ_0} and P_{θ_1} provided $\mathrm{TV}(P_{\theta_0}, P_{\theta_1}) = \Omega\left(\frac{1}{\varepsilon n}\right)$, as required. The following is the main technical lemma in this section, it is proved in Section B.6.

Lemma 4.11. *For all positive constants ζ and b, and $\Omega\left(\frac{\ln n}{n}\right) \leq \varepsilon_n \leq O\left(\frac{1}{\sqrt{n}}\right)$, there exists constants $N \in \mathbb{N}$ and $C > 0$ such for all measures μ, $\theta_0 \in \Phi$ and $n \in \mathbb{N}$ such that $n \geq N$, if*

1. $\max_{\theta \in \Phi(\theta_0)} \frac{\mathbb{E}_{P_\theta}[(X - A'(\theta_0))^4]}{A''(\theta)^2} \leq \zeta.$

2. $|\hat{t} - A'(\theta_0)| \leq b\left(\sqrt{\frac{A''(\theta_0)}{n}} + \frac{\sqrt{A''(\theta_0)\ln(n)}}{\varepsilon n}\right)$

3. $\kappa(\theta_0) \geq J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{8\sqrt{2}}{\varepsilon_n n}\right)$

then for all $\theta_1 \in \Phi(\theta_0)$ such that $|\theta_0 - \theta_1| \geq C J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$, then there exists a threshold τ such that the test

$$\begin{cases} P_{\theta_0} & \text{if } f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right) \leq \tau \\ P_{\theta_1} & \text{if } f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right) \geq \tau \end{cases}$$

distinguishes between θ_0 and θ_1. Furthermore, $\mathbb{E}_{P_{\theta_0}}[f_{\hat{t}}(X)] \leq \tau \leq \mathbb{E}_{P_{\theta_1}}[f_{\hat{t}}(X)]$.

The following corollary translates from the above testing result to an estimation result. The intuition for this conversion is that if the test statistic $f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)$ can distinguish θ_0 from θ_1, then the estimation algorithm $\texttt{Count}(X, \hat{t})$ is unlikely to output θ_1 when the data is drawn from P_{θ_0}.

Corollary 4.12. *For all positive constants ζ and b, and $\Omega\left(\frac{\ln n}{n}\right) \leq \varepsilon \leq O\left(\frac{1}{\sqrt{n}}\right)$, there exists constants $N \in \mathbb{N}$ and $C > 0$ such for all measures μ, $\theta_0 \in \Phi$ and $n \in \mathbb{N}$ such that $n \geq N$, if*

1. $\max_{\theta \in \Phi(\theta_0)} \frac{\mathbb{E}_{P_\theta}[(X - A'(\theta_0))^4]}{A''(\theta)^2} \leq \zeta,$

2. $|\hat{t} - A'(\theta_0)| \leq b\left(\sqrt{\frac{A''(\theta_0)}{n}} + \frac{\sqrt{A''(\theta_0)\ln(n)}}{\varepsilon n}\right),$

3. $\kappa(\theta_0) \geq J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right),$

20

then we have with probability 0.75,

$$|\texttt{Count}(X, \hat{t}) - \theta_0| \leq C \cdot J_{\mathrm{TV}, \theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right) = O\left(\frac{1}{n\varepsilon_n \sqrt{A''(\theta_0)}}\right).$$

Proof. Let C and N be as in Lemma 4.11. Let $X \sim P_{\theta_0}^n$ and suppose for sake of contradiction that $\mathcal{A}_{\mathrm{high}}(X) = \theta_1$ where $|\theta_1 - \theta_0| \geq C J_{\mathrm{TV}, \theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$. Then by definition, $f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right) = \mathbb{E}_{Y \sim P_{\theta_1}}[f_{\hat{t}}(Y)]$. Therefore, the test in Lemma 4.11 would have rejected P_{θ_0}, which only happens with probability 0.25. $\square$

The proof of Proposition 4.10 follows directly from Corollary 4.12 and Theorem 8. Note that the conditions in Proposition 4.10 are sufficient to ensure that the mean estimator from Theorem 8 is accurate enough that $\hat{t}$ satisfies the required condition in Corollary 4.12 with high probability.

4.6 Low Privacy Regime

In the low privacy regime, $\varepsilon = \Omega\left(\frac{1}{\sqrt{n}}\right)$, we claim that the optimal local estimation rate for privacy is asymptotically the same as the non-private rate. The estimator that achieves the optimal local estimation rate is derived from the noisy clamped log-likelihood test outlined in Section 3, the optimal algorithm for privately distinguishing two distributions.

4.6.1 Characterising the Optimal Local Estimation Rate in Low Privacy Regime

A main component of this claim is that for exponential families, the modulus of continuity of the non-private sample complexity SC is equal to the modulus of continuity of the private sample complexity, $\mathrm{SC}_{\varepsilon_n}$ in this parameter regime. The proof of the following proposition is found in Section B.7.

Proposition 4.13. *For all $\zeta > 0$, there exists positive constants k, C_1, C_2 and N such that for all measures μ, $\theta \in \Phi$ and $n \in \mathbb{N}$ such that $n \geq N$, if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$,

- $\max_{\theta' \in \Phi(\theta)} \frac{\mathbb{E}_{\theta'}((x - A'(\theta'))^4)}{A''(\theta')^2} \leq \zeta$

- $\kappa(\theta) \geq \frac{C_2}{\sqrt{nA''(\theta)}}$

then,

$$\frac{C_1}{\sqrt{nA''(\theta)}} \leq \omega_n(P_\theta, \mathcal{F}_\varepsilon^{\mathrm{test}}) \leq \frac{C_2}{\sqrt{nA''(\theta)}}.$$

This implies $\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\mathrm{test}}) = \Theta\big(\omega_n(P_\theta, \mathcal{F}^{\mathrm{test}})\big)$.

The following corollary follows immediately from Theorem 2.1 and Proposition 4.13.

Corollary 4.14 (Lower Bound in the Low Privacy Regime). *For all $\zeta > 0$, there exists positive constants k and constants C_1, C_2 and N such that under the same conditions as Proposition 4.13,*

$$\mathcal{E}_n^{\mathrm{loc}}(P_{\theta_0}, \mathcal{P}_\mu, \mathcal{F}_\varepsilon, \theta) \in \left[\frac{C_1}{2\sqrt{nA''(\theta_0)}}, \frac{C_2}{2\sqrt{nA''(\theta_0)}}\right].$$

21

4.6.2 Uniform Achievability in the Low Privacy Regime

Our estimator in the low privacy regime is based on the ncLLR test described in Section 3.2. Recall that the probability density function of P_θ has the form $p_\theta(x) = e^{\theta x - A(\theta)}$. Intuitively, the idea is that if we appropriately clamp the test statistic $x - \hat{t}$ then it contains roughly the same amount of information as the clamped log-likelihood ratio. Suppose $|\theta_1 - \theta_0| = \omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})$ then

$$
\begin{aligned}
(\theta_1 - \theta_0)[x - \hat{t}]_{-\varepsilon/\omega_n(P_{\theta_0},\mathcal{F}_\varepsilon^{\text{test}})}^{\varepsilon/\omega_n(P_{\theta_0},\mathcal{F}_\varepsilon^{\text{test}})} &= [(\theta_1 - \theta_0)x - (\theta_1 - \theta_0)\hat{t}]_{-\varepsilon}^{\varepsilon} \\
&= \left[\ln\frac{P_{\theta_1}(x)}{P_{\theta_0}(x)} + A(\theta_1) - A(\theta_0) - (\theta_1 - \theta_0)\hat{t}\right]_{-\varepsilon}^{\varepsilon} \\
&\approx \left[\ln\frac{P_{\theta_1}(x)}{P_{\theta_0}(x)}\right]_{-\varepsilon}^{\varepsilon}
\end{aligned}
$$

where the final approximation holds since $\hat{t} \approx \mathbb{E}_{\theta_0}[x] = A'(\theta)$. The final term is the optimal test for distinguishing the two distributions. The main technical difficulty is in showing that the approximations do not affect the sample complexity of the test too much. As in the high privacy setting, we need an initial mean estimator to estimate $\hat{t} \approx \mathbb{E}_{\theta_0}[x]$ and $\tilde{\alpha} \approx \omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})$. Again the estimator from Theorem 8 will suffice. For ease of notation, let $\alpha_n(\theta) = \frac{1}{\sqrt{nA''(\theta)}}$.

Algorithm 3 nCLLRE

Require: Sample $X_1 = (x_1, \cdots, x_n) \sim P_{\theta_0}^n$, $\hat{t}$, C, ε

1: $\tilde{\alpha} = \alpha_n(A'^{-1}(\hat{t}))$.
2: $f_{\tilde{\alpha}}(X_1) = \frac{1}{n}\sum_{i=1}^{n}[x_i - \hat{t}]_{\frac{-\varepsilon}{C\tilde{\alpha}}}^{\frac{\varepsilon}{C\tilde{\alpha}}}$
3: $\widehat{f} = f_{\tilde{\alpha}}(X_1) + \text{Lap}\left(\frac{2}{\varepsilon nC\tilde{\alpha}}\right))$
4: **return** $\hat{\theta} = \arg\min_\theta\left|\widehat{f} - \mathbb{E}_{X \sim P_\theta^n}(f_{\tilde{\alpha}}(X))\right|$

Algorithm 4 $\mathcal{A}_{\text{low}}$, estimating exponential family parameters

---

Require: Sample $X_1 \sim P_{\theta_0}^n$, $X_2 \sim P_{\theta_0}^n$, and an $(\varepsilon/2, \delta)$-DP mean estimator $\mathcal{M}$, C

1: $\hat{t} = \mathcal{M}(X_2)$.
2: **return** $\hat{\theta} = \text{nCLLRE}(X_1, \hat{t}, C, \varepsilon/2)$

Proposition 4.15. [Uniform Achievability in Low Privacy Regime] *For any $\varepsilon_n > 0, \delta_n \in [0,1], C > 0$, $\mathcal{A}_{low}$ is $(\varepsilon_n, \delta_n)$-DP. Further, for all $\zeta > 0$, $C > 0$, $\varepsilon_n > 0$, and $\delta_n \in [0,1]$, there exists an initial estimator $\mathcal{M}_{\zeta,C}$ such that there exists constants $N \in \mathbb{N}$, $k \geq 0$, $C' > 0$, $D > 0$ such that for all exponential families (i.e. any measure μ) and $\theta_0 \in \Phi$ if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((x - \mathbb{E}_\theta(x))^4)}{A''(\theta)^2} \leq \zeta$,

- $\kappa(\theta) \geq \frac{1}{C}\frac{1}{\sqrt{A''(\theta)}}$

then for all $n \in \mathbb{N}$ such that $n \geq \max\left\{N, \frac{c\zeta^2\ln(1/\delta_n)}{\varepsilon_n}, \frac{2\ln\left(\frac{1}{(DA''(\theta_0))^2}\right)}{(DA''(\theta_0))^2}\right\}$ then we have

$$
\mathfrak{R}_n(\theta_0, \mathcal{A}_{\text{low},C,\zeta}) \leq C'\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}}),
$$

where $\mathcal{A}_{\text{low},C,\zeta}$ is $\mathcal{A}_{\text{low}}$ with initial mean estimator $\mathcal{M}_{\zeta,C}$.

The main technical part of our proof is the following lemma, whose proof we defer to Section B.9.

Lemma 4.16. *For all $\zeta > 0$, there exists constants $N \in \mathbb{N}$, $k \geq 0$, $C > 0$, $b > 0$ and $D > 0$ such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((x - \mathbb{E}_\theta(x))^4)}{A''(\theta)^2} \leq \zeta,$

- $|\theta_0 - A'^{-1}(\hat{t})| \leq \min\{\kappa(\theta_0), b\varepsilon_n\sqrt{nA''(\theta_0)}\},$

then for all $n \in \mathbb{N}$ such that $n \geq \max\{N, \frac{D}{A''(\theta_0)(\kappa(\theta_0))^2}\}$ and all θ_1 such that $|\theta_1 - \theta_0| \geq C\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})$, there exists a threshold τ such that the test

$$
\begin{cases}
P_{\theta_0} & \text{if } f_{\tilde{\alpha}}(X) + \text{Lap}\left(\frac{2}{\varepsilon_n n C\tilde{\alpha}}\right) \leq \tau \\
P_{\theta_1} & \text{if } f_{\tilde{\alpha}}(X) + \text{Lap}\left(\frac{2}{\varepsilon_n n} C\tilde{\alpha}\right) \geq \tau
\end{cases}
\tag{6}
$$

distinguishes between P_{θ_0} and P_{θ_1} with n samples. Furthermore, τ can be chosen so $|\mathbb{E}_{X \sim P_{\theta_0}^n}[\hat{f}_{\tilde{\alpha}}(X)] - \tau| \leq |\mathbb{E}_{X \sim P_{\theta_1}^n}[\hat{f}_{\tilde{\alpha}}(X)] - \tau|.$

As in the previous section, in the translation from testing result in Lemma 4.16 to the bound on the estimation rate we argue that algorithm nCLLRE is unlikely to return θ such that $|\theta_0 - \theta| \geq D\omega_{n,\text{SC}_\varepsilon}(\theta_0)$ since this would result in the induced test failing, which is unlikely to occur.

Corollary 4.17. *For all $\zeta > 0$, there exists constants $N \in \mathbb{N}$, $k \geq 0$, $C > 0$, $b > 0$ and $D > 0$ such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((x - \mathbb{E}_{\theta_0}(x))^4)}{A''(\theta_0)^2} \leq \zeta,$

- $|\theta_0 - A'^{-1}(\hat{t})| \leq \min\{\kappa(\theta_0), b\varepsilon_n\sqrt{nA''(\theta_0)}\},$

then for all $n \in \mathbb{N}$ such that $n \geq \max\{N, \frac{D}{A''(\theta_0)(\kappa(\theta_0))^2}\}$ we have with probability 0.75

$$|\text{nCLLRE}(X, \hat{t}, C) - \theta_0| \leq C\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\text{test}})$$

Proof of Theorem 4.15. First, we claim that the estimator from Theorem 8 is sufficient to satisfy the conditions of Lemma 4.16. Suppose that $\mathcal{A}_{\text{low}}(X) = \theta_1$ and $|\theta_1 - \theta_0| \geq C\omega_{n,\text{SC}_\varepsilon}(\theta_0)$. Then, $|\hat{f} - \mathbb{E}_{\theta_1}(f_{\tilde{\alpha}})| \leq |\hat{f} - \mathbb{E}_{\theta_0}(f_{\tilde{\alpha}})|$, which implies that the test in Lemma 4.16 would have rejected θ_0, which only happens with probability 0.25. □

Finally, as in the previous section, Proposition 4.15 follows by showing that the estimator from Corollary 4.7 satisfies the conditions of Corollary 4.17.

