
Optimal Compression of Locally Differentially Private Mechanisms

Abhin Shah

Massachusetts Institute of Technology
abhin@mit.edu

Wei-Ning Chen

Stanford University
wnchen@stanford.edu

Johannes Balle

Google Research
jballe@google.com

Peter Kairouz

Google Research
kairouz@google.com

Lucas Theis

Google Research
theis@google.com

Abstract

Compressing the output of ϵ -locally differentially private (LDP) randomizers naively leads to suboptimal utility. In this work, we demonstrate the benefits of using schemes that jointly compress and privatize the data using shared randomness. In particular, we investigate a family of schemes based on Minimal Random Coding (Havasi et al., 2019) and prove that they offer optimal privacy-accuracy-communication tradeoffs. Our theoretical and empirical findings show that our approach can compress PrivUnit_2 (Bhowmick et al., 2018) and Subset Selection (Ye and Barg, 2018), the best known LDP algorithms for mean and frequency estimation, to the order of ϵ bits of communication while preserving their privacy and accuracy guarantees.

1 INTRODUCTION

Machine learning and data analytics are critical tools for designing better products and services. So far, these

Work done while A.S and W.C were interns at Google. J.B. provided the initial idea. P.K. and L.T. gave the conceptual and theoretical framework. A.S. and L.T. designed the algorithm. A.S., W.C., and L.T., devised the proofs. A.S. designed and performed the simulations with support from J.B.. A.S., W.C., P.K., and L.T. wrote the manuscript. J.B., P.K., and L.T., are listed alphabetically.

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

tools have been predominantly applied in datacenters on data that was curated from millions of users. However, centralized data collection and processing can expose individuals to privacy risks and organizations to legal risks if data is not properly managed. Indeed, increasing privacy concerns are fueling the demand for distributed learning and analytics systems that ensure that the underlying data remains private and secure. This is evident from the recent surge of interest in federated learning and analytics (e.g., Ramage and Mazzocchi, 2020; Kairouz et al., 2021).

Designing private and efficient distributed learning and analytics systems involves addressing three main challenges: (a) preserving the privacy of the user’s local data, (b) communicating the privatized data efficiently to a central server, and (c) achieving high accuracy on a task (e.g., mean or frequency estimation). Privacy is often achieved by enforcing ϵ -local differential privacy (ϵ -LDP) (Warner, 1965; Evfimievski et al., 2003; Dwork et al., 2006; Kasiviswanathan et al., 2011), which guarantees that the outcome from a privatization mechanism will not release too much individual information statistically. Efficient communication, on the other hand, is achieved via compression and dimensionality reduction techniques (Suresh et al., 2017; Alistarh et al., 2017; Wen et al., 2017; Wang et al., 2018; Han et al., 2018a,b; Agarwal et al., 2018; Gandikota et al., 2019; Barnes et al., 2020; Chen et al., 2021).

Most existing works focus on addressing two of the three above-mentioned challenges, such as achieving good privacy-accuracy or good communication-accuracy tradeoffs separately. However, doing so can lead to suboptimal performance where all three desiderata are concerned. It is thus important to investigate the joint privacy-communication-accuracy tradeoffs when designing communication-efficient and private distributed algorithms. Under ϵ -LDP constraints,

Chen et al. (2020) presents minimax order-optimal mechanisms for frequency and mean estimation that require only ε bits (independent of the underlying dimensionality of the problem) by using shared randomness¹. However, as noted by Feldman and Talwar (2021), the algorithms of Chen et al. (2020) are not competitive in terms of accuracy with the best known schemes – **Subset Selection** for frequency estimation (Ye and Barg, 2018) and **PrivUnit₂** for mean estimation (Bhowmick et al., 2018). Motivated by this fact, the present work addresses the following fundamental question: *Can we attain the best known accuracy under ε -LDP while only using on the order of ε bits of communication?* We answer this question affirmatively by leveraging a technique based on importance sampling called Minimal Random Coding (Havasi et al., 2019; Cuff, 2008; Song et al., 2016).

1.1 Our Contributions

We first demonstrate that Minimal Random Coding (MRC) can compress any ε -LDP mechanism in a near-lossless fashion using only on the order of ε bits of communication (see Theorem 3.1). We also prove that the resulting compressed mechanism is 2ε -LDP (see Theorem 3.2). Thus, to achieve ε -LDP, one has to simulate an $\varepsilon/2$ mechanism and pay the corresponding penalty in accuracy. Similar to Chen et al. (2020), this approach can achieve the order optimal privacy-accuracy tradeoffs with about ε bits of communication but is not competitive with the best known LDP schemes. However, we show that this approach is optimal if one is willing to accept approximate LDP with a small δ (see Theorem 3.3).

To overcome the limitations of MRC in the pure LDP case, we present a modified version (MMRC) such that the resulting compressed mechanism is ε -LDP (see Theorem 3.4). We show that MMRC can simulate a large class of LDP mechanisms in a near-lossless fashion using only on the order of ε bits of communication (see Theorem 3.5 in conjunction with Theorem 3.1).

While the class of LDP mechanisms MMRC can simulate includes the best-known schemes for mean and frequency estimation, MMRC (similar to MRC) is biased for a fixed number of bits of communication. We show that MMRC simulating **PrivUnit₂** and **Subset Selection** can be debiased (see Lemma 4.1 and Lemma 5.1), while preserving the corresponding accuracy guarantees (see Theorem 4.1 and Theorem 5.1).

Finally, we empirically demonstrate that MMRC achieves an accuracy comparable to **PrivUnit₂** and

Subset Selection (see Section 4.2 and Section 5.2)² while only using about ε bits.

We discuss interesting open problems in Section 6 and defer all additional results and experiments to the accompanying supplementary material.

1.2 Related Work

Recent works have examined approaches for compressing LDP schemes in the presence of shared randomness. When $\varepsilon \leq 1$, for frequency estimation, Bassily and Smith (2015) showed that a single bit is enough to simulate any LDP randomizer with (almost) no impact on its utility although with a large amount of shared randomness. Their result was improved upon, in terms of the amount of shared randomness required, by Bassily et al. (2017), Bun et al. (2019), and Acharya and Sun (2019).

Chen et al. (2020) generalized these methods to arbitrary ε 's, and provided order-optimal schemes for both frequency and mean estimation that only use on the order of ε bits. However, their method is only order-optimal and cannot achieve the accuracy of the best known schemes: **PrivUnit₂** (Bhowmick et al., 2018) for mean estimation and **Subset Selection** (Ye and Barg, 2018)³ for frequency estimation. We show how one can achieve the accuracy of these schemes with on the order of ε bits of communication (when $\varepsilon \geq 1$). While we don't advocate large ε (our methods work for $\varepsilon = 1$ as well), we note that larger ε are both of theoretical and practical interest since *amplification via shuffling* can convert a local $\varepsilon > 1$ to a small central ε (Erlingsson et al., 2019; Balle et al., 2019; Erlingsson et al., 2020).

In the absence of shared randomness, Girgis et al. (2021b), Girgis et al. (2021a), Chen et al. (2020) provided order-optimal mechanisms for frequency and mean estimation but their mechanisms do not achieve the best known accuracy. Feldman and Talwar (2021) presented an approach for compressing ε -LDP schemes in a lossless fashion using a pseudorandom generator (PRG). Their approach, which relies on cryptographic hardness of the PRG, can compress **Subset Selection** to $O(\ln d)$ bits and **PrivUnit₂** to $O(\varepsilon + \ln d)$ bits, where d is the dimension of the underlying problem. Their approach, similar to ours, can achieve the privacy vs accuracy tradeoffs of the best known schemes, i.e., **Subset Selection** and

²The source code of our implementation is available at <https://tinyurl.com/rcc-dp>.

³**Subset Selection** is similar to asymmetric RAPPOR (Erlingsson et al., 2014) in the sense that both have the same marginal distribution. Here, we focus on simulating **Subset Selection**.

¹We assume that the encoder and the decoder can depend on a random quantity that both the server and the user have access to. See Section 2.2 for details.

PrivUnit_2 . Nevertheless, their approach is designed to work without shared randomness, therefore requiring more bits than necessary if shared randomness is available, as in our work.

Unlike previous work, our technique of compressing generic LDP schemes relies on Minimal Random Coding (MRC), which was designed to simulate noisy channels. Several papers in information theory and related fields have studied the problem of efficiently simulating noisy channels over digital channels (e.g., Bennett and Shor, 2002; Harsha et al., 2007; Li and El Gamal, 2018) and proposed general solutions. In particular, these papers showed that any noisy channel can be simulated at a bit-rate which is close to the mutual information between the information available to the sender and the receiver. However, this result only holds if a shared source of randomness is available. Without such a source, the achievable rate has been shown to be close to Wyner’s common information (Wyner, 1975; Cuff, 2008), which can be significantly larger than the mutual information (Xu et al., 2011). While promising as a recipe for simulating arbitrary differentially private mechanisms, the general coding schemes discussed in these papers have not been analyzed for their effect on differential privacy guarantees. MRC (Havasi et al., 2019), which we analyze and build upon here, is one of these schemes and is also known as likelihood encoder in information theory (Cuff, 2008; Song et al., 2016).

Finally, mean and frequency estimation under LDP constraints, two canonical problems in distributed learning and analytics, have been widely studied (Duchi et al., 2013; Nguyen et al., 2016; Bhowmick et al., 2018; Wang et al., 2019; Gandikota et al., 2019; Erlingsson et al., 2014; Bassily and Smith, 2015; Kairouz et al., 2016; Ye and Barg, 2018; Acharya et al., 2019).

