

# Locally Private Histograms in All Privacy Regimes

Clément L. Canonne   

School of Computer Science, University of Sydney, Australia

Abigail Gentle   

School of Computer Science, University of Sydney, Australia

---

## Abstract

Frequency estimation, a.k.a. histograms, is a workhorse of data analysis, and as such has been thoroughly studied under differential privacy. In particular, computing histograms in the *local* model of privacy has been the focus of a fruitful recent line of work, and various algorithms have been proposed, achieving the order-optimal  $\ell_\infty$  error in the high-privacy (small  $\epsilon$ ) regime while balancing other considerations such as time- and communication-efficiency. However, to the best of our knowledge, the picture is much less clear when it comes to the medium- or low-privacy regime (large  $\epsilon$ ), despite its increased relevance in practice. In this paper, we investigate locally private histograms, and the very related distribution learning task, in this medium-to-low privacy regime, and establish near-tight (and somewhat unexpected) bounds on the  $\ell_\infty$  error achievable. As a direct corollary of our results, we obtain a protocol for histograms in the *shuffle* model of differential privacy, with accuracy matching previous algorithms but significantly better message and communication complexity.

Our theoretical findings emerge from a novel analysis, which appears to improve bounds across the board for the locally private histogram problem. We back our theoretical findings by an empirical comparison of existing algorithms in all privacy regimes, to assess their typical performance and behaviour beyond the worst-case setting.

**2012 ACM Subject Classification** Security and privacy; Security and privacy  $\rightarrow$  Usability in security and privacy; Security and privacy  $\rightarrow$  Privacy protections; Theory of computation  $\rightarrow$  Theory of database privacy and security

**Keywords and phrases** Differential Privacy, Local Differential Privacy, Histograms, Frequency Estimation, Lower Bounds, Maximum Error

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2025.25

**Related Version** *Full Version*: <https://arxiv.org/abs/2408.04888>

**Funding** *Clément L. Canonne*: Supported by an ARC DECRA (DE230101329).

**Acknowledgements** The authors would like to thank Guy Blanc for the proof of Lemma 17, and Albert Cheu for insightful discussions regarding the use of amplification by shuffling (Theorem 25). This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing.

## 1 Introduction

Frequency estimation is a fundamental problem of statistics: besides its use for basic surveying, it is also used as a building block in distribution learning, identifying heavy hitters in sparse domains, and regression, correlation or covariance analysis. As such, frequency estimation (and the closely related problem of heavy hitters) have been thoroughly studied in a variety of settings, ranging from streaming to privacy-preserving analytics. We here focus on the latter, and specifically on the *local* model of differential privacy (LDP), where the data is distributed across a large number of users, and each datum is subject to stringent differential privacy requirements. This question has, over the past decade, received a lot of attention, starting with [8, 28]; as we elaborate in Section 1.1, much is known about locally private frequency



© Clément L. Canonne and Abigail Gentle;

licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 25; pp. 25:1–25:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

estimation, to the point that it may seem the question has been, already fully resolved – both in theory and practice. However, most recent large-scale implementations of local privacy (see, e.g., [5, 18]) must balance many efficiency objectives, including the bandwidth and computational requirements (as, among others, a proxy for energy consumption), and, above all, the estimation accuracy. Keeping this accuracy in check has, in turn, led to the common use of relatively large values of the “privacy parameter”  $\varepsilon$ , far from the original rule of thumb of  $\varepsilon \ll 1$ .<sup>1</sup>

In view of the relevance of frequency estimation to privacy-preserving algorithms and deployments, it is crucial that theory be informed by, and applicable to practice; and that constant factors, bounds on achievability, and the behaviour of estimation error in all parameter regimes are well understood. Yet, to the best of our knowledge, most if not all of the previous theoretical work on locally private frequency estimation focuses on the high-privacy  $\varepsilon \ll 1$  regime, leaving the low-privacy regime by and large uncharted. Addressing this gap in our theoretical understanding is the focus of our work.

### The question and setting

We will focus on examining the standard worst-case estimation error, or the  $\ell_\infty$ -error as it is most relevant to the problem, with special importance for controlling error rates in heavy hitters. Specifically, when estimating the empirical frequencies  $q = (q_1, \dots, q_k)$  over a universe  $\mathcal{X}$  of size  $k$  from  $n$  users, the (expected) error is given by

$$\mathbb{E}[\|\hat{q} - q\|_\infty] = \mathbb{E}\left[\max_{1 \leq i \leq k} |\hat{q}_i - q_i|\right] \quad (1)$$

where the expectation is taken over the (possibly randomised) algorithm given as input the users’ data, and  $\hat{q}$  is the vector of estimated frequencies. We are interested in the worst-case error over all possible datasets, that is, the supremum of the above quantity over all inputs  $X := (X_1, \dots, X_n) \in \mathcal{X}^n$  (equivalently, all  $q$ ). We will seek algorithms (“protocols”) to minimize this worst-case error under the constraint of (pure) local privacy (see Section 2), parameterized by the privacy parameter  $\varepsilon > 0$ . As in most distributed settings, one can consider several variants depending on whether a common random seed is available to the users (public-coin protocols) or not (private-coin protocols). It is known that for locally private frequency estimation allowing public-coin protocols can provably reduce the communication requirements [3];<sup>2</sup> similarly, one can allow some back and forth between users and the central server (interactive protocols). However, the public-coin or interactive settings come at an increased deployment costs, and are often less easy (or even impossible) to implement as they require either broadcasting from the center or sustained two-way communication between the parties. As our focus is chiefly on analyzing the most versatile setting, we hereafter restrict ourselves for our algorithms to the private-coin (non-interactive) setting; however, we mention that our lower bounds apply to the public-coin setting as well.

A closely related question is that of *distribution learning* under  $\ell_\infty$  loss, which differs from frequency estimation in that the dataset  $X$  is not arbitrary, but instead assumed to be drawn i.i.d. from an unknown probability distribution  $q$  over  $\mathcal{X}$ : in this sense, distribution learning

<sup>1</sup> See, for instance, the list of real-case use of differential privacy maintained here, which includes use cases of LDP along with the correspond stated values of  $\varepsilon$ : <https://desfontain.es/blog/real-world-differential-privacy.html>.

<sup>2</sup> Specifically, the cited paper establishes a separation between public- and private- coin protocols, showing that the former can achieve vanishing error even with constant-length communication per user, a setting where the latter must incur  $\Omega(1)$  error. To complement this, it is known that using public randomness one can always reduce the per-player communication to  $\lceil \varepsilon \rceil$  bits [8].

is an “easier” problem than frequency estimation, and an algorithm for the latter implies one for the former.<sup>3</sup> Formal definitions, as well as the relation between these two questions, can be found in Section 2. In this paper, we therefore focus on frequency estimation, and will point out the implications for distribution learning as corollaries.

### Connection to shuffle privacy

To conclude this subsection, we mention that besides the use of large values of  $\varepsilon$  in practical deployments of locally private algorithms, which calls for a better understanding of the tasks in this parameter regime, another motivation for studying the “low-privacy regime” comes from the emergence of another model of differential privacy, *shuffle privacy* [15, 21], and its increasingly widespread adoption. Indeed, it is known that one generic way to obtain shuffle private algorithms is via the so-called “amplification-by-shuffling” technique (see, e.g., [23]), whereby a locally private algorithm with high  $\varepsilon$  (low privacy) yields a shuffle private protocol with small  $\varepsilon$  (high privacy), and the same number of messages and communication cost. This makes characterizing the low-privacy regime of locally private algorithms a very consequential question, especially for fundamental tasks such as frequency estimation and the related distribution learning.

### 1.1 Prior work

Many innovative algorithms have been proposed in recent years for locally private histograms and distribution estimation [22, 3, 14, 34, 33, 25, 24], with a subset focusing on  $\ell_\infty$ -error.<sup>4</sup> A lower bound of  $\ell_\infty = \Omega(\sqrt{\log k/\varepsilon^2 n})$  for frequency estimation under local differential privacy was established by [8] for most regimes of  $k$ ,  $n$ , and small privacy parameter  $\varepsilon$ ; and various LDP protocols asymptotically achieving this bound have been proposed (see Table 1).

Most recently, upper bounds on  $\ell_\infty$  for frequency estimation were derived in [14], and evaluated empirically in [24]. For large  $\varepsilon$ , i.e., the “low-privacy regime,” to the best of our knowledge the best current results are (1) a bound of  $O(\sqrt{\log k/\varepsilon n})$  on the  $\ell_\infty$  error rate established in [14], and, (2) an upper bound of  $O(\sqrt{\log k/(e^{\varepsilon/2} n)} + (\log k)/n)$  from [29] (Theorem 2), based on CountSketch: note that this error does not vanish as  $\varepsilon$  grows.

In [25], it was stated that under most commonly used metrics ( $\ell_\infty, \ell_2^2, \ell_1$ ), the error is primarily driven by the variance of an algorithm. We show that at least for  $\ell_\infty$  a close look at sub-Gaussian, and sub-gamma behaviour can better capture the error in low-privacy regimes. Additionally, [14] raise the question of whether the upper bound of  $O(\sqrt{\log k/\varepsilon n})$ , is tight in the low privacy regime. We show in Theorem 14 that, quite surprisingly, it is not. Histograms have also seen detailed analysis in the shuffle model as both a practical question [26, 27, 15, 16, 9] and as a benchmark for reasoning about the power of the shuffle model [6]. As two points of comparison we highlight [26] which introduced a multiple-message protocol for shuffled histograms achieving expected  $\ell_\infty$ -error  $O(\sqrt{\log k \log(1/\delta)}/(\varepsilon n))$  with  $O(k^{1/100})$  rounds of communication, and [6] which demonstrated expected maximum error of  $O(\log(1/\delta)/(\varepsilon^2 n))$  with  $k + 1$  rounds. While our results do not achieve the latter bound, we show surprisingly that the former can be achieved in some fairly permissive parameter regimes with only a single message.

<sup>3</sup> There are some subtleties here, but it is worth noting, as a baseline, that absent privacy constraints it is well-known that the error of learning to  $\ell_\infty$  scales as  $1/\sqrt{n}$  with no dependence on the alphabet size  $k$ ; while frequency estimation can be done with zero error, as each user can simply send their data.

<sup>4</sup> For distribution estimation, a common error measure in the literature is  $\ell_1$ , corresponding to total variation distance.

