

583 A Distributed Frequency Estimation

584 In this section, we consider the frequency estimation problem for federated analytics. Recall that
 585 for the frequency estimation task, each client's private data $x_i \in \{0, 1\}^d$ satisfies $\|x_i\|_0 = 1$, and the
 586 goal is to estimate $\pi := \frac{1}{n} \sum_i x_i$ by minimizing the ℓ_2 (or ℓ_1, ℓ_∞) error $\mathbb{E} \left[\|\pi - \hat{\pi}(Y^n)\|_2^2 \right]$ subject
 587 to communication and (ε, δ) -DP constraints. When the context is clear, we sometimes use x_i to
 588 denote, by abuse of notation, the index of the item, i.e., $x_i \in [d]$.

589 To fully make use of the ℓ_0 structure of the problem, a standard technique is applying a Hadamard
 590 transform to convert the ℓ_0 geometry to an ℓ_∞ one and then leveraging the recursive structure of
 591 Hadamard matrices to efficiently compress local messages.

Specifically, for a given b -bit constraint, we partition each local item x_i into 2^{b-1} chunks
 $x_i^{(1)}, \dots, x_i^{(2^{b-1})} \in \{0, 1\}^B$, where $B := d/2^{b-1}$ and $x_i^{(j)} = x_i[B \cdot (j-1) : B \cdot j - 1]$. Note
 that since x_i is one-hot, only one chunk of $x_i^{(j)}$ is non-zero. Then, client i performs the following
 Hadamard transform for each chunk: $y_i^{(\ell)} = H_B \cdot x_i^{(\ell)}$, where H_B is defined recursively as follows:

$$H_{2^n} = \frac{1}{\sqrt{2}} \begin{bmatrix} H_{2^{n-1}} & H_{2^{n-1}} \\ H_{2^{n-1}} & -H_{2^{n-1}} \end{bmatrix}, \text{ and } H_0 = [1].$$

592 Each client then generates a sampling vector $Z_{ij} \stackrel{\text{i.i.d.}}{\sim} \text{Bern}\left(\frac{1}{B}\right)$ via shared randomness that
 593 is also known by the server, and commits $(y_i^{(1)}(j), \dots, y_i^{(2^{b-1})}(j))$ as its local report. Since
 594 $(y_i^{(1)}(j), \dots, y_i^{(2^{b-1})}(j))$ only contains a single non-zero entry that can be $\frac{1}{\sqrt{B}}$ or $-\frac{1}{\sqrt{B}}$, the local
 595 report can be represented in b bits ($b-1$ bits for the location of the non-zero entry and 1 bit for its
 596 sign).

597 From the local reports, the server can compute an unbiased estimator by summing them together
 598 (with proper normalization) and performing an inverse Hadamard transform. Moreover, with an
 599 adequate injection of Gaussian noise, the frequency estimator satisfies (ε, δ) -DP.

600 The idea has been used in previous literature under local DP [19, 6, 3, 32], but in order to obtain the
 601 order-optimal trade-off under *central*-DP, one has to combine Hadamard transform with a random
 602 subsampling step and incorporate the privacy amplification due to random compression in the
 603 analysis. In Algorithm 3, we provide a summary of the resultant scheme which builds on the
 604 Recursive Hadamard Response (RHR) mechanism from [32], which was originally designed for
 605 communication-efficient frequency estimation under *local* DP.

606 In the following theorem, we control the ℓ_∞ error of Algorithm 3.

607 **Theorem A.1.** *Let $\hat{\pi}(x^n)$ be the output of Algorithm 3. Then it holds that for all $j \in [d]$,*

$$\mathbb{E} [|\pi(j) - \hat{\pi}(j)|] \leq \sqrt{\frac{\sum_i \mathbb{1}_{\{x_i \in [B \cdot (j-1) : B \cdot j - 1]\}}}{n^2}} + \frac{\sigma^2}{B}, \quad (7)$$

608 and the ℓ_2^2 and ℓ_1 errors are bounded by

$$\mathbb{E} \left[\|\pi - \hat{\pi}\|_2^2 \right] \leq \frac{B}{n} + \frac{d\sigma^2}{B}, \text{ and} \quad (8)$$

609

$$\mathbb{E} [\|\pi - \hat{\pi}\|_1] \leq \sqrt{\frac{dB}{n} + \frac{d^2\sigma^2}{B}}. \quad (9)$$

Theorem A.2. *For any $\varepsilon, \delta > 0$, Algorithm 3 is (ε, δ) -DP, if*

$$\sigma^2 \geq O \left(\frac{B^2 \log(B/\delta)}{n^2} + \frac{B(\log(1/\delta) + \varepsilon) \log(B/\delta)}{n^2 \varepsilon^2} \right).$$

610 By combining Theorem A.1 and Theorem A.2, we conclude that Algorithm 3 achieves (ε, δ) -DP
 611 with ℓ_2^2 error

$$\begin{aligned} & O \left(\frac{B}{n} + \frac{dB \log(B/\delta)}{n^2} + \frac{d(\log(1/\delta) + \varepsilon) \log(B/\delta)}{n^2 \varepsilon^2} \right) \\ & = O \left(\frac{d}{n2^b} + \frac{d^2 \log(d/\delta)}{n^2 2^b} + \frac{d(\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2} \right). \end{aligned}$$

Algorithm 3 Subsampled Recursive Hadamard Response

Input: user data $x_1, \dots, x_n \in \{0, 1\}^d$ (where d is a power of two), DP parameters (ε, δ) , communication budget b .