Proof of Proposition 4.15. The proposition follows from a combination of Corollary 4.17 and Corollary 4.7. Let N, k, C_1, D_1 and D_2 be as in Corollary 4.17. Note that since $\frac{D_2}{A''(\theta_0)(\kappa(\theta_0))^2} \leq \frac{D_2}{B^2}$ we can assume that $N \geq \frac{D_2}{A''(\theta_0)(\kappa(\theta_0))^2}$. By Corollary 4.7, there exists a constant C_2 such that

$$|A'^{-1}(\mathcal{M}_{\zeta, C}(X)) - \theta_0| \leq C_2\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right).$$

Again since $\kappa(\theta_0)\sqrt{A''(\theta_0)} \geq B$, there exists constants N_1 and D such that for all $n > N_1$ if $\frac{\sqrt{\ln n}}{\sqrt{n}} \leq DA''(\theta_0)$ then,

$$C_2\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right) \leq \min\{\kappa(\theta_0), D_1\varepsilon\sqrt{nA''(\theta_0)}.$$

Thus, the estimator from Corollary 4.7 satisfies the requirements of Corollary 4.17 and so we are done. $\square$

5 Nonparametric Estimation of Functionals

In the previous section, the statistic of interest fully characterised the distribution. In this section, we will study the problem of estimating a statistic that does not characterise the distribution. While one can still define the local estimation rate as in Equation (1) in this setting, we will follow the standard set by Donoho and Liu (1991) by focusing on a slightly different notion of local estimation rate and modulus of continuity. Given a family of distributions $\mathcal{P}$ and a statistic $\theta : \mathcal{P} \to \mathbb{R}$, we define the *modulus of continuity with respect to θ* at any value $t \in \mathbb{R}$ as

$$\omega_n(t, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta) = \sup\left\{|\theta(P) - \theta(Q)| \mid \text{SC}_{\mathcal{F}^{\text{test}}}(P, Q) \geq n, \theta(P) = t, P, Q \in \mathcal{P}\right\}$$
$$= \max_{P \in \mathcal{P}, \theta(P) = t} \omega_n(P, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta). \tag{7}$$

The quantity is the worst case modulus of any distribution P in the family $\mathcal{P}$ such that $\theta(P) = t$. Donoho and Liu (1991) showed that for some estimation problems, in the non-private setting, one can design an algorithm that can universally achieve error, $\max_{\theta(P)=t} \mathfrak{R}_n(P, \hat{\theta}) \leq \omega_n(t, \mathcal{F}, \mathcal{P}, \theta)$.

Rather than using of simple hypothesis tests, Donoho and Liu (1991) turn to the problem of distinguishing

$$\mathcal{P}_{\leq t} = \{f | \theta(f) \leq t\} \quad \text{and} \quad \mathcal{P}_{\geq t+\Delta} = \{f | \theta(f) \geq t + \Delta\}. \tag{8}$$

Using this test, given the promise that $t \in [t_{\min}, t_{\max}]$ and setting $\Delta = |t_{\max} - t_{\min}|/3$, we can rule out the true parameter lying in either $[t_{\min}, t_{\min} + \Delta]$ or $[t_{\max} - \Delta, t_{\max}]$, reducing the search space by a factor of 2/3 for the next round. We will call this ternary search. If we run this algorithm for $\lceil\frac{\log(|t_1 - t_0|/\omega)}{\log(3/2)}\rceil$ steps, then the resulting error on the final estimate is at most ω. The total sample complexity of the estimator is the sum of the sample complexities of the tests performed at each step. This is at most a logarithmic factor times the sample complexity of the most stringent (final) test; however, in many cases it is quite a bit lower than that, since the exponential decrease in Δ can mean that the sample complexity of the final test dominates the overall the sample complexity of the estimator.

Algorithm 5 describes the algorithm that uses ternary search to estimate $\theta(P)$. In it, $T^*_{t,\Delta}$ denotes the optimal test for distinguishing the hypotheses in Eq. (8). The algorithm works for essentially

---

Algorithm 5 Ternary Search for Functional Estimation, $\hat{\theta}$

Require: Sample oracle for sample from P, $t_0 < t_1$, number of rounds k
1: $t_{\min} = t_0$, $t_{\max} = t_1$, and $\Delta_1 = |t_1 - t_0|/3$
2: **for** $i \in [k]$ **do**
3:      Let X_i be a sample from P large enough to distinguish between $\mathcal{P}_{\leq t_{\min} + \Delta_i}$ and $\mathcal{P}_{\geq t_{\min} + 2\Delta_i}$
     with probability $1 - \frac{1}{3k}$.
4:      **if** $T^*_{t_{\min} + \Delta_i, \Delta_i}(X_i) = \mathcal{P}_{\leq t_{\min} + \Delta_i}$ **then**
5:          $t_{\max} = t_{\max} - \Delta_i$
6:      **else**
7:          $t_{\min} = t_{\min} + \Delta_i$
8:      $\Delta_{i+1} = |t_{\max} - t_{\min}|/3$
9: **return** $t_{\min}$

---

any one-dimensional estimation problem. Under some conditions, we can compare its error on a given total sample size to the modulus of continuity (Equation 7 above). The main condition is that the difficulty of distinguishing between the two compound hypotheses in Eq. 8 should be captured by the difficulty of a *simple* hypothesis test: that is, for $t \in [t_0, t_1]$ and $\Delta > 0$, there should exist distributions P_0^* and P_1^* such that the number of samples needed to distinguish $\mathcal{P}_{\leq t}$ and $\mathcal{P}_{\geq t + \Delta}$ is equal to the number of samples needed to distinguish between P_0^* and P_1^*. Suppose we would like a test which competes with the modulus of continuity at some target sample size m. Given this condition, if we run Algorithm 5 for $\left\lceil \frac{\log\left(\omega_m(t, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta)/|t_1 - t_0|\right)}{\log(2/3)} \right\rceil$ rounds then it achieves error at most $\omega_m(t, \mathcal{F}^{\text{test}}, \mathcal{P}, \theta)$. The final, most stringent test requires a sample of size m, by definition.[4] The overall sample size is thus at most a logarithmic factor larger than m, and in some cases even closer to m.

Algorithm 5 can be adapted to the private setting by letting $T^*_{t, \Delta}$ be the optimal ε-DP test. In the following section, we adapt an example of this framework from Donoho and Liu (1991) to the private setting. The key difficulty is showing that the sample complexity of distinguishing $\mathcal{P}_{\leq t}$ and $\mathcal{P}_{\geq t + \Delta}$ is characterised by the hardest simple test in the private setting.

5.1 Tail Rates

In this section, we will consider estimation of the tail decay rate of a certain class of distributions. We will show that we can use the framework of Algorithm 5 to design an algorithm for tail decay rate estimation that is instance optimal up to a logarithmic factor. This section is a private analogue of Section 5 of Donoho and Liu (1991). The fact that much of Donoho and Liu's argument can be immediately adapted to the private setting is largely due to the fact that the optimal private test is similar to the optimal non-private test, which the optimal local estimator is built from. However, Donoho and Liu (1991) are able to alter Algorithm 5 to eliminate the logarithmic factor increase in the sample complexity. This alteration can not be adapted to the private setting. As a result, unlike in the non-private setting, where the resulting algorithm is instance optimal, our private analogue will be instance optimal up to a logarithmic factor.

As in Donoho and Liu, rather than estimate the rate at which the tail of the density approaches 0 as $x \to \infty$, we will consider a transformation of the problem to observations $Y_i = 1/X_i$. This leads us to estimating the rate at which a density approaches 0 as $x \to 0^+$. Let

---

[4]This discussion elides the dependency on the tests' error probability, which must be set sufficiently low to ensure that the decisions made at every round are correct; see Algorithm 5.

$\mathcal{P} = \textbf{Tails}(C_-, C_+, \delta, t_0, t_1, \gamma, p)$ be the set of distributions defined on $[0, \infty)$ with densities satisfying:

$$f(x) = Cx^t(1 + h(x)) \ \ 0 \le x \le \delta < 1,$$

where

$$0 < t_0 \le t \le t_1 < \infty, \ \ 0 < C_- \le C \le C_+ < \infty, \ \ \text{and} \ \ |h(x)| \le \gamma x^p.$$

For such a density function, the *tail rate* is given by $\theta(f) = t$. The statistic $\theta(f)$ is one dimensional and lies in the bounded interval $[t_0, t_1]$.

Theorem 9. *For every positive integer n:*

1. *For every $t \in [t_0, t_1]$ there exists a density f with $\theta(f) = t$ such that every differentially private estimator has error at least $\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)$ on some distribution in a neighborhood of f. That is,*

$$\mathcal{E}_n^{\text{loc}}(f, \mathcal{F}_\varepsilon, \mathcal{P}, \theta) \ge \frac{1}{2}\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)$$

2. *There exists an ε-differentially private estimator $\hat{\theta}$ with the following property. For all $t \in [t_0, t_1]$, if $k^*(n) = \lceil \log_{3/2}(\frac{|t_1 - t_0|}{\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)}) \rceil$ and $N = n \cdot k^*(n) \cdot \lceil \log k^*(n) \rceil$, then $\hat{\theta}$ has error at most $\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)$ when run on N samples. That is,*

$$\mathfrak{R}_N(f, \hat{\theta}) \le \omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)$$

Typically, we expect N is $O(n \cdot \log(n|t_1 - t_0|) \cdot \log\log(n|t_1 - t_0|))$ for $\varepsilon \le 1$ when $\omega_{n/\log k^*(n)}(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta) \approx 1/poly(n)$. As in Donoho et al., the first aspect of the proof is to show that the corresponding testing problem has the property that the difficulty of distinguishing between two intervals in $[t_0, t_1]$ is captured by the difficulty of the hardest two-point testing problem. That is, for $t \in [t_0, t_1]$ and $\Delta > 0$, there exists distributions f_0^* and f_1^* such that the number of samples needed to distinguish $\mathcal{P}_{\le t}$ and $\mathcal{P}_{\ge t+\Delta}$ is equal to the number of samples needed to distinguish between f_0^* and f_1^*. Donoho and Liu (1991) showed that in the non-private case, the distributions f_0^* and f_1^* satisfy the following conditions:

$$f_0^*(x) = C_- x^t(1 - \gamma x^p), \quad x \le a_1(t, \Delta),$$
$$f_1^*(x) = C_+ x^{t+\Delta}(1 + \gamma x^p), \quad x \le a_1(t, \Delta),$$

and

$$\frac{f_0^*(x)}{f_1^*(x)} = \frac{f_0^*(a_1(t, \Delta))}{f_1^*(a_1(t, \Delta))}, \quad x > a_1(t, \Delta),$$

and

$$\frac{f_0^*(a_1(t, \Delta))}{f_1^*(a_1(t, \Delta))} = \frac{1 - \int_0^{a_1(t, \Delta)} f_0^*(v)dv}{1 - \int_0^{a_1(t, \Delta)} f_1^*(v)dv}.$$

In the following lemma we mirror their proof to show that the same distributions also satisfy this condition in the private case.

We say a real random variable X is stochastically less than a random variable Y, denoted $X \preceq Y$, if

$$\mathbb{P}(X > x) \le \mathbb{P}(Y > x) \quad \text{for all } x \in (-\infty, \infty).$$

Lemma 5.1. *If u is non-decreasing and $X \preceq Y$ then $u(X) \preceq u(Y)$. If $X_i \preceq Y_i$ for $i = 1, \cdots, n$ then $\sum_{i=1}^{n} X_i \preceq \sum_{i=1}^{n} Y_i$.*

Lemma 5.2. *For the distribution f_0^* and f_1^* described above, we have*

$$\mathrm{SC}_\varepsilon(\mathcal{P}_{\leq t}, \mathcal{P}_{\geq t+\Delta}) = \mathrm{SC}_\varepsilon(f_0^*, f_1^*).$$

Furthermore, the test statistic for distinguishing $\mathcal{P}_{\leq t}$ and $\mathcal{P}_{\geq t+\Delta}$ is the clamped log-likelihood ratio between f_0^* and f_1^*.

Proof. To simplify notation, we will let $a = a_1(t, \Delta)$. Recall from Proposition 3.1, the optimal test statistic for distinguishing between f_0^* and f_1^* is given by the noisy clamped log-likelihood ratio. Now, let

$$L_{t,\Delta}(x) = \begin{cases} \ln\left(\frac{C_+}{C_-} x^\Delta \frac{1+\gamma x^p}{1-\gamma x^p}\right), & 0 < x < a, \\ \ln\left(\frac{C_+}{C_-} a^\Delta \frac{1+\gamma a^p}{1-\gamma a^p}\right), & x \geq a, \end{cases}$$

be the log likelihood ratio and let $L_{t,\Delta}^\varepsilon(x) = [L_{t,\Delta}(x)]_{-\varepsilon}^\varepsilon$ be the clamped log likelihood ratio (where $L_{t,\Delta}(x)$ is projected onto the domain $[-\varepsilon, \varepsilon]$). Donoho and Liu (1991), showed that the distribution of $L_{t,\Delta}(x)$ where x is distributed according to f_0^* is stochastically larger than $L_{t,\Delta}(x)$ where x is distributed according to f for any $f \in \mathcal{P}_{\leq t}$. Since clamping is a monotone function, this implies that the same property holds for $L_{t,\Delta}^\varepsilon$. Now, for $X = (x_1, \cdots, x_n)$,

$$L_{n,t,\Delta}^\varepsilon(X) = \sum_{i=1}^{n} L_{t,\Delta}^\varepsilon(x_i)$$

and $\widetilde{L_{n,t,\Delta}^\varepsilon}(X) = L_{n,t,\Delta}^\varepsilon(X) + \mathrm{Lap}(2)$. Again by Lemma 5.1, we have that among all the distributions in $\mathcal{P}_{\leq t}$, $\widetilde{L_{n,t,\Delta}^\varepsilon}(X)$ is stochastically largest at f_0^*. Similarly, we can show that among all the distributions in $\mathcal{P}_{\geq t+\Delta}$, $\widetilde{L_{n,t,\Delta}^\varepsilon}(X)$ is stochastically smallest at f_1^*. Therefore, if we design a test which `accepts` (corresponding to choosing $\mathcal{P}_{\leq t}$) if $\widetilde{L_{n,t,\Delta}}(X) \leq 1$ and `rejects` otherwise. Then we must have

$$\sup_{f \in \mathcal{P}_{\leq t}} \mathbb{P}_f(\texttt{rejects}) = \mathbb{P}_{f_0^*}(\texttt{rejects}), \text{ and } \sup_{f \in \mathcal{P}_{\geq t+\Delta}} \mathbb{P}_f(\texttt{accepts}) = \mathbb{P}_{f_1^*}(\texttt{accepts}).$$

It follows that if $n \geq \mathrm{SC}_\varepsilon(f_0^*, f_1^*)$ then $\mathbb{P}_{f_0^*}(\texttt{rejects}) \leq 1/3$ and $\mathbb{P}_{f_1^*}(\texttt{accepts}) \leq 1/3$ since the clamped log-likelihood test is optimal for distinguishing f_0^* and f_1^*. Thus this distinguishes $\mathcal{P}_{\leq t}$ and $\mathcal{P}_{\geq t+\Delta}$ with the same number of samples. $\qquad\square$

Lemma 5.2 gives us the tools we need to use Algorithm 5 to obtain a near instance optimal (up to logarithmic factors) differentially private algorithm for tail bound estimation.