■ **Table 1** Selection of best known upper bounds for communication,  $\ell_\infty$  and  $\ell_2^2$ -error, where the shaded cells are our results ( $\dagger$  indicates results not established in the corresponding paper, but that we derive for completeness.) As discussed above, “private-coin” refers to the fact that the users do not require access to a common random seed (typically easier to implement than “public-coin” protocols, where they do). We note that all public-coin protocols mentioned here can be made private-coin at the cost of an  $O(\log k)$  blowup factor in the communication cost.

Protocol	Private-coin	Communication	$\ell_\infty$ error	$\ell_2^2$ error
RAPPOR	✓	$k$	$\sqrt{\frac{\log k}{n \min(1, \varepsilon^2)}}$ $\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}$ $\sqrt{\frac{\log k}{n e^{\varepsilon/2}} + \frac{\log k}{n \varepsilon} \cdot \log n}$	$\frac{k e^{\varepsilon/2}}{n(e^{\varepsilon/2}-1)^2}$
Subset Selection [33]	✓	$\frac{k}{e^\varepsilon} \max(1, \varepsilon)$	?	$\frac{k e^\varepsilon}{n(e^\varepsilon-1)^2}$
(G)HR [3, Theorem 7]	✓	$\log k$	?	$\frac{k e^\varepsilon}{n(e^\varepsilon-1)^2}$
RHR [14, Theorem 3.1]	×	$\ell$	?	$\frac{k}{n \min((e^{\varepsilon/2}-1)^2, e^\varepsilon, 2^\ell, k)}$
No name [14, Theorem 3.1]	×	$\ell$	$\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2, \ell)}}$	$\frac{k e^{2\varepsilon/\ell}}{n \ell (e^{\varepsilon/\ell}-1)^2}$ ( $\dagger$ )
CountSketch-based [29]	×	$\max(\varepsilon, \log \frac{1}{\varepsilon})$	$\sqrt{\frac{\log k}{n(e^{\varepsilon/2}-1)^2}} + \frac{\log k}{n}$	$\frac{k e^{\varepsilon/2}}{n(e^{\varepsilon/2}-1)^2}$
PGR [24]	✓	$\log k$	$\sqrt{\frac{\log k}{n e^\varepsilon}} + \frac{\log k}{n \varepsilon} \cdot \log n$	$\frac{k e^\varepsilon}{n(e^\varepsilon-1)^2}$

## 1.2 Overview of results

We here summarise our main results, and briefly discuss the underlying techniques. While we focus on two protocols in particular, the techniques themselves are broadly applicable. Our first set of results concerns the  $\ell_\infty$  error rate achievable by LDP protocols for frequency estimation. We first recall, as a baseline, a general transformation which converts any LDP protocol with optimal error in the high-privacy regime into another LDP protocol with *reasonable error* in the low-privacy regime, at the cost of a blowup in communication:

► **Proposition 1** (Informal; see Proposition 9). *Let  $A$  be any locally private protocol for frequency estimation with expected  $\ell_\infty$  error  $O\left(\sqrt{\log k/(n\varepsilon^2)}\right)$  for  $\varepsilon \leq 1$ , using  $\ell$  bits of communication per user. Then there is a locally private protocol  $A'$  achieving error*

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

*for all  $\varepsilon > 0$ , using  $\ell \lceil \varepsilon \rceil$  bits of communication per user.*

We emphasise that this transformation is not new, and mimics an argument found in, e.g., [14]. This general transformation is quite appealing, as it provides (in theory) good performance in the low-privacy regime “for free” given any good enough LDP protocol for the high-privacy one. However, this comes at a price: first, the communication blowup, which could be impractical; second, a loss in constant factors, which while relatively small might still be prohibitive; and, perhaps more importantly, *this requires changing the existing algorithm* (and as a result the data analysis pipeline), which is often a significant hurdle. Our

second result, focusing on one of the earliest, versatile, and (at least in its “vanilla” version) conceptually simple LDP protocols for frequency estimation, RAPPOR [22], shows that this transformation is not actually necessary, and that RAPPOR actually achieves this improved bound *without* modification:

► **Theorem 2** (Informal; see Theorem 13). *The (simple version of) RAPPOR achieves expected  $\ell_\infty$  error*

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

for all  $\varepsilon > 0$ , using  $k$  bits of communication per user.

In comparison, using the generic transformation above on RAPPOR would require  $\lceil \varepsilon \rceil$   $k$ -bit messages per user, and, in terms of worst-case theoretical bounds, an expected error worse by a factor  $\simeq 2.04$ . As such, our first result can be summarized as saying that *analyzing (again) an existing algorithm can be better than modifying it* – and, quite importantly, that it may not be necessary to change an existing algorithmic pipeline to inherit better guarantees.

The proof of the above result relies on a careful analysis of the expected maximum of sums of Bernoulli random variables, and specifically on a fine-grained analysis of their subgaussian behaviour in the “highly biased” regime. While RAPPOR is particularly amenable to this analysis, we believe that this technique is highly general and applicable to a broad range of LDP protocols, for example those following the general “scheme template” of [4].

Yet, trying to establish optimality of this  $\min(\varepsilon, \varepsilon^2)$  scaling turned out to be very challenging. And indeed, *this is for a good reason*: as we show, it is actually possible to achieve significantly better error rate in the low-privacy regime – and, surprisingly, this much better error is again attained by RAPPOR, out-of-the-box:

► **Theorem 3** (Informal; see Theorem 14). *The (simple version of) RAPPOR achieves expected  $\ell_\infty$  error*

$$\tilde{O}\left(\max\left(\sqrt{\frac{\log k}{ne^{\varepsilon/2}}}, \frac{\log k}{n\varepsilon}\right)\right)$$

for all  $\varepsilon \geq 1$ , using  $k$  bits of communication per user.

The  $\tilde{O}$  notation here only hides a single  $\log n$  factor in the second term. As the  $\ell_\infty$  error is always at most 1, up to this  $\log n$  factor this is always better than Theorem 2.

The proof of this result requires going beyond the sub-gaussian concentration behaviour alluded to before, and instead analyse the maximum of these sums of Bernoulli random variables as *sub-gamma* random variables. More precisely, we draw upon very recent work by [17] (see also [10]) on “local Glivenko–Cantelli” results, which provide refined concentration bounds for mean estimation of high-dimensional product distributions – in our case, the distributions over  $\{0, 1\}^k$  arising from the use of RAPPOR.

The above result is appealing, in that it not only yields better error rate than previously known (or, in many cases, believed to be possible) in the  $\varepsilon \gg 1$  regime; but also in that it is achieved “for free” by an existing and widely used algorithm. However, it does have one negative aspect: namely, that RAPPOR is, in its standard version, very inefficient from a communication point of view, as it require one  $k$ -bit message from each user – far from the ideal  $O(\log k)$  bits one could hope for. Luckily, as mentioned above, the analysis underlying Theorem 3 is quite general, and applies to a broad range of locally private

## 25:6 Locally Private Histograms in All Privacy Regimes

estimation algorithms. While it does not lead to the improved bound for *all* such protocols, we show that it applies, for instance, to the recently proposed Projective Geometry Response protocol of Feldman, Nelson, Nguyen, and Talwar [24], whose performance was originally only analyzed for  $\ell_2$  error:

► **Theorem 4** (Informal; see Theorem 18). *Projective Geometry Response (PGR) achieves expected  $\ell_\infty$  error*

$$\tilde{O}\left(\max\left(\sqrt{\frac{\log k}{ne^\varepsilon}}, \frac{\log k}{n\varepsilon}\right)\right)$$

for all  $\varepsilon \geq 1$ , using  $O(\log k)$  bits of communication per user.

Note that again this is achieved “out-of-the-box” by an existing algorithm, without any modification! As a result, this inherits all of the features of PGR its authors originally established: crucially, its computational efficiency. We also point out that the stated guarantee is even slightly better than that of RAPPOR, as the first term now features an  $e^{-\varepsilon}$  dependence (instead of the  $e^{-\varepsilon/2}$  of Theorem 3).

We then turn to prove optimality of this error rate, and show that this is, up to a logarithmic factor, optimal for any LDP protocol:

► **Theorem 5** (Informal; see Theorem 24). *Any LDP protocol for frequency estimation must have, in the worst case, expected  $\ell_\infty$  error*

$$\Omega\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$$

for  $\varepsilon \in (0, 1]$ , and

$$\Omega\left(\max\left(\sqrt{\frac{\log k}{ne^\varepsilon}}, \frac{\log k}{n\varepsilon}\right)\right)$$

for  $\varepsilon \geq 1$ .

We remark that this lower bound also applies to the “easier” question of distribution learning, as we will state momentarily. The first part of this lower bound, as mentioned before, was already known; the second part is new, and, while not necessarily difficult to show in hindsight, does require significant care in combining inequalities between various information-theoretic quantities to avoid ending up with a vacuous bound. In this sense, our main contribution for the lower bound is to establish it in a self-contained, streamlined fashion, by drawing on the “chi-square contraction” framework of [2]; and – importantly – that it matches the upper bound obtained earlier in *both* the high- and low-privacy regimes. Another interesting aspect of this lower bound is that all three terms are derived separately, but from the same family of hard instances: a dataset where almost all users hold the same element.

### Implications for shuffle privacy

By combining our new results on LDP protocols for histograms in the low-privacy regime with known “amplification-by-shuffling” results, we are able to obtain a simple, robust shuffle (approximate) DP protocol using only *one* message per user of logarithmic length, while achieving state-of-the-art:

► **Theorem 6** (Informal; see Theorem 27). *Shuffling the Projective Geometry Response protocol achieves expected  $\ell_\infty$  error*

$$O\left(\frac{\sqrt{\log(k) \log(1/\delta)}}{n\varepsilon}\right)$$

for all  $\varepsilon \in [\Omega(1/\sqrt{n}), 1]$ , with  $O(\log k)$  bits of communication and a single message per user.

This matches that of the best known shuffle DP protocols, but with a much lower message and communication complexity. It was, to the best of our knowledge, still open whether achieving these “best of three worlds” guarantees was possible.

We conclude by providing the corollary of our results for distribution learning, using the known connection to frequency estimation (along with the standard lower bound of  $\Omega(1/\sqrt{n})$  on the error rate for the question, absent privacy constraints):

► **Corollary 7.** *For distribution learning under local privacy constraints, the minmax  $\ell_\infty$  error achievable is at most*

$$\tilde{O}\left(\max\left(\sqrt{\frac{\log k}{ne^\varepsilon}}, \frac{\log k}{n\varepsilon}, \frac{1}{\sqrt{n}}\right)\right)$$

for all  $\varepsilon \geq 1$ , and

$$O\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$$

for  $\varepsilon \in (0, 1]$ ; and this is tight, up to a logarithmic factor (in  $n$ ). Moreover, the upper bound is attained by PGR, using  $O(\log k)$  bits of communication per user.