2 PRELIMINARIES

2.1 Locally Differentially Private (LDP)

Suppose $\mathbf{x} \in \mathcal{X}$ is some user’s data that must remain private. A privatization mechanism q is a randomized mapping that maps $\mathbf{x} \in \mathcal{X}$ to $\mathbf{z} \in \mathcal{Z}$ with probability $q(\mathbf{z}|\mathbf{x})$ where \mathcal{Z} can be arbitrary. The user transmits $\mathbf{z} \sim q(\cdot|\mathbf{x})$, i.e., a privatized version of \mathbf{x} to the server. Further, q is ε -LDP if

$$\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}, q(\mathbf{z}|\mathbf{x}) \leq e^\varepsilon q(\mathbf{z}|\mathbf{x}') \quad (1)$$

and q is (ε, δ) -LDP if $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \mathcal{Z} \subseteq \mathcal{Z}$,

$$\sum_{\mathbf{z} \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}) \leq e^\varepsilon \sum_{\mathbf{z} \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}') + \delta.$$

Here, we focus on ε -LDP mechanisms where $\varepsilon \geq 1$.

2.2 Shared Randomness

Here, we allow ε -LDP mechanisms to use *shared randomness*. That is, q can depend on a random variable $\mathbf{u} \in \mathcal{U}$ that is known to both the user and the server (but \mathbf{u} is independent of \mathbf{x}). The corresponding ε -LDP constraint is

$$\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}, \mathbf{u} \in \mathcal{U}, q(\mathbf{z}|\mathbf{x}, \mathbf{u}) \leq \exp(\varepsilon) q(\mathbf{z}|\mathbf{x}', \mathbf{u}).$$

The server wishes to reconstruct \mathbf{x} from \mathbf{z} and the corresponding estimator is allowed to implicitly depend on \mathbf{u} . However, for simplicity, we suppress the dependence on \mathbf{u} in our notation. In practice, shared randomness can be achieved via downlink communication, that is, the server generates \mathbf{u} (e.g., a random seed) and communicates it to the user. Further, we note that such shared randomness can be established well before the advent of any private data⁴.

2.3 PrivUnit₂

The PrivUnit_2 mechanism q^{pu} , proposed by Bhowmick et al. (2018), is an ε -LDP sampling scheme when the input alphabet \mathcal{X} is the d -dimensional unit ℓ_2 sphere \mathbb{S}^{d-1} . Formally, given a vector $\mathbf{x} \in \mathbb{S}^{d-1}$, PrivUnit_2 draws a random vector \mathbf{z} from a spherical cap $\{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\}$ with probability p_0 or from its complement $\{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle < \gamma\}$ with probability $1 - p_0$, where $\gamma \in [0, 1]$ and $p_0 \geq 1/2$ are parameters (depending on ε and d) that trade accuracy and privacy (see Appendix D). In other words, q^{pu} is as follows:

$$q^{\text{pu}}(\mathbf{z}|\mathbf{x}) = \begin{cases} \frac{2p_0}{A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})} & \text{if } \langle \mathbf{x}, \mathbf{z} \rangle \geq \gamma \\ \frac{2(1-p_0)}{2A(1, d) - A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})} & \text{otherwise} \end{cases}$$

where $A(1, d)$ denotes the area of \mathbb{S}^{d-1} and $I_x(a, b)$ denotes the regularized incomplete beta function. The estimator of the PrivUnit_2 mechanism (denoted by $\hat{\mathbf{x}}^{\text{pu}}$) is obtained by dividing every coordinate of \mathbf{z} by m_{pu} i.e., $\hat{\mathbf{x}}^{\text{pu}} := \mathbf{z}/m_{\text{pu}}$ where

$$m_{\text{pu}} := \frac{(1-\gamma^2)^\alpha}{2^{d-2}(d-1)} \left[\frac{p_0}{B(1; \alpha, \alpha) - B(\tau; \alpha, \alpha)} - \frac{1-p_0}{B(\tau; \alpha, \alpha)} \right] \quad (2)$$

with $\alpha = (d-1)/2$, $\tau = (1+\gamma)/2$, and $B(x; \alpha, \beta)$ denoting the incomplete beta function. The estimator $\hat{\mathbf{x}}^{\text{pu}}$ is (a) unbiased i.e., $\mathbb{E}[\hat{\mathbf{x}}^{\text{pu}}|\mathbf{x}] = \mathbf{x}$, (b) has order-optimal utility i.e., $\mathbb{E}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] = \Theta\left(\frac{d}{\min(\varepsilon, (\varepsilon-1)^2, d)}\right)$, and (c) achieves the best known constants for mean estimation. See Appendix D for more details on PrivUnit_2 .

⁴Quantifying the amount of such shared randomness required remains an open question. See Section 6.

2.4 Subset Selection

The **Subset Selection** mechanism q^{ss} , proposed by [Ye and Barg \(2018\)](#), is an ε -LDP sampling scheme when the input alphabet \mathcal{X} can take d different values. Without loss of generality, let $\mathcal{X} := \{e_1, e_2, \dots, e_d\}$, where $e_j \in \{0, 1\}^d$ is the j^{th} standard unit vector, i.e., the one-hot encoding of j . The output alphabet \mathcal{Z} is the set of all d -bit binary strings with Hamming weight $s := \lceil \frac{d}{1+e^\varepsilon} \rceil$, i.e.,

$$\mathcal{Z} = \left\{ \mathbf{z} = (z^{(1)}, \dots, z^{(d)}) \in \{0, 1\}^d : \sum_{i=1}^d z^{(i)} = s \right\}.$$

Given $\mathbf{x} \in \mathcal{X}$, **Subset Selection** maps it to $\mathbf{z} \in \mathcal{Z}$ with the following conditional probability:

$$q^{\text{ss}}(\mathbf{z}|\mathbf{x}) := \begin{cases} \frac{e^\varepsilon}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}} & \text{if } \mathbf{z} \in \mathcal{Z}_{\mathbf{x}} \\ \frac{1}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}} & \text{if } \mathbf{z} \in \mathcal{Z} \setminus \mathcal{Z}_{\mathbf{x}} \end{cases} \quad (3)$$

where $\mathcal{Z}_{\mathbf{x}} = \{\mathbf{z} = (z^{(1)}, \dots, z^{(d)}) \in \mathcal{Z} : z^{(x)} = 1\}$ is the set of elements in \mathcal{Z} with 1 in the x^{th} location. The estimator of the **Subset Selection** mechanism (denoted by $\hat{\mathbf{x}}^{\text{ss}}$) is obtained by subtracting b_{ss} from every component of \mathbf{z} and dividing every component of the result by m_{ss} i.e., $\hat{\mathbf{x}}^{\text{ss}} := (\mathbf{z} - b_{\text{ss}})/m_{\text{ss}}$ where

$$m_{\text{ss}} := \frac{s(d-s)(e^\varepsilon - 1)}{(d-1)(s(e^\varepsilon - 1) + d)}, \quad b_{\text{ss}} := \frac{s((s-1)e^\varepsilon + (d-s))}{(d-1)(s(e^\varepsilon - 1) + d)}. \quad (4)$$

The estimator $\hat{\mathbf{x}}^{\text{ss}}$ is (a) unbiased i.e., $\mathbb{E}[\hat{\mathbf{x}}^{\text{ss}}|\mathbf{x}] = \mathbf{x}$, (b) has optimal utility i.e., $\mathbb{E}[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2] = \Theta\left(\frac{d}{\min(e^\varepsilon, (e^\varepsilon - 1)^2, d)}\right)$ and (c) achieves the best known constants for frequency estimation. See [Appendix G](#) for more details on **Subset Selection**.

3 MAIN RESULTS

In this section, first, we describe the Minimal Random Coding algorithm for compressing any ε -LDP mechanism and prove its order-optimal privacy-accuracy-communication tradeoffs. Then, we propose the Modified Minimal Random Coding algorithm for compressing any ε -LDP *cap-based mechanism*⁵ and prove that it achieves optimal privacy-accuracy-communication tradeoffs.

3.1 Minimal Random Coding (MRC)

Consider an ε -LDP mechanism $q(\cdot|\mathbf{x})$ that we wish to compress. Under MRC, first, a number of candidates $\mathbf{z}_1, \dots, \mathbf{z}_N$ are drawn from a fixed reference distribution $p(\cdot)$ (known to both the user and the

server). This can be achieved via a pseudorandom number generator with a known seed. Next, the user transmits an index $K \in [N]$ to the server where K is drawn according to some distribution $\pi^{\text{mrc}}(\cdot)$ such that $\mathbf{z}_K \sim q(\cdot|\mathbf{x})$ approximately. The distribution π^{mrc} is such that, $\forall k \in [N]$, $\pi^{\text{mrc}}(k) \propto w(k)$ where $w(k) := q(\mathbf{z}_k|\mathbf{x})/p(\mathbf{z}_k)$ are the importance weights⁶ (see [Algorithm 1](#)). To communicate the index K of MRC, $\log N$ bits are required.