Theorem 10. *For all $n \in \mathbb{N}$, let $k^*(n) = \lceil \log_{3/2}(\frac{|t_1 - t_0|}{\omega_n(t, \mathcal{F}_\varepsilon^{\mathrm{test}}, \mathcal{P}, \theta)}) \rceil$, $N = n \cdot k^*(n) \cdot \lceil \log k^*(n) \rceil$ and $\hat{\theta}$ be the output of Algorithm 5 run for $k^*(n)$ rounds. Then*

$$\mathfrak{R}_N(f, \hat{\theta}) \leq \omega_n(t, \mathcal{F}_\varepsilon^{\mathrm{test}}, \mathcal{P}, \theta).$$

Proof of Theorem 10. We can think of Algorithm 5 as at each step dividing the distance between $t_{\min}$ and $t_{\max}$ by $2/3$ and concluding that the true value t^* lies between $t_{\min}$ and $t_{\max}$. Thus, in order to show that $\mathfrak{R}_N(f, \hat{\theta}) \leq \omega_n(t, \mathcal{F}_\varepsilon^{\mathrm{test}}, \mathcal{P}, \theta)$, it suffices to show that in order to run for

$k^*(n) = \left\lceil \log_{\frac{3}{2}} \left(\frac{|t_1 - t_0|}{\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)} \right) \right\rceil$ iterations, it suffices to have at least $N = n \cdot \lceil \log k^*(n) \rceil \cdot k^*(n)$ samples. In order to make the correct decision with probability $1/3k^*(n)$, it suffices for the last iteration to use $n \cdot \lceil \log k^*(n) \rceil$ samples. Since the hypothesis test at the last step has the largest sample size, $n \cdot \lceil \log k^*(n) \rceil \cdot k^*(n)$ samples is sufficient to run $k^*(n)$ rounds. $\qquad \square$

Donoho and Liu (1991) are able to alter Algorithm 5 to remove the logarithmic factor. Given $\Delta > 0$, they define the following estimator

$$\theta_{n,\Delta}^*(X) = \frac{\Delta}{2} + \sup\{t \mid L_{n,t,\Delta}(X) \le 0\}, \tag{9}$$

which outputs the largest t for such the hypothesis $\mathcal{P}_{\le t}$ would be accepted. This estimator is well-defined since $L_{n,t,\Delta}$ is deterministic, a crucial distinction when we move to the private setting. Donoho et al. show that if Δ is sufficiently small then $L_{n,t,\Delta}$ is monotonically decreasing in t for a given x, which implies that given input distribution f such that $\theta(f) = t$, the estimation algorithm $T_{n,\Delta}^*$ has error rate $\omega_n(t, \mathcal{F}_\varepsilon^{\text{test}}, \mathcal{P}, \theta)$.

The estimator $\theta_{n,\Delta}^*(X)$ in eqn (9) can be viewed as performing the test $L_{n,t,\Delta}(X)$ on every t value and outputting the threshold point, where the test flips from `accept` to `reject`. We can not replicate this directly in the private setting both because the private test is stochastic (so there is likely to be some false negatives and false positives), and because performing the test on every t value would result in an unreasonably large privacy cost.

Acknowledgments

References

J. Acharya, Z. Sun, and H. Zhang. Differentially private testing of identity and closeness of discrete distributions. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 6879–6891, Red Hook, NY, USA, 2018. Curran Associates Inc.

N. Alon, R. Livni, M. Malliaris, and S. Moran. Private pac learning implies finite littlestone dimension. In *ACM Symposium on Theory of Computing*, STOC '19, pages 852–860, 2019.

Apple Differential Privacy Team. Learning with privacy at scale. *Apple Machine Learning Journal*, 1(8), 2017. https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialprivacysystem.pdf.

H. Asi and J. C. Duchi. Near instance-optimality in differential privacy. *arXiv preprint arXiv:2005.10630*, 2020.

B. Balle, G. Barthe, and M. Gaboardi. Privacy profiles and amplification by subsampling. *Journal of Privacy and Confidentiality*, 10(1), Jan. 2020.

R. F. Barber and J. C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.

O. Barndorff-Nielsen. *Introductory Theory of Exponential Families*, chapter 8. John Wiley & Sons, Ltd, 1978. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118857281.ch8.

A. Beimel, K. Nissim, and E. Omri. Distributed private data analysis: Simultaneously solving how and what. In *International Cryptology Conference*, CRYPTO '08, pages 451–468, Santa Barbara, CA, USA, 2008. https://arxiv.org/abs/1103.2626.

J. Blocki, A. Blum, A. Datta, and O. Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *4th ACM Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 87–96, Berkeley, CA, USA, 2013. ACM.

M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. volume 9985, pages 635–658, 11 2016. ISBN 978-3-662-53640-7.

M. Bun, J. Ullman, and S. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *ACM Symposium on the Theory of Computing*, STOC '14, pages 1–10, New York, NY, USA, 2014. https://arxiv.org/abs/1311.3158.

M. Bun, K. Nissim, and U. Stemmer. Simultaneous private learning of multiple concepts. 11 2015. doi: 10.1145/2840728.2840747.

C. Canonne, G. Kamath, A. McMillan, A. Smith, and J. Ullman. The structure of optimal private tests for simple hypotheses. In *51st Annual ACM Symposium on Theory of Computing*, pages 310–321, 06 2019. doi: 10.1145/3313276.3316336.

T.-H. H. Chan, E. Shi, and D. Song. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)*, 14(3):26, 2011.

S. Chen and S. Zhou. Recursive mechanism: Towards node differential privacy and unrestricted joins. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, page 653–664, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320375.

D. L. Donoho and R. C. Liu. Geometrizing rates of convergence, I. *Technical Report 137*, 1987.

D. L. Donoho and R. C. Liu. Geometrizing rates of convergence, II. *Ann. Statist.*, 19(2):633–667, 06 1991. doi: 10.1214/aos/1176348114. URL https://doi.org/10.1214/aos/1176348114.

J. Duchi and F. Ruan. The right complexity measure in locally private estimation: It is not the fisher information. *arXiv:1806.05756*, 06 2018.

J. Duchi, M. Jordan, and M. Wainwright. Local privacy and statistical minimax rates. In *IEEE Symposium on Foundations of Computer Science*, FOCS '13, pages 429–438, Berkeley, CA, USA, 2013. https://arxiv.org/abs/1302.3203.

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

C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, STOC '09, pages 371–380. ACM, 2009.

C. Dwork and G. N. Rothblum. Concentrated differential privacy. *ArXiv*, abs/1603.01887, 2016.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Conference on Theory of Cryptography*, TCC '06, pages 265–284, New York, NY, USA, 2006.

C. Dwork, A. Smith, T. Steinke, J. Ullman, and S. Vadhan. Robust traceability from trace amounts. In *IEEE Symposium on Foundations of Computer Science*, FOCS '15, 2015.

A. Edmonds, A. Nikolov, and J. Ullman. The power of factorization meisms in local and central differential privacy. In *ACM Symposium on the Theory of Computing*, STOC '20, pages 425–438, Chicago, IL, USA, 2020. https://arxiv.org/abs/1911.08339.

G. Kamath and J. Ullman. A primer on private statistics. *arXiv preprint arXiv:2005.00010*, 2020.

G. Kamath, J. Li, V. Singhal, and J. Ullman. Privately learning high-dimensional distributions. In *Annual Conference on Learning Theory*, COLT '19. JMLR.org, 2019.

G. Kamath, V. Singhal, and J. Ullman. Private mean estimation of heavy-tailed distributions. *https://arxiv.org/abs/2002.09464*, 2020.

V. Karwa and S. P. Vadhan. Finite sample differentially private confidence intervals. volume abs/1711.03908 of *Innovations in Theoretical Computer Science '18*, 2018.

S. P. Kasiviswanathan and A. D. Smith. On the 'semantics' of differential privacy: A bayesian formulation. *CoRR*, abs/0803.3946, 2008.

S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? In *IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 531–540, Philadelphia, PA, USA, 2008. https://arxiv.org/abs/0803.0924.

S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. D. Smith. Analyzing graphs with node differential privacy. In *10th IACR Theory of Cryptography Conference*, TCC '13, pages 457–476, Tokyo, Japan, 2013. Springer.

B. McMahan and A. Thakurta. Blog post: Federated Learning with Formal Differential Privacy Guarantees, Feb. 2022. https://ai.googleblog.com/2022/02/federated-learning-with-formal.html. Last accessed: June 2022.

S. Narayanan. Private high-dimensional hypothesis testing. *ArXiv*, abs/2203.01537, 2022.

K. Nissim, S. Raskhodnikova, and A. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 30th annual ACM Symposium on Theory of Computing, STOC*, pages 75–84, 2007.

A. Rohde and L. Steinberger. Geometrizing rates of convergence under local differential privacy constraints. *arXiv:1805.01422*, 2018.

A. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing*, STOC '11, pages 813–822, New York, NY, USA, 2011. ACM.

T. Steinke and J. Ullman. Tight lower bounds for differentially private selection. In *IEEE Symposium on Foundations of Computer Science*, FOCS '17, 2017.

S. Vadhan. *The Complexity of Differential Privacy*, pages 347–450. 04 2017. ISBN 978-3-319-57047-1.

A Proofs for Section 3.2

A.1 Proof of Proposition 3.1

Proposition 3.1. *If $\varepsilon = O(1)$ then for all $a = \Theta(\varepsilon)$ and $b = \Theta(\varepsilon)$, there exists constants C_1 and C_2 such that for all distributions P and Q,*

$$\mathrm{SC}_{\mathrm{ncLLR}_a^b}(P,Q) \in [C_1 \cdot \mathrm{SC}_\varepsilon(P,Q), C_2 \cdot \mathrm{SC}_\varepsilon(P,Q)].$$

As in Canonne et al. (2019), define

$$\tau = \tau(P,Q) = \max\left\{\int \max\{P(x) - e^\varepsilon Q(x), 0\}dx, \int \max\{Q(x) - e^\varepsilon P(x), 0\}dx\right\}$$

and assume without loss of generality that $\tau = \int \max\{P(x) - e^\varepsilon Q(x), 0\}dx$. Let $0 \leq \varepsilon' \leq \varepsilon$ be the smallest value such that $\tau = \int \max\{Q(x) - e^{\varepsilon'} P(x), 0\}dx$. Define $P' = \frac{1}{1-\tau}\min\{P, e^\varepsilon Q\}$ and $Q' = \frac{1}{1-\tau}\min\{Q, e^{\varepsilon'} P\}$.

Lemma A.1. *For any $\varepsilon \in [0,1]$, and distributions P and Q with the same support,*

$$\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon'}^\varepsilon}(P,Q) = \Theta(\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon}^\varepsilon}(P,Q)).$$

Proof. First note that since $\mathrm{ncLLR}_{-\varepsilon'}^\varepsilon$ is an optimal test up to a constant factor (Canonne et al., 2019),

$$\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon'}^\varepsilon}(P,Q) = O(\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon}^\varepsilon}(P,Q)).$$

Canonne et al. (2019) show that the following two inequalities are sufficient to prove that $\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon'}^\varepsilon}(P,Q) = \Theta\left(\frac{1}{\tau\varepsilon + (1-\tau)H^2(P',Q')}\right)$:

$$\mathbb{E}_{P^n}[\mathrm{cLLR}_{-\varepsilon'}^\varepsilon] - \mathbb{E}_{Q^n}[\mathrm{cLLR}_{-\varepsilon'}^\varepsilon] \geq \Omega(n(\tau\varepsilon + (1-\tau)H^2(P',Q')))$$

and

$$\max\left\{\mathbb{E}_P\left[\left[\ln \frac{P(x)}{Q(x)}\right]_{-\varepsilon'}^\varepsilon\right]^2, \mathbb{E}_Q\left[\left[\ln \frac{P(x)}{Q(x)}\right]_{-\varepsilon'}^\varepsilon\right]^2\right\} \leq O(\tau\varepsilon + (1-\tau)H^2(P',Q'))$$

We first note that the gap between the expectations increases when we move from $\mathrm{cLLR}_{-\varepsilon'}^\varepsilon$ to $\mathrm{cLLR}_{-\varepsilon}^\varepsilon$. If $P(x) < Q(x)$, then $\mathrm{cLLR}_{-\varepsilon}^\varepsilon(x) < \mathrm{cLLR}_{-\varepsilon'}^\varepsilon(x)$, and otherwise, $\mathrm{cLLR}_{-\varepsilon}^\varepsilon(x) = \mathrm{cLLR}_{-\varepsilon'}^\varepsilon(x)$. Thus, $(P(x) - Q(x))\mathrm{cLLR}_{-\varepsilon}^\varepsilon(x) \geq (P(x) - Q(x))\mathrm{cLLR}_{-\varepsilon'}^\varepsilon(x)$ and

$$\begin{aligned}
\mathbb{E}_{P^n}[\mathrm{cLLR}_{-\varepsilon}^\varepsilon] - \mathbb{E}_{Q^n}[\mathrm{cLLR}_{-\varepsilon}^\varepsilon] &= n\int (P(x) - Q(x))\,\mathrm{cLLR}_{-\varepsilon}^\varepsilon(x)dx \\
&\geq n\int (P(x) - Q(x))\,\mathrm{cLLR}_{-\varepsilon'}^\varepsilon(x)dx \\
&= \mathbb{E}_P[\mathrm{cLLR}_{-\varepsilon'}^\varepsilon] - \mathbb{E}_Q[\mathrm{cLLR}_{-\varepsilon'}^\varepsilon] \\
&\geq \Omega(n(\tau\varepsilon + (1-\tau)H^2(P',Q')))
\end{aligned}$$

The next step is to bound the second moment of $[\ln \frac{P(x)}{Q(x)}]_{-\varepsilon}^\varepsilon$ under both P and Q. Define the following three regions:

$$A = \left\{x \mid \ln \frac{P(x)}{Q(x)} \geq -\varepsilon'\right\} \quad B = \left\{x \mid \ln \frac{P(x)}{Q(x)} \in [-\varepsilon, -\varepsilon']\right\} \quad \text{and} \quad C = \left\{x \mid \ln \frac{P(x)}{Q(x)} \leq -\varepsilon\right\},$$

so $\text{cLLR}^\varepsilon_{-\varepsilon}(x) = \text{cLLR}^\varepsilon_{-\varepsilon'}(x) + g(x)$ where

$$g(x) = \begin{cases} 0 & \text{if } x \in A \\ \ln \frac{P(x)}{Q(x)} + \varepsilon' & \text{if } x \in B \\ -\varepsilon + \varepsilon' & \text{if } x \in C \end{cases}$$

Let us first bound the expectation of $-g(x)$ under P. Note that since $\ln x$ is concave, $\ln x \leq e^{-\varepsilon'} x + \varepsilon' - 1$, since the right hand side is the tangent at $x = e^{\varepsilon'}$.