## Organisation

After providing some background and setting notation, we establish our theoretical results (upper bounds on the error rate) in Section 3, followed by information-theoretic lower bounds in Section 4 and applications to shuffle privacy in Section 5. Finally, we discuss our results and potential future work in Section 6. Several proofs and a discussion of empirical analyses – omitted for clarity of exposition – can be found in the full version.

## 2 Preliminaries and notation

### (Local) differential privacy

We first recall the definition of local differential privacy (LDP):

► **Definition 8** (Locally private randomiser). *An algorithm  $Q: \mathcal{X} \rightarrow \mathcal{Y}$  satisfies  $\varepsilon$ -differential privacy if for all pairs of inputs  $x, x' \in \mathcal{X}$ , all sets of outputs  $S \subseteq \mathcal{Y}$  are  $\varepsilon$ -close,*

$$\Pr[Q(x) \in S] \leq e^\varepsilon \Pr[Q(x') \in S].$$

A *locally private protocol* is typically a pair of algorithms  $Q: \mathcal{X} \rightarrow \mathcal{Y}$ , executed by the user on their data, and  $A: \mathcal{Y}^n \rightarrow \mathcal{X}^k$ , executed by the server to transform the randomised outputs into an unbiased estimator for the true quantity. It is well known (see, e.g., [19]) that all LDP Frequency Estimation algorithms have a stochastic matrix representation which maps every element in the input alphabet to a probability distribution over the output

alphabet. Furthermore, [30] demonstrated that all optimal mechanisms obey a “binary” property, in that their stochastic matrix only contains probabilities weighted by either  $e^\varepsilon$  or 1 (appropriately normalised). Recently another “family of LDP protocols” was introduced by [4] and used in [24], where each input is associated with a set of high probability outputs determined by some useful set system.

### Frequency Estimation

Under frequency estimation, the aim is to estimate the empirical frequency of the observations. We denote the dataset of  $n$  observations as  $X^n := \{x_1, \dots, x_n\}$  and their empirical frequencies as the probability vector  $q := q(X^n)$ , where  $q_i := \frac{1}{n} \sum_{j=1}^n \mathbb{1}_{\{x_j=i\}}$ . That is,  $q$  is an element of the probability simplex,

$$\Delta_k := \left\{ p \in \mathbb{R}_{\geq 0}^k : \sum_{i=1}^k p_i = 1 \right\}$$

A *frequency estimator* over an alphabet  $[k]$  is a (possibly randomized) function  $\mathcal{A}: [k]^n \rightarrow \Delta_k$ , which approximates the true frequencies sufficiently well. The quality of the resulting approximation, denoted  $\hat{q}$ , is measured through a suitably chosen loss function. Typical choices of distance are the  $\ell_r$  norms, for  $r \in [1, \infty]$ ; in this paper, we will be mostly concerned with the  $\ell_\infty$  distance, where  $d(q, \hat{q}) = \|q - \hat{q}\|_\infty = \max_{1 \leq i \leq k} |q_i - \hat{q}_i|$ , and consider the *expected loss*  $\mathbb{E}[\|q - \hat{q}\|_\infty]$  (where the expectation is over the randomness of the frequency estimator).

### Distribution learning (estimation)

Given independent and identically distributed (i.i.d.) samples from an unknown probability distribution  $p$ , the goal of distribution learning (density estimation) is to find a distribution  $\hat{p}$  that approximates  $p$  sufficiently well. As before, the quality of the approximation is typically measured through a suitably chosen loss function, quantifying the distance between the true distribution  $p$  and its approximation  $\hat{p}$ . From the above, one can see that the key difference between frequency and distribution (under the same loss function) is that in the former, observations are arbitrary, while in the latter they are assumed i.i.d. from some common underlying probability distribution. As such, the expected loss also takes into account the randomness of the samples themselves.

### Frequency Estimation implies Distribution Estimation

For any given distance measure  $d$  we denote the expected sampling error (under loss  $d$ ) between the empirical histogram  $\hat{p}_n$  from  $n$  i.i.d. samples and the true underlying distribution  $p$  as  $\Phi(d, k, n)$ :

$$\Phi(d, k, n) = \sup_{p \in \Delta_k} \mathbb{E}[d(p, \hat{p}_n)]$$

Further, denote the optimal worst-case (over  $p$ ) expected error under loss  $d$  from  $n$  i.i.d. samples under  $\varepsilon$ -LDP constraints as  $\Phi^{\text{priv}} := \Phi^{\text{priv}}(d, \varepsilon, k, n)$ . The “optimal” here refers to quantifying over all locally private distribution estimators. We can analogously express the optimal expected loss when learning the empirical frequencies as  $\varphi^{\text{priv}} := \varphi^{\text{priv}}(d, \varepsilon, k, n)$ .

► **Fact 1.** *Accurate frequency estimation implies accurate distribution estimation: namely,  $\Phi^{\text{priv}}(d, \varepsilon, k, n) \leq \Phi(d, k, n) + \varphi^{\text{priv}}(d, \varepsilon, k, n)$ .*



For specific norms this allows us to derive (1) a lower bound on  $\varphi^{\text{priv}}$  (locally private frequency estimation) from a lower bound on  $\Phi^{\text{priv}}$  (locally private distribution estimation), and (2) an upper bound on  $\Phi^{\text{priv}}$  from an upper bound on  $\varphi^{\text{priv}}$ , recalling that (see, for instance, [31, 13])

$$\Phi(d, k, n) = \begin{cases} \Theta\left(\sqrt{\frac{k}{n}}\right) & \text{when } d = \ell_1 \\ \Theta\left(\frac{1}{\sqrt{n}}\right) & \text{when } d = \ell_2, \ell_\infty \end{cases}$$

### 3 Better algorithms in the low-privacy regime

Although the order-optimal  $\ell_\infty$  error rates for LDP frequency estimation and distribution learning are well-understood by now in the high-privacy regime, with many distinct algorithms achieving the tight

$$O\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$$

expected error bound, much less is understood about the best achievable error for large, or even medium, values of  $\varepsilon$ . In this section, we first revisit (and slightly generalize) an idea from [14], which shows how to convert any protocol “optimal in the high-privacy regime” to a related protocol “good enough in the low-privacy regime as well”, at the price of a blowup in communication (Section 3.1). We then focus on the specific example of RAPPOR, showing that – perhaps surprisingly – this simple protocol does already achieve the same bound by *without any modification* nor communication blowup (Section 3.2). We finally show that for distribution learning, this very same RAPPOR in fact achieves an *even better* error rate than this, leveraging very recent results on concentration of the empirical mean for high-dimensional distributions, due to [17, 10] (Section 3.3).

#### 3.1 A generic transformation, and a baseline

Here we prove the following generic statement:

► **Proposition 9.** *Suppose there exists a symmetric<sup>5</sup> LDP protocol  $A$  for frequency estimation achieving expected  $\ell_\infty$  error*

$$\mathcal{E}(n, k, \varepsilon) = O\left(\sqrt{\frac{\log k}{n \min(1, \varepsilon^2)}}\right)$$

for  $\varepsilon \leq 1$ , with  $m$  messages and  $\ell$  bits of communication per user. Then, for every integer  $L \geq 1$ , there exists an LDP protocol  $A'$  for frequency estimation achieving expected  $\ell_\infty$  error

$$\mathcal{E}\left(n \cdot \min(\lceil \varepsilon \rceil, L), k, \frac{\varepsilon}{\min(\lceil \varepsilon \rceil, L)}\right) = O\left(\sqrt{\frac{\log k}{n \min(L, \varepsilon, \varepsilon^2)}}\right)$$

for all  $\varepsilon > 0$ , with  $m \min(\lceil \varepsilon \rceil, L)$  messages and  $\ell \min(\lceil \varepsilon \rceil, L)$  bits of communication per user. Further, if  $A$  is a public-coin (resp. private-coin) protocol, then so is  $A'$ .

<sup>5</sup> i.e., where all users use the same randomiser.

## 25:10 Locally Private Histograms in All Privacy Regimes

The idea behind this result is not new, and is used (in a slightly less general form) in [14, Appendix E.4]. We restate and prove it in this paper in a self-contained form, as we believe it to be of broader interest.

As a direct corollary (setting  $L = \lceil \varepsilon \rceil$ ), applying the above to Hadamard Response and RAPPOR, for instance, we obtain the following bounds:

► **Corollary 10.** *For every  $\varepsilon > 0$ , there exists a private-coin  $\varepsilon$ -LDP protocol (namely, a modification of Hadamard Response) for frequency estimation with expected  $\ell_\infty$  error*

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

using  $\lceil \varepsilon \rceil$  messages per user, and  $\log k + O(1)$  bits per message.

► **Corollary 11.** *For every  $\varepsilon > 0$ , there exists a private-coin  $\varepsilon$ -LDP protocol (namely, a modification of RAPPOR) for frequency estimation with expected  $\ell_\infty$  error*

$$O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right)$$

using  $\lceil \varepsilon \rceil$  messages per user, and  $k$  bits per message.

### 3.2 Tighter analysis for RAPPOR

To do so, we first recall some facts and notation about RAPPOR, which will help with the analysis of its performance. A common simplification of the RAPPOR protocol [4, 3] is to parameterise it by  $\varepsilon$  as follows.

1. Given input  $x_i \in [k]$ , one-hot encode it onto a  $k$ -bit vector  $e_x$  s.t. only the  $x$ 'th bit is 1.
2. Flip each bit independently with probability  $\frac{1}{e^{\varepsilon/2} + 1}$  and send the resulting noisy bit-vector  $Y_i$  to the server.<sup>6</sup>
3. The server then receives  $n$  noisy bit vectors, computes  $\bar{Y} \leftarrow \frac{1}{n} \sum_{i=1}^n Y_i$ , and estimates

$$\hat{q} = \frac{e^{\varepsilon/2} + 1}{e^{\varepsilon/2} - 1} \bar{Y} - \frac{1}{e^{\varepsilon/2} - 1} \mathbf{1}_k \quad (2)$$

where  $\mathbf{1}$  is the 1-vector of size  $k$ .