Algorithm 1: MRC

Input: ε -LDP mechanism $q(\cdot|\mathbf{x})$, reference distribution $p(\cdot)$, number of candidates N
 Draw samples $\mathbf{z}_1, \dots, \mathbf{z}_N$ from $p(\mathbf{z})$ // **Using the shared source of randomness**
for $k \in \{1, \dots, N\}$ **do**
 $w(k) \leftarrow q(\mathbf{z}_k|\mathbf{x})/p(\mathbf{z}_k)$
 $\pi^{\text{mrc}}(\cdot) \leftarrow w(\cdot)/\sum_k w(k)$
Output: $\pi^{\text{mrc}}(\cdot), \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$

Let q^{mrc} denote the distribution of \mathbf{z}_K where $K \sim \pi^{\text{mrc}}(\cdot)$. The following theorem shows that when the number of candidates is exponential in ε , samples drawn from q^{mrc} will be similar to samples drawn from $q(\cdot|\mathbf{x})$ in terms of ℓ_2 error. In other words, q^{mrc} can compress $q(\cdot|\mathbf{x})$ to the order of ε bits of communication as well as simulate it in a near-lossless fashion. A proof can be found in [Appendix B.1](#).

Theorem 3.1 (Utility of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP mechanism $q(\cdot|\mathbf{x})$. Consider any reference distribution $p(\cdot)$ such that $|\ln(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \forall \mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$.⁷ Let the number of candidates be $N = 2^{(\log e + 4c)\varepsilon}$ for some constant $c \geq 0$. Then, for $\alpha \in [0, 1/2]$, q^{mrc} is such that*

$$\left| \mathbb{E}_{q^{\text{mrc}}}[\|\mathbf{z} - \mathbf{x}\|_2^2] - \mathbb{E}_q[\|\mathbf{z} - \mathbf{x}\|_2^2] \right| \leq \frac{2\alpha\sqrt{\mathbb{E}_q[\|\mathbf{z} - \mathbf{x}\|_2^4]}}{1-\alpha} \quad (5)$$

holds with probability at least $1 - 2\alpha$, with c and α related by the following: $\alpha = \sqrt{2^{-c\varepsilon} + 2^{-c^2/\log e + 1}}$.

In general, $\mathbb{E}_q[\|\mathbf{z} - \mathbf{x}\|_2^4]$ in (5) can be well-controlled. See [Remark B.1](#) in [Appendix B.1](#) for more details.

In the next Theorem, we show that π^{mrc} is 2ε -LDP. Hence, the compressed mechanism q^{mrc} is 2ε -LDP.

Theorem 3.2 (Pure DP guarantee of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , and data $\mathbf{x} \in \mathcal{X}$. Consider any ε -LDP mechanism $q(\cdot|\mathbf{x})$, reference distribution $p(\cdot)$, and number of candidates $N \geq 1$. Then, $\pi^{\text{mrc}}(\cdot)$ obtained from [Algorithm 1](#) is a 2ε -LDP mechanism.*

⁶We suppress dependence of π^{mrc} & w on \mathbf{x} for simplicity.

⁷Note that this condition holds for many reference distributions $p(\cdot)$. For example, one can simply choose $p(\cdot) = q(\cdot|\mathbf{x}^*)$ for some $\mathbf{x}^* \in \mathcal{X}$.

⁵The family of cap-based mechanisms includes [PrivUnit₂](#) and **Subset Selection**. See [Definition 3.1](#).

A proof is provided in Appendix B.2.1 and it relies on fact that the following ratio can be bounded by $e^{2\varepsilon}$:

$$\frac{\pi_{\mathbf{x}}^{\text{mrc}}(k)}{\pi_{\mathbf{x}'}^{\text{mrc}}(k)} = \frac{q(\mathbf{z}_k|\mathbf{x})}{q(\mathbf{z}_k|\mathbf{x}')} \cdot \frac{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x}')/p(\mathbf{z}_{k'})}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})}.$$

In the following Theorem, we show that π^{mrc} is $(\varepsilon + \varepsilon_0, \delta)$ -LDP implying that the compressed mechanism q^{mrc} is $(\varepsilon + \varepsilon_0, \delta)$ -LDP where $\varepsilon_0 > 0$ and $\delta \leq 1$ are free parameters. This Theorem can be viewed complementary to Theorem 3.2 where a stronger privacy parameter can be achieved (i.e., $\varepsilon + \varepsilon_0$ which can get arbitrarily close to ε as opposed to 2ε) albeit at the cost of trading pure privacy for approximate privacy. A proof is provided in Appendix B.2.2.

Theorem 3.3 (Approximate DP guarantee of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP mechanism $q(\cdot|\mathbf{x})$. Consider any reference distribution $p(\cdot)$ such that $|\ln(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \forall \mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$.⁵ Let $c_0 \geq 0$ be some constant and let the number of candidates $N = \exp(2\varepsilon + 2c_0)$. Then, for any $\delta \leq 1$, $\pi^{\text{mrc}}(\cdot)$ obtained from Algorithm 1 is $(\varepsilon + \varepsilon_0, \delta)$ -LDP mechanism where*

$$\varepsilon_0 := \ln \frac{1 + a_0}{1 - a_0} \quad \text{and} \quad a_0 := \exp(-c_0) \sqrt{\frac{1}{2} \ln \frac{2}{\delta}}.$$

3.2 Modified Minimal Random Coding (MMRC)

While the results regarding MRC in Section 3.1 are general and offer order optimal privacy-accuracy tradeoffs with about ε bits of communication, the resulting compressed mechanism is not exactly ε -LDP. More specifically, Theorem 3.2 introduces an additional factor of 2 in the LDP guarantee and Theorem 3.3 provides an approximate privacy guarantee instead of a pure privacy guarantee. To address these limitations, we focus on a class of ε -LDP mechanisms which we call *cap-based* mechanisms and propose a modification to MRC such that the resulting compressed mechanism is ε -LDP. Further, like MRC, MMRC can simulate the underlying ε -LDP mechanism in a near-lossless fashion while using only on the order of ε bits.

We start with the definition of cap-based mechanism which is inspired from the structure of `PrivUnit2` and `Subset Selection`.

Definition 3.1 (Cap-based Mechanisms). *An ε -LDP mechanism $q(\mathbf{z}|\mathbf{x})$ with input alphabet \mathcal{X} and output alphabet \mathcal{Z} is a cap-based mechanism if it can be written in the following way:*

$$q(\mathbf{z}|\mathbf{x}) = \begin{cases} c_1(\varepsilon, d) & \text{if } \mathbf{z} \in \text{Cap}_{\mathbf{x}} \\ c_2(\varepsilon, d) & \text{if } \mathbf{z} \notin \text{Cap}_{\mathbf{x}} \end{cases} \quad (6)$$

where (a) $c_1(\varepsilon, d)$ and $c_2(\varepsilon, d)$ are constants with respect to \mathbf{x} and \mathbf{z} such that $c_1(\varepsilon, d) \geq c_2(\varepsilon, d)$, and (b) $\text{Cap}_{\mathbf{x}} \subseteq \mathcal{Z}$ such that $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}})$ is independent of \mathbf{x} and is at least $c_2(\varepsilon, d)/2c_1(\varepsilon, d)$.

In words, a cap-based ε -LDP mechanism samples uniformly either from $\text{Cap}_{\mathbf{x}}$ or from $\mathcal{Z} \setminus \text{Cap}_{\mathbf{x}}$ where $\text{Cap}_{\mathbf{x}} \subseteq \mathcal{Z}$ is such that if \mathbf{z} is sampled uniformly from \mathcal{Z} , it will belong to $\text{Cap}_{\mathbf{x}}$ with probability at least $c_2(\varepsilon, d)/2c_1(\varepsilon, d)$. It is easy to see that q^{ss} defined in (3) is a cap-based mechanism with $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$, $c_1(\varepsilon, d) = \frac{e^\varepsilon}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$, and $c_2(\varepsilon, d) = \frac{1}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$. See Appendix G where we evaluate $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \mathcal{Z}_{\mathbf{x}})$ and show that it is at least $1/2e^\varepsilon$. In Appendix D, we show that q^{pu} is a cap-based mechanism.

For a cap-based ε -LDP mechanism $q(\mathbf{z}|\mathbf{x})$ and a uniform reference distribution $p(\cdot)$, the distribution π^{mrc} obtained from Algorithm 1 takes a special form:

$$\pi^{\text{mrc}}(k) = \begin{cases} \frac{1}{N} \times \frac{c_1(\varepsilon, d)}{\theta c_1(\varepsilon, d) + (1-\theta)c_2(\varepsilon, d)} & \text{if } \mathbf{z}_k \in \text{Cap}_{\mathbf{x}} \\ \frac{1}{N} \times \frac{c_2(\varepsilon, d)}{\theta c_1(\varepsilon, d) + (1-\theta)c_2(\varepsilon, d)} & \text{if } \mathbf{z}_k \notin \text{Cap}_{\mathbf{x}} \end{cases} \quad (7)$$

where θ is the fraction of candidates inside the $\text{Cap}_{\mathbf{x}}$, i.e., $\theta = \frac{1}{N} \sum_k \mathbb{1}(\mathbf{z}_k \in \text{Cap}_{\mathbf{x}})$. As is, π^{mrc} in (7) is not necessarily ε -LDP because θ can be different for \mathbf{x} and \mathbf{x}' . However, as $N \rightarrow \infty$, $\theta \rightarrow \mathbb{E}[\theta] = \mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}})$, which is not a function of \mathbf{x} , implying that $\pi_{\mathbf{x}}^{\text{mrc}}(k)/\pi_{\mathbf{x}'}^{\text{mrc}}(k) \leq c_1(\varepsilon, d)/c_2(\varepsilon, d) \leq \exp(\varepsilon)$ ⁸. This shows that π^{mrc} is ε -LDP when $N \rightarrow \infty$. This motivates us to modify π^{mrc} to π^{mmrc} such that π^{mmrc} is ε -LDP irrespective of N . Further, when N is large enough, the modification is not by much, i.e., a sample from π^{mrc} is similar to a sample from π^{mmrc} .