$$\int_{x \in B} P(x)(-g(x)) dx = \int_{x \in B} P(x)\left(\ln \frac{Q(x)}{P(x)} - \varepsilon'\right) dx$$

$$\leq \int_{x \in B} P(x)\left(e^{-\varepsilon'} \frac{Q(x)}{P(x)} + \varepsilon' - 1 - \varepsilon'\right) dx$$

$$= \int_{x \in B} e^{-\varepsilon'} Q(x) - P(x) dx$$

$$= e^{-\varepsilon'} \int_{x \in B} (Q(x) - e^{\varepsilon'} P(x)) dx$$

$$\leq \int_{x \in B} (Q(x) - e^{\varepsilon'} P(x)) dx.$$

Note also that since $\varepsilon > \varepsilon' > 0$, and $\varepsilon - \varepsilon' \leq e^\varepsilon - e^{\varepsilon'}$ so

$$\int_{x \in C} P(x)(-g(x)) dx = \int_{x \in C} P(x)(\varepsilon - \varepsilon') dx$$

$$\leq \int_{x \in C} P(x)(e^\varepsilon - e^{\varepsilon'}) dx$$

$$\leq \int_{x \in C} Q(x) - e^{\varepsilon'} P(x) dx.$$

where the last inequality follows from the fact that $e^\varepsilon P(x) < Q(x)$ for $x \in C$. Therefore, $\mathbb{E}_P[-g(x)] \leq \int_{x \in B \cup C}(Q(x) - e^{\varepsilon'} P(x)) dx = \tau$. Now,

$$\mathbb{E}_P\left[\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon}\right]^2 = \int P(x)\left[\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon'}\right]^2 dx + 2\int P(x)g(x)\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon'} dx$$

$$+ \int P(x)g(x)^2 dx.$$

Now, whenever $g(x) \neq 0$, $\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon'} \in [-\varepsilon, -\varepsilon']$ so

$$\int P(x)g(x)\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon'} dx \leq \varepsilon \int P(x)(-g(x)) dx.$$

Thus,

$$\mathbb{E}_P\left[\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon}\right]^2 \leq \mathbb{E}_P\left[\left[\ln \frac{P(x)}{Q(x)}\right]^\varepsilon_{-\varepsilon'}\right]^2 + 2\varepsilon \int P(x)(-g(x)) dx + \int P(x)g(x)^2 dx$$

$$\leq O(\tau\varepsilon + (1-\tau)H^2(P', Q')) + 2\varepsilon \mathbb{E}_P[-g(x)] + (\varepsilon - \varepsilon')\mathbb{E}_P[-g(x)]$$

$$= O(\tau\varepsilon + (1-\tau)H^2(P', Q'))$$

Also on the support of $g(x)$ $(B \cup C)$, $Q(x) = e^{\varepsilon'} P(x) + F(x)$ where $\int_{B \cup C} F(x) = \tau$. So,

$$\mathbb{E}_Q[-g(x)] = \int (e^{\varepsilon'} P(x) + F(x))(-g(x))dx \leq e^{\varepsilon'} \mathbb{E}_P[-g(x)] + \varepsilon \int_{B \cup C} F(x)dx = O(\tau),$$

since $\varepsilon = O(1)$. Then, by the same argument as above, $\mathbb{E}_Q\left[\left[\log \frac{P(x)}{Q(x)}\right]_{-\varepsilon}^{\varepsilon}\right]^2 = O(\tau\varepsilon + (1 - \tau)H^2(P', Q'))$, and so we are done.

$\square$

Lemma A.2. *For all constants C_1, C_2, there exists constants C_3, C_4 such that for all $\varepsilon \in [0, 1]$, distributions P and Q, if $-a, b \in [C_1\varepsilon, C_2\varepsilon]$ then*

$$\mathrm{SC}_{\mathrm{ncLLR}_a^b}(P, Q) \in [C_3 \mathrm{SC}_\varepsilon(P, Q), C_4 \mathrm{SC}_\varepsilon(P, Q)].$$

Proof. Define $f(\varepsilon) = \mathrm{SC}_\varepsilon(P, Q)$. Let C be a constant, and let us first establish some simple bounds on $f(C\varepsilon)$. Suppose $C \leq 1$. Then since $f(\varepsilon)$ is the optimal sample complexity, we have $f(\varepsilon) \leq f(C\varepsilon)$. Also, by the secrecy-of-the-sample lemma for DP algorithms (Balle et al., 2020), we have that $f(C\varepsilon) \leq \frac{1}{C}f(\varepsilon)$. If $C \geq 1$ then by the same arguments we have $f(C\varepsilon) \in [\frac{1}{C}f(\varepsilon), f(\varepsilon)]$.

Now, note that by Lemma A.1 and the proof of Theorem 2.5 in Canonne et al. (2019), there exist constants D_1, D_2, D_3, D_4 such that for all ε and P, Q, $\mathrm{SC}_{\mathrm{ncLLR}_{-\varepsilon}^\varepsilon}(P, Q) \in [D_1 \mathrm{SC}_\varepsilon(P, Q), D_2 \mathrm{SC}_\varepsilon(P, Q)]$ and

$$|\mathbb{E}_P^n(\mathrm{cLLR}_{-\varepsilon}^\varepsilon(X)) - \mathbb{E}_Q^n(\mathrm{cLLR}_{-\varepsilon}^\varepsilon(X))| \geq D_3 n \frac{1}{f(\varepsilon)}$$

and

$$\mathrm{var}_P(\mathrm{cLLR}_\varepsilon^\varepsilon(X)) \leq D_4 n \frac{1}{f(\varepsilon)}.$$

Note that

$$|\mathbb{E}_P^n(\mathrm{cLLR}_a^b(X)) - \mathbb{E}_Q^n(\mathrm{cLLR}_a^b(X))| \geq |\mathbb{E}_P^n(\mathrm{cLLR}_{-C_1\varepsilon}^{C_1\varepsilon}(X)) - \mathbb{E}_Q^n(\mathrm{cLLR}_{-C_1\varepsilon}^{C_1\varepsilon}(X))|$$

$$\geq D_3 n \frac{1}{f(C_1\varepsilon)}$$

$$\geq D_3 \min\{1, C_1\} n \frac{1}{f(\varepsilon)}$$

where the first inequality follows by definition. Also,

$$\mathrm{var}_P(\mathrm{cLLR}_a^b(X)) \leq \mathrm{var}_P(\mathrm{cLLR}_{-C_2\varepsilon}^{C_2\varepsilon}(X)) \leq D_4 n \frac{1}{f(C_2\varepsilon)} \leq D_4 \max\{1, C_2\} n \frac{1}{f(\varepsilon)}$$

where again the first inequality follows by definition. Therefore, there exist constants C_3, C_4, depending on $C_1, C_2, D_1, D_2, D_3, D_4$ such that

$$\mathrm{SC}_{\mathrm{ncLLR}_a^b}(P, Q) \in [C_3 f(\varepsilon), C_4 f(\varepsilon)]$$

$\square$

B Proofs about Exponential Families

B.1 Proof of Corollary 4.2

Corollary 4.2. *Assume that A'' is continuous. Then for all θ, the function $h \mapsto \mathrm{TV}(P_\theta, P_{\theta+h})$ is continuous and monotonically increasing on $h \geq 0$.*

Proof of Corollary 4.2. Recall from Lemma 4.1 that one-parameter exponential families have monotone likelihood ratio so for any $\theta < \theta'$ and any t, $\mathbb{P}_\theta(x > t) < \mathbb{P}_{\theta'}(x > t)$. Suppose $\theta_1 \leq \theta_2 \leq \theta_3$ then

$$
\begin{aligned}
\mathrm{TV}(P_{\theta_1}, P_{\theta_2}) &= \mathbb{P}_{\theta_2}\left(x \geq \frac{A(\theta_2) - A(\theta_1)}{\theta_2 - \theta_1} \right) - \mathbb{P}_{\theta_1}\left(x \geq \frac{A(\theta_2) - A(\theta_1)}{\theta_2 - \theta_1} \right) \\
&\leq \mathbb{P}_{\theta_3}\left(x \geq \frac{A(\theta_2) - A(\theta_1)}{\theta_2 - \theta_1} \right) - \mathbb{P}_{\theta_1}\left(x \geq \frac{A(\theta_2) - A(\theta_1)}{\theta_2 - \theta_1} \right) \\
&\leq \mathrm{TV}(P_{\theta_1}, P_{\theta_3}).
\end{aligned}
$$

Thus, monotonicity holds. To prove continuity, note that

$$
\begin{aligned}
\mathrm{KL}(P_\theta \| P_{\theta+h}) &= \int p_\theta(x) \ln \frac{p_\theta(x)}{p_{\theta+h}(x)} d\mu \\
&= \int p_\theta(x)(\theta x - A(\theta) - (\theta + h)x + A(\theta + h)) d\mu \\
&= -h\mathbb{E}_\theta[x] + A(\theta + h) - A(\theta) \\
&= A(\theta + h) - A(\theta) - hA'(\theta) \\
&\leq \max_{\theta' \in [\theta, \theta+h]} A''(\theta')h^2.
\end{aligned}
$$

Therefore, by Pinsker's inequality,

$$
\mathrm{TV}(P_\theta, P_{\theta+h}) \leq |h| \sqrt{\max_{\theta' \in [\theta, \theta+h]} A''(\theta')}. \tag{10}
$$

Therefore,

$$
\mathrm{TV}(P_\theta, P_{\theta+h_1}) - \mathrm{TV}(P_\theta, P_{\theta+h_2}) \leq \mathrm{TV}(P_{\theta+h_1}, P_{\theta+h_2}) \leq |h_1 - h_2| \sqrt{\max_{\theta' \in [\theta+h_1, \theta+h_2]} A''(\theta')}.
$$

Since A'' is continuous, there exists γ such that if $|h_1 - h_2| \leq \gamma$ then $\sqrt{\max_{\theta' \in [\theta+h_1, \theta+h_2]} A''(\theta')} \leq 2\sqrt{A''(\theta + h_1)}$. Thus for any $\rho \geq 0$, if $|h_1 - h_2| \leq \min\left\{ \gamma, \frac{\rho}{2\sqrt{A''(\theta+h_1)}} \right\}$ then $\mathrm{TV}(P_\theta, P_{\theta+h_1}) - \mathrm{TV}(P_\theta, P_{\theta+h_2}) \leq \rho$. Therefore, $h \to \mathrm{TV}(P_\theta, P_{\theta+h})$ is continuous and monotone. $\square$

B.2 Proof of Lemma 4.4

Lemma 4.4. [Concentration Inequality for Exponential Families] *For all measures μ, $\theta \in \Phi$, and $\beta \in [0, 1]$,*

$$
\mathbb{P}_{x \sim P_\theta}\left(|x - A'(\theta)| \geq 2\sqrt{A''(\theta)}\sqrt{\ln(2/\beta)} + \frac{\ln(2/\beta)}{\kappa(\theta)} \right) \leq \beta.
$$

In particular, if $\kappa(\theta) \geq \frac{1}{C} \cdot \frac{\sqrt{\log(2/\beta)}}{\sqrt{A''(\theta)}}$,

$$\mathbb{P}_{x \sim P_\theta}\left(|x - A'(\theta)| \geq (2 + C)\sqrt{A''(\theta)}\sqrt{\ln(2/\beta)}\right) \leq \beta.$$

Proof of Lemma 4.4. Recall that $\int e^{\theta x} d\mu = e^{A(\theta)}$. Let $\lambda \leq \kappa(\theta)$ then

$$
\begin{aligned}
\mathbb{E}_{P_\theta}\left[e^{\lambda x}\right] &= \int e^{\lambda x} e^{\theta x - A(\theta)} d\mu \\
&= e^{-A(\theta)} \int e^{(\lambda + \theta)x} d\mu \\
&= e^{A(\theta + \lambda) - A(\theta)}.
\end{aligned}
$$

Now,

$$
\begin{aligned}
\mathbb{E}_{P_\theta}\left[e^{|\lambda(x - A'(\theta))|}\right] &= \mathbb{E}_{P_\theta}\left[e^{\lambda(x - A'(\theta))}\mathbf{1}_{\lambda(x - A'(\theta)) \geq 0}\right] + \mathbb{E}_{P_\theta}\left[e^{-\lambda(x - A'(\theta))}\mathbf{1}_{\lambda(x - A'(\theta)) \leq 0}\right] \\
&\leq \mathbb{E}_{P_\theta}\left[e^{\lambda x - \lambda A'(\theta)}\right] + \mathbb{E}_{P_\theta}\left[e^{-\lambda x + \lambda A'(\theta)}\right] \\
&= e^{A(\theta + \lambda) - A(\theta) - \lambda A'(\theta)} + e^{A(\theta - \lambda) - A(\theta) + \lambda A'(\theta)} \\
&\leq e^{\frac{\lambda^2}{2} \max_{\theta' \in [\theta, \theta + \lambda]} A''(\theta')} + e^{\frac{\lambda^2}{2} \max_{\theta' \in [\theta - \lambda, \theta]} A''(\theta')} \\
&\leq 2 e^{\lambda^2 A''(\theta)}
\end{aligned}
$$

where the last inequality follows since $\lambda \leq \kappa(\theta)$. Therefore for any $u > 0$,

$$
\begin{aligned}
\mathbb{P}_{P_\theta}(|x - A'(\theta)| \geq u) &= \mathbb{P}_{P_\theta}(e^{\lambda|x - A'(\theta)|} \geq e^{\lambda u}) \\
&\leq \frac{\mathbb{E}_{P_\theta}\left[e^{|\lambda(x - A'(\theta))|}\right]}{e^{\lambda u}} \\
&\leq \frac{2 e^{\lambda^2 A''(\theta)}}{e^{\lambda u}}
\end{aligned}
$$

where the first inequality follows from Markov's inequality. Let $u = 2\sqrt{A''(\theta)}\sqrt{\ln(2/\beta)} + \frac{\ln(2/\beta)}{\kappa(\theta)}$ and $\lambda = \min\left\{\kappa(\theta), \frac{\sqrt{\ln(2/\beta)}}{\sqrt{A''(\theta)}}\right\}$, so $u \geq \lambda A''(\theta) + \frac{\ln(2/\beta)}{\lambda}$ and

$$
\begin{aligned}
\mathbb{P}_{P_\theta}\left(|x - A'(\theta)| \geq 2\sqrt{A''(\theta)}\sqrt{\log(2/\beta)} + \frac{\ln(2/\beta)}{\kappa(\theta)}\right) &\leq \mathbb{P}_{P_\theta}\left(|x - A'(\theta)| \geq \lambda A''(\theta) + \frac{\ln(2/\beta)}{\lambda}\right) \\
&\leq 2 e^{\lambda^2 A''(\theta) - \lambda^2 A''(\theta) - \ln(2/\beta)} \\
&= 2 e^{-\ln(2/\beta)} \\
&= \beta.
\end{aligned}
$$

The second statement follows immediately.