A standard fact, which we recall here for completeness, is that  $\hat{q}$  defined above is an unbiased estimator for the true vector of frequencies  $q$ :

► **Lemma 12** (Expectation of  $\hat{q}$ ). *We have  $\mathbb{E}[\hat{q}] = q$ .*

We then have  $\mathbb{E}[\|\hat{q} - q\|_\infty] = \frac{e^{\varepsilon/2} + 1}{e^{\varepsilon/2} - 1} \mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty]$ , and so to bound the expected  $\ell_\infty$  error it suffices to bound that of the expected maximum deviation of the  $\bar{Y}_j$ 's from their mean. Now, each  $\bar{Y}_j$  is the (normalised) sum of  $n$  independent Bernoulli random variables: the standard way to analyse this maximum is to recall that a sum of  $n$  independent Bernoullis is a sub-gaussian random variable with parameter<sup>7</sup> at most  $\frac{n}{4}$ . Standard results (see, for example [32, Chapter 2]) on the maximum of  $k$  (not necessarily independent) sub-Gaussian random variables then give

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \leq \frac{e^{\varepsilon/2} + 1}{e^{\varepsilon/2} - 1} \sqrt{\frac{\log k}{2n}}$$

<sup>6</sup> This is just the ‘‘Permanent Randomised Response’’ step of RAPPOR parameterised by  $\varepsilon$ .

<sup>7</sup> Importantly, note that this may not coincide with the variance of  $X$ , although it does in some important cases (e.g., for a Gaussian r.v.).

which indeed behaves as  $O(\sqrt{\log k/(n\varepsilon^2)})$  for small  $\varepsilon$ . However, the bound quickly degrades for large  $\varepsilon$ , and only yields  $O(\sqrt{\log k/n})$  (no dependence on the privacy parameter at all!) as  $\varepsilon$  grows.

To remedy this, we will need a tighter analysis of the subgaussian parameter of Bernoulli random variables in the “very biased” regime (which is the one we have to handle for large  $\varepsilon$ , as then  $\frac{1}{e^{\varepsilon/2}+1} \approx 0$ ). Specifically, we will rely on the following characterisation of the sub-gaussian norm  $\sigma^2(p)$  of a (centered) Bern( $p$ ) random variable, known as the Kearns–Saul inequality, and which can be found in, e.g., [12]:

$$\sigma^2(p) = \begin{cases} 0, & p \in \{0, 1\} \\ \frac{1}{4}, & p = \frac{1}{2} \\ \frac{2p-1}{2 \log \frac{p}{1-p}}, & \end{cases} \quad (3)$$

Importantly, this expression is symmetric:  $\sigma^2(p) = \sigma^2(1-p)$  for all  $p \in [0, 1]$ . In our setting, each  $Y_j$  is the sum of  $n$  independent Bernoulli random variables with one of two symmetric parameters,  $p := \frac{1}{e^{\varepsilon/2}+1}$  or  $1-p = \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1}$ . The sub-Gaussian norm at play is then  $\sigma^2(p) = \frac{e^{\varepsilon/2}-1}{(e^{\varepsilon/2}+1)\varepsilon}$ . Using sub-additivity of the sub-gaussian parameter for independent random variables, we can bound the sub-Gaussian norm of  $Y_j$  as

$$\sigma^2(Y_j) \leq \frac{(e^{\varepsilon/2} - 1)}{(e^{\varepsilon/2} + 1)n\varepsilon} \quad (4)$$

We can use this to bound the maximum error of the estimator.

► **Theorem 13** (Expected maximum error of RAPPOR). *The expected maximum of  $\hat{q}$  is given by*

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \leq \sqrt{\frac{2(e^{\varepsilon/2} + 1) \log k}{n(e^{\varepsilon/2} - 1)\varepsilon}} \in O\left(\sqrt{\frac{\log k}{n \min(\varepsilon, \varepsilon^2)}}\right) \quad (5)$$

proving the main result of this subsection.

### 3.3 Optimal analysis for RAPPOR

It is natural to wonder whether this  $\min(\varepsilon, \varepsilon^2)$  dependence on the privacy parameter is order-optimal; especially as it appears in other locally private estimation tasks, such as mean estimation for high-dimensional Gaussians or product distributions [20, 1]. Quite surprisingly, we will show that for frequency estimation the  $\ell_\infty$  error rate given in Equation (5) is *not* optimal for large values of  $\varepsilon$ , and that even the simple RAPPOR algorithm can achieve significantly better.

► **Theorem 14.** *For  $\varepsilon \geq 1$ , the expected  $\ell_\infty$  error of RAPPOR for frequency estimation satisfies*

$$\mathbb{E}[\|\hat{q} - q\|_\infty] = O\left(\sqrt{\frac{\log k}{n\varepsilon^{3/2}}} + \frac{\log k}{n\varepsilon} \cdot \log n\right).$$

Importantly, this is better than the bound in Equation (5) whenever  $n = \tilde{\Omega}(\log k/\varepsilon)$ , which is the regime of interest (small constant, or vanishing, error). To see why this better bound may hold for large  $\varepsilon$ , recall that the bound given in Equation (5) relies on analyzing the

## 25:12 Locally Private Histograms in All Privacy Regimes

expected maximum of  $k$  (centered) Binomials random variables using their *sub-gaussian* behavior. This is good when the parameters of the Binomials are not too skewed; however, in the low privacy regimes, the parameters of the Bernoulli summands become very close to 0 (or 1): in that case, to analyze the expected maximum of the Binomials it is tighter to see them as having a *sub-gamma* behavior. Details follow.

**Proof of Theorem 14.** As alluded to above, we want to analyze the expected behavior of the maximum of Binomial random variables *beyond* the sub-gaussian regime. Invoking generic bounds for sub-gamma random variables such as [11, Corollary 2.6], unfortunately, does not lend itself to the order-optimal bounds either. Instead, we rely on the “local Glivenko–Cantelli” bounds recently obtained by [17, 10], which provide a more refined upper bound: to introduce the result we will invoke, we first need some notation.

► **Definition 15.** Given  $\mu \in [0, 1]^{\mathbb{N}}$ , denote by  $\tilde{\mu} \in [0, 1/2]^{\mathbb{N}}$  the sequence defined by  $\tilde{\mu}_i = \min(\mu, 1 - \mu)$  for all  $i \geq 1$ , and by  $\tilde{\mu}^\downarrow$  its non-increasing rearrangement. Finally, let  $p_\mu$  denote the product distribution over  $\{0, 1\}^{\mathbb{N}}$  with mean vector  $\mu$ .

With this in hand, the main result of [17] can be restated as follows:

► **Theorem 16** ([17, Theorem 3]). Let  $n \geq 21$ , and suppose that  $\mu \in [0, 1]^{\mathbb{N}}$  is such that  $\lim_{i \rightarrow \infty} \tilde{\mu}_i^\downarrow = 0$ . Let  $\hat{\mu}_n = \frac{1}{n} \sum_{j=1}^n X_j \in [0, 1]^{\mathbb{N}}$  denote the empirical estimator for  $\mu$  given  $n$  i.i.d. samples  $X_1, \dots, X_n$  from  $p_\mu$ . Then the following holds:

$$\mathbb{E}[\|\hat{\mu}_n - \mu\|_\infty] \lesssim \sup_{i \geq 1} \sqrt{\frac{\tilde{\mu}_i^\downarrow \log(i+1)}{n}} + \frac{\log n}{n} \sup_{i \geq 1} \frac{\log(i+1)}{\log \frac{1}{\tilde{\mu}_i^\downarrow}}$$

(Note that [10] recently improved on this results by up to a  $\log n$  factor in the second term. For simplicity, we use the slightly weaker result of [17], as it is easier to manipulate.)

We want to apply this result to bounding  $\mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty]$ . However, we face one obstacle in doing so, as we do not have a product distribution over  $\{0, 1\}^k$ : while  $\bar{Y}_1, \dots, \bar{Y}_k$  are indeed independent, and each of the form

$$\bar{Y}_i = \frac{1}{n} \sum_{j=1}^n Y_{i,j}$$

where  $Y_{i,1}, \dots, Y_{i,n}$  are independent Bernoulli random variables, these Bernoullis are not identically distributed: exactly  $n_i = nq_i$  of them have parameter  $1 - p = \frac{e^{\epsilon/2}}{e^{\epsilon/2} + 1}$ , and the remaining  $n - n_i$  have parameter  $p = \frac{1}{e^{\epsilon/2} + 1}$ . Of course, we do not know the  $n_i$ 's, as this is what we are trying to estimate; all we have is that

$$\sum_{i=1}^k n_i = n.$$

To circumvent this issue, let us write, for each  $1 \leq i \leq k$ ,

$$\bar{Y}_i = \bar{Y}_i^+ + \bar{Y}_i^-$$

where  $n\bar{Y}_i^+ \sim \text{Bin}(n_i, 1 - p)$ ,  $n\bar{Y}_i^- \sim \text{Bin}(n - n_i, p)$  are independent. We can then express

$$\begin{aligned}
\mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty] &= \mathbb{E}\left[\max_{1 \leq i \leq k} |\bar{Y}_i - \mathbb{E}[\bar{Y}_i]|\right] \\
&= \mathbb{E}\left[\max_{1 \leq i \leq k} (|\bar{Y}_i^+ - \mathbb{E}[\bar{Y}_i^+]| + |\bar{Y}_i^- - \mathbb{E}[\bar{Y}_i^-]|)\right] \\
&\leq \mathbb{E}\left[\max_{1 \leq i \leq k} |\bar{Y}_i^+ - \mathbb{E}[\bar{Y}_i^+]| \right] + \mathbb{E}\left[\max_{1 \leq i \leq k} |\bar{Y}_i^- - \mathbb{E}[\bar{Y}_i^-]| \right] \tag{6}
\end{aligned}$$

Now, instead of taking the maximum of  $k$  sums of Bernoullis with different parameters but same number of summands ( $n$  summands), we take the maximum of  $k$  sums of Bernoullis with the same parameter (i.e., Binomials) but different number of summands (at most  $n$ ). This does not necessarily seem like an improvement, and still does not let us apply Theorem 16 to either of the two expectations. However, *if* we could argue that “adding summands to each Binomial” cannot decrease the expected maximum, then we would be in good shape: that is, we want to upper bound

$$\mathbb{E}\left[\max_{1 \leq i \leq k} |\bar{Y}_i^+ - \mathbb{E}[\bar{Y}_i^+]| \right]$$

by

$$\mathbb{E}\left[\max_{1 \leq i \leq k} |Z_i^+ - \mathbb{E}[Z_i^+]| \right]$$

where  $nZ_i^+ \sim \text{Bin}(n, 1-p)$  (instead of  $\text{Bin}(n_i, 1-p)$ ). Intuitively, this seems reasonable, as adding independent summands should make the Binomial more likely to deviate from its expectation. The next lemma makes this intuition rigorous:

► **Lemma 17.** *Fix  $n_1, \dots, n_k, m_1, \dots, m_k \in \mathbb{N}$  and  $p_1, \dots, p_k \in [0, 1]$ , with  $n_i \leq m_i$  for all  $i$ . Let  $N_1, \dots, N_k$  and  $M_1, \dots, M_k$  be (not necessarily independent) random variables with  $N_i \sim \text{Bin}(n_i, p_i)$  and  $M_i \sim \text{Bin}(m_i, p_i)$ . Then*

$$\mathbb{E}\left[\max_{1 \leq i \leq k} |N_i - \mathbb{E}[N_i]| \right] \leq \mathbb{E}\left[\max_{1 \leq i \leq k} |M_i - \mathbb{E}[M_i]| \right].$$

**Proof.** Set  $\tilde{N}_i := N_i - \mathbb{E}[N_i]$  and  $\tilde{M}_i := M_i - \mathbb{E}[M_i]$  for all  $1 \leq i \leq k$ . We give a coupling of  $\tilde{N}_i, \tilde{M}_i$  such that

$$\mathbb{E}[\tilde{M}_i \mid \tilde{N}_i] = \tilde{N}_i.$$

Such a coupling can be obtained by setting  $\tilde{M}_i = \tilde{N}_i + \Delta_i - \mathbb{E}[\Delta_i]$ , where  $\Delta_i \sim \text{Bin}(m_i - n_i, p_i)$  is independent of  $\tilde{N}_i$ . Then it is easy to check that  $\tilde{M}_i$  has the right distribution, since  $N_i + \Delta_i \sim \text{Bin}(m_i, p_i)$ ; and the the conditional expectation is indeed as claimed. Using this coupling, we obtain

$$\begin{aligned}
\mathbb{E}\left[\max_{1 \leq i \leq k} |\tilde{N}_i|\right] &= \mathbb{E}\left[\max_{1 \leq i \leq k} |\mathbb{E}[\tilde{M}_i \mid \tilde{N}_i]| \right] \\
&\leq \mathbb{E}\left[\mathbb{E}\left[\max_{1 \leq i \leq k} |\tilde{M}_i| \mid \tilde{N}_i\right]\right] && \text{(Jensen's inequality)} \\
&= \mathbb{E}\left[\max_{1 \leq i \leq k} |\tilde{M}_i|\right]
\end{aligned}$$

establishing the lemma.<sup>8</sup> ◀

<sup>8</sup> More generally, via the existence of this coupling the argument shows that  $(\tilde{N}_1, \dots, \tilde{N}_k) \preceq_{\text{cx}} (\tilde{M}_1, \dots, \tilde{M}_k)$  (domination in the convex order), which in turn is equivalent to having  $\mathbb{E}[\phi(\tilde{N}_1, \dots, \tilde{N}_k)] \leq \mathbb{E}[\phi(\tilde{M}_1, \dots, \tilde{M}_k)]$  for every convex function  $\phi$ .

## 25:14 Locally Private Histograms in All Privacy Regimes

Let  $Z_1^+, \dots, Z_k^+$  (resp.  $Z_1^-, \dots, Z_k^-$ ) be i.i.d. with  $nZ_i^+ \sim \text{Bin}(n, 1-p)$  random variables (resp.  $nZ_i^- \sim \text{Bin}(n, p)$ ). Invoking Lemma 17 with  $m_1 = \dots = m_k = n$  separately on the two expectations of Equation (6), we get

$$\mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty] \leq \mathbb{E}\left[\max_{1 \leq i \leq k} |Z_i^+ - \mathbb{E}[Z_i^+]| \right] + \mathbb{E}\left[\max_{1 \leq i \leq k} |Z_i^- - \mathbb{E}[Z_i^-]| \right]$$

Both of the terms in the RHS now fit the setting of Theorem 16. Moreover, since their parameters are  $p$  (for the first expectation) and  $1-p$  (for the second), in both case the corresponding  $\tilde{m}u_i = \min(p, 1-p) = p = \frac{1}{e^{\varepsilon/2} + 1}$  is the same, and applying the theorem will give the same upper bound for both expectations. Thus, Theorem 16 yields

$$\mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty] \lesssim \sqrt{\frac{p \log(k+1)}{n}} + \frac{\log n}{n} \cdot \frac{\log(k+1)}{\log \frac{1}{p}}$$

Recalling the setting of  $p$ , along with the fact that  $\mathbb{E}[\|\hat{q} - q\|_\infty] = \frac{e^{\varepsilon/2} + 1}{e^{\varepsilon/2} - 1} \mathbb{E}[\|\bar{Y} - \mathbb{E}[\bar{Y}]\|_\infty]$ , finally gives

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \lesssim \sqrt{\frac{\log(k+1)}{n} \cdot \frac{e^{\varepsilon/2} + 1}{(e^{\varepsilon/2} - 1)^2}} + \frac{\log n}{n} \cdot \frac{e^{\varepsilon/2} + 1}{e^{\varepsilon/2} - 1} \cdot \frac{\log(k+1)}{\log(e^{\varepsilon/2} + 1)} \quad (7)$$

To conclude, observe that the right hand side is  $O\left(\sqrt{\frac{\log k}{n\varepsilon^2}} + \frac{\log k}{n\varepsilon} \cdot \log n\right)$  for  $\varepsilon \leq 1$ , and  $O\left(\sqrt{\frac{\log k}{ne^{\varepsilon/2}}} + \frac{\log k}{n\varepsilon} \cdot \log n\right)$  for  $\varepsilon \geq 1$ .  $\blacktriangleleft$

### 3.4 Upper bound for Projective Geometry Response

Projective geometry response (PGR), introduced by [24] achieves optimal rates for  $\ell_2^2$  error, communication, and near-optimal processing time. The protocol is based on the general template established in [4]: specifically, PGR relies on a set structure defined by *projective planes*, as detailed next. For a prime power  $d$  and  $k = \frac{d^t - 1}{d - 1}$ , the authors define a  $t$ -dimensional vector space  $\mathbb{F}_d^t$ , where each element  $x \in [k]$  is represented by one of the canonical basis vectors. These basis vectors in turn each uniquely determine a projective plane  $S(x)$ , such that there are  $s = |S(x)| = \frac{d^{t-1} - 1}{d - 1}$  “high probability” elements. Every one of these sets in turn has an intersection with every other set of size  $c = |S(x) \cap S(x')| = \frac{d^{t-2} - 1}{d - 1}$ . By choosing a prime power  $d \approx e^\varepsilon + 1$ , optimal error is achieved.

► **Theorem 18.** *Projective geometry response (PGR) [24] achieves the optimal rate for  $\ell_\infty$  error. More specifically, the expected  $\ell_\infty$  error of Projective Geometry Response for frequency estimation satisfies*

$$\begin{aligned} \mathbb{E}[\|\hat{q} - q\|_\infty] &\leq \sqrt{\frac{16(2e^\varepsilon + 1)^2 \log(k+1)}{e^\varepsilon (e^\varepsilon - 1)^2 n}} + \frac{4(2e^\varepsilon + 1) \log(k+1)}{(e^\varepsilon - 1)\varepsilon n} \log n \\ &\in O\left(\sqrt{\frac{\log k}{ne^\varepsilon}} + \frac{\log k}{n\varepsilon} \log n\right) \end{aligned}$$

We will make use of the size of each subset  $s$  and the size of each intersection  $c$  as described above, as well as the probability of returning any element  $y$  from the output alphabet, given  $x$  as an input:

$$Q(Y = y \mid X = x) = \begin{cases} \frac{e^\varepsilon}{se^\varepsilon + k - s} & \text{if } y \in S(x) \\ \frac{s - c}{se^\varepsilon + k - s} & \text{otherwise.} \end{cases} \quad (8)$$

The estimate  $\hat{q}$  of the frequency vector is then given by

$$\hat{q}_x := \alpha \cdot \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{Y_i \in S(x)} + \beta, \quad x \in [k] \quad (9)$$

where

$$\alpha = \frac{(e^\varepsilon - 1)s + k}{(e^\varepsilon - 1)(s - c)}, \quad \beta = -\frac{(e^\varepsilon - 1)c + s}{(e^\varepsilon - 1)(s - c)} \quad (10)$$

so that  $\hat{q}$  is an unbiased estimator of the true frequency vector  $q$ . Since a prime power can be found within a factor 2 of any number, we can choose  $d$  such that

$$e^\varepsilon + 1 \leq d \leq 2(e^\varepsilon + 1) \quad (11)$$

(While we could instead choose  $d < e^\varepsilon + 1$  and set the inequality to be an undershoot by a factor two, an investigation in that direction led to worse constants and trickier analysis.)

Recall that, for every integer  $\ell$ ,

$$d^\ell - 1 = (d - 1)(1 + d + \dots + d^{\ell-1}), \quad (12)$$

an identity we will rely on extensively in the rest of the section.

► **Remark 19.** Note that, as introduced above, we must have  $k = \frac{d^t - 1}{d - 1}$  for a prime power  $d$  satisfying (11). Other values of  $k$  must be rounded up, losing up to a factor  $O(e^\varepsilon)$  in the domain size (i.e., working instead with a domain size  $k' = O(e^\varepsilon k)$ ). We hereafter ignore this detail, which does not affect the final bound of Theorem 18 unless  $\varepsilon \gg \log k$ ; and assume that  $k$  is of the form stated above. For the same reason, we additionally can assume  $t \geq 3$ .

**Proof of Theorem 18.** We start with two lemmas which will be useful in proving the optimal error rate.

► **Lemma 20.** *The common size  $s$  of every subset  $S(x)$ ,  $x \in [k]$ , satisfies  $s \geq e^\varepsilon + 2$ .*

**Proof.** We can rewrite  $s$  as

$$s = \frac{d^{t-1} - 1}{d - 1} = \sum_{i=0}^{t-2} d^i$$

As  $t \geq 3$ , we have, by applying Equations (11) and (12)

$$s = \frac{d^{t-1} - 1}{d - 1} = \sum_{i=0}^{t-2} d^i \geq 1 + d \geq e^\varepsilon + 2. \quad \blacktriangleleft$$

Now remembering that whatever expected error we compute will be multiplied by the normalising constant  $\alpha$ , we would like to bound that in terms of  $\varepsilon$ .