To that end, define an upper threshold $t_u = \frac{1}{N} \times \frac{c_1(\varepsilon, d)}{\mathbb{E}[\theta]c_1(\varepsilon, d) + (1-\mathbb{E}[\theta])c_2(\varepsilon, d)}$ and a lower threshold $t_l = \frac{1}{N} \times \frac{c_2(\varepsilon, d)}{\mathbb{E}[\theta]c_1(\varepsilon, d) + (1-\mathbb{E}[\theta])c_2(\varepsilon, d)}$, and initialize π^{mmrc} to be equal to π^{mrc} . We want to modify π^{mmrc} so as to ensure:

$$t_l \leq \pi^{\text{mmrc}}(k) \leq t_u \quad \forall k \in [N], \quad (8)$$

which, as argued above, guarantees ε -LDP irrespective of the choice of N . First, it is easy to see that $\theta c_1(\varepsilon, d) + (1-\theta)c_2(\varepsilon, d)$ is an increasing function of θ . Next, we will look at 3 cases depending on the relationship between θ and $\mathbb{E}[\theta]$: (A) If $\theta = \mathbb{E}[\theta]$, then π^{mmrc} already satisfies (8); (B) If $\theta < \mathbb{E}[\theta]$, then only the upper threshold is violated and we set $\pi^{\text{mmrc}}(k) = t_u \forall k : \mathbf{z}_k \in \text{Cap}_{\mathbf{x}}$ and re-normalize the remaining $\pi^{\text{mmrc}}(k)$; (C) If $\theta > \mathbb{E}[\theta]$, then only the lower threshold is violated, we set $\pi^{\text{mmrc}}(k) = t_l \forall k : \mathbf{z}_k \notin \text{Cap}_{\mathbf{x}}$ and re-normalize the remaining $\pi^{\text{mmrc}}(k)$. The re-normalization step does not violate (8). We provide pseudo-code to calculate π^{mmrc} in Algorithm 2.

⁸This follows from (1) and (6) because $q(\cdot|\mathbf{x})$ is ε -LDP.

Algorithm 2: MMRC

Input: ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$, the associated $\text{Cap}_{\mathbf{x}}$, reference distribution $p(\cdot)$, number of candidates N , lower threshold t_l , upper threshold t_u

$\pi^{\text{mrc}}(\cdot), \{z_1, \dots, z_N\} \leftarrow \text{MRC}(p(\cdot), q(\cdot|\mathbf{x}), N)$

$\theta \leftarrow \frac{1}{N} \sum_k \mathbb{1}(z_k \in \text{Cap}_{\mathbf{x}})$ // Compute the fraction of candidates inside the cap

Initialization: $\pi^{\text{mmrc}}(\cdot) \leftarrow \pi^{\text{mrc}}(\cdot)$

if $\max_k \pi^{\text{mmrc}}(k) > t_u$ **then**

// Upper threshold is violated

$\pi^{\text{mmrc}}(k) \leftarrow t_u, \forall k : z_k \in \text{Cap}_{\mathbf{x}}$

$\pi^{\text{mmrc}}(k) \leftarrow \frac{1-N\theta t_u}{N(1-\theta)}, \forall k : z_k \notin \text{Cap}_{\mathbf{x}}$

else if $\min_k \pi^{\text{mmrc}}(k) < t_l$ **then**

// Lower threshold is violated

$\pi^{\text{mmrc}}(k) \leftarrow t_l, \forall k : z_k \notin \text{Cap}_{\mathbf{x}}$

$\pi^{\text{mmrc}}(k) \leftarrow \frac{1-N(1-\theta)t_l}{N\theta}, \forall k : z_k \in \text{Cap}_{\mathbf{x}}$

Output: $\pi^{\text{mmrc}}(\cdot), \{z_1, \dots, z_N\}$

Let q^{mmrc} denote the distribution of z_K where $K \sim \pi^{\text{mmrc}}(\cdot)$. In the following Theorem, we show that π^{mmrc} is ε -LDP implying that the compressed mechanism q^{mmrc} is ε -LDP. The proof follows from (8) and can be found in Appendix C.1.

Theorem 3.4 (DP guarantee of MMRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$. Let the reference distribution $p(\cdot)$ be the uniform distribution on \mathcal{Z} . Consider any number of candidates $N \geq 1$. Then, $\pi^{\text{mmrc}}(\cdot)$ obtained from Algorithm 2 is an ε -LDP mechanism.*

The following Theorem shows that, with number of candidates exponential in ε , samples drawn from q^{mmrc} will be similar to the samples drawn from q^{mrc} in terms of ℓ_2 error. A proof can be found in Appendix C.3.

Theorem 3.5 (Utility of MMRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$. Let the reference distribution $p(\cdot)$ be the uniform distribution on \mathcal{Z} . Let N denote the number of candidates. Then, q^{mmrc} is such that*

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] \leq \mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] + \sqrt{\frac{\rho(1+\varepsilon)}{2}} \max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2 \quad (9)$$

where $\rho \in (0, 1)$ is such that $N = \frac{2(\exp(\varepsilon)-1)^2}{\rho^2} \ln \frac{2}{\rho}$.

For bounded mechanisms, $\max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2$ in (9) can be well-controlled. See Remark C.1 in Appendix C.3 for a discussion.

In conjunction with Theorem 3.1, Theorem 3.5 implies that q^{mmrc} can compress $q(\cdot|\mathbf{x})$ to the order of ε bits of communication and simulate it in a near-lossless fashion. This is stated formally and proved in Appendix C.3.

4 MEAN ESTIMATION

In this section, we focus on the mean estimation problem, which is a canonical statistical task in distributed estimation with applications in distributed stochastic gradient descent, federated learning, etc. Let the input space \mathcal{X} be the d -dimensional unit ℓ_2 sphere, i.e., $\mathcal{X} = \mathbb{S}^{d-1}$. Consider n users where user i has some data $\mathbf{x}_i \in \mathcal{X}$. For every $i \in [n]$, let \mathbf{x}_i be privatized using an ε -LDP mechanism $q(\cdot|\mathbf{x}_i)$ and potentially post-processed to obtain an estimate $\hat{\mathbf{x}}_i$ of \mathbf{x}_i . We are interested in estimating the empirical mean $\boldsymbol{\mu} \triangleq \frac{1}{n} \sum_i \mathbf{x}_i$ using $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n$ such that the mean estimation error defined below is minimized

$$r_{\text{ME}}(\hat{\boldsymbol{\mu}}, q) \triangleq \max_{\mathbf{x}^n \in \mathcal{X}^n} \mathbb{E} \left[\|\hat{\boldsymbol{\mu}}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n) - \boldsymbol{\mu}\|_2^2 \right], \quad (10)$$

where $\hat{\boldsymbol{\mu}}$ is an estimate of $\boldsymbol{\mu}$ and the expectation is with respect to $q(\cdot|\mathbf{x}_i)$ as well as all (possibly shared) randomness used by $q(\cdot|\mathbf{x}_i) \forall i \in [n]$.

Bhowmick et al. (2018) show that PrivUnit_2 achieves the order-optimal privacy-accuracy trade-off for mean estimation, i.e., $r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{\text{pu}}, q^{\text{pu}}) = \Theta\left(\frac{d}{\min(\varepsilon, (\varepsilon^\varepsilon - 1)^2, d)}\right)$ where $\hat{\boldsymbol{\mu}}^{\text{pu}} := \frac{1}{n} \sum_i \hat{\mathbf{x}}_i^{\text{pu}}$. Moreover, compared to other (order-optimal) ε -LDP mean estimation mechanisms, PrivUnit_2 admits the best constants and gives the smallest ℓ_2 error in practice (see Feldman and Talwar (2021)). However, PrivUnit_2 requires each user to send a d -dimensional real vector, so without any compression, the communication needed is $\Theta(d)$ bits, which can be an issue in many practical scenarios.

To compress and simulate PrivUnit_2 , one can directly apply the generic MMRC mechanism defined in Section 3.2. However, for a fixed number of candidates N , MMRC yields a biased estimate of \mathbf{x} and hence cannot get the correct (optimal) order of estimation error in (10), i.e., the error would not decay with n ⁹. Fortunately, we show (in Section 4.1) that the bias can be corrected by appropriately scaling the privatized version of \mathbf{x} , i.e., by using an estimator which is slightly different compared to the original estimator of PrivUnit_2 . Further, we also show (in Section 4.2) that the resulting unbiased estimator for mean estimation ($\hat{\boldsymbol{\mu}}^{\text{mmrc}}$) can simulate PrivUnit_2 closely while only using on the order of ε bits of communication.

⁹We note that this does not undermine the significance of Theorem 3.1 and Theorem 3.5. These are useful in single-user settings (i.e., $n = 1$) and are generic as they can compress (near-losslessly) any ε -LDP and ε -LDP cap-based mechanism, respectively.

4.1 Debiasing MMRC to simulate PrivUnit₂

Let us focus on a single user and consider some data $\mathbf{x} \in \mathcal{X}$. Recall the PrivUnit₂ ε -LDP mechanism q^{pu} described in Section 2 with parameters p_0 and γ . PrivUnit₂ is a cap-based mechanism with $\text{Cap}_{\mathbf{x}} = \{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\}$ (see Appendix D for details). Let π^{mmrc} be the distribution and $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ be the candidates obtained from Algorithm 2 when the reference distribution is $\text{Unif}(\mathbb{S}^{d-1})$. Let $K \sim \pi^{\text{mmrc}}(\cdot)$. Therefore, \mathbf{z}_K is the privatized version of \mathbf{x} using MMRC.