Output: frequency estimate $\hat{\pi}$

Set $B := d/2^{b-1}$ and partition each one-hot vector x_i into 2^{b-1} chunks: $x_i^{(1)}, \dots, x_i^{(2^{b-1})} \in \{0, 1\}^B$.

for user $i \in [n]$ **do**

 Compute the Hadamard transform of each chunk: $y_i^{(\ell)} = H_B \cdot x_i^{(\ell)}$.

for coordinate $j \in [B]$ **do**

 Draw $Z_{i,j} \stackrel{\text{i.i.d.}}{\sim} \text{Bern}\left(\frac{1}{B}\right)$

if $Z_{i,j} = 1$ **then**

 Send $(y_i^{(1)}(j), \dots, y_i^{(2^{b-1})}(j))$ to the server.

end if

end for

end for

Server computes the average: $\forall \ell \in [2^{b-1}], j \in [B]$,

$$\hat{y}^{(\ell)}(j) := \frac{B}{n} \sum_{i: Z_{i,j}=1} y_i^{(\ell)}(j) + N(0, \sigma^2),$$

where σ^2 is computed according to Theorem [A.2](#)

Server performs the inverse Hadamard transform $\hat{\pi}^{(\ell)} = H_B \cdot \hat{y}^{(\ell)}$, for $\ell = 1, \dots, B$.

Return: $\hat{\pi} = \left((\hat{\pi}^{(1)})^\top, \dots, (\hat{\pi}^{(2^{b-1})})^\top \right)$.

Notice that when $n = \tilde{\Omega}(d)$, the error can be simplified to

$$O\left(\frac{d}{n2^b} + \frac{d(\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right),$$

612 which matches the order-optimal estimation error (up to a log d factor) subject to a b -bit constraint
613 [\[54, 3, 2\]](#) and (ε, δ) -DP constraint [\[15, 7\]](#).

614 **B Proof of Theorem 4.1**

615 It is trivial to see that the average communication cost is $d \cdot \gamma = b$ bits. To compute the ℓ_2^2 estimation
616 error, observe that

$$\begin{aligned}
& \mathbb{E} \left[\|\hat{\mu}_{x^n} - \mu_{x^n}\|_2^2 \right] \\
&= \sum_{j=1}^d \mathbb{E} \left[\left(\frac{1}{n\gamma} \sum_i x_i(j) \cdot Z_{i,j} + N(0, \sigma^2) - \frac{1}{n} \sum_i x_i(j) \right)^2 \right] \\
&= \sum_{j=1}^d \frac{1}{n^2} \mathbb{E} \left[\left(\frac{1}{\gamma} \sum_i x_i(j) \cdot Z_{i,j} - \sum_i x_i(j) \right)^2 \right] + d\sigma^2 \\
&= \sum_{j=1}^d \frac{1}{n^2} \mathbb{E} \left[\left(\frac{1}{\gamma} \sum_i x_i(j) \cdot Z_{i,j} \right)^2 \right] - \frac{1}{n^2} \left(\sum_i x_i(j) \right)^2 + d\sigma^2 \\
&= \sum_{j=1}^d \frac{1}{n^2} \mathbb{E} \left[\frac{1}{\gamma^2} \sum_i x_i^2(j) \cdot Z_{i,j}^2 + \frac{1}{\gamma^2} \sum_{i \neq i'} x_i(j)x_{i'}(j)Z_{i,j}Z_{i',j} \right] - \frac{1}{n^2} \left(\sum_i x_i(j) \right)^2 + d\sigma^2 \\
&= \sum_{j=1}^d \frac{1}{n^2} \left(\frac{1}{\gamma} \sum_i x_i^2(j) + \sum_{i \neq i'} x_i(j)x_{i'}(j) \right) - \frac{1}{n^2} \left(\sum_i x_i(j) \right)^2 + d\sigma^2 \\
&= \sum_{j=1}^d \frac{1}{n^2} \left(\frac{1}{\gamma} - 1 \right) \left(\sum_i x_i^2(j) \right) + d\sigma^2 \\
&\leq \frac{dc^2}{n\gamma} + d\sigma^2,
\end{aligned}$$

617 which yields the inequality of (2). Next, we analyze the privacy of Algorithm 1. We first the following
618 two lemmas for subsampling and the Gaussian mechanism:

619 **Lemma B.1** ([65, 81]). *If \mathcal{M} is (ϵ, δ) -DP, then \mathcal{M}' that applies $\mathcal{M} \circ \text{PoissonSample}$ satisfies
620 (ϵ', δ') -DP with $\epsilon' = \log(1 + \gamma(e^\epsilon - 1))$ and $\delta' = \gamma\delta$.*