► **Lemma 21.** *The value  $\alpha$  defined in (10) satisfies  $\alpha \leq 2 + \frac{2+d}{e^\varepsilon - 1} \leq 2 + \frac{4+2e^\varepsilon}{e^\varepsilon - 1}$ .*

**Proof.** First note that

$$\alpha = \frac{(e^\varepsilon - 1)s + k}{(e^\varepsilon - 1)(s - c)} = \frac{s}{(s - c)} + \frac{k}{(e^\varepsilon - 1)(s - c)}$$

## 25:16 Locally Private Histograms in All Privacy Regimes

and note that by simple application of Equation (12) we get  $s - c = d^{t-2}$ . Applying the identity again we get:

$$\frac{s}{d^{t-2}} = \sum_{i=0}^{t-2} \frac{1}{d^{i-t+2}} = \frac{1}{d^{t-2}} + \frac{1}{d^{t-1}} + \dots + 1 \leq 2 \quad (13)$$

where the final inequality comes from the fact that with  $d \geq 2$  this is bounded by the geometric series. For  $k$  we have the same series, with the addition of a single  $d$  term.

$$\frac{k}{d^{t-2}} = \sum_{i=0}^{t-1} \frac{1}{d^{i-t+2}} = \frac{1}{d^{t-2}} + \frac{1}{d^{t-1}} + \dots + 1 + d$$

The result is then at most  $2 + d \leq 2 + 2(e^\varepsilon + 1)$ , by applying Equation (11). ◀

We are now ready to apply Theorem 16. As we did for RAPPOR, we will break the expected maximum into the sum of error over two vectors of binomials  $Z^+ = \text{Bin}(n, p^+)$  and  $Z^- = \text{Bin}(n, p^-)$  where, recalling (8),

$$p^+ = \frac{e^\varepsilon}{se^\varepsilon + k - s}, \quad p^- = \frac{s - c}{se^\varepsilon + k - s} \quad (14)$$

In this case we do not have that  $p^+ = 1 - p^-$  so we will need to bound the following

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \leq \alpha \left( \sqrt{\frac{\log(k+1)}{n}} (\sqrt{p^+} + \sqrt{p^-}) + \frac{\log(k+1)}{n} \log(n) \left( \frac{1}{\log(1/p^+)} + \frac{1}{\log(1/p^-)} \right) \right).$$

We briefly note that while RAPPOR has clear independence between coordinates, the result of Theorem 16 due to [17] does not require independence; therefore it applies to PGR, and is likely suitable for the analysis of other subset-based protocols.

First we will upper bound  $\sqrt{p^+} + \sqrt{p^-}$ :

► **Lemma 22.** *For  $p^+, p^-$  defined in (14), we have  $p^+ + p^- \leq 2/e^\varepsilon$ , and so  $\sqrt{p^+} + \sqrt{p^-} \leq \sqrt{4/e^\varepsilon}$ .*

**Proof.** First,  $p^+$  is clearly less than  $\frac{1}{s}$  when we remove  $k - s$  from the denominator, and so is at most  $1/(e^\varepsilon + 2) \leq 1/e^\varepsilon$  by Lemma 20. For the other term, notice that by Equation (13) we have  $s/(s - c) \geq 1$ , from which,

$$p^- \leq \frac{s - c}{se^\varepsilon} \leq \frac{1}{e^\varepsilon}$$

Overall, we get that  $p^+ + p^- \leq 2/e^\varepsilon$ . The conclusion follows from the AM-GM inequality, as  $\sqrt{p^+} + \sqrt{p^-} \leq \sqrt{2(p^+ + p^-)}$ . ◀

Next we bound  $\log^{-1}(1/p^+) + \log^{-1}(1/p^-)$ :

► **Lemma 23.** *We have  $\frac{1}{\log(1/p^+)} + \frac{1}{\log(1/p^-)} \leq \frac{2}{\varepsilon}$ .*

**Proof.** We will proceed in both cases by lower-bounding the denominators. First,

$$\log(1/p^-) = \log\left(\frac{se^\varepsilon + k - s}{s - c}\right) \geq \log\left(\frac{se^\varepsilon}{s - c}\right) \geq \log(e^\varepsilon) = \varepsilon.$$



Next,

$$\begin{aligned} \log(1/p^+) &= \log\left(\frac{se^\varepsilon + k - s}{e^\varepsilon}\right) \geq \log(s) \\ &\geq \log(e^\varepsilon + 2) && \text{(Lemma 20)} \\ &\geq \varepsilon \end{aligned}$$

As such adding the reciprocals of both terms give a bound of  $2/\varepsilon$ .  $\blacktriangleleft$

The only step left to bound the  $\ell_\infty$  error of PGR is to take the product of these bounds. Proceeding to do so, we arrive at:

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \leq \sqrt{\frac{16(2e^\varepsilon + 1)^2 \log(k+1)}{e^\varepsilon(e^\varepsilon - 1)^2 n}} + \frac{4(2e^\varepsilon + 1) \log(k+1)}{(e^\varepsilon - 1)\varepsilon n} \log n$$

which in the low privacy regime gives:

$$\mathbb{E}[\|\hat{q} - q\|_\infty] \in O\left(\sqrt{\frac{\log k}{ne^\varepsilon}} + \frac{\log k}{n\varepsilon} \log n\right)$$

as claimed.  $\blacktriangleleft$

#### 4 Worst-case, information-theoretic lower bounds

We will follow the “chi-squared lower bound” framework of [2] to obtain our information-theoretic lower bounds against non-interactive locally private protocols:

► **Theorem 24.** *Fix any  $\varepsilon > 0$ . Any non-interactive (public- or private-coin) protocol  $\Pi$  for distribution estimation from  $n$  users must have minmax expected  $\ell_\infty$  error*

$$\Omega\left(\max\left(\sqrt{\frac{\log k}{n(e^\varepsilon - 1)^2}}, \sqrt{\frac{\log k}{ne^\varepsilon}}, \frac{\log k}{n\varepsilon}\right)\right)$$

**Proof.** Suppose there exists a (non-interactive, public- or private-coin)  $\varepsilon$ -LDP protocol  $\Pi$  for  $n$  users which learns any  $\mathbf{p}$  to expected  $\ell_\infty$  error  $\alpha$  when each user gets an independent sample from  $\mathbf{p}$  as input:

$$\mathbb{E}[\|\mathbf{p} - \hat{\mathbf{p}}\|_\infty] \leq \alpha \tag{15}$$

where  $\hat{\mathbf{p}}$  is the output of  $\Pi$  when run on  $X_1, \dots, X_n \sim \mathbf{p}$ .

Now, consider the family  $\mathcal{P}_\alpha = \{\mathbf{p}_z\}_{z \in [k]}$  of probability distributions over  $[k]$ , where, for  $z \in [k]$ ,  $\mathbf{p}_z$  is defined by

$$\mathbf{p}_z(x) = \frac{1 - 4\alpha}{k} + 4\alpha \mathbf{1}_{\{z=x\}}, \quad x \in [k] \tag{16}$$

(that is,  $\mathbf{p}_z$  is a mixture of the uniform distribution  $\mathbf{u}_k$  and a point mass on  $z$ , with mixture coefficients  $1 - 4\alpha$  and  $4\alpha$ ). Consider the following process: first, we select  $Z$  from  $[k]$  uniformly at random, then generate  $n$  i.i.d. samples  $X_1, \dots, X_n$  from  $\mathbf{p}_Z$  and run  $\Pi$  on these samples, obtaining a hypothesis  $\hat{\mathbf{p}}$ . We finally set  $\hat{Z}$  to be the element of the domain with the largest probability under  $\hat{\mathbf{p}}$ , i.e.,

$$\hat{Z} = \operatorname{argmax}_{i \in [k]} \hat{\mathbf{p}}(i)$$

## 25:18 Locally Private Histograms in All Privacy Regimes

(breaking ties arbitrarily). The first claim is that doing so allows one to guess the correct value of  $Z$  with high probability: indeed, since the gap between the highest and second highest element of  $\mathbf{p}_Z$  is  $4\alpha$ , we have

$$\Pr\left[\hat{Z} = Z\right] \geq \Pr\left[\|\mathbf{p} - \hat{\mathbf{p}}\|_\infty < 2\alpha\right] \geq 1 - \frac{\mathbb{E}[\|\mathbf{p}_Z - \hat{\mathbf{p}}\|_\infty]}{2\alpha} \geq \frac{1}{2} \quad (17)$$

the second-to-last inequality being Markov's. However, by Fano's inequality applied to the Markov chain  $Z - \mathbf{p}_Z - X^n - (Y^n, R) - \hat{\mathbf{p}} - \hat{Z}$  (where  $X^n$  is the tuple of i.i.d. samples, and  $(Y^n, R)$  denotes the tuple of  $n$  messages, along with the public randomness, resulting from the protocol  $\Pi$ ), we get, recalling that  $Z$  is chosen uniformly in  $[k]$ , that

$$\Pr\left[\hat{Z} = Z\right] \leq \frac{I(Z \wedge (Y^n, R)) + \log 2}{\log k} \quad (18)$$

and so, putting Equations (17) and (18) together, we get

$$I(Z \wedge (Y^n, R)) \geq \frac{1}{2} \log \frac{k}{4} = \Omega(\log k). \quad (19)$$

This gives us the first ingredient: a lower bound on  $I(Z \wedge (Y^n, R))$ . For the second, we need to obtain an upper bound on this same mutual information as a function of  $\varepsilon, n, k$ , and  $\alpha$ . To do so, observe first that, by the chain rule

$$I(Z \wedge (Y^n, R)) = I(Z \wedge Y^n | R) + I(Z \wedge R) = I(Z \wedge Y^n | R)$$

the second equality follows from the independence of the public randomness  $R$  from  $Z$ . This is convenient, as the messages  $Y^n = (Y_1, \dots, Y_n)$  are independent conditioned on  $R$ , and so we get

$$I(Z \wedge (Y^n, R)) \leq \sum_{j=1}^n I(Z \wedge Y_j | R) \leq n \cdot \max_{1 \leq j \leq n} I(Z \wedge Y_j | R) \quad (20)$$

Consider any user  $j$ , using locally private randomizer  $Q = Q_{j,R}: [k] \rightarrow \mathcal{Y}$  (for notational simplicity, we drop afterwards the dependence on  $j$  and  $R$ ). Let  $\bar{\mathbf{p}} = \mathbb{E}_Z[\mathbf{p}_Z]$  denote the average input distribution (over  $Z$ ), i.e., the uniform mixture of all  $\mathbf{p}_z$ 's. Then we can rewrite the mutual information as

$$I(Z \wedge Y_j | R) = \mathbb{E}_Z[\text{KL}(Q^{\mathbf{p}_Z} \parallel Q^{\bar{\mathbf{p}}})] \quad (21)$$

where, for a given input distribution  $\mathbf{p}$ ,  $Q^{\mathbf{p}}(\cdot) = \mathbb{E}_{X \sim \mathbf{p}}[Q(\cdot | X)]$  denotes the output distribution (over  $\mathcal{Y}$ ) induced by the randomizer  $Q$  on input  $X \sim \mathbf{p}$ .