Define $p_{\text{mmrc}} := \mathbb{P}(\mathbf{z}_K \in \text{Cap}_{\mathbf{x}})$ to be the probability with which the sampled candidate \mathbf{z}_K belongs to the spherical cap associated with PrivUnit₂. Define m_{mmrc} as the scaling factor in (2) when p_0 in (2) is replaced by p_{mmrc} . Define $\hat{\mathbf{x}}^{\text{mmrc}} := \mathbf{z}_K / m_{\text{mmrc}}$ as the estimator of the MMRC mechanism simulating PrivUnit₂. The following Lemma shows that $\hat{\mathbf{x}}^{\text{mmrc}}$ is an unbiased estimator. See Appendix F.1 for a proof.

Lemma 4.1. *Let $\hat{\mathbf{x}}^{\text{mmrc}}$ be the estimator of the MMRC mechanism simulating PrivUnit₂ as defined above. Then, $\mathbb{E}_{q^{\text{mmrc}}}[\hat{\mathbf{x}}^{\text{mmrc}}] = \mathbf{x}$.*

4.2 Simulating PrivUnit₂ using MMRC

Finally, we consider estimating the empirical mean $\boldsymbol{\mu}$ defined earlier using the MMRC scheme simulating PrivUnit₂. To that end, consider n users and let $\hat{\mathbf{x}}_i^{\text{mmrc}}$ be the unbiased estimator of \mathbf{x}_i at the i^{th} user. Let the (unbiased) estimate of $\boldsymbol{\mu}$ be $\hat{\boldsymbol{\mu}}^{\text{mmrc}} := \frac{1}{n} \sum_i \hat{\mathbf{x}}_i^{\text{mmrc}}$.

The following Theorem shows that, for mean estimation, MMRC can simulate PrivUnit₂ in a near-lossless manner (when n is large and λ is small) while only using on the order of ε bits of communication. A proof can be found in Appendix F.2. The key idea in the proof is to show that when the number of candidates N is exponential in ε , the scaling factor m_{mmrc} is close to the scaling parameter associated with PrivUnit₂ (i.e., m_{pu} defined in (2)).

Theorem 4.1. *Let $r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{pu}, q^{pu})$ and $r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{\text{mmrc}}, q^{\text{mmrc}})$ be the empirical mean estimation error for PrivUnit₂ with parameter p_0 and MMRC simulating PrivUnit₂ with N candidates respectively. Consider any $\lambda > 0$. Then,*

$$r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{\text{mmrc}}, q^{\text{mmrc}}) \leq (1 + \lambda)^2 r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{pu}, q^{pu}) + 2(1 + \lambda)(2 + \lambda) \sqrt{\frac{r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{pu}, q^{pu})}{n}} + \frac{(2 + \lambda)^2}{n}.$$

as long as

$$N \geq \frac{e^{2\varepsilon}}{2} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)} \right). \quad (11)$$

We note that while a specific value of λ can be chosen (say 0.1 or smaller) in (11), in practice, the number of

bits could be fixed (see Section 4.3), determining the value of λ .

4.3 Empirical Comparisons

Next, we empirically demonstrate the privacy-accuracy-communication tradeoffs of MMRC simulating PrivUnit₂. Along with PrivUnit₂, we compare against the SQKR algorithm of Chen et al. (2020) which offers order-optimal privacy-accuracy tradeoffs while requiring only ε bits. Following Chen et al. (2020), we generate data independently but non-identically to capture the distribution-free setting as well as ensure that the data non-central, i.e. $\boldsymbol{\mu} \neq 0$. More specifically, we set $\mathbf{x}_1, \dots, \mathbf{x}_{n/2} \stackrel{\text{i.i.d.}}{\sim} N(1, 1)^{\otimes d}$ and $\mathbf{x}_{n/2+1}, \dots, \mathbf{x}_n \stackrel{\text{i.i.d.}}{\sim} N(10, 1)^{\otimes d}$. Further, to ensure that each data lies on \mathbb{S}^{d-1} , we normalize each \mathbf{x}_i by setting $\mathbf{x}_i \leftarrow \mathbf{x}_i / \|\mathbf{x}_i\|_2$. We report the average ℓ_2 estimation error over 10 runs. See more variations in Appendix F.3.

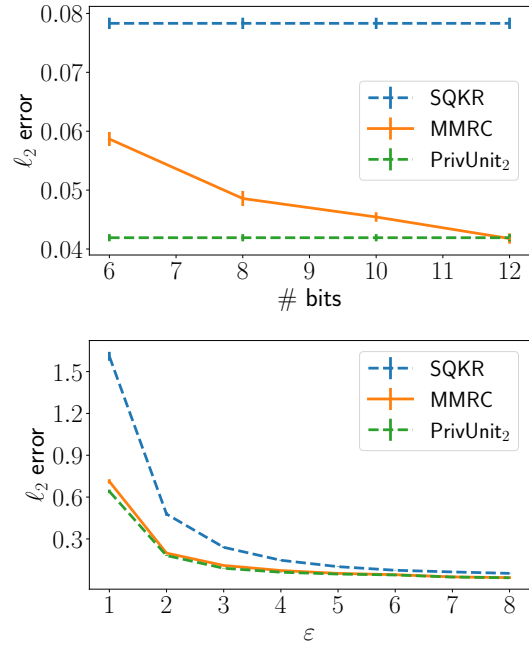


Figure 1: Comparing PrivUnit₂, MMRC simulating PrivUnit₂ and SQKR for mean estimation with $d = 500$ and $n = 5000$. **Top:** ℓ_2 error vs #bits for $\varepsilon = 0.04$. **Bottom:** ℓ_2 error vs ε for #bits = $\max\{\lceil \varepsilon / \ln 2 \rceil + 2, 8\}$. SQKR uses #-bits = ε for both as it leads to a poor performance if #-bits $> \varepsilon$.

In Figure 1 (top), we show the communication-accuracy tradeoffs. We see that with correct order of bits, the accuracy of MMRC simulating PrivUnit₂ converges to the accuracy of the uncompressed PrivUnit₂. In Figure 1 (bottom), we show the privacy-accuracy tradeoffs. We see that MMRC simulating PrivUnit₂ can

attain accuracy of the uncompressed `PrivUnit2` for the range of ε 's typically considered by LDP mechanisms while only using $\max\{\lceil(\varepsilon/\ln 2)\rceil + 2, 8\}$ bits.

5 FREQUENCY ESTIMATION

In this section, we study the frequency estimation problem, which is another canonical statistical task in distributed distribution estimation, with application to federated analytics (Ramage and Mazzocchi, 2020).

Let \mathcal{X} be a set of d distinct symbols and without loss of generality $\mathcal{X} := \{e_1, e_2, \dots, e_d\}$, where $e_j \in \{0, 1\}^d$ is the j^{th} standard unit vector i.e., e_j is the one-hot encoding of j . Consider n users where user i has some data $\mathbf{x}_i \in \mathcal{X}$. For every $i \in [n]$, let \mathbf{x}_i be privatized using an ε -LDP mechanism $q(\cdot|\mathbf{x}_i)$ and potentially post-processed to obtain an estimate $\hat{\mathbf{x}}_i$ of \mathbf{x}_i . We are interested in estimating the empirical distribution of $\mathbf{x}_1, \dots, \mathbf{x}_n$ defined as $\Pi \triangleq \frac{1}{n} \sum_i \mathbf{x}_i$ using $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n$ such that the estimation error defined below is minimized:

$$r_{\text{FE}}(\hat{\Pi}, q, \ell) \triangleq \max_{\mathbf{x}^n \in \mathcal{X}^n} \mathbb{E} \left[\ell \left(\hat{\Pi}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n), \Pi \right) \right], \quad (12)$$

where $\ell = \|\cdot\|_1$ or $\|\cdot\|_2^2$, $\hat{\Pi}$ is an estimate of Π and the expectation is with respect to $q(\cdot|\mathbf{x}_i)$ as well as all (possibly shared) randomness used by $q(\cdot|\mathbf{x}_i) \forall i \in [n]$. For simplicity, we only focus on ℓ_2 error i.e., $\ell = \|\cdot\|_2^2$.

Ye and Barg (2018) show that the `Subset Selection` achieves the order-optimal privacy-accuracy trade-off for frequency estimation i.e., $r_{\text{FE}}(\hat{\Pi}^{\text{ss}}, q^{\text{ss}}) = \Theta\left(\frac{d}{\min(e^\varepsilon, (e^\varepsilon - 1)^2, d)}\right)$ (where $\hat{\Pi}^{\text{ss}} := \frac{1}{n} \sum_i \hat{\mathbf{x}}_i^{\text{ss}}$). Like `PrivUnit2`, compared to other (order-optimal) ε -LDP frequency estimation mechanisms, `Subset Selection` admits the best constants and gives the smallest ℓ_2 error in practice (see Chen et al. (2020)). However, the communication cost associated with `Subset Selection` is $O\left(\frac{d}{e^\varepsilon + 1}\right)$ bits per user, which which can be an issue for small and moderate ε .

Similar to `PrivUnit2`, one could apply the generic MMRC scheme defined in Section 3 to compress and simulate `Subset Selection`. However, for a fixed number of candidates N , it yields a biased estimate of \mathbf{x} and hence cannot get the correct (optimal) order of estimation error in (12) i.e., the error would not decay with n . Fortunately, similar to `PrivUnit2`, we show (in Section 5.1) that the bias can be corrected by appropriately translating and scaling the privatized version of \mathbf{x} i.e., by using an estimator which is slightly different compared to the original estimator of `Subset Selection`. Further, we also show (in Section 5.2) that the resulting unbiased estimator for frequency estimation ($\hat{\Pi}^{\text{mmrc}}$) can simulate

`Subset Selection` closely while only using on the order of ε -bits communication.