$\square$

B.3 Initial Estimator - Proof of Theorem 8

Theorem 8. *There exists constants $c > 0$ and $b > 0$ such that for all $\varepsilon > 0$, $\delta \in [0,1]$, $\zeta > 0$, and $C > 0$, there exists an (ε, δ)-DP algorithm, $\mathcal{M}_{\zeta, C}$, such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\dfrac{\mathbb{E}_{\theta_0}[|x - A'(\theta_0)|^3]}{\sqrt{A''(\theta_0)}^3} \leq \zeta$

- $\kappa(\theta_0) \geq \dfrac{1}{C} \dfrac{1}{\sqrt{A''(\theta_0)}}$

then for all $n \in \mathbb{N}$ such that $n \geq \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon}$, if $X \sim P_{\theta_0}^n$, then with probability 0.8,

$$\left| \mathcal{M}_{\zeta, C}(X) - \frac{1}{n} \sum_{x \in X} x \right| \leq b(6 + C) \left(\frac{\sqrt{A''(\theta_0)}}{n\varepsilon} \sqrt{\ln(n)} \right).$$

In this section we slightly generalise the algorithm and analysis from Karwa and Vadhan (2018) beyond Gaussian distributions. We will show that their algorithm provides accurate estimates of the mean of sufficiently nice exponential families. This algorithm first estimates the variance of the distribution, then estimates the mean. Both steps of the estimation are performed using differentially private histogram queries.

Let $\rho = \mathbb{E}_P[|X - \mathbb{E}_P(x)|^3]$ be the absolute third moment of P and σ be the standard deviation. Since the algorithm of Karwa and Vadhan (2018) is designed for Gaussian distributions we will use the following lemma that describes the rate of convergence of the central limit theorem.

Lemma B.1 (Berry-Esseen theorem)**.** *Let $n \in \mathbb{N}$ and $X_1, \cdots, X_n$ be iid samples from a distribution P, and $\rho = \mathbb{E}_P[|X - \mathbb{E}_P(x)|^3]$. Set $S_n = \frac{1}{n}\sum_{j=1}^n X_j$, $\mu = \mathbb{E}_P[x]$ and $\sigma^2 = var(P)$, and let $Y \sim \mathcal{N}(\mu, \frac{\sigma^2}{n})$ then for some absolute constant $\nu > 0$,*

- *(Uniform) For all $a > 0$,*
$$|\mathbb{P}[S_n \leq a] - \mathbb{P}[Y \leq a]| \leq \frac{\nu\rho}{\sigma^3\sqrt{n}}$$

- *(Non-uniform) For all $a > 0$,*
$$|\mathbb{P}[S_n \leq a] - \mathbb{P}[Y \leq a]| \leq \frac{\nu\rho}{(1 + |a|)^3\sigma^3\sqrt{n}}.$$

Lemma B.2 (Histogram Learner Dwork et al. (2006); Bun et al. (2015); Vadhan (2017))**.** *For all $K \in \mathbb{N}$ and domain Ω, for any collection of disjoint bins $B_1, \cdots, B_K$ defined on Ω, $n \in \mathbb{N}$, $\varepsilon \geq 0$, $\delta \in (0, 1/n)$, $\lambda > 0$ and $\beta \in (0,1)$ there exists an (ε, δ)-DP algorithm $M : \Omega^n \to \mathbb{R}^K$ such that for every distribution D on Ω, if*

1. *$X_1, \cdots, X_N \sim D$ and $p_k = \mathbb{P}(X_i \in B_k)$*

2. *$(\tilde{p_1}, \cdots, \tilde{p_K}) = M(X_1, \cdots, X_n)$ and*

3.
$$n \geq \max\left\{ \min\left\{ \frac{8}{\varepsilon\lambda} \ln\left(\frac{2K}{\beta}\right), \frac{8}{\varepsilon\lambda} \ln\left(\frac{4}{\beta\delta}\right) \right\}, \frac{1}{2\lambda^2} \ln\left(\frac{4}{\beta}\right) \right\}$$

then,

$$\mathbb{P}_{X \sim D, M}(\max_k |\tilde{p}_k - p_k| \leq \lambda) \geq 1 - \beta \quad \text{and},$$

$$\mathbb{P}_{X \sim D, M}(\arg\max_k \tilde{p}_k = j) \leq \begin{cases} np_j + 2e^{-(\varepsilon n/8)\cdot(\max_k p_k)} & \text{if } K < 2/\delta \\ np_j & \text{if } K \geq 2/\delta \end{cases}$$

where the probability is taken over the randomness of M and the data $X_1, \cdots, X_n$.

Algorithm 6 Variance estimator

Require: Sample $X = (x_1, \ldots, x_n) \sim P, \varepsilon, \delta, \sigma_{\min}, \sigma_{\max}, \beta, \zeta$.
1: Let $\phi = \lceil (600\nu\zeta)^2 \rceil$, where ν is the absolute constant from Lemma B.1.
2: If

$$n < c\phi \min\left\{ \frac{1}{\varepsilon} \ln\left(\frac{\ln\left(\frac{\sigma_{\max}}{\sigma_{\min}}\right)}{\beta} \right), \frac{1}{\varepsilon} \ln\left(\frac{1}{\delta\beta} \right) \right\},$$

output $\perp$, where c is an absolute constant whose existence is ensured by Lemma B.2.
3: Divide $[\sigma_{\min}, \sigma_{\max}]$ into bins of exponentially increasing length. The bins are of the form $B_j = (2^j, 2^{j+1}]$ for $j = j_{\min}, \cdots, j_{\max}$, where $j_{\max} = \lceil \ln_2 \frac{\sigma_{\max}}{\sqrt{\phi}} \rceil + 1$ and $j_{\min} = \lfloor \ln_2 \frac{\sigma_{\min}}{\sqrt{\phi}} \rfloor - 2$.
4: Let $Z_i = \frac{1}{\phi} \sum_{j=1}^{\phi} x_{(i-1)\phi+j}$ for $i = 1, \cdots, \lfloor n/\phi \rfloor$.
5: Let $Y_i = Z_{2i} - Z_{2i-1}$ for $i = 1, \cdots, \lfloor n/2\phi \rfloor$
6: Run the histogram learner of Lemma B.2 with privacy parameters (ε, δ) and bins $B_{j_{\min}}, \cdots, B_{j_{\max}}$ on input $|Y_1|, \cdots, |Y_n|$ to obtain noisy estimates $\tilde{p_{j_{\min}}}, \cdots, \tilde{p_{j_{\max}}}$. Let

$$\hat{l} = \arg\max \tilde{p}_j$$

7: Output $\hat{\sigma} = 2^{\hat{l}+2}\sqrt{\phi}$.

Note in particular that the use of approximate (ε, δ)-DP allows us to set the $K = \infty$, while the sample complexity remains finite. The following lemma states that provided ρ/σ^3 is bounded, Algorithm 6 can estimate the standard deviation up to a multiplicative constant.

Lemma B.3. *For all $n \in \mathbb{N}$, $\sigma_{\min} < \sigma_{\max} \in [0, \infty], \varepsilon > 0, \delta \in (0, \frac{1}{n}], \beta \in (0, 1/2), \zeta > 0$, Algorithm 6 is (ε, δ)-DP and satisfies that if $X_1, \cdots, X_n$ are iid draws from P, where P has standard deviation $\sigma \in [\sigma_{\min}, \sigma_{\max}]$ and $\frac{\rho}{\sigma^3} \leq \zeta$ then if*

$$n \geq c\zeta^2 \min\left\{ \frac{1}{\varepsilon} \ln\left(\frac{\ln\left(\frac{\sigma_{\max}}{\sigma_{\min}}\right)}{\beta} \right), \frac{1}{\varepsilon} \ln\left(\frac{1}{\delta\beta} \right) \right\},$$

(where c is a universal constant), we have

$$\mathbb{P}_{X \sim P, M}(\sigma \leq \hat{\sigma} \leq 8\sigma) \geq 1 - \beta.$$

Proof of Lemma B.3. This proof follows almost directly from Theorem 3.2 of Karwa and Vadhan (2018). Note that each Y_i is sampled from a distribution with mean 0 and variance $\frac{2\sigma^2}{\phi}$, and in addition is the sum of ϕ independent random variables. As in Karwa and Vadhan (2018), there exists a bin B_l with label $l \in (\lfloor \ln_2 \frac{\sigma_{\min}}{\sqrt{\phi}} \rfloor - 1, \lceil \ln_2 \frac{\sigma_{\max}}{\sqrt{\phi}} \rceil)$ such that $\frac{\sigma}{\sqrt{\phi}} \in (2^l, 2^{l+1}] = B_l$. Define,

$$p_j = \mathbb{P}(|Y_i| \in B_j).$$

Sort the p_j's as $p_{(1)} \geq p_{(2)} \geq \cdots$ and let $j_{(1)}, j_{(2)}, \cdots$ be the corresponding bins. Then, the following two facts imply the result (as in Karwa and Vadhan (2018)).

Fact 1: The bins corresponding to the largest and second largest mass $p_{(1)}, p_{(2)}$ are $(j_{(1)}, j_{(2)}) \in \{(l, l-1), (l, l+1), (l+1, l)\}$.

Fact 2: $p_{(1)} - p_{(3)} > 1/300$.

Now, let $W_i \sim N(0, 2\frac{\sigma^2}{\phi})$ and let $q_i, q_{(i)}$ be the corresponding probabilities for W_i. Then Karwa and Vadhan (2018) showed that:

- The bins corresponding to the largest and second largest mass $q_{(1)}, q_{(2)}$ are $(j_{(1)}, j_{(2)}) \in \{(l, l-1), (l, l+1), (l+1, l)\}$.

- $q_{(1)} - q_{(3)} > 1/100$.

By Lemma B.1, since $\phi = \lceil (600\nu\zeta)^2 \rceil$, for all j, $|p_j - q_j| \leq 1/300$. Therefore, $\{p_{(1)}, p_{(2)}\} = \{q_{(1)}, q_{(2)}\}$, which implies both Fact 1 and Fact 2. $\square$

Algorithm 7 Range estimator

Require: Sample $X = (x_1, \ldots, x_n) \sim P, \varepsilon > 0, \delta \in [0,1], \beta \in (0, 1/2), R \in (0, \infty), \sigma > 0, B > 0$.
1: Let $r = \lceil \frac{R}{2\sigma} \rceil$. Divide $[-R - \sigma/2, R + \sigma/2]$ into $2r + 1$ bins of length at most 2σ each in the following manner - bin B_j equals $(2(j - 0.5)\sigma, 2(j + 0.5)\sigma]$, for $j \in \{-r, \cdots, r\}$.
2: Run the histogram learner of Lemma B.2 with privacy parameters (ε, δ) and bins $B_{-r}, \cdots, B_r$ on input $x_1, \cdots, x_n$ to obtain noisy estimates $\tilde{p}_{-r}, \cdots, \tilde{p}_r$. Let

$$\hat{l} = \arg\max_{j=-r,\cdots,r} \tilde{p}_j.$$

3: Output $(x_{\min}, x_{\max})$, where

$$x_{\min} = 2\sigma\hat{l} - \sigma(6 + C)\sqrt{\ln(4n/\beta)}, \quad x_{\max} = 2\sigma\hat{l} + \sigma(6 + C)\sqrt{\ln(4n/\beta)}.$$

Theorem 11. *For all $n \in \mathbb{N}, \sigma > 0, \varepsilon > 0, \delta \in [0,1], \beta \in (0, 1/2), R \in (0, \infty], C \geq 0$, Algorithm 7 is (ε, δ)-DP. For all measures μ and $\theta \in \Phi$, if $x_1, \cdots, x_n$ are sampled from P_θ where $A'(\theta) \in (-R, R)$, $A''(\theta) \leq \sigma^2$, $\kappa(\theta) \geq \frac{1}{C} \frac{\sqrt{\log(2/\beta)}}{\sqrt{A''(\theta)}}$, and*

$$n \geq c \min\left\{ \frac{1}{\varepsilon} \ln\left(\frac{R}{\sigma\beta} \right), \frac{1}{\varepsilon} \ln\left(\frac{1}{\delta\beta} \right) \right\},$$

(where c is a universal constant), then we have

$$\mathbb{P}_{x \sim P_\theta, M}(\forall i, \ x_{min} \leq x_i \leq x_{max}) \geq 1 - \beta$$

and ,

$$|x_{\max} - x_{\min}| = \sigma(6 + C)\sqrt{\ln(4n/\beta)}.$$

Proof. By Lemma 4.4 and a union bound, with probability $1 - \beta/2$, we have

$$\forall i \ : \ |x_i - \mu| \leq \sigma(2 + C)\sqrt{\ln(4n/\beta)}.$$

Next, as in the proof of Theorem 3.1 from Karwa and Vadhan (2018), we want to show that with probability $1 - \beta/2$, we have
$$|\mu - \hat{l}\sigma| \leq 2\sigma.$$

Note that by Chebyshev's inequality $\mathbb{P}[|x - \mu| \leq 2\sigma] \geq 3/4$ so there exists a pair of neighbouring bins B_j, B_{j+1} such that $\mathbb{P}[x \in B_j \cup B_{j+1}] \geq 3/4$ and $\mu \in B_j \cup B_{j+1}$. Also, for all $i \notin \{j, j+1\}$, $\mathbb{P}[x \in B_i] \leq 1/4$. Let $j^* = \arg\max_{k=j,j+1} \mathbb{P}[x \in B_k]$. Then $\mathbb{P}[x \in B_{j^*}] \geq 3/8$, and $\mathbb{P}[x \in B_{j^*}] - \mathbb{P}[x \in B_i] \geq 1/8$ for all $i \notin \{j, j+1\}$. Then by Lemma B.2, setting $\lambda = 1/8$, n is large enough that with probability $1 - \beta/2$, $\hat{l} \in \{j, j+1\}$. Therefore, $|\mu - 2\hat{l}\sigma| \leq 4\sigma$. Therefore, with probability $1 - \beta$, for all i,
$$|x_i - 2\hat{l}\sigma| \leq |x_i - \mu| + |\mu - 2\hat{l}\sigma| \leq \sigma(2 + C)\sqrt{\ln(4n/\beta)} + 4\sigma \leq \sigma(6 + C)\sqrt{\ln(4n/\beta)}.$$

$\square$

As in Karwa and Vadhan (2018) combining these two algorithms gives us an estimator of the range with unknown variance. Since this range contains all the data points with high probability, we can clamp the data to this range, and add noise proportional to the width of the range. Note that we can remove the dependence on the range $[0, R]$ and $[\sigma_{\min}, \sigma_{\max}]$ in the sample complexity since $n \geq \frac{c\zeta}{\varepsilon} \ln(\frac{1}{\delta\beta})$ is sufficient to ensure that the bounds required in both Theorem 11 and Lemma B.3 hold. This completes the proof of Theorem 8.