621 **Lemma B.2** ([15]). *For any $\epsilon, \delta \in (0, 1)$, the Gaussian output perturbation mechanism with
622 $\sigma^2 := \frac{\Delta^2 2 \log(1.25/\delta)}{\epsilon^2}$ satisfies (ϵ, δ) -DP, where Δ is the ℓ_2 sensitivity of the target function.*

Now, we use the above two lemmas to analyze the per-coordinate privacy leakage of Algorithm 1. For
simplicity, we analyze the sum of $x_i(j)$'s instead (and normalized it in the last step). Let $S_j(x^n) :=$
 $\sum_{i=1}^n x_i(j)$, then clearly the sensitivity of $S_j(x^n)$ is c , so Lemma B.2 implies $S_j(x^n) + N(0, \sigma_1^2)$
satisfies (ϵ_1, δ_1) -DP if we set $\sigma_1^2 = \frac{2c^2 \log(1.25/\delta_1)}{\epsilon_1^2}$ (assuming $\epsilon_1 < 1$). Next, if applying subsampling
before computing the sum, i.e.,

$$S_j \circ \text{PoissonSample}_\gamma(x^n) := \sum_{i=1}^n x_i(j)Z_{i,j},$$

where $Z_{i,j} \stackrel{\text{i.i.d.}}{\sim} \text{Bern}(1/\gamma)$ as defined in Algorithm 1, then by Lemma B.1

$$S_j \circ \text{PoissonSample}_\gamma(x^n) + N(0, \sigma_1^2)$$

623 satisfies (ϵ_2, δ_2) -DP with $\epsilon_2 := \log(1 + \gamma(e^{\epsilon_1} - 1)) = C_1\gamma\epsilon_1$ (since we assume $\epsilon_1 < 1$) and
624 $\delta_2 := \gamma\delta_1$. Equivalently, we have

$$\begin{cases} \epsilon_1 = \tilde{C}_1 \frac{1}{\gamma} \epsilon_2 \\ \delta_1 = \frac{1}{\gamma} \delta_2. \end{cases} \quad (10)$$

625 Now, since we have established the per-coordinate privacy leakage, we apply the following composi-
626 tion theorem to account for the total privacy budgets.

Theorem B.3. For any $\varepsilon > 0$, $\delta \in [0, 1]$ and $\tilde{\delta} \in (0, 1]$, the class of (ε, δ) -DP mechanisms satisfies $(\tilde{\varepsilon}_{\tilde{\delta}}, d\delta + \tilde{\delta})$ -DP under d -fold adaptive composition, for

$$\tilde{\varepsilon}_{\tilde{\delta}} = d\varepsilon(e^\varepsilon - 1) + \varepsilon\sqrt{2d\log(1/\tilde{\delta})}.$$

627 According to Theorem B.3, Algorithm 1 satisfies (ε, δ) -DP for

$$\varepsilon = d\varepsilon_2(e^{\varepsilon_2} - 1) + \varepsilon_2\sqrt{2d\log(1/\tilde{\delta})}, \quad (11)$$

628 and $\delta = d\delta_2 + \tilde{\delta}$ (where $\tilde{\delta}$ is a free parameter that we can optimize).

Consequently, for a pre-specified (total) privacy budget (ε, δ) , we set parameters as follows. Let $\tilde{\delta} = \frac{\delta}{2}$ and $\delta_1 = \frac{1}{\gamma}\delta_2 = \frac{1}{2d\gamma}\delta$. Let $\varepsilon_2 \leq 1$ so that $e^{\varepsilon_2} - 1 \leq 2\varepsilon_2$ holds. Then (11) implies Algorithm 1 is

$$\varepsilon = 2d\varepsilon_2^2 + \varepsilon_2\sqrt{2d\log(1/\tilde{\delta})} \geq d\varepsilon_2(e^{\varepsilon_2} - 1) + \varepsilon_2\sqrt{2d\log(1/\tilde{\delta})}.$$

Solving the above quadratic (in-)equality for ε_2 , it yields that

$$\varepsilon_2 = \min\left(1, \frac{-\sqrt{2d\log(2/\tilde{\delta})} + \sqrt{2d\log(2/\tilde{\delta}) + 8\varepsilon d}}{4d}\right) = O\left(\min\left(1, \frac{\varepsilon}{\sqrt{d(\log(1/\tilde{\delta}) + \varepsilon)}}\right)\right).$$

629 Consequently, we set $\varepsilon_1 = \frac{\tilde{c}_1}{\gamma}\varepsilon_2 = O\left(\min\left(1, \frac{\varepsilon}{\gamma\sqrt{d(\log(1/\tilde{\delta}) + \varepsilon)}}\right)\right)$ (note that we require $\varepsilon_1 = O(1)$)
630 so that (10) holds).