5.1 Debiasing MMRC to simulate Subset Selection

Let us focus on a single user and consider some data $\mathbf{x} \in \mathcal{X}$. Recall the `Subset Selection` ε -LDP mechanism q^{ss} described in Section 2 with $s := \lceil \frac{d}{1+\varepsilon} \rceil$. `Subset Selection` is cap-based mechanism as discussed in Section 3 and Appendix G with $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}}) = s/d$. Similar to Section 4.1, let \mathbf{z}_K be the privatized version of \mathbf{x} using MMRC. We define $\hat{\mathbf{x}}^{\text{mmrc}} := (\mathbf{z}_K - b_{\text{mmrc}})/m_{\text{mmrc}}$ as the estimator of the MMRC mechanism simulating `Subset Selection` where m_{mmrc} and b_{mmrc} (defined in Appendix I.1) are translation and scaling factor analogous to m_{ss} and b_{ss} in (4). The following Lemma shows that $\hat{\mathbf{x}}^{\text{mmrc}}$ is an unbiased estimator. See Appendix I.1 for a proof.

Lemma 5.1. *Let $\hat{\mathbf{x}}^{\text{mmrc}}$ be the estimator of the MMRC mechanism simulating `Subset Selection` as defined above. Then, $\mathbb{E}[\hat{\mathbf{x}}^{\text{mmrc}}] = \mathbf{x}$.*

5.2 Simulating Subset Selection using MMRC

Finally, we consider estimating the empirical frequency Π defined earlier using the MMRC scheme simulating `Subset Selection`. To that end, consider n users and let $\hat{\mathbf{x}}_i^{\text{mmrc}}$ be the unbiased estimator of \mathbf{x}_i at the i^{th} user. Let the (unbiased) estimate of Π be $\hat{\Pi}^{\text{mmrc}} := \frac{1}{n} \sum_i \hat{\mathbf{x}}_i^{\text{mmrc}}$. The following Theorem shows that, for frequency estimation, MMRC can simulate `Subset Selection` in a near-lossless manner (when λ is small) while only using on the order of ε bits of communication. A proof can be found in Appendix I.2. Similar to `PrivUnit2`, the key idea in the proof is to show that when the number of candidates N is exponential in ε , the scaling factor m_{mmrc} is close to the scaling parameter associated with q^{ss} (i.e., m_{ss} defined in (4)).

Theorem 5.1. *Let $r_{\text{FE}}(\hat{\Pi}^{\text{ss}}, q^{\text{ss}})$ and $r_{\text{FE}}(\hat{\Pi}^{\text{mmrc}}, q^{\text{mmrc}})$ be the empirical mean estimation error for `Subset Selection` and MMRC simulating `Subset Selection` with N candidates respectively. Consider any $\lambda > 0$. Then*

$$r_{\text{FE}}(\hat{\Pi}^{\text{mmrc}}, q^{\text{mmrc}}) \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) r_{\text{FE}}(\hat{\Pi}^{\text{ss}}, q^{\text{ss}}),$$

as long as

$$N \geq \frac{2(e^\varepsilon + 1)^2(1 + \lambda)^2}{0.24^2 \lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right). \quad (13)$$

Similar to mean estimation, while a specific value of λ can be chosen in (13), in practice, the number of bits

could be fixed (see Section 5.3), determining the value of λ .

5.3 Empirical Comparisons.

Next, we empirically demonstrate the privacy-accuracy-communication tradeoffs of MMRC simulating Subset Selection. Along with Subset Selection, we compare against the RHR algorithm of Chen et al. (2020) which offers order-optimal privacy-accuracy tradeoffs while requiring only ε bits. Following Acharya et al. (2019), we generate $\mathbf{x}_1, \dots, \mathbf{x}_n$ from the Zipf distribution with degree 1. We report the average ℓ_2 estimation error over 10 runs. See more variations in Appendix I.3.




Figure 2: Comparing Subset Selection, MMRC simulating Subset Selection and RHR for frequency estimation with $d = 500$ and $n = 5000$. **Top:** ℓ_2 error vs #bits for $\varepsilon = 6$. **Bottom:** ℓ_2 error vs ε for #bits = $\max\{\lceil \varepsilon / \ln 2 \rceil + 3, 8\}$. RHR uses #bits = ε for both as it leads to a poor performance if #bits $> \varepsilon$.

In Figure 2 (top), we show the communication-accuracy tradeoffs. We see that with correct order of bits, the accuracy of MMRC simulating Subset Selection converges to the accuracy of the uncompressed Subset Selection. In Figure 2 (bottom), we show the privacy-accuracy tradeoffs. More specifically, MMRC simulating Subset Selection can attain the accuracy of the uncompressed Subset Selection for the range of ε 's typically considered by LDP mechanisms while only using $\max\{\lceil \varepsilon / \ln 2 \rceil + 3, 8\}$ bits.

6 CONCLUSION AND OPEN PROBLEMS

We demonstrated how Minimal Random Coding can be used to simulate a class of ε -LDP mechanisms in a manner which is communication efficient while preserving accuracy and differential privacy guarantees. Further, for mean and frequency estimation, we proposed unbiased versions of our schemes (relying only on translation and scaling) that attain the privacy-accuracy tradeoffs of the best known schemes i.e., PrivUnit₂ and Subset Selection, while requiring on the order of ε bits of communication.

We now briefly discuss a few non-trivial and interesting open questions.

Computational Cost. The computational cost of our approach, similar to Feldman and Talwar (2021) grows linearly in d and exponentially in ε (as we need $N = \exp(O(\varepsilon))$ candidates to properly simulate the optimal mechanisms). An important question for future research is therefore how to increase the computational efficiency of MRC and MMRC with respect to ε .

Privacy Amplification via Shuffling. As mentioned in Section 1.2, privacy amplification via shuffling techniques ensure a central $\varepsilon \approx 1$ even when the local $\varepsilon > 1$. While our method could be combined with these amplification techniques in principle, we leave the analysis of the privacy, utility, and communication guarantees of the resulting scheme as a question for future research.

Other schemes to simulate noisy channels. MRC is only one of several channel simulation schemes studied in information theory which could be considered for compression of ε -LDP mechanisms. Similar to MRC, other channel simulation schemes, e.g., rejection sampling (Harsha et al., 2007) or schemes based on the Poisson functional representation (Li and El Gamal, 2018), can also compress noisy signals to a number of bits which is close to the information contained in the signal (which decreases as noise increases). Analyzing these schemes for their effect on differential privacy guarantees is an interesting open question.

Shared Randomness. Finally, here we assumed the existence of a shared source of randomness. We further assumed that each user is using a different source of shared randomness. While shared randomness only adds to the cost of downlink and not uplink communication (which is usually the bottleneck in settings like federated learning), a question left for future research is how much communication is required to establish and select these sources of randomness.

Acknowledgements

We thank the anonymous reviewers of AISTATS 2022 for their comments and suggestions. We sincerely thank Jakub Konečný and Wennan Zhu for helpful discussions. We also thank Zachary Charles for support with the software infrastructure.

References

- J. Acharya and Z. Sun. Communication complexity in locally private distribution estimation and heavy hitters. In *International Conference on Machine Learning*, pages 51–60, 2019.
- J. Acharya, Z. Sun, and H. Zhang. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1120–1129. PMLR, 2019.
- N. Agarwal, A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan. cpsgd: Communication-efficient and differentially-private distributed sgd. In *Advances in Neural Information Processing Systems*, pages 7564–7575, 2018.
- D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 1709–1720. Curran Associates, Inc., 2017.
- B. Balle, J. Bell, A. Gascón, and K. Nissim. The privacy blanket of the shuffle model. In *Annual International Cryptology Conference*, pages 638–667. Springer, 2019.
- L. P. Barnes, H. A. Inan, B. Isik, and A. Ozgur. rtop-k: A statistical estimation approach to distributed sgd, 2020.
- R. Bassily and A. Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 127–135, 2015.
- R. Bassily, K. Nissim, U. Stemmer, and A. Thakurta. Practical locally private heavy hitters. *arXiv preprint arXiv:1707.04982*, 2017.
- C. H. Bennett and P. W. Shor. Entanglement-Assisted Capacity of a Quantum Channel and the Reverse Shannon Theorem. *IEEE Trans. Info. Theory*, 48(10), 2002.
- A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv preprint arXiv:1812.00984*, 2018.
- M. Bun, J. Nelson, and U. Stemmer. Heavy hitters and the structure of local privacy. *ACM Transactions on Algorithms (TALG)*, 15(4):1–40, 2019.
- W.-N. Chen, P. Kairouz, and A. Özgür. Breaking the communication-privacy-accuracy trilemma. *arXiv preprint arXiv:2007.11707*, 2020.
- W.-N. Chen, P. Kairouz, and A. Özgür. Breaking the dimension dependence in sparse distribution estimation under communication constraints. *arXiv preprint arXiv:2106.08597*, 2021.
- P. Cuff. Communication requirements for generating correlated random variables. In *2008 IEEE International Symposium on Information Theory*, pages 1393–1397, 2008.
- J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 429–438. IEEE, 2013.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- U. Erlingsson, V. Pihur, and A. Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 21st ACM Conference on Computer and Communications Security*, Scottsdale, Arizona, 2014. URL <https://arxiv.org/abs/1407.6981>.
- Ú. Erlingsson, V. Feldman, I. Mironov, A. Raghunathan, K. Talwar, and A. Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2468–2479. SIAM, 2019.
- Ú. Erlingsson, V. Feldman, I. Mironov, A. Raghunathan, S. Song, K. Talwar, and A. Thakurta. Encode, shuffle, analyze privacy revisited: Formalizations and empirical evaluation. *arXiv preprint arXiv:2001.03618*, 2020.
- A. Evfimievski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 211–222, 2003.
- V. Feldman and K. Talwar. Lossless compression of efficient private local randomizers. *arXiv preprint arXiv:2102.12099*, 2021.
- V. Gandikota, D. Kane, R. K. Maity, and A. Mazumdar. vqsgd: Vector quantized stochastic gradient descent, 2019.
- A. Girgis, D. Data, S. Diggavi, P. Kairouz, and A. T. Suresh. Shuffled model of differential privacy in

- federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2521–2529. PMLR, 2021a.
- A. M. Girgis, D. Data, S. Diggavi, P. Kairouz, and A. T. Suresh. Shuffled model of federated learning: Privacy, accuracy and communication trade-offs. *IEEE Journal on Selected Areas in Information Theory*, 2(1):464–478, 2021b.
- Y. Han, P. Mukherjee, A. Ozgur, and T. Weissman. Distributed statistical estimation of high-dimensional and nonparametric distributions. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 506–510. IEEE, 2018a.
- Y. Han, A. Özgür, and T. Weissman. Geometric lower bounds for distributed parameter estimation under communication constraints. *arXiv preprint arXiv:1802.08417*, 2018b.
- P. Harsha, R. Jain, D. McAllester, and J. Radhakrishnan. The communication complexity of correlation. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 10–23. IEEE, 2007.
- M. Havasi, R. Peharz, and J. M. Hernández-Lobato. Minimal Random Code Learning: Getting Bits Back from Compressed Model Parameters. In *International Conference on Learning Representations*, 2019.
- P. Kairouz, K. Bonawitz, and D. Ramage. Discrete distribution estimation under local privacy. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 2436–2444, New York, New York, USA, 20–22 Jun 2016.
- P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, R. G. L. D’Oliveira, H. Eichner, S. E. Rouayheb, D. Evans, J. Gardner, Z. Garrett, A. Gascón, B. Ghazi, P. B. Gibbons, M. Gruteser, Z. Harchaoui, C. He, L. He, Z. Huo, B. Hutchinson, J. Hsu, M. Jaggi, T. Javidi, G. Joshi, M. Khodak, J. Konečný, A. Korolova, F. Koushanfar, S. Koyejo, T. Lepoint, Y. Liu, P. Mittal, M. Mohri, R. Nock, A. Özgür, R. Pagh, H. Qi, D. Ramage, R. Raskar, M. Raykova, D. Song, W. Song, S. U. Stich, Z. Sun, A. T. Suresh, F. Tramèr, P. Vepakomma, J. Wang, L. Xiong, Z. Xu, Q. Yang, F. X. Yu, H. Yu, and S. Zhao. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021. ISSN 1935-8237. doi: 10.1561/22000000083. URL <http://dx.doi.org/10.1561/22000000083>.
- S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- C. T. Li and A. El Gamal. Strong functional representation lemma and applications to coding theorems. *IEEE Transactions on Information Theory*, 64(11):6967–6978, 2018.
- T. T. Nguyễn, X. Xiao, Y. Yang, S. C. Hui, H. Shin, and J. Shin. Collecting and analyzing data from smart device users with local differential privacy, 2016.
- D. Ramage and S. Mazzocchi. Federated analytics: Collaborative data science without data collection. <https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html>, 2020.
- E. C. Song, P. Cuff, and H. V. Poor. The likelihood encoder for lossy compression. *IEEE Transactions on Information Theory*, 62(4):1836–1849, 2016. doi: 10.1109/TIT.2016.2529657.
- A. T. Suresh, F. X. Yu, S. Kumar, and H. B. McMahan. Distributed mean estimation with limited communication. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 3329–3337. JMLR.org, 2017.
- H. Wang, S. Sievert, S. Liu, Z. Charles, D. Papailiopoulos, and S. Wright. Atomo: Communication-efficient learning via atomic sparsification. In *Advances in Neural Information Processing Systems*, pages 9850–9861, 2018.
- T. Wang, J. Zhao, X. Yang, and X. Ren. Locally differentially private data collection and analysis. *arXiv preprint arXiv:1906.01777*, 2019.
- S. L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- W. Wen, C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in neural information processing systems*, pages 1509–1519, 2017.
- A. Wyner. The common information of two dependent random variables. *IEEE Transactions on Information Theory*, 21(2):163–179, 1975. doi: 10.1109/TIT.1975.1055346.
- G. Xu, W. Liu, and B. Chen. Wyners common information for continuous random variables - a lossy source coding interpretation. In *45th Annual Conference on Information Sciences and Systems*, pages 1–6, 2011. doi: 10.1109/CISS.2011.5766249.

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

Supplementary Material: Optimal Compression of Locally Differentially Private Mechanisms

Organization. The Appendix is organized as follows. In Appendix A, we discuss the societal impact associated with our work. In Appendix B, we focus on MRC and provide the proofs of Theorem 3.1, Theorem 3.2, and Theorem 3.3. In Appendix C, we focus on MMRC and provide the proofs of Theorem 3.4 and Theorem 3.5. Further, we also provide Theorem C.1 where we show that MMRC can simulate any ε -LDP cap-based mechanism in a nearly lossless fashion with about ε bits of communication. In Appendix D, we provide additional preliminary on PrivUnit_2 and also show that PrivUnit_2 is a cap-based mechanism (Definition 3.1). In Appendix E, we show how PrivUnit_2 can be simulated using MRC analogous to how we simulated PrivUnit_2 using MMRC in Section 4. Along with the theoretical guarantees, we also provide some empirical comparisons between MRC simulating PrivUnit_2 and PrivUnit_2 . In Appendix F, we provide the proofs of Lemma 4.1 and Theorem 4.1 as well as some additional empirical comparisons between MMRC simulating PrivUnit_2 and PrivUnit_2 . In Appendix G, we provide additional preliminary on **Subset Selection** and also show that **Subset Selection** is a cap-based mechanism (Definition 3.1). In Appendix H, we show how **Subset Selection** can be simulated using MRC analogous to how we simulated **Subset Selection** using MMRC in Section 5. Along with the theoretical guarantees, we also provide some empirical comparisons between MRC simulating **Subset Selection** and **Subset Selection**. In Appendix I, we provide the proofs of Lemma 5.1 and Theorem 5.1 as well as some additional empirical comparisons between MMRC simulating **Subset Selection** and **Subset Selection**.

A SOCIETAL IMPACT

Collecting large datasets allows us to build better machine learning models which can facilitate our lives in many different ways. However, harnessing data from devices can expose their users to privacy risks. Research into differential privacy can help to minimize these risks. At present, our work is mostly theoretical in nature as there are a few unsolved questions. In particular, for large ε the computational complexity of our approach may be too expensive to be practical.

B MINIMAL RANDOM CODING

Let $q(\mathbf{z}|\mathbf{x})$ be an ε -LDP mechanism for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$. Let $p(\mathbf{z})$ be the fixed reference distribution over \mathcal{Z} and let $\{\mathbf{z}_k\}_{k=1}^N$ be N candidates drawn from $p(\mathbf{z})$. From Algorithm 1, the distribution over the indices $k \in [N]$ under minimal random coding (MRC) is as follows:

$$\pi_{\mathbf{x}}^{\text{mrc}}(k) := \frac{q(\mathbf{z}_k|\mathbf{x})/p(\mathbf{z}_k)}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})}$$

π^{mrc} can be viewed as a function that maps \mathbf{x} and $(\mathbf{z}_1, \dots, \mathbf{z}_N)$ to a distribution in $[N]$. However for notational convenience, when the context is clear, we will omit the dependence on \mathbf{x} and $(\mathbf{z}_1, \dots, \mathbf{z}_N)$.

Let q^{mrc} denote the distribution of \mathbf{z}_K where $K \sim \pi^{\text{mrc}}(\cdot)$ i.e., with $\delta(\cdot)$ denoting the Dirac delta function:

$$q^{\text{mrc}}(\mathbf{z}|\mathbf{x}) := \sum_k \pi^{\text{mrc}}(k) \delta(\mathbf{z} - \mathbf{z}_k).$$

B.1 Utility of MRC

In this section, we prove Theorem 3.1 i.e., we show that MRC can simulate any ε -LDP mechanism in a nearly lossless fashion with about ε bits of communication

Theorem 3.1 (Utility of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP mechanism $q(\cdot|\mathbf{x})$. Consider any reference distribution $p(\cdot)$ such that $|\ln(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \forall \mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$.¹⁰ Let the number of candidates be $N = 2^{(\log e + 4c)\varepsilon}$ for some constant $c \geq 0$. Then, for $\alpha \in [0, 1/2]$, q^{mrc} is such that*

$$\left| \mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|^2] - \mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|^2] \right| \leq \frac{2\alpha \sqrt{\mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|^4]}}{1 - \alpha} \quad (5)$$

holds with probability at least $1 - 2\alpha$, with c and α related by the following: $\alpha = \sqrt{2^{-c\varepsilon} + 2^{-c^2/\log e + 1}}$.

Proof. In order to prove this theorem, we invoke Theorem 3.2 of Havasi et al. (2019).