Algorithm 8 Initial mean estimator, $\mathcal{M}_{\zeta, C}$

Require: $X_1, \cdots, X_n, \beta, \varepsilon, \delta, \zeta, B$.
 1: If
$$n < \frac{c \max\{\zeta^2, 1\}}{\varepsilon} \ln\left(\frac{1}{\delta\beta}\right),$$

 output 0.
 2: Run Algorithm 6 to obtain an estimate $\hat{\sigma}$ of the variance with privacy parameters (ε, δ), $\sigma_{\min} = 0$ and $\sigma_{\max} = \infty$.
 3: Run Algorithm 7 with privacy parameters (ε, δ), $R = \infty$, and standard deviation $\hat{\sigma}$ to obtain a range $[X_{\min}, X_{\max}]$.
 4: Let
$$Y_i = \begin{cases} X_i & \text{if } X_i \in [X_{\min}, X_{\max}] \\ X_{\max} & \text{if } X_i > X_{\max} \\ X_{\min} & \text{if } X_i > X_{\min} \end{cases}$$

 5: Let Z be a Laplace random variable with mean 0 and scale parameter $\frac{X_{\max} - X_{\min}}{\varepsilon n}$.
 6: Output
$$\frac{\sum_{i=1}^n Y_i}{n} + Z.$$

B.4 Proof of Corollary 4.7

Corollary 4.7. *There exists a constant $c > 0$ such that for all $\varepsilon_n = \Omega(\frac{\ln n}{n})$, $\delta \in [0, 1]$, $\zeta > 0$, $C > 0$ there exists an (ε_n, δ)-DP algorithm, $\mathcal{M}_{\zeta, C}$ and constants $D > 0$ and $N \in \mathbb{N}$ such that for all measures μ and $\theta_0 \in \Phi$ if*

- $\dfrac{\mathbb{E}_{\theta_0}[|x-A'(\theta_0)|^3]}{\sqrt{A''(\theta_0)}^3} \leq \zeta$

- $\kappa(\theta_0) \geq \dfrac{1}{C}\dfrac{1}{\sqrt{A''(\theta_0)}}$

then for all $n \in \mathbb{N}$ such that $n \geq \max\{N, \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon n}\}$, with probability at least 0.8,

$$|A'^{-1}(\mathcal{M}_{\zeta,C}(X)) - \theta_0| \leq D\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon_n\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right).$$

Proof. By Theorem 4.5 and Theorem 8, there exists a constant $C > 0$ such that for all $\theta_0 \in \Phi$ satisfying the two conditions and $n \geq \frac{c\zeta^2 \ln(1/\delta)}{\varepsilon}$ with probability 0.8 we have,

$$|\mathcal{M}_{\zeta,C}(X) - A'(\theta_0)| = \frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}}{n\varepsilon}\sqrt{\ln(n)}\right).$$

Now, since $\kappa(\theta) \geq \frac{1}{C}\frac{\sqrt{\log(2/\beta)}}{\sqrt{A''(\theta)}}$ and $\varepsilon = \Omega(\frac{\ln n}{n})$, there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}\sqrt{\ln n}}{\varepsilon n}\right) \leq \frac{1}{2}A''(\theta_0)\kappa(\theta_0)$$

combined with Lemma 4.3 implies that $A'^{-1}(\mathcal{M}_{\zeta,C}(X)) \in \Phi(\theta_0)$. Therefore,

$$\begin{aligned}
|(A')^{-1}(\mathcal{M}_{\zeta,C}(X)) - \theta_0| &\leq \max_{t\in[\mathcal{M}_{\zeta,C}(X),A'(\theta_0)]}(A')^{-1})'(t)|\mathcal{M}_{\zeta,C}(X) - \theta_0| \\
&= \max_{t\in[\mathcal{M}_{\zeta,C}(X),A'(\theta_0)]}\frac{1}{A''(A'^{-1}(t))}|\mathcal{M}_{\zeta,C}(X) - \theta_0| \\
&\leq 2\frac{1}{A''(\theta_0)}\frac{C}{2}\left(\frac{\sqrt{A''(\theta_0)}}{\sqrt{n}} + \frac{\sqrt{A''(\theta_0)}}{n\varepsilon}\sqrt{\ln(n)}\right) \\
&= C\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right)
\end{aligned}$$

$\square$

B.5 Proof of Lemma 4.8

Lemma 4.8. *For all $\zeta > 0$, there exists a constant C such that for all measures μ, $\theta_0 \in \Phi$, and $\beta \in [0, 1]$, if $\frac{\mathbb{E}_{P_\theta}(X-A'(\theta_0))^4}{A''(\theta_0)^2} \leq \zeta \leq \frac{9}{128\beta}$ and $\kappa(\theta) \geq J_{\mathrm{TV},\theta}^{-1}(\beta)$ then,*

$$J_{\mathrm{TV},\theta}^{-1}(\beta) \in \left[\frac{1}{\sqrt{2}}\frac{\beta}{\sqrt{A''(\theta)}}, C\frac{\beta}{\sqrt{A''(\theta)}}\right].$$

Proof. **Lower Bound:** Recall from the proof of Corollary 4.2 that $\mathrm{TV}(P_\theta, P_{\theta+h}) \leq |h|\sqrt{\max_{\theta'\in[\theta,\theta+h]} A''(\theta')}$. Now, let h be such that $\mathrm{TV}(P_\theta, P_{\theta+h}) = \beta$, so

$$|h| \geq \frac{\beta}{\sqrt{\max_{\theta'\in[\theta,\theta+h]} A''(\theta')}}.$$

By assumption, $\kappa(\theta) \geq J_{\text{TV},\theta}^{-1}(\beta) \geq |h|$ and $\max_{\theta' \in [\theta, \theta+h]} A''(\theta') \leq 2A''(\theta_0)$. Therefore $|h| \geq \frac{\beta}{\sqrt{2A''(\theta)}}$ so $J_{\text{TV},\theta}^{-1}(\beta) \geq \frac{\beta}{\sqrt{2A''(\theta)}}$.

Upper Bound: Since $\beta < \frac{9}{128\zeta}$, there exists a constant C such that $\frac{16\zeta}{9} \leq C \leq \frac{1}{8\beta}$. Let $h = \frac{C\beta}{\sqrt{A''(\theta)}}$. If $h \geq \kappa(\theta)$ then we are done since $J_{\text{TV},\theta}^{-1}(\beta) \leq \kappa(\theta) \leq h$, so assume that $h \leq \kappa(\theta)$. It suffices to prove that $\text{TV}(P_\theta, P_{\theta+h}) \geq \beta$ since then again by monotonicity and continuity of $h \to \text{TV}(P_\theta, P_{\theta+h})$, we are done. By the Paley-Zygmund inequality,

$$\mathbb{P}_{P_\theta}\left(X \geq A'(\theta) + \frac{1}{2}\sqrt{A''(\theta)}\right) \geq \left(1 - \frac{1}{4}\right)^2 \frac{(\mathbb{E}(X - \mathbb{E}(X))^2)^2}{\mathbb{E}[(X - \mathbb{E}(X))^4]} \geq \frac{9}{16\zeta}.$$

Then,

$$
\begin{aligned}
TV(P_\theta, P_{\theta+h}) &= \int_{\frac{A(\theta+h)-A(\theta)}{h}}^{\infty} \left(e^{(\theta+h)x - A(\theta+h)} - e^{\theta x - A(\theta)}\right) d\mu \\
&\geq \int_{\frac{A(\theta+h)-A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)}}^{\infty} \left(e^{(\theta+h)x - A(\theta+h)} - e^{\theta x - A(\theta)}\right) d\mu \\
&= \int_{\frac{A(\theta+h)-A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)}}^{\infty} \left(e^{hx - (A(\theta+h)-A(\theta))} - 1\right) P_\theta(x) d\mu \\
&\geq \left(e^{\frac{1}{4}h\sqrt{A''(\theta)}} - 1\right) \mathbb{P}_{P_\theta}\left(X \geq \frac{A(\theta+h) - A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)}\right).
\end{aligned}
$$

Now,

$$
\begin{aligned}
\frac{A(\theta+h) - A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)} &\leq A'(\theta) + h \max_{\theta' \in [\theta, \theta+h]} A''(\theta) + \frac{1}{4}\sqrt{A''(\theta)} \\
&\leq A'(\theta) + \frac{1}{8\sqrt{A''(\theta)}} 2A''(\theta) + \frac{1}{4}\sqrt{A''(\theta)} \\
&\leq A'(\theta) + \frac{1}{2}\sqrt{A''(\theta)},
\end{aligned}
$$

where the second inequality holds since $h \geq \frac{1}{8}\frac{1}{\sqrt{A''(\theta)}}$. Thus,

$$
\begin{aligned}
TV(P_\theta, P_{\theta+h}) &\geq \left(e^{\frac{1}{4}h\sqrt{A''(\theta)}} - 1\right) \mathbb{P}_{P_\theta}\left(X \geq A'(\theta) + \frac{1}{2}\sqrt{A''(\theta)}\right) \\
&\geq \frac{1}{4}\frac{9}{4\zeta} h\sqrt{A''(\theta)} \\
&\geq \beta,
\end{aligned}
$$

where the second inequality follows from the fact that $e^x - 1 \geq x$ for all $x > 0$, and the final inequality follows from the definition of h, and the assumptions on C. $\qquad\square$

B.6 Proof of Lemma 4.11

Lemma 4.11. *For all positive constants ζ and b, and $\Omega\left(\frac{\ln n}{n}\right) \leq \varepsilon_n \leq O\left(\frac{1}{\sqrt{n}}\right)$, there exists constants $N \in \mathbb{N}$ and $C > 0$ such for all measures μ, $\theta_0 \in \Phi$ and $n \in \mathbb{N}$ such that $n \geq N$, if*

1. $\max_{\theta \in \Phi(\theta_0)} \frac{\mathbb{E}_{P_\theta}[(X - A'(\theta_0))^4]}{A''(\theta)^2} \leq \zeta.$

2. $|\hat{t} - A'(\theta_0)| \le b\left(\sqrt{\frac{A''(\theta_0)}{n}} + \frac{\sqrt{A''(\theta_0)\ln(n)}}{\varepsilon n}\right)$

3. $\kappa(\theta_0) \ge J_{\text{TV},\theta_0}^{-1}\left(\frac{8\sqrt{2}}{\varepsilon_n n}\right)$

then for all $\theta_1 \in \Phi(\theta_0)$ such that $|\theta_0 - \theta_1| \ge C J_{\text{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$, then there exists a threshold τ such that the test

$$\begin{cases} P_{\theta_0} & \text{if } f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right) \le \tau \\ P_{\theta_1} & \text{if } f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right) \ge \tau \end{cases}$$

distinguishes between θ_0 and θ_1. Furthermore, $\mathbb{E}_{P_{\theta_0}}[f_{\hat{t}}(X)] \le \tau \le \mathbb{E}_{P_{\theta_1}}[f_{\hat{t}}(X)]$.

Proof of Lemma 4.11. By assumption there exists constant A_1 and A_2 such that $A_1 \frac{\ln n}{n} \le \varepsilon \le A_2 \frac{1}{\sqrt{n}}$. Set $N = \frac{8\sqrt{2}}{\varepsilon_n} \max\{\frac{1}{J_{TV,\theta_0}(\kappa(\theta_0))}, \frac{128\zeta}{9}\}$ then $n \ge N$ implies $\frac{8\sqrt{2}}{\varepsilon n} \le \frac{9}{128\zeta}$ and $\kappa(\theta_0) \ge J_{\text{TV},\theta}^{-1}(\frac{8\sqrt{2}}{\varepsilon n})$. Combined with the first assumption and Lemma 4.8, this implies that that there exists constants C_1 and C_2 (depending only on ζ) such that for all $\beta \le \frac{8\sqrt{2}}{\varepsilon n}$,

$$J_{\text{TV},\theta_0}^{-1}(\beta) \in \left[\frac{C_1\beta}{\sqrt{A''(\theta_0)}}, \frac{C_2\beta}{\sqrt{A''(\theta_0)}}\right].$$

Let $C = \frac{C_2 8\sqrt{2}}{C_1}$. Then

$$|\theta_0 - \theta_1| \ge C J_{\text{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon n}\right) \ge \frac{C_2 8\sqrt{2}}{C_1} \frac{C_1}{\varepsilon n \sqrt{A''(\theta_0)}} \ge J_{\text{TV},\theta_0}^{-1}\left(\frac{8\sqrt{2}}{\varepsilon n}\right).$$

Thus, $\text{TV}(P_{\theta_0}, P_{\theta_1}) \ge \frac{8\sqrt{2}}{\varepsilon n}$.

Next, assume that $|\theta_0 - \theta_1| = C J_{\text{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$. Assume without loss of generality that $\theta_0 \le \theta_1$. Note that by Markov's inequality, it is sufficient to show that

$$\mathbb{E}_{\theta_1}\left[f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right]$$

$$\ge \frac{1}{4} \min\left\{\sqrt{\text{var}_{\theta_0}\left(f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right)}, \sqrt{\text{var}_{\theta_1}\left(f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right)}\right\}$$

Let us first analyze the gap in expectations in the test statistic $f_{\hat{t}}(\cdot) + \text{Lap}(\frac{1}{\varepsilon n})$. Note,

$$\mathbb{P}_{\theta_1}\left[T(X) > \frac{A(\theta_0) - A(\theta_1)}{\theta_0 - \theta_1}\right] - \mathbb{P}_{\theta_0}\left[T(X) > \frac{A(\theta_0) - A(\theta_1)}{\theta_0 - \theta_1}\right] = TV(P_{\theta_0}, P_{\theta_1}).$$

Therefore,

$$\mathbb{E}_{\theta_1}\left[f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \text{Lap}\left(\frac{1}{\varepsilon n}\right)\right]$$

$$= TV(P_{\theta_0}, P_{\theta_1}) + \mathbb{P}_{\theta_1}[T(x) \in I] - \mathbb{P}_{\theta_0}[T(x) \in I],$$

where I has endpoints $\hat{t}$ and $\frac{A(\theta_0)-A(\theta_1)}{\theta_0-\theta_1}$. Assume, for ease of notation, that $\hat{t} \geq \frac{A(\theta_0)-A(\theta_1)}{\theta_0-\theta_1}$ so $\hat{t} = \frac{A(\theta_0)-A(\theta_1)}{\theta_0-\theta_1} + \Gamma$ where

$$\Gamma = \hat{t} - \frac{A(\theta_0) - A(\theta_1)}{\theta_0 - \theta_1}$$

$$= \hat{t} - A'(\theta_0) + A'(\theta_0) - \frac{A(\theta_0) - A(\theta_1)}{\theta_0 - \theta_1}$$

$$\leq b\left(\sqrt{\frac{A''(\theta_0)}{n}} + \frac{\sqrt{A''(\theta_0)\ln(n)}}{\varepsilon n}\right) + \max_{\theta' \in [\theta_0,\theta_1]} A''(\theta')|\theta_0 - \theta_1|.$$