Plug in $(\varepsilon_1, \delta_1)$ into σ_1^2 , we have

$$\sigma_1^2 := \frac{2c^2\log(1.25/\delta_1)}{\varepsilon_1^2} = \Omega\left(\max\left(c^2\log(d/\delta), \frac{\gamma^2c^2d(\log(1/\tilde{\delta}) + \varepsilon)\log(d/\delta)}{\varepsilon^2}\right)\right).$$

Finally, as we are interested in estimating the (subsampled) mean instead of the sum, we will normalize the private sum by

$$\hat{\mu}_j(x^n) = \frac{1}{n\gamma} (S_j \circ \text{PoissonSample}_\gamma(x^n) + N(0, \sigma_1^2)) = \frac{1}{n\gamma} S_j \circ \text{PoissonSample}_\gamma(x^n) + N(0, \sigma^2),$$

where

$$\sigma^2 = O\left(\max\left(\frac{c^2\log(d/\delta)}{n^2\gamma^2}, \frac{c^2d(\log(1/\tilde{\delta}) + \varepsilon)\log(d/\delta)}{n^2\varepsilon^2}\right)\right).$$

631 Plugging in σ^2 above and $\gamma = b/d$ yields the desired accuracy in Theorem 4.1 □

632 Since we will reuse the above result, we summarize it into the following lemma:

Lemma B.4. Let $f_i : \mathbb{R}^{d \times m} \mapsto \mathbb{R}^D$ for $i = 1, \dots, B$ be n functions with sensitivity bounded by Δ (where the number of inputs m can be a random variable). Then

$$(f_1 \circ \text{PoissonSample}_\gamma(x^n) + N(0, \sigma^2), \dots, f_B \circ \text{PoissonSample}_\gamma(x^n) + N(0, \sigma^2))$$

satisfies (ε, δ) -DP, if

$$\sigma^2 \geq O\left(\max\left(\Delta^2\log(B/\delta), \frac{\gamma^2\Delta^2B(\log(1/\tilde{\delta}) + \varepsilon)\log(B/\delta)}{\varepsilon^2}\right)\right).$$

633 C Omitted details of dimension-free communication cost

634 C.1 Proof of Theorem 4.4

635 To prove Theorem 4.4, it suffices to prove the following ℓ_∞ version:

636 **Theorem C.1.** Let $x_1, \dots, x_n \in \{-c, c\}^d$, $d' = \min\left(nb, \frac{n^2 \varepsilon^2}{(\log(1/\delta) + \varepsilon) \log(d/\delta)}\right)$, and

$$\sigma^2 = O\left(\frac{c^2 \log(1/\delta)}{n^2 \gamma^2} + \frac{c^2 d' (\log(d'/\delta) + \varepsilon) \log(d'/\delta)}{n^2 \varepsilon^2}\right). \quad (12)$$

637 Then Algorithm 2 is (ε, δ) -DP and yields an unbiased estimator on μ . In addition, the (average)
638 per-client communication cost is $\gamma d' = b$ bits, and the ℓ_2^2 estimation error is at most

$$O\left(c^2 d^2 \log\left(\frac{d}{\delta}\right) \max\left(\frac{1}{nb}, \frac{(\log(1/\delta) + \varepsilon)}{n^2 \varepsilon^2}\right)\right). \quad (13)$$

With a slight abuse of notation, we let $\mu_{\mathcal{J}} \in \mathbb{R}^d$ be such that

$$\mu_{\mathcal{J}}(j) = \begin{cases} 0, & \text{if } j \notin \mathcal{J} \\ \frac{d\mu_j}{d'}, & \text{else.} \end{cases}$$

639 Note that $\mu_{\mathcal{J}}$ is an unbiased estimate of μ if \mathcal{J} is selected uniformly at random. Then the ℓ_2^2 error
640 can be controlled by

$$\begin{aligned} \mathbb{E}\left[\|\mu - \hat{\mu}\|_2^2\right] &\stackrel{(a)}{=} \mathbb{E}\left[\|\mu - \mu_{\mathcal{J}}\|_2^2\right] + \mathbb{E}\left[\|\mu_{\mathcal{J}} - \hat{\mu}\|_2^2\right] \\ &\stackrel{(b)}{\leq} \mathbb{E}\left[\|\mu - \mu_{\mathcal{J}}\|_2^2\right] + \frac{d^2}{d'^2} O\left(\max\left(\frac{d^2 c^2}{nb}, \frac{d'^3 c^2 \log(d/\delta)}{n^2 b^2}, \frac{c^2 d'^2 (\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right)\right) \\ &= \mathbb{E}\left[\|\mu - \mu_{\mathcal{J}}\|_2^2\right] + O\left(\max\left(\frac{d^2 c^2}{nb}, \frac{d^2 d' c^2 \log(d/\delta)}{n^2 b^2}, \frac{c^2 d^2 (\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right)\right) \\ &\stackrel{(c)}{\leq} \frac{d^2 c^2}{d'} + O\left(\max\left(\frac{d^2 c^2}{nb}, \frac{d^2 d' c^2 \log(d/\delta)}{n^2 b^2}, \frac{c^2 d^2 (\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right)\right), \end{aligned}$$

641 where (a) holds since $\mu_{\mathcal{J}}$ is an unbiased estimate of μ and conditioned on \mathcal{J} , $\hat{\mu}$ is an unbiased
642 estimate of $\mu_{\mathcal{J}}$; (b) follows from Theorem 4.1; (c) holds due to the following fact:

$$\mathbb{E}\left[\|\mu - \mu_{\mathcal{J}}\|_2^2\right] \leq \sum_{j \in \mathcal{J}} \mu_{\mathcal{J}}(j)^2 + \sum_{j \in [d]} \mu_j^2 \leq \frac{d^2 c^2}{d'} + dc^2 \leq \frac{2d^2 c^2}{d'}.$$