**First lower bound (good for small  $\varepsilon$ ).** We then proceed by upperbounding the KL divergence by the  $\chi^2$  one and unrolling the latter's definition, getting

$$\begin{aligned} I(Z \wedge Y_j | R) &\leq \mathbb{E}_Z[\chi^2(Q^{\mathbf{p}_Z} \parallel Q^{\bar{\mathbf{p}}})] = \mathbb{E}_Z\left[\sum_{y \in \mathcal{Y}} \frac{(\mathbb{E}_{X \sim \mathbf{p}_Z}[Q(y | X)] - \mathbb{E}_{X \sim \bar{\mathbf{p}}}[Q(y | X)])^2}{\mathbb{E}_{X \sim \bar{\mathbf{p}}}[Q(y | X)]}\right] \\ &= \mathbb{E}_Z\left[\sum_{y \in \mathcal{Y}} \frac{\left(\sum_{x \in [k]} Q(y | x)(\mathbf{p}_Z(x) - \bar{\mathbf{p}}(x))\right)^2}{\sum_{x \in [k]} Q(y | x)\bar{\mathbf{p}}(x)}\right] \end{aligned}$$

Note, observing that, for our choice of  $\mathcal{P}_\alpha$ ,  $\bar{\mathbf{p}}$  is simply the uniform distribution over  $[k]$ , and that for all  $x \in [k]$  we then have  $\mathbf{p}_Z(x) - \bar{\mathbf{p}}(x) = 4\alpha(\mathbf{1}_{\{Z=x\}} - \frac{1}{k})$ . Then, letting  $\Phi(x) = e_x - \frac{1}{k}\mathbf{1}_k \in \mathbb{R}^k$  for  $x \in [k]$ , we get

$$\begin{aligned}
I(Z \wedge Y_j \mid R) &\leq 16\alpha^2 k^2 \cdot \mathbb{E}_Z \left[ \sum_{y \in \mathcal{Y}} \frac{(\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)(\mathbf{1}_{\{Z=X\}} - \frac{1}{k})])^2}{\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)]} \right] \\
&= 16\alpha^2 k^2 \cdot \sum_{y \in \mathcal{Y}} \frac{\mathbb{E}_Z \left[ (\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)(\mathbf{1}_{\{Z=X\}} - \frac{1}{k})])^2 \right]}{\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)]} \\
&= 16\alpha^2 k \cdot \sum_{y \in \mathcal{Y}} \frac{\sum_{i=1}^k (\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)(\mathbf{1}_{\{i=X\}} - \frac{1}{k})])^2}{\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)]} \\
&= 16\alpha^2 k \cdot \sum_{y \in \mathcal{Y}} \frac{\sum_{i=1}^k \mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)\Phi(X)_i]^2}{\mathbb{E}_{X \sim \mathbf{u}}[Q(y \mid X)]} \tag{22}
\end{aligned}$$

Note that, for any fixed  $1 \leq i < j \leq k$ , we have

$$\mathbb{E}_{X \sim \mathbf{u}}[\Phi_i(X)] = 0, \quad \mathbb{E}_{X \sim \mathbf{u}}[\Phi_i(X)^2] = \frac{1}{k} \left(1 - \frac{1}{k}\right), \quad \mathbb{E}_{X \sim \mathbf{u}}[|\Phi_i(X)|] = \frac{2}{k} \left(1 - \frac{1}{k}\right),$$

For simplicity, we write  $\mathbb{E}_X[\cdot]$  for  $\mathbb{E}_{X \sim \mathbf{u}}[\cdot]$ . We will use the fact that,  $Q$  being  $\varepsilon$ -LDP, we have

$$|Q(y \mid x) - \mathbb{E}_X[Q(y \mid X)]| \leq (e^\varepsilon - 1)\mathbb{E}_X[Q(y \mid X)] \tag{23}$$

for every  $x \in [k]$  and  $y \in \mathcal{Y}$ . With this in hand, starting from (22), we can write

$$I(Z \wedge Y_j \mid R) \leq 16\alpha^2 k \sum_{y \in \mathcal{Y}} \frac{\sum_{i=1}^k \mathbb{E}_X[Q(y \mid X)\Phi(X)_i]^2}{\mathbb{E}_X[Q(y \mid X)]} \leq 64\alpha^2 (e^\varepsilon - 1)^2.$$

where the derivation relies on the fact that  $\mathbb{E}_X[\Phi_i(X)] = 0$ , (23), and  $\mathbb{E}_{X \sim \mathbf{u}}[|\Phi_i(X)|] \leq \frac{2}{k}$  (details can be found in the full version). Using this last bound along with Equations (19) and (20), we get that

$$\frac{1}{2} \log \frac{k}{4} \leq I(Z \wedge (Y^n, R)) \leq 64\alpha^2 n (e^\varepsilon - 1)^2, \tag{24}$$

i.e.,

$$\alpha \geq \frac{1}{8\sqrt{2}} \sqrt{\frac{\log \frac{k}{4}}{n(e^\varepsilon - 1)^2}}, \tag{25}$$

showing the  $\Omega\left(\sqrt{\frac{\log k}{n\varepsilon^2}}\right)$  lower bound for small  $\varepsilon$ .

The second and third lower bounds can be established in an analogous fashion: the details can be found in the full version.  $\blacktriangleleft$

## 5 Amplification by shuffling

As mentioned in the Introduction, one of the key motivations for studying locally private histogram estimation in the low privacy (high- $\varepsilon$ ) regime is the implication for histogram estimation in the *shuffle* model of privacy (see, e.g., [27, 26, 7, 16]), in light of the “plug-and-play” amplification-by-shuffling results allowing to “translate” the former into the latter. Specifically, we will use the following result of Feldman, McMillan, and Talwar:

## 25:20 Locally Private Histograms in All Privacy Regimes

► **Theorem 25** ([23, Theorem 3.1]). For any domain  $\mathcal{X}$ , let  $R: \mathcal{X} \rightarrow \mathcal{Y}$  be an  $\varepsilon_L$ -DP local randomiser; and let  $S$  be the algorithm that given a tuple of  $n$  messages  $\vec{y} \in \mathcal{Y}^n$ , samples a uniform random permutation  $\pi$  over  $[n]$  and outputs  $(\vec{y}_{\pi(1)}, \dots, \vec{y}_{\pi(n)})$ . Then for any  $\delta \in (0, 1]$  such that  $\varepsilon_L \leq \log \frac{n}{16 \log(2/\delta)}$ ,  $S \circ R^n$  is  $(\varepsilon, \delta)$ -DP, where

$$\varepsilon \leq \log \left( 1 + 8 \frac{e^{\varepsilon_L} - 1}{e^{\varepsilon_L} + 1} \left( \sqrt{\frac{e^{\varepsilon_L} \log(4/\delta)}{n}} + \frac{e^{\varepsilon_L}}{n} \right) \right).$$

In particular, if  $\varepsilon_L \geq 1$  then  $\varepsilon = O(\sqrt{e^{\varepsilon_L} \log(1/\delta)/n})$ , and if  $\varepsilon_L < 1$  then  $\varepsilon = O(\varepsilon_L \sqrt{\log(1/\delta)/n})$ .

This implies the following:

► **Lemma 26** (Amplification by shuffling). Fix any  $\delta \in (0, 1]$ ,  $\varepsilon \in (0, 1]$ , and  $n$  such that  $\varepsilon > 16\sqrt{\log(4/\delta)/n}$ . Then, for

$$\varepsilon_L := \log \frac{\varepsilon^2 n}{256 \log(4/\delta)} = \Theta \left( \log \frac{\varepsilon^2 n}{\log(1/\delta)} \right),$$

shuffling the messages of  $n$  users using the same  $\varepsilon_L$ -LDP randomizer satisfies (robust)  $(\varepsilon, \delta)$ -shuffle differential privacy.

**Proof.** Note that for  $\varepsilon, \delta$  as in the statement and  $\varepsilon_L$  as defined, we have  $0 < \varepsilon_L \leq \log \frac{n}{16 \log(2/\delta)}$ . Applying Theorem 25, we get  $(\varepsilon', \delta)$  privacy for

$$\varepsilon' = \log \left( 1 + 8 \underbrace{\frac{e^{\varepsilon_L} - 1}{e^{\varepsilon_L} + 1}}_{\leq 1} \left( \frac{\varepsilon}{16} + \underbrace{\frac{\varepsilon^2}{256 \log(4/\delta)}}_{\leq \varepsilon/16} \right) \right) \leq \varepsilon$$

which proves the statement. ◀

► **Theorem 27.** For  $n = \Omega\left(\frac{\log(1/\delta)}{\varepsilon^2}\right)$  and  $\varepsilon \in (0, 1]$ , Shuffled Projective Geometry Response achieves maximum error  $O\left(\frac{\sqrt{\log(k) \log(1/\delta)}}{n\varepsilon}\right)$ , with  $O(\log k)$  bits of communication (and one single message) per user.

**Proof.** For  $n \geq \frac{500}{\varepsilon^2} \log \frac{4}{\delta}$ , the restriction on  $\varepsilon$  from Lemma 26 holds, and, setting  $\varepsilon_L$  as in the lemma, we also have  $\varepsilon_L = \Omega(1)$ . We then invoke the bound of Theorem 18, focusing on first part of the bound, which dominates when  $\frac{\log k}{n} \cdot \log^2 n \leq \frac{\varepsilon_L^2}{e^{\varepsilon_L}}$ . This leads to an upper bound on the error of the order

$$\sqrt{\frac{\log k}{n e^{\varepsilon_L}}} = \sqrt{\frac{\log(k)}{n \varepsilon^2 \frac{n}{\log(1/\delta)}}} = \frac{\sqrt{\log(k) \log(1/\delta)}}{n \varepsilon}$$

as desired. It only remains to argue that the first term of the bound did, indeed, dominate the error. As mentioned above, this is the case whenever

$$\frac{e^{\varepsilon_L}}{\varepsilon_L^2} \cdot \log k \ll \frac{n}{\log^2 n}$$

for which a weaker, sufficient condition is  $\frac{e^{\varepsilon_L}}{\varepsilon_L^2} \ll \frac{n}{\log^2 n}$ , that is,  $n \gg e^{\varepsilon_L}$ . But this follows from our setting of  $\varepsilon_L$ , such that  $e^{\varepsilon_L} = \underbrace{\frac{\varepsilon^2}{256 \log(4/\delta)}}_{\ll 1} \cdot n$ . ◀

Comparing this result with the summary of local, shuffle, and central histogram error bounds available in [16, Table 1] shows that with shuffled PGR achieves the best error of any protocol which sends a constant number of messages.