Now,

$$\mathbb{P}_{\theta_1}[T(x) \in I] - \mathbb{P}_{\theta_0}[T(x) \in I] = \int_{a \in I} e^{a\theta_1 - A(\theta_1)} d\mu - \int_{a \in I} e^{a\theta_0 - A(\theta_0)} d\mu$$

$$= \int_{a \in I} e^{a(\theta_1 - \theta_0) + A(\theta_0) - A(\theta_1)} e^{a\theta_0 - A(\theta_0)} d\mu - \int_{a \in I} e^{a\theta_0 - A(\theta_0)} d\mu$$

$$\leq \max_{a \in I}(e^{a(\theta_1 - \theta_0) + A(\theta_0) - A(\theta_1)} - 1)\mathbb{P}_{\theta_0}[T(X) \in I]$$

$$\leq \max_{a \in I}(e^{a(\theta_1 - \theta_0) + A(\theta_0) - A(\theta_1)} - 1)$$

$$= e^{\hat{t}(\theta_1 - \theta_0) + A(\theta_0) - A(\theta_1)} - 1$$

$$= e^{\Gamma(\theta_1 - \theta_0)} - 1.$$

Now $|\theta_0 - \theta_1| = CJ_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right) \leq \frac{C_2 C}{\varepsilon n \sqrt{A''(\theta_0)}}$ and thus there exists a constant C_3 (depending on ζ, C_1, C_2, A_1 and A_2) such that

$$\mathbb{P}_{\theta_0}[T(x) \in I] - \mathbb{P}_{\theta_1}[T(x) \in I] \leq e^{\left[b\left(\sqrt{\frac{A''(\theta_0)}{n}} + \frac{\sqrt{A''(\theta_0)\ln(n)}}{\varepsilon n}\right) + \max_{\theta' \in [\theta_0,\theta_1]} A''(\theta')|\theta_0 - \theta_1|\right](\theta_1 - \theta_0)} - 1$$

$$\leq e^{\frac{bC_2 C}{\varepsilon n^{1.5}} + \frac{bC_2 C\sqrt{\ln n}}{\varepsilon^2 n^2} + \frac{2C_2^2 C^2}{\varepsilon^2 n^2}} - 1$$

$$= e^{\frac{C_3\sqrt{\ln n}}{\varepsilon^2 n^2}} - 1$$

where the second inequality follows from $\varepsilon = O(1/\sqrt{n})$, which implies that $\varepsilon n^{1.5} = \Omega(\varepsilon^2 n^2)$. Now $\varepsilon = \Omega(\ln n/n)$ implies $e^{\frac{C_3\sqrt{\ln n}}{\varepsilon^2 n^2}} - 1 = o(\frac{1}{\varepsilon n})$, thus since $\mathrm{TV}(P_{\theta_0}, P_{\theta_1}) \geq \frac{8\sqrt{2}}{\varepsilon n}$, there exists N (depending on C, B, A_1 and A_2) such that if $n \geq \max\{N, \frac{8\sqrt{2}}{\varepsilon n J_{TV,\theta_0}(\kappa(\theta_0))}, \frac{8\sqrt{2}}{\varepsilon n}\frac{128}{9\zeta}\}$ then

$$\mathbb{E}_{\theta_1}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right] \geq \frac{1}{2}(\mathrm{TV}(P_{\theta_0}, P_{\theta_1})).$$

Also,

$$\min\left\{\mathrm{var}_{\theta_0}\left(f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right), \mathrm{var}_{\theta_1}\left(f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right)\right\} \leq \frac{1}{n} + \frac{1}{\varepsilon^2 n^2} \leq 2\left(\frac{1}{\varepsilon^2 n^2}\right),$$

where the second inequality holds since $\varepsilon \leq 1/\sqrt{n}$. Thus, since $\mathrm{TV}(P_{\theta_0}, P_{\theta_1})] \geq \frac{8\sqrt{2}}{\varepsilon n}$, we have

$$\sqrt{\mathrm{var}_{\theta_0}\left(f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right)} \leq \frac{1}{4}\left(\mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right]\right)$$

Thus, the test distinguishes between P_{θ_0} and P_{θ_1}.

Next, assume that $|\theta_0 - \theta_1| \geq C J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$. Let θ_1' be such that $|\theta_0 - \theta_1'| = C J_{\mathrm{TV},\theta_0}^{-1}\left(\frac{1}{\varepsilon_n n}\right)$ and let $n > \max\{N, \frac{8\sqrt{2}}{\varepsilon_n J_{TV,\theta_0}(\kappa(\theta_0))}, \frac{8\sqrt{2}}{\varepsilon_n}\frac{128}{9\zeta}\}$. Then by Lemma 4.1,

$$\mathbb{E}_{\theta_1}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right]$$

$$\geq \mathbb{E}_{\theta_1'}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right] - \mathbb{E}_{\theta_0}\left[f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right]$$

$$\geq \frac{1}{2}\mathrm{TV}\left(P_{\theta_0}, P_{\theta_1'}\right)$$

and as above

$$\min\left\{\mathrm{var}_{\theta_0}\left(f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right), \mathrm{var}_{\theta_1}\left(f_{\hat{t}}(X) + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)\right)\right\} \leq \frac{1}{n} + \frac{1}{\varepsilon^2 n^2}$$

so we are done. □

B.7 Proof of Proposition 4.13

Proposition 4.13. *For all $\zeta > 0$, there exists positive constants k, C_1, C_2 and N such that for all measures μ, $\theta \in \Phi$ and $n \in \mathbb{N}$ such that $n \geq N$, if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$,

- $\max_{\theta' \in \Phi(\theta)} \frac{\mathbb{E}_{\theta'}((x - A'(\theta'))^4)}{A''(\theta')^2} \leq \zeta$

- $\kappa(\theta) \geq \frac{C_2}{\sqrt{nA''(\theta)}}$

then,

$$\frac{C_1}{\sqrt{nA''(\theta)}} \leq \omega_n(P_\theta, \mathcal{F}_\varepsilon^{\mathrm{test}}) \leq \frac{C_2}{\sqrt{nA''(\theta)}}.$$

This implies $\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\mathrm{test}}) = \Theta\big(\omega_n(P_\theta, \mathcal{F}^{\mathrm{test}})\big)$.

Lemma B.4. *For all constants $\zeta > 0$, there exists constants $k, C, N > 0$ such that for all $\theta_0 \in \Phi$ if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((X - A'(\theta))^4)}{A''(\theta)^2} \leq \zeta$

then for all $h \geq \frac{C}{\sqrt{nA''(\theta)}}$ and n such that $\kappa(\theta) \geq \frac{C}{\sqrt{nA''(\theta)}}$ and $n \geq N$, we have

$$\mathrm{SC}_{\varepsilon_n}(P_\theta, P_{\theta+h}) \leq n.$$

Proof. Let C and N be constants to be specified later with the relationship that $\frac{C}{\sqrt{N}} \leq \frac{1}{8}$. Let $n \geq \max\{N, \frac{C^2}{A''(\theta)(\kappa(\theta))^2}\}$. First, assume that $h = \frac{C}{\sqrt{nA''(\theta)}}$. Let

$$\mathrm{cLLR}(X) = \sum_{i=1}^{n}\left[\ln \frac{P_{\theta+h}(x_i)}{P_\theta(x_i)}\right]_{-\varepsilon}^{\varepsilon}$$

and
$$\text{ncLLR}_\varepsilon(X) = \text{cLLR}(X) + \text{Lap}(2).$$

Our goal is to show that under the conditions outlined in the lemma statement,
$$\sqrt{\text{var}_\theta(\text{ncLLR})} \leq (1/4)|\mathbb{E}_{\theta+h}(\text{ncLLR}) - \mathbb{E}_\theta(\text{ncLLR})|$$
and
$$\sqrt{\text{var}_{\theta+h}(\text{ncLLR})} \leq (1/4)|\mathbb{E}_{\theta+h}(\text{ncLLR}) - \mathbb{E}_\theta(\text{ncLLR})|.$$

Firstly, note that since for any x and y,
$$\left| \left[\ln \frac{P_{\theta+h}(x)}{P_\theta(x)} \right]_{-\varepsilon}^{\varepsilon} - \left[\ln \frac{P_{\theta+h}(y)}{P_\theta(y)} \right]_{-\varepsilon}^{\varepsilon} \right| \leq \left| \ln \frac{P_{\theta+h}(x)}{P_\theta(x)} - \ln \frac{P_{\theta+h}(y)}{P_\theta(y)} \right|.$$

Further, $\text{var}_\theta(LLR) \leq 18nH^2(P_\theta, P_{\theta+h}) \leq 18n\text{TV}(P_\theta, P_{\theta+h}) \leq 36nh^2 A''(\theta)$, where the first inequality follows from a simple manipulation (see the proof of Theorem 2.5 in Canonne et al. (2019) for details), the second is a standard inequality, and the third follows from eqn (10). Thus, we have that $\text{var}_\theta(\text{ncLLR}) \leq \text{var}_\theta(\text{LLR}) + 2 = 36nh^2 A''(\theta) + 2$ and $\text{var}_{\theta+h}(\text{ncLLR}) \leq \text{var}_{\theta+h}(\text{LLR}) + 2 = nh^2 A''(\theta + h) + 2 \leq 72nh^2 A''(\theta) + 2$. Now, also as in the proof of Lemma 4.8, since $|h| \leq \frac{1}{8\sqrt{A''(\theta)}}$ and $|h| \leq \kappa(\theta_0)$, we have

$$\mathbb{P}_{P_\theta}\left(x \geq \frac{A(\theta + h) - A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)} \right) \geq \mathbb{P}_{P_\theta}\left(x \geq A'(\theta) + h \max_{\theta' \in [\theta, \theta+h]} A''(\theta') + \frac{1}{4}\sqrt{A''(\theta)} \right)$$
$$\geq \mathbb{P}_{P_\theta}\left(x \geq A'(\theta) + \frac{1}{2}\sqrt{A''(\theta)} \right)$$
$$\geq \frac{9}{16\zeta}$$

Let $C_2 = \sqrt{\frac{32\zeta}{9} + \frac{1}{16}}$, by Chebyshev's's inequality we have

$$\mathbb{P}_{P_\theta}\left(x \geq \frac{A(\theta + h) - A(\theta)}{h} + C_2\sqrt{A''(\theta)} \right) \leq \mathbb{P}_{P_\theta}\left(x \geq A'(\theta) - h \max_{\theta' \in [\theta, \theta+h]} A''(\theta) + C_2\sqrt{A''(\theta)} \right)$$
$$\leq \mathbb{P}_{P_\theta}\left(x \geq A'(\theta) + \left(C_2 - \frac{1}{16} \right)\sqrt{A''(\theta)} \right)$$
$$\leq \frac{1}{(C_2 - \frac{1}{16})^2}$$
$$\leq \frac{9}{32\zeta}$$

where the second inequality follows since $h \leq \frac{1}{8\sqrt{A''(\theta)}}$. Therefore,

$$\mathbb{P}_{P_\theta}\left(\frac{A(\theta + h) - A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)} \leq x \leq \frac{A(\theta + h) - A(\theta)}{h} + C_2\sqrt{A''(\theta)} \right) \geq \frac{9}{32\zeta}.$$

45

Now,

$$\frac{1}{n}(\mathbb{E}_{\theta+h}(\text{ncLLR}) - \mathbb{E}_\theta(\text{ncLLR})) = \int_{-\infty}^{\infty} \left[\ln \frac{P_{\theta+h}(x_i)}{P_\theta(x_i)}\right]_{-\varepsilon}^{\varepsilon} (P_{\theta+h}(x) - P_\theta(x))d\mu(x)$$

$$= \int_{-\infty}^{\infty} [hx - A(\theta+h) + A(\theta)]_{-\varepsilon}^{\varepsilon} (P_{\theta+h}(x) - P_\theta(x))d\mu(x)$$

$$\geq \int_{\frac{A(\theta+h)-A(\theta)}{h}+\frac{1}{4}\sqrt{A''(\theta)}}^{\frac{A(\theta+h)-A(\theta)}{h}+C_2\sqrt{A''(\theta)}} [hx - A(\theta+h) + A(\theta)]_{-\varepsilon}^{\varepsilon} (P_{\theta+h}(x) - P_\theta(x))d\mu(x)$$

Now, if $x \in \left[\frac{A(\theta+h)-A(\theta)}{h} + \frac{1}{4}\sqrt{A''(\theta)}, \frac{A(\theta+h)-A(\theta)}{h} + C_2\sqrt{A''(\theta)}\right]$ then

$$hx - A(\theta+h) + A(\theta) \in \left[\frac{1}{4}h\sqrt{A''(\theta)}, hC_2\sqrt{A''(\theta)}\right] \subset \left[0, \frac{C_2 C}{\sqrt{n}}\right],$$

where the subset relationship holds since $h = \frac{C}{\sqrt{nA''(\theta)}}$. Thus, letting $k = CC_2$, if $\varepsilon \geq \frac{k}{\sqrt{n}}$ then the truncation has no impact inside this region and

$$\frac{1}{n}(\mathbb{E}_{\theta+h}(\text{ncLLR}) - \mathbb{E}_\theta(\text{ncLLR}))$$

$$\geq \int_{\frac{A(\theta+h)-A(\theta)}{h}+\frac{1}{4}\sqrt{A''(\theta)}}^{\frac{A(\theta+h)-A(\theta)}{h}+C_2\sqrt{A''(\theta)}} (hx - A(\theta+h) + A(\theta))(P_{\theta+h}(x) - P_\theta(x))d\mu(x)$$

$$= \int_{\frac{A(\theta+h)-A(\theta)}{h}+\frac{1}{4}\sqrt{A''(\theta)}}^{\frac{A(\theta+h)-A(\theta)}{h}+C_2\sqrt{A''(\theta)}} (hx - A(\theta+h) + A(\theta))P_\theta(x)\left(e^{hx-A(\theta+h)+A(\theta)} - 1\right)d\mu(x)$$

$$\geq \frac{1}{4}h\sqrt{A''(\theta)}\left(e^{\frac{1}{4}h\sqrt{A''(\theta)}} - 1\right)\mathbb{P}_{P_\theta}\left(\frac{1}{4}\sqrt{A''(\theta)} \leq x - \frac{A(\theta+h)-A(\theta)}{h} \leq C_2\sqrt{A''(\theta)}\right)$$

$$\geq \frac{1}{16}\frac{9}{32\zeta}h^2 A''(\theta).$$

Therefore, recalling that $h = \frac{C}{\sqrt{nA''(\theta)}}$ and $\max\{\text{var}_\theta(\text{ncLLR}), \text{var}_{\theta+h}(\text{ncLLR})\} \leq 72nh^2 A''(\theta)$, there exists a large enough C (where large enuogh depends only on ζ) such that $\max\left\{\sqrt{\text{var}_\theta(\text{ncLLR})}, \sqrt{\text{var}_{\theta+h}(\text{ncLLR})}\right\} \leq (1/4)(\mathbb{E}_{\theta+h}(\text{ncLLR}) - \mathbb{E}_\theta(\text{ncLLR}))$, as required.