643 Therefore, by setting $d' = \min\left(nb, \frac{n^2 \varepsilon^2}{(\log(1/\delta) + \varepsilon) \log(d/\delta)}\right)$ we ensure the first term in (c) is always
644 smaller than the second term, and the second term can be simplified as follows:

$$\begin{aligned} &O\left(c^2 d^2 \max\left(\frac{1}{nb}, \frac{d' \log(d/\delta)}{n^2 b^2}, \frac{(\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right)\right) \\ &\leq O\left(c^2 d^2 \max\left(\frac{1}{nb}, \frac{nb \log(d/\delta)}{n^2 b^2}, \frac{(\log(1/\delta) + \varepsilon) \log(d/\delta)}{n^2 \varepsilon^2}\right)\right) \\ &\leq O\left(c^2 d^2 \log(d/\delta) \max\left(\frac{1}{nb}, \frac{(\log(1/\delta) + \varepsilon)}{n^2 \varepsilon^2}\right)\right). \end{aligned}$$

645 Finally, applying the same trick of Kashin's representation, we can transform the ℓ_∞ geometry to ℓ_2
646 (similar to Corollary 4.3), hence proving Theorem 4.4. \square

647 D Proof of Theorem A.1

648 Let $\pi := \frac{1}{n} \sum_i x_i$ and $\pi^{(\ell)}$ be defined in the same way as $x_i^{(\ell)}$ for $\ell \in [B]$. Then our goal is to bound
649 $|\pi^{(\ell)}(j) - \hat{\pi}^{(\ell)}(j)|$, for all $\ell \in [2^{b-1}]$ and $j \in [B]$.

650 To this end, let $y^{(\ell)} := H_B \cdot \pi^{(\ell)}$ (so it holds that $\pi^{(\ell)} = \frac{1}{B} H_B \cdot y^{(\ell)}$). Then we have

$$\begin{aligned} \mathbb{E}\left[|\pi^{(\ell)}(j) - \hat{\pi}^{(\ell)}(j)|\right] &\stackrel{(a)}{\leq} \sqrt{\mathbb{E}\left[\left(\pi^{(\ell)}(j) - \hat{\pi}^{(\ell)}(j)\right)^2\right]} \\ &= \sqrt{\mathbb{E}\left[\left(\frac{1}{B} H_B \cdot (y^{(\ell)} - \hat{y}^{(\ell)})(j)\right)^2\right]}. \end{aligned} \quad (14)$$

Next, observe that due to the subsampling step, for all $\ell \in [2^{b-1}]$ and $j \in [B]$,

$$\hat{y}^{(\ell)}(j) = \frac{B}{n} \sum_{i=1}^n \langle (H_B)_j, x_i^{(\ell)} \rangle \cdot Z_{ij} + N(0, \sigma^2),$$

651 where recall that $Z_{ij} \stackrel{\text{i.i.d.}}{\sim} \text{Ber}(1/B)$. Therefore, $\hat{y}^{(\ell)}(j)$ is an unbiased estimator of $y^{(\ell)}(j)$. In
652 addition, since we choose Z_{ij} independently in Algorithm 3, $\hat{y}^{(\ell)}(j)$'s are independent for different
653 j 's, so we have

$$\begin{aligned} \mathbb{E} \left[\left(\hat{y}^{(\ell)}(j) - y^{(\ell)}(j) \right)^2 \right] &= \text{Var} \left(\hat{y}^{(\ell)}(j) \right) \\ &= \sigma^2 + \frac{B^2}{n^2} \sum_{i=1}^n \langle (H_B)_j, x_i^{(\ell)} \rangle^2 \text{Var} (Z_{ij}) \\ &\leq \sigma^2 + \frac{B}{n^2} \sum_{i=1}^n \langle (H_B)_j, x_i^{(\ell)} \rangle^2 \\ &= \sigma^2 + \frac{B}{n^2} \underbrace{\sum_{i=1}^n \mathbb{1}_{\{x_i \in \ell\text{-th chunk}\}}}_{:= C_\ell}, \end{aligned} \quad (15)$$

654 and for all $j \neq j'$

$$\mathbb{E} \left[\left(\hat{y}^{(\ell)}(j) - y^{(\ell)}(j) \right) \cdot \left(\hat{y}^{(\ell)}(j') - y^{(\ell)}(j') \right) \right] = 0. \quad (16)$$