More specifically, focusing on 3 representative known protocols: the only known one-message-per-user protocol, due to [16], achieves much worse error,<sup>9</sup> and requires  $k$  bits of communication per user; while the protocol of [26], which achieves the same error as Theorem 27, uses  $k^{\Omega(1)}$  ( $\log k$ )-bit messages per user. Finally, a protocol of [6] does achieve better error, but at the cost of performing  $k + 1$  rounds, each with  $\log k$  communication per user. Thus, our bounds demonstrates that shuffled PGR achieve state-of-the-art  $\ell_\infty$ -error with only *one* message of  $\log k$  bits.

## 6 Discussion and future work

**The logarithmic factor in the upper bound.** We prove our upper bounds using the result of [17], which is where the  $\log n$  term appears. The authors conjectured that this term could be removed in general, but follow-up work [10] demonstrated by counter-example that this is not the case. They do, however, show that there is a distribution-dependent interpolation between the term being  $T(n)/n$  and  $T(n)\log n/n$ . This logarithmic factor has implications for understanding which error regime is dominating, given the parameters of the algorithm  $(\varepsilon, k, n)$ . While it is possible to imagine that a careful application of the tools in [10] could resolve this matter in general or for a specific use-case, in the meantime we emphasise that empirical analysis is still crucial.

**Tighter analysis of other protocols.** While we analyse RAPPOR which admits a simple analysis due to the independence of coordinates, and Projective Geometry Response which represents the state of the art in low-communication LDP protocols, we believe the tools introduced in this paper are applicable in a very general way to most LDP protocols. Of particular interest would be Subset Selection, which has optimal mutual information between inputs and outputs of the local randomiser, and the two Hadamard Response protocols that we include in our empirical analysis.

**Histograms in the shuffle model.** While our shuffle DP result (Theorem 27) yields better error, communication, and number of messages, it does have a limitation on the parameter range, namely  $n = \Omega(\log(1/\delta)/\varepsilon^2)$ . It would be interesting to weaken this requirement to match the central DP one (where the dependence on  $\varepsilon$  is only linear). Moreover, one could hope to further improve the resulting error to the central DP bound (i.e., a  $\min(\log k, \log(1/\delta))$  dependence instead of  $\sqrt{\log k \log(1/\delta)}$ ), or, alternatively, prove a matching lower bound separating the two models.

**Broader applicability of lower bound techniques.** While we prove lower bounds for the specific case of  $\ell_\infty$  frequency estimation, the tools used are extremely general and we imagine that their application could provide new lower bounds for a variety of problems in the LDP setting, especially with the low-privacy regime in mind.

---

<sup>9</sup> And has the same restriction  $n = \Omega(\log(1/\delta)/\varepsilon^2)$  on the parameters.

## References

- 1 Jayadev Acharya, Clément L. Canonne, Ziteng Sun, and Himanshu Tyagi. Unified lower bounds for interactive high-dimensional estimation under information constraints. In *NeurIPS*, 2023.
- 2 Jayadev Acharya, Clément L. Canonne, and Himanshu Tyagi. Inference under information constraints I: Lower bounds from chi-square contraction. *Institute of Electrical and Electronics Engineers*, 66(12):7835–7855, 2020. doi:10.1109/TIT.2020.3028440.
- 3 Jayadev Acharya and Ziteng Sun. Communication complexity in locally private distribution estimation and heavy hitters. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 51–60, Long Beach, California, USA, 2019. PMLR. URL: <http://proceedings.mlr.press/v97/acharya19c.html>.
- 4 Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1120–1129. PMLR, 2019. URL: <http://proceedings.mlr.press/v89/acharya19a.html>.
- 5 Apple Privacy Team. Learning with privacy at scale, 2017. URL: <https://machinelearning.apple.com/research/learning-with-privacy-at-scale>.
- 6 Victor Balcer and Albert Cheu. Separating Local & Shuffled Differential Privacy via Histograms, April 2020. arXiv:1911.06879, doi:10.48550/arXiv.1911.06879.
- 7 Victor Balcer and Albert Cheu. Separating local & shuffled differential privacy via histograms. In *ITC*, volume 163 of *LIPICs*, pages 1:1–1:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ITC.2020.1.
- 8 Raef Bassily and Adam Smith. Local, Private, Efficient Protocols for Succinct Histograms. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, pages 127–135, June 2015. doi:10.1145/2746539.2746632.
- 9 Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnés, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. *CoRR*, abs/1710.00901, 2017. arXiv:1710.00901.
- 10 Moïse Blanchard and Vaclav Voracek. Tight bounds for local glivenko-cantelli. In Claire Vernade and Daniel Hsu, editors, *Proceedings of The 35th International Conference on Algorithmic Learning Theory*, volume 237 of *Proceedings of Machine Learning Research*, pages 179–220. PMLR, February 2024. URL: <https://proceedings.mlr.press/v237/blanchard24a.html>.
- 11 Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013.
- 12 V. V. Buldygin and K. K. Moskvichova. The sub-Gaussian norm of a binary random variable. *Theory of Probability and Mathematical Statistics*, 86:33–49, August 2013. doi:10.1090/S0094-9000-2013-00887-4.
- 13 Clément L. Canonne. A short note on learning discrete distributions, 2020. arXiv:2002.11457.
- 14 Wei-Ning Chen, Peter Kairouz, and Ayfer Özgür. Breaking the communication-privacy-accuracy trilemma. *IEEE Trans. Inf. Theory*, 69(2):1261–1281, 2023. doi:10.1109/TIT.2022.3218772.
- 15 Albert Cheu, Adam D. Smith, Jonathan R. Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *EUROCRYPT (1)*, volume 11476 of *Lecture Notes in Computer Science*, pages 375–403. Springer, 2019. doi:10.1007/978-3-030-17653-2\_13.
- 16 Albert Cheu and Maxim Zhilyaev. Differentially private histograms in the shuffle model from fake users. In *SP*, pages 440–457. IEEE, 2022. doi:10.1109/SP46214.2022.9833614.
- 17 Doron Cohen and Aryeh Kontorovich. Local glivenko-cantelli. In *COLT*, volume 195 of *Proceedings of Machine Learning Research*, page 715. PMLR, 2023. URL: <https://proceedings.mlr.press/v195/cohen23a.html>.

- 18 Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In *SIGMOD Conference*, pages 1655–1658. ACM, 2018. doi:10.1145/3183713.3197390.
- 19 John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013*, pages 429–438. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.53.
- 20 John C. Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In *COLT*, volume 99 of *Proceedings of Machine Learning Research*, pages 1161–1191. PMLR, 2019. URL: <http://proceedings.mlr.press/v99/duchi19a.html>.
- 21 Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Shuang Song, Kunal Talwar, and Abhradeep Thakurta. Encode, Shuffle, Analyze Privacy Revisited: Formalizations and Empirical Evaluation. *arXiv:2001.03618 [cs]*, January 2020. arXiv:2001.03618.
- 22 Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security, CCS '14*, pages 1054–1067, New York, NY, USA, 2014. ACM. doi:10.1145/2660267.2660348.
- 23 Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 954–964. IEEE, 2021. doi:10.1109/FOCS52979.2021.00096.
- 24 Vitaly Feldman, Jelani Nelson, Huy Nguyen, and Kunal Talwar. Private frequency estimation via projective geometry. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 6418–6433. PMLR, July 2022. URL: <https://proceedings.mlr.press/v162/feldman22a.html>.
- 25 Vitaly Feldman and Kunal Talwar. Lossless compression of efficient private local randomizers. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3208–3219. PMLR, July 2021. URL: <https://proceedings.mlr.press/v139/feldman21a.html>.
- 26 Badih Ghazi, Noah Golowich, Ravi Kumar, Rasmus Pagh, and Ameya Velingker. On the power of multiple anonymous messages: Frequency estimation and selection in the shuffle model of differential privacy. In *EUROCRYPT (3)*, volume 12698 of *Lecture Notes in Computer Science*, pages 463–488. Springer, 2021. doi:10.1007/978-3-030-77883-5\_16.
- 27 Badih Ghazi, Ravi Kumar, Pasin Manurangsi, and Rasmus Pagh. Private counting from anonymous messages: Near-optimal accuracy with vanishing communication overhead. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 3505–3514. PMLR, 2020. URL: <http://proceedings.mlr.press/v119/ghazi20a.html>.
- 28 Justin Hsu, Sanjeev Khanna, and Aaron Roth. Distributed private heavy hitters. In *ICALP (1)*, volume 7391 of *Lecture Notes in Computer Science*, pages 461–472. Springer, 2012. doi:10.1007/978-3-642-31594-7\_39.
- 29 Ziyue Huang, Yuan Qiu, Ke Yi, and Graham Cormode. Frequency estimation under multiparty differential privacy: One-shot and streaming. *Proc. VLDB Endow.*, 15(10):2058–2070, 2022. doi:10.14778/3547305.3547312.
- 30 Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. *Journal of Machine Learning Research*, 17(17):1–51, 2016. URL: <http://jmlr.org/papers/v17/15-135.html>.
- 31 Sudeep Kamath, Alon Orlitsky, Dheeraj Pichapati, and Ananda Theertha Suresh. On Learning Distributions from their Samples. In *Proceedings of The 28th Conference on Learning Theory*, pages 1066–1100. PMLR, June 2015. URL: <http://proceedings.mlr.press/v40/Kamath15.html>.
- 32 Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

## 25:24 Locally Private Histograms in All Privacy Regimes

- 33 Shaowei Wang, Liusheng Huang, Pengzhan Wang, Yiwen Nie, Hongli Xu, Wei Yang, Xiang-Yang Li, and Chunming Qiao. Mutual information optimally local private discrete distribution estimation. *ArXiv*, abs/1607.08025, 2016. [arXiv:1607.08025](https://arxiv.org/abs/1607.08025).
- 34 Min Ye and Alexander Barg. Optimal schemes for discrete distribution estimation under locally differential privacy. *Institute of Electrical and Electronics Engineers*, 64(8):5662–5676, 2018. [doi:10.1109/TIT.2018.2809790](https://doi.org/10.1109/TIT.2018.2809790).