Next, assume that $h' \geq \frac{C}{\sqrt{nA''(\theta)}}$. Note that $\left[\ln \frac{P_{\theta+h}(x)}{P_\theta(x)}\right]_{-\varepsilon}^{\varepsilon}$ is monotone increasing in x, so for any threshold τ, $\mathbb{P}_{X \sim P_{\theta+h'}^n}(\text{ncLLR}(X) \geq \tau) \geq \mathbb{P}_{X \sim P_{\theta+h}^n}(\text{ncLLR}(X) \geq \tau)$. Therefore, if ncLLR distinguishes θ and $\theta + h$ with n samples then it also distinguishes θ and $\theta + h'$ with n samples. $\square$

Proof of Proposition 4.13. First note that by definition $\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}}) \geq \omega_n(P_\theta, \mathcal{F}^{\text{test}})$ and $\omega_n(P_\theta, \mathcal{F}^{\text{test}}) = \Theta\left(\frac{1}{\sqrt{nA''(\theta)}}\right)$. By Lemma B.4, there exists $k, C_2 \geq 0$ such that if n is sufficiently large, $\varepsilon \geq \frac{k}{\sqrt{n}}$ and $h \geq \frac{C_2}{\sqrt{nA''(\theta)}}$ then $\text{SC}_\varepsilon(P_{\theta+h}, P_\theta) \leq n$. Therefore, $\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}}) \leq \frac{C_2}{\sqrt{nA''(\theta)}}$. $\square$

B.8 Proof of Proposition 4.15

Proposition 4.15. [Uniform Achievability in Low Privacy Regime] *For any $\varepsilon_n > 0, \delta_n \in [0,1], C > 0$, $\mathcal{A}_{low}$ is $(\varepsilon_n, \delta_n)$-DP. Further, for all $\zeta > 0, C > 0, \varepsilon_n > 0$, and $\delta_n \in [0,1]$, there exists an initial*

estimator $\mathcal{M}_{\zeta,C}$ *such that there exists constants* $N \in \mathbb{N}$, $k \geq 0$, $C' > 0$, $D > 0$ *such that for all exponential families (i.e. any measure* μ*) and* $\theta_0 \in \Phi$ *if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((x - \mathbb{E}_\theta(x))^4)}{A''(\theta)^2} \leq \zeta$,

- $\kappa(\theta) \geq \frac{1}{C} \frac{1}{\sqrt{A''(\theta)}}$

then for all $n \in \mathbb{N}$ *such that* $n \geq \max\left\{N, \frac{c\zeta^2 \ln(1/\delta_n)}{\varepsilon_n}, \frac{2\ln\left(\frac{1}{(DA''(\theta_0))^2}\right)}{(DA''(\theta_0))^2}\right\}$ *then we have*

$$\mathfrak{R}_n(\theta_0, \mathcal{A}_{\mathrm{low},C,\zeta}) \leq C' \omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\mathrm{test}}),$$

where $\mathcal{A}_{\mathrm{low},C,\zeta}$ *is* $\mathcal{A}_{\mathrm{low}}$ *with initial mean estimator* $\mathcal{M}_{\zeta,C}$.

Proof. The proposition follows from a combination of Corollary 4.17 and Corollary 4.7. Let N, k, C_1, D_1 and D_2 be as in Corollary 4.17. Note that since $\frac{D}{A''(\theta_0)(\kappa(\theta_0))^2} \leq DC^2 = O(1)$ we can set N large enough that $N \geq \frac{D_2}{A''(\theta_0)(\kappa(\theta_0))^2}$. By Corollary 4.7, there exists a constant C_2 such that

$$|A'^{-1}(\mathcal{M}_{\zeta,C}(X)) - \theta_0| \leq C_2\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right).$$

Again since $\kappa(\theta_0)\sqrt{A''(\theta_0)} \geq \frac{1}{C}$, there exists constants N_1 and D such that for all $n > N_1$ if $\frac{\sqrt{\ln n}}{\sqrt{n}} \leq DA''(\theta_0)$ then,

$$C_2\left(\frac{1}{\sqrt{nA''(\theta_0)}} + \frac{1}{n\varepsilon\sqrt{A''(\theta_0)}}\sqrt{\ln(n)}\right) \leq \min\{\kappa(\theta_0), D_1\varepsilon\sqrt{nA''(\theta_0)}.$$

The condition $\frac{\sqrt{\ln n}}{\sqrt{n}} \leq DA''(\theta_0)$ is implied by the assumption that $n \geq \frac{2\ln(1/(DA''(\theta_0))^2)}{(DA''(\theta_0))^2}$. Thus, the estimator from Corollary 4.7 satisfies the requirements of Corollary 4.17 and so we are done. $\qquad\square$

B.9   Proof of Lemma 4.16

Lemma 4.16. *For all* $\zeta > 0$, *there exists constants* $N \in \mathbb{N}$, $k \geq 0$, $C > 0$, $b > 0$ *and* $D > 0$ *such that for all measures* μ *and* $\theta_0 \in \Phi$ *if*

- $\varepsilon_n \geq \frac{k}{\sqrt{n}}$

- $\frac{\mathbb{E}_\theta((x - \mathbb{E}_\theta(x))^4)}{A''(\theta)^2} \leq \zeta$,

- $|\theta_0 - A'^{-1}(\hat{t})| \leq \min\{\kappa(\theta_0), b\varepsilon_n\sqrt{nA''(\theta_0)}\}$,

then for all $n \in \mathbb{N}$ *such that* $n \geq \max\{N, \frac{D}{A''(\theta_0)(\kappa(\theta_0))^2}\}$ *and all* θ_1 *such that* $|\theta_1 - \theta_0| \geq C\omega_n(P_{\theta_0}, \mathcal{F}_\varepsilon^{\mathrm{test}})$, *there exists a threshold* τ *such that the test*

$$\begin{cases} P_{\theta_0} & \textit{if } f_{\tilde{\alpha}}(X) + \mathrm{Lap}\left(\frac{2}{\varepsilon_n nC\tilde{\alpha}}\right) \leq \tau \\ P_{\theta_1} & \textit{if } f_{\tilde{\alpha}}(X) + \mathrm{Lap}\left(\frac{2}{\varepsilon_n n}C\tilde{\alpha}\right) \geq \tau \end{cases} \tag{6}$$

distinguishes between P_{θ_0} *and* P_{θ_1} *with* n *samples. Furthermore,* τ *can be chosen so* $|\mathbb{E}_{X \sim P_{\theta_0}^n}[\hat{f}_{\tilde{\alpha}}(X)] - \tau| \leq |\mathbb{E}_{X \sim P_{\theta_1}^n}[\hat{f}_{\tilde{\alpha}}(X)] - \tau|$.

Proof of Lemma 4.16. Let $\tilde{\theta} = A'^{-1}(\hat{t})$ and recall that for ease of notation we let $\alpha_n(\theta) = \frac{1}{\sqrt{nA''(\theta)}}$. Since $|\theta_0 - A'^{-1}(\hat{t})| \leq \kappa(\theta_0)$, we have $\alpha_n(\theta_0) \in [\frac{1}{2}\tilde{\alpha}, 2\tilde{\alpha}]$. Thus by Proposition 4.13, there exists constants k, C_1, and C_2 (depending only on ζ) such that $\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}}) \in [C_1\alpha_n(\theta_0), C_2\alpha_n(\theta_0)]$. Further, by Lemma A.2, there exist a constant C_3 (depending on C_1 and C_2) such that if $a, b \in [C_1/4\varepsilon, 4C_2\varepsilon]$ then

$$\text{SC}_{\text{ncLLR}_a^b}(P, Q) \leq C_3 \cdot \text{SC}_\varepsilon(P, Q).$$

Now, set $C = \sqrt{C_3}\frac{C_2}{C_1}$ and assume first that $|\theta_0 - \theta_1| = C \cdot \omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}})$. Then

$$
\begin{aligned}
|\theta_0 - \theta_1| &= C\omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}}) \\
&\geq \frac{CC_1}{\sqrt{nA''(\theta_0)}} \\
&= \frac{C_2}{\sqrt{(C_2/C_1C)^2 nA''(\theta_0)}} \\
&\geq \omega_{(C_2/C_1C)^2 n, \text{SC}_\varepsilon}(P_{\theta_0}).
\end{aligned}
$$

Therefore, $\text{SC}_\varepsilon(P_{\theta_0}, P_{\theta_1}) \leq (\frac{C_2}{C_1 C})^2 n$. Thus, if $a, b \in [C_1/4, 4C_2]$ then

$$\text{SC}_{\text{ncLLR}_a^b}(P_{\theta_0}, P_{\theta_1}) \leq C_3\left(\frac{C_2}{C_1 C}\right)^2 n \leq n. \tag{11}$$

Now, recall that $P_\theta(x) = e^{\theta x - A(\theta)}$ and $|\theta_0 - \theta_1| = C \cdot \omega_n(P_\theta, \mathcal{F}_\varepsilon^{\text{test}})$ so

$$
\begin{aligned}
(\theta_1 - \theta_0)[x - \hat{t}]_{-\varepsilon/C\tilde{\alpha}}^{\varepsilon/C\tilde{\alpha}} &= \left[(\theta_1 - \theta_0)(x - \hat{t})\right]_{-\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}}}^{\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}}} \\
&= \left[\ln\frac{P_{\theta_1}(x)}{P_{\theta_0}(x)} + (A(\theta_1) - A(\theta_0)) - (\theta_1 - \theta_0)\hat{t}\right]_{-\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}}}^{\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}}} \\
&= \left[\ln\frac{P_{\theta_1}(x)}{P_{\theta_0}(x)}\right]_{\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}} - \Gamma}^{\varepsilon\frac{\omega_{n,\text{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}} - \Gamma} + \Gamma
\end{aligned}
$$

where

$$
\begin{aligned}
\Gamma &= A(\theta_1) - A(\theta_0) - (\theta_1 - \theta_0)\hat{t} \\
&\leq A(\theta_1) - A(\theta_0) - (\theta_1 - \theta_0)A'(\theta_0) + (\theta_1 - \theta_0)(A'(\theta_0) - \hat{t}) \\
&\leq \left(\max_{\theta' \in [\theta_1, \theta_0]} A''(\theta')\right)(\theta_1 - \theta_0)^2 + |\theta_1 - \theta_0||A'(\theta_0) - \hat{t}|
\end{aligned}
$$

Now, if we let $D_2 = C_2^2 C^2$ then $n \geq \frac{D_2}{A''(\theta_0)(\kappa(\theta_0))^2}$ implies that

$$|\theta_0 - \theta_1| = C\omega_{n,\text{SC}_\varepsilon}(\theta_0) \leq C\frac{C_2}{\sqrt{nA''(\theta_0)}} = \frac{\sqrt{D_2}}{\sqrt{nA''(\theta_0)}} \leq \kappa(\theta_0).$$

Therefore,

$$\Gamma \leq 2A''(\theta_0)(\theta_1 - \theta_0)^2 + |\theta_1 - \theta_0||A'(\theta_0) - \hat{t}| \leq 2C^2 C_2^2\frac{1}{n} + C_2 C\alpha_n(\theta_0)|A'(\theta_0) - \hat{t}|.$$

48

Thus, noting that $1/n \ll \varepsilon$, and setting $b = \frac{C_1}{8CC_2}$ then if $|\theta_0 - A'^{-1}(\hat{t})| \leq \frac{b\varepsilon}{\alpha_n(\theta_0)} = b\varepsilon\sqrt{nA''(\theta_0)}$ then there exists N such that for all $n \geq \max\{N, \frac{D_2}{A''(\theta_0)(\kappa(\theta_0))^2}\}$, $\Gamma \leq \frac{C_1\varepsilon}{4}$. Therefore, for all $n \geq N$, the truncation parameters $\varepsilon\frac{\omega_{n,\mathrm{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}} - \beta, \varepsilon\frac{\omega_{n,\mathrm{SC}_\varepsilon}(\theta_0)}{\tilde{\alpha}} + \beta \in [C_1/4, 4C_2]$. So, by eqn (11), $\mathrm{SC}_{\mathrm{ncLLR}^b_{-a}}(P_{\theta_0}, P_{\theta_1}) \leq n$ which implies n samples are sufficient for the test statistic $[x - \hat{t}]^{\varepsilon/C\tilde{\alpha}}_{-\varepsilon/C\tilde{\alpha}}$ to distinguish between P_{θ_0} and P_{θ_1}. The threshold τ can be chosen as the midpoint between $\mathbb{E}_{X\sim P^n_{\theta_0}}[\hat{f}_{\tilde{\alpha}}(X)]$ and $\mathbb{E}_{X\sim P^n_{\theta_1}}[\hat{f}_{\tilde{\alpha}}(X)]$.

Now, assume that $|\theta_0 - \theta_1| \geq C \cdot \omega_n(P_\theta, \mathcal{F}^{\mathrm{test}}_\varepsilon)$. Let θ'_1 be such that $\theta_0 < \theta'_1 < \theta_1$ and $|\theta_0 - \theta_1| = C \cdot \omega_n(P_\theta, \mathcal{F}^{\mathrm{test}}_\varepsilon)$. Then, by the previous argument, there exists a threshold τ such that n samples are sufficient for the test statistic $[x - \tilde{t}]^{\varepsilon/C\tilde{\alpha}}_{-\varepsilon/C\tilde{\alpha}}$ to distinguish between P_{θ_0} and $P_{\theta'_1}$. Noting that this test statistic is monotone in x, we have by Lemma 4.1 (the fact that P_{θ_1} stochastically dominates $P_{\theta'_1}$) that this test statistic also distinguishes between P_{θ_0} and P_{θ_1} with n samples. Additionally, since $\mathbb{E}_{\theta'_1}[\mathrm{ncLLR}(X)] \geq \mathbb{E}_{\theta_1}[\mathrm{ncLLR}(X)] \geq \mathbb{E}_{\theta_0}[\mathrm{ncLLR}(X)]$, we maintain that $|\mathbb{E}_{X\sim P^n_{\theta_0}}[\hat{f}_{\tilde{\alpha}}(X)] - \tau| \leq |\mathbb{E}_{X\sim P^n_{\theta_1}}[\hat{f}_{\tilde{\alpha}}(X)] - \tau|$. $\qquad\square$

C   Proofs for Section 5.1

C.1   Proof of Theorem 10

Theorem 10. *For all $n \in \mathbb{N}$, let $k^*(n) = \lceil\log_{3/2}(\frac{|t_1-t_0|}{\omega_n(t,\mathcal{F}^{\mathrm{test}}_\varepsilon,\mathcal{P},\theta)})\rceil$, $N = n \cdot k^*(n) \cdot \lceil\log k^*(n)\rceil$ and $\hat{\theta}$ be the output of Algorithm 5 run for $k^*(n)$ rounds. Then*

$$\mathfrak{R}_N(f, \hat{\theta}) \leq \omega_n(t, \mathcal{F}^{\mathrm{test}}_\varepsilon, \mathcal{P}, \theta).$$

Proof of Theorem 10. We can think of Algorithm 5 as at each step dividing the distance between $t_{\min}$ and $t_{\max}$ by $2/3$ and concluding that the true value t^* lies between $t_{\min}$ and $t_{\max}$. Thus, in order to show that $\mathfrak{R}_N(f, \hat{\theta}) \leq \omega_n(t, \mathcal{F}^{\mathrm{test}}_\varepsilon, \mathcal{P}, \theta)$, it suffices to show that it is possible to run for $k^*(n) = \lceil\log_{\frac{3}{2}}\left(\frac{|t_1-t_0|}{\omega_n(t,\mathcal{F}^{\mathrm{test}}_\varepsilon,\mathcal{P},\theta)}\right)\rceil$ iterations with $N = n \cdot \lceil\log k^*(n)\rceil \cdot k^*(n)$ samples. In order to make the correct decision with probability $1/3k^*(n)$, it suffices for the last iteration to use $n \cdot \lceil\log k^*(n)\rceil$ samples. Since the hypothesis test at the last step has the largest sample size, $n \cdot \lceil\log k^*(n)\rceil \cdot k^*(n)$ samples is sufficient to run $k^*(n)$ rounds. $\qquad\square$