655 Therefore, we continue bounding (14) as follows:

$$\begin{aligned} \sqrt{\mathbb{E} \left[\left(\frac{1}{B} H_B \cdot (y^{(\ell)} - \hat{y}^{(\ell)})(j) \right)^2 \right]} &= \sqrt{\frac{1}{B^2} \mathbb{E} [\langle (H_B)_j, (y^{(\ell)} - \hat{y}^{(\ell)}) \rangle^2]} \\ &= \sqrt{\frac{1}{B^2} \mathbb{E} \left[\left(\sum_{k=1}^B (H_B)_{jk} \cdot (y^{(\ell)}(k) - \hat{y}^{(\ell)}(k)) \right)^2 \right]} \\ &\stackrel{(a)}{=} \sqrt{\frac{1}{B^2} \mathbb{E} \left[\sum_{k=1}^B (y^{(\ell)}(k) - \hat{y}^{(\ell)}(k))^2 \right]} \\ &\stackrel{(b)}{=} \sqrt{\frac{C_\ell}{n^2} + \frac{\sigma^2}{B}} \\ &\stackrel{(c)}{\leq} \sqrt{\frac{1}{n} + \frac{\sigma^2}{B}}, \end{aligned}$$

656 where (a) holds since each entry of H_B takes value in $\{-1, 1\}$ and by (16), (b) holds due to (15), and
657 (c) holds because $C_\ell \leq n$ for all ℓ .

Finally, to bound the ℓ_2^2 error, observe that the above analysis ensures that

$$\mathbb{E} \left[\left(\pi^{(\ell)}(j) - \hat{\pi}^{(\ell)}(j) \right)^2 \right] \leq \frac{C_{\ell(j)}}{n^2} + \frac{\sigma^2}{B},$$

where $\ell(j) \in [2^{b-1}]$ is the index of the chunk containing j . Therefore, summing over $j \in [d]$, we must have

$$\mathbb{E} \left[\left\| \pi^{(\ell)} - \hat{\pi}^{(\ell)} \right\|_2^2 \right] \leq \sum_{j=1}^d \frac{C_{\ell(j)}}{n^2} + \frac{d\sigma^2}{B} = \frac{B}{n} + \frac{d\sigma^2}{B},$$

since

$$\sum_j C_{\ell(j)} = \sum_{\ell=1}^{2^{b-1}} \sum_{j' \in \ell\text{-th chunk}} \sum_{i=1}^n \mathbb{1}_{\{i \in \ell\text{-th chunk}\}} = B \sum_{\ell=1}^{2^{b-1}} \sum_{i=1}^n \mathbb{1}_{\{i \in \ell\text{-th chunk}\}} = B \cdot n.$$

658 □

659 **E Proof of Theorem [A.2](#)**

660 Let $f_j(x^n) := (\pi^{(1)}(j), \dots, \pi^{(2^{b-1})}(j))$, for $j = 1, \dots, B$. Then the ℓ_2 sensitivity of f_j is $\Delta = \frac{B}{n}$. Set
 661 the sampling rate $\gamma = \frac{1}{B}$ and the proof is complete by Lemma [B.4](#). \square

662 **F Algorithm of Shuffled SQKR**

Algorithm 4 Shuffled SQKR

Input: users' data x_1, \dots, x_n , local-DP parameter ε_0 , communication parameters b_0, T
Output: mean estimator $\hat{\mu}$
for round $k \in [T]$ **do**
 for user $i \in [n]$ **do**
 Sample $s(i, 1), \dots, s(i, b_0) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}[d]$
 Sample $Z \sim \text{Bern}\left(\frac{e^{\varepsilon_0}}{e^{\varepsilon_0} + 2^{b_0} - 1}\right)$
 if $Z=1$ **then**
 Set $Y(i, 1), \dots, Y(i, b_0) \leftarrow x_i(s(i, 1)), \dots, x_i(s(i, b_0))$
 else
 Sample $Y(i, 1), \dots, Y(i, b_0) \stackrel{\text{i.i.d.}}{\sim} \text{Unif}\{-c, c\}$
 end if
 Send $Y(i, 1), \dots, Y(i, b_0)$ and $s(i, 1), \dots, s(i, b_0)$ to shuffler
 end for
 Shuffler samples a permutation $\pi \sim \text{Unif}\{f : [n] \rightarrow [n] \text{ bijective}\}$
 for $j \in [b_0]$ **do**
 Shuffler sends $Y(\pi(1), j), \dots, Y(\pi(n), j)$ and $s(\pi(1), j), \dots, s(\pi(n), j)$ to server
 end for
 $\hat{\mu}^{(k)} \leftarrow \frac{d}{nb_0} \frac{e^{\varepsilon_0} + 2^{b_0} - 1}{e^{\varepsilon_0} - 1} \sum_{i=1}^n \sum_{j=1}^{b_0} Y(\pi(i), j) e_{s(\pi(i), j)}$
end for
 Return $\hat{\mu} := \frac{1}{T} \sum_{k=1}^T \hat{\mu}^{(k)}$

663 **G Proof of Theorem [5.3](#)**

664 Each round $x^n \mapsto \hat{\mu}^{(k)}$ of Algorithm [4](#) implements the private-coin SQKR scheme of [\[32\]](#), achieving
 665 the communication cost and error as stated in Lemma [5.2](#)

666 **Lemma G.1** (SQKR [\[32\]](#)). *For all $\varepsilon_0 > 0, b_0 > 0$, the random mapping*
 667 $x_i \mapsto Y(i, 1), \dots, Y(i, b_0), s(i, 1), \dots, s(i, b_0)$ *in Algorithm [4](#) is $(\varepsilon_0, 0)$ -LDP and has*
 668 *output that can be communicated in $b_0 \log(d)$ bits, and the $\hat{\mu}^{(k)}$ computed from*
 669 $Y(i, 1), \dots, Y(i, b_0), s(i, 1), \dots, s(i, b_0)$ *is an unbiased estimator of μ satisfying*

$$\max_{x^n} \mathbb{E} \left[\left\| \mu(x^n) - \hat{\mu}^{(k)}(x^n) \right\|_2^2 \right] = O \left(\frac{c^2 d}{n \min(\varepsilon_0^2, \varepsilon_0, b_0)} \right). \quad (17)$$

670 We now characterize the error performance of Algorithm [4](#) for general choices of parameters that
 671 satisfy privacy and communication constraints.

672 **Proposition G.2.** *For all $\varepsilon > 0, b > 0, n > 0$, with any arbitrary choice of*

$$\delta_1 \in \left(e^{-n/16e}, 1 \right] \quad (18)$$

$$\delta_2 \in (0, 1], \quad (19)$$

673 *there exists a choice of parameters ε_0, b_0, T such that Algorithm [4](#) is $(\varepsilon, T\delta_1 + \delta_2)$ -DP, uses no more*
 674 *than b bits of communication, and produces $\hat{\mu}$ such that*

$$\max_{x^n} \mathbb{E} \left[\left\| \mu - \hat{\mu} \right\|_2^2 \right] = O \left(\max \left(\frac{c^2 d \log(d) b_0}{nb}, \frac{c^2 d \log(1/\delta_1) (\log(1/\delta_2) + \varepsilon)}{n^2 \varepsilon^2} \right) \right). \quad (20)$$

675 *Proof.* For arbitrary choice of

$$b_0 < \log\left(\frac{n}{16 \log(2)}\right), \quad (21)$$

676 it suffices to choose

$$T = \left\lfloor \frac{b}{(\log_2(d) + 1)b_0} \right\rfloor \quad (22)$$

$$\varepsilon_0 = O\left(\min\left(1, \frac{\varepsilon\sqrt{n}}{\sqrt{T \log(1/\delta_1) (\log(1/\delta_2) + \varepsilon)}}\right)\right). \quad (23)$$

677 The fact that Algorithm 4 uses less than b bits is immediate from the choice of T .

678 Applying Lemma G.1 by construction the mapping from each x_i to $Y(i, 1), \dots, Y(i, b_0)$ is $(\varepsilon_0, 0)$ -
679 LDP. By assumption

$$\delta_1 > e^{-n/16e}, \quad (24)$$

680 the inequality

$$1 < \log\left(\frac{n}{16 \log(2/\delta_1)}\right) \quad (25)$$

681 is satisfied. Then the choice of

$$\varepsilon_0 \leq 1 \quad (26)$$

682 also satisfies $\varepsilon_0 \leq \log\left(\frac{n}{16 \log(2/\delta_1)}\right)$, so by Lemma 5.1 the mapping $x^n \mapsto \hat{\mu}^{(k)}$ is $(\varepsilon_1, \delta_1)$ -DP. where

$$\varepsilon_1 = O\left(\frac{\varepsilon_0 \sqrt{\log(1/\delta_1)}}{\sqrt{n}}\right). \quad (27)$$

683 Since the output of Algorithm 4 is a function of $(\hat{\mu}^{(1)}, \dots, \hat{\mu}^{(T)})$, by B.3 it suffices to have

$$\varepsilon_1 = O\left(\min\left(1, \frac{\varepsilon}{\sqrt{T(\log(1/\delta_2) + \varepsilon)}}\right)\right) \quad (28)$$

684 for Algorithm 4 to be $(\varepsilon, T\delta_1 + \delta_2)$ -DP. The first inequality follows from the assumption of $\delta_1 >$
685 $e^{-n/16e}$ and choice of $\varepsilon_0 = O(1)$, and the second from choice of

$$\varepsilon_0 = O\left(\frac{\varepsilon\sqrt{n}}{\sqrt{T \log(1/\delta_1) (\log(1/\delta_2) + \varepsilon)}}\right). \quad (29)$$

686 Since $\varepsilon_0 \leq 1 \leq b$, we have $\min(\varepsilon_0^2, \varepsilon_0, b) = \varepsilon_0^2$. Applying Lemma G.1

$$\max_{x^n} \mathbb{E} \left[\|\mu - \hat{\mu}\|_2^2 \right] = \frac{1}{T} \max_{x^n} \mathbb{E} \left[\left\| \mu - \hat{\mu}^{(1)} \right\|_2^2 \right] \quad (30)$$

$$= O\left(\frac{d}{Tn\varepsilon_0^2}\right) \quad (31)$$

$$= O\left(\max\left(\frac{d}{Tn}, \frac{d \log(1/\delta_1) (\log(1/\delta_2) + \varepsilon)}{n^2\varepsilon^2}\right)\right). \quad (32)$$

687 Substituting the choice of T gives the desired result. \square

688 To show Theorem 5.3, it suffices to choose

$$b_0 = 1 \quad (33)$$

$$\delta_1 = \frac{\delta}{2T} \quad (34)$$

$$\delta_2 = \frac{\delta}{2}, \quad (35)$$

689 which requires $n > 16e \log(2) \approx 30.14$ due to (21), and apply the previous proposition.

690 H Rényi-DP for Shuffled SQKR

691 In this section we restate some results for RDP which are useful for privacy accounting in experiments.
 692 Following the proof of Corollary 4.3 in [44], applying Theorem 4.1 in the same paper yields the
 693 following.

694 **Lemma H.1.** *Let \mathcal{M}_i be an independent $(\varepsilon_0, 0)$ -LDP mechanism for each $i \in [n]$ with $\varepsilon_0 \leq 1$ and π
 695 be a random permutation of $[n]$. Then for any $\alpha < \frac{n}{16\varepsilon_0 \exp(\varepsilon_0)}$, the mechanism*

$$\mathcal{S} : (x_1, \dots, x_n) \mapsto (\mathcal{M}_1(x_{\pi(1)}), \dots, \mathcal{M}_n(x_{\pi(n)}))$$

696 is $(\varepsilon_1(\alpha), \delta)$ -RDP with

$$\varepsilon_1(\alpha) = \frac{\log\left(e^{2\alpha^2\sigma^2} + 4\delta_{\min}e^{\alpha\varepsilon_0}\right)}{\alpha - 1}, \quad (36)$$

697 where

$$\sigma = 8\sqrt{\frac{e^{\varepsilon_0}}{n}} \quad (37)$$

$$\delta_{\min} = e^{-\frac{n}{8(e^{\varepsilon_0}+1)}}. \quad (38)$$

698 For small ε_0 , the result below is useful.

699 **Lemma H.2** ([40]). *Under the same assumptions as Lemma [H.1] \mathcal{S} is $(\varepsilon(\alpha), \delta)$ -RDP*

$$\varepsilon_1(\alpha) = 2\alpha e^{4\varepsilon_0} (e^{\varepsilon_0} - 1)^2 / n. \quad (39)$$

700 Applying Lemma [G.1] by construction the mapping from each x_i to $y(i, 1), \dots, y(i, b_0)$ is $(\varepsilon_0, 0)$ -
 701 LDP. By Lemma [H.1] respectively Lemma [H.2] the mapping $x^n \mapsto \hat{\mu}^{(k)}$ is $(\varepsilon_1(\alpha), \alpha)$ -RDP where
 702 $\varepsilon_1(\alpha)$ is given by [36], respectively [39]. By composition, Algorithm [4] is $(T\varepsilon(\alpha), \alpha)$ -RDP.

703 We can convert this bound back to (ε, δ) -DP using Proposition 12 from [30].

704 **Proposition H.3.** *For all $\delta > 0$, Algorithm [4] is (ε, δ) -DP where*

$$\varepsilon = \inf_{\alpha \in (1, \infty)} T\varepsilon_1(\alpha) + \frac{\log(1/\delta) + (\alpha - 1)\log(1 - 1/\alpha) - \log(\alpha)}{\alpha - 1}, \quad (40)$$

705 where

$$\varepsilon_1(\alpha) = \min\left(2\alpha e^{4\varepsilon_0} (e^{\varepsilon_0} - 1)^2 / n, \frac{\log\left(e^{2\alpha^2\sigma^2} + 4\delta_{\min}e^{\alpha\varepsilon_0}\right)}{\alpha - 1}\right) \quad (41)$$

706 and σ, δ_{\min} are given by [37], [38] respectively.

707 I Additional Experiments

708 Here experiments are done with the same setup as in Section [6] with local vectors $X_i(j) \stackrel{\text{i.i.d.}}{\sim}$
 709 $\frac{1}{\sqrt{d}}(2 \cdot \text{Ber}(0.8) - 1)$. We set $\delta = 10^{-6}$.

710 Figure [2] illustrates separation between Algorithm [4] and LDP schemes. Algorithm [4] achieves error
 711 decreasing quadratically with n as guaranteed by Theorem [5.3]. With only one round of shuffling,
 712 there is separation from the LDP scheme only when n is sufficiently large, and thus order-optimal
 713 error performance only occurs for large n (or equivalently small ε). This problem is avoided with
 714 multiple rounds of shuffling.

715 Figure [3] compares the performance of CSGM with and without coordinate pre-selection. In this
 716 regime coordinate pre-selection improves performance for all b . As predicted by Corollary [4.3] and
 717 Corollary [4.5], the MSE decreases with b but is effectively constant for sufficiently high b where the
 718 privacy term dominates. We can determine the communication cost needed for order-optimal central
 719 DP error performance to be the b at which the MSE is within some fixed constant factor away from
 720 the limiting value. We see that the communication cost increases with dimension d with the vanilla
 721 CSGM scheme, but a dimension-free communication cost is achieved with coordinate pre-selection.




Figure 2: Comparison of MSE vs. number of clients n for LDP scheme (SQKR) and shuffled SQKR. For shuffled SQKR, we set $b_0 = 1$ and choose ε_0 using results in Section [H](#). Communication cost is $\lceil (\log_2(2000) + 1) \rceil = 12$ bits per round.




Figure 3: CSGM with and without coordinate pre-selection using $d^l = 833$.