Recall Theorem 3.2 (Havasi et al., 2019): Let q' and p be distributions over \mathcal{Z} . Let $t \geq 0$ be some constant and let $N' = 2^{(D_{\text{KL}}(q'(\mathbf{z}) \| p(\mathbf{z})) + t)}$. Let \tilde{q} be a discrete distribution constructed by drawing N' samples $\{\mathbf{z}_k\}_{k=1}^{N'}$ from p and defining

$$\tilde{q}(\mathbf{z}) := \sum_{k=1}^{N'} \frac{q(\mathbf{z}_k)/p(\mathbf{z}_k)}{\sum_{k'} q(\mathbf{z}_{k'})/p(\mathbf{z}_{k'})} \delta(\mathbf{z} - \mathbf{z}_k).$$

Furthermore, let f be a measurable function and $\|f\|_{q'} = \sqrt{\mathbb{E}_{q'(\mathbf{z})}[f^2(\mathbf{z})]}$ be its 2-norm under q' . Then it holds that

$$\mathbb{P}\left(\left|\mathbb{E}_{\tilde{q}(\mathbf{z})}[f(\mathbf{z})] - \mathbb{E}_{q'(\mathbf{z})}[f(\mathbf{z})]\right| \geq \frac{2\|f\|_{q'}\alpha'}{1 - \alpha'}\right) \leq 2\alpha'$$

where

$$\alpha' = \sqrt{2^{-t/4} + 2\sqrt{\mathbb{P}(\log(q'(\mathbf{z})/p(\mathbf{z})) > D_{\text{KL}}(q'(\mathbf{z}) \| p(\mathbf{z})) + t/2)}}.$$

We apply Theorem 3.2 (Havasi et al., 2019) to $q'(\mathbf{z}) := q(\mathbf{z}|\mathbf{x})$ and $f(\mathbf{z}) := \|\mathbf{z} - \mathbf{x}\|^2$. We identify $\tilde{q}(\mathbf{z}) = q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ and $N' = N$. To prove Theorem 3.1, it suffices to show that $\alpha \geq \alpha'$. Note that

$$D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) \stackrel{(a)}{=} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \left(\frac{q(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} \right) \right] \stackrel{(b)}{\leq} \varepsilon \log e,$$

where (a) follows the definition of KL-divergence and (b) follows since $|\log(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \log e$ by the assumption on p . We therefore have

$$t = (\log e + 4c)\varepsilon - D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) \geq 4c\varepsilon.$$

It follows that

$$\begin{aligned} \mathbb{P}(\log(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) > D_{\text{KL}}(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) + t/2) &\leq \mathbb{P}(\log(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) > \mathbb{E}[\log(q(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z}))] + 2c\varepsilon) \\ &\stackrel{(b)}{\leq} \exp(-2c^2/(\log e)^2) = 2^{-2c^2/\log e}. \end{aligned}$$

where (b) follows from Hoeffding's inequality since $|\log(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \log e$ by the assumption on p . Therefore,

$$\alpha' \leq \sqrt{2^{-c\varepsilon} + 2\sqrt{2^{-2c^2/\log e}}} = \sqrt{2^{-c\varepsilon} + 2^{-c^2/\log e + 1}} = \alpha.$$

□

Remark B.1. *For most ε -LDP mechanisms $q(\cdot|\mathbf{x})$, the term $\mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|^4]$ in (5) can be well-controlled. For instance, for Subset Selection and PrivUnit₂, the output spaces are bounded, and therefore, $\sqrt{\mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|^4]}$ is of the same order as $\mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|^2]$. Therefore, by making α small enough (in Theorem 3.1) i.e. by increasing c , the estimation error of MRC can be arbitrarily close to the estimation error of the underlying scheme it is simulating.*

¹⁰Note that this condition holds for many reference distributions $p(\cdot)$. For example, one can simply choose $p(\cdot) = q(\cdot|\mathbf{x}^*)$ for some $\mathbf{x}^* \in \mathcal{X}$.

B.2 Privacy of MRC

B.2.1 Pure Privacy of MRC

In this section, we prove Theorem 3.2 i.e., we show that π^{mrc} is a 2ε -LDP mechanism.

Theorem 3.2 (Pure DP guarantee of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , and data $\mathbf{x} \in \mathcal{X}$. Consider any ε -LDP mechanism $q(\cdot|\mathbf{x})$, reference distribution $p(\cdot)$, and number of candidates $N \geq 1$. Then, $\pi^{\text{mrc}}(\cdot)$ obtained from Algorithm 1 is a 2ε -LDP mechanism.*

Proof. For any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$, using the definition of an ε -LDP mechanism, we have

$$q(\mathbf{z}|\mathbf{x}) \leq \exp(\varepsilon)q(\mathbf{z}|\mathbf{x}'). \quad (14)$$

For any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}, \{\mathbf{z}_k\}_{k=1}^N \in \mathcal{Z}^N$ and $k \in [N]$, we have

$$\begin{aligned} \frac{\pi_{\mathbf{x}}^{\text{mrc}}(k)}{\pi_{\mathbf{x}'}^{\text{mrc}}(k)} &\stackrel{(a)}{=} \frac{q(\mathbf{z}_k|\mathbf{x})}{q(\mathbf{z}_k|\mathbf{x}')} \times \frac{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x}')/p(\mathbf{z}_{k'})}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})} \\ &\stackrel{(b)}{\leq} \exp(\varepsilon) \times \frac{\sum_{k'} \exp(\varepsilon)q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})} \\ &= \exp(\varepsilon) \times \frac{\exp(\varepsilon) \sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})} \\ &= \exp(2\varepsilon). \end{aligned}$$

where (a) follows from the definition of π^{mrc} and (b) follows from (14). \square

B.2.2 Approximate Privacy of MRC

In this section, we prove Theorem 3.3 i.e., we provide the approximate DP guarantee of π^{mrc} .

Theorem 3.3 (Approximate DP guarantee of MRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP mechanism $q(\cdot|\mathbf{x})$. Consider any reference distribution $p(\cdot)$ such that $|\ln(q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}))| \leq \varepsilon \forall \mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$.⁵ Let $c_0 \geq 0$ be some constant and let the number of candidates $N = \exp(2\varepsilon + 2c_0)$. Then, for any $\delta \leq 1$, $\pi^{\text{mrc}}(\cdot)$ obtained from Algorithm 1 is $(\varepsilon + \varepsilon_0, \delta)$ -LDP mechanism where*

$$\varepsilon_0 := \ln \frac{1 + a_0}{1 - a_0} \quad \text{and} \quad a_0 := \exp(-c_0) \sqrt{\frac{1}{2} \ln \frac{2}{\delta}}.$$

Proof. Fix any $\mathbf{x} \in \mathcal{X}$. Let us define the following random variable:

$$w(\mathbf{z}|\mathbf{x}) = q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}). \quad (15)$$

Assuming $\mathbf{z} \sim p(\cdot)$, the expected value of the random variable $w(\mathbf{z}|\mathbf{x})$ is

$$\mathbb{E}_p[w(\mathbf{z}|\mathbf{x})] = \mathbb{E}_p[q(\mathbf{z}|\mathbf{x})/p(\mathbf{z})] = \int_{\mathbf{z} \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}) = 1.$$

Further, the random variable $w(\mathbf{z}|\mathbf{x})$ can be bounded as follows:

$$|w(\mathbf{z}|\mathbf{x})| = |q(\mathbf{z}|\mathbf{x})/p(\mathbf{z})| \stackrel{(a)}{\leq} \exp(\varepsilon).$$

where (a) follows from the assumption on $p(\cdot)$. Therefore, we have

$$\mathbb{P}\left(\left|\frac{1}{N} \sum_{k=1}^N w(\mathbf{z}_k|\mathbf{x}) - 1\right| \geq a_0\right) \stackrel{(a)}{\leq} 2 \exp\left(\frac{-2Na_0^2}{(\exp(\varepsilon) - \exp(-\varepsilon))^2}\right) \leq 2 \exp\left(\frac{-2Na_0^2}{\exp(2\varepsilon)}\right) \stackrel{(b)}{=} \delta \quad (16)$$

where (a) follows from Hoeffding's inequality and (b) follows from the definition of a_0 and N . Now, for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}, \{\mathbf{z}_k\}_{k=1}^N \in \mathcal{Z}^N$ and $k \in [N]$, we have

$$\frac{\pi_{\mathbf{x}}^{\text{mrc}}(k)}{\pi_{\mathbf{x}'}^{\text{mrc}}(k)} \stackrel{(a)}{=} \frac{q(\mathbf{z}_k|\mathbf{x})}{q(\mathbf{z}_k|\mathbf{x}')} \times \frac{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x}')/p(\mathbf{z}_{k'})}{\sum_{k'} q(\mathbf{z}_{k'}|\mathbf{x})/p(\mathbf{z}_{k'})}$$

$$\begin{aligned}
 & \stackrel{(b)}{=} \frac{q(\mathbf{z}_k|\mathbf{x})}{q(\mathbf{z}_k|\mathbf{x}')} \times \frac{\sum_{k'} w(\mathbf{z}_{k'}|\mathbf{x}')}{\sum_{k'} w(\mathbf{z}_{k'}|\mathbf{x})} \\
 & \stackrel{(c)}{\leq} \exp(\varepsilon) \times \frac{\frac{1}{N} \sum_{k'} w(\mathbf{z}_{k'}|\mathbf{x}')}{\frac{1}{N} \sum_{k'} w(\mathbf{z}_{k'}|\mathbf{x})}
 \end{aligned} \tag{17}$$

where (a) follows from the definition of π^{mrc} , (b) follows from (15) and (c) follows from (14). Now, using (16) in (17), we have with probability at least $1 - \delta$:

$$\frac{\pi_{\mathbf{x}}^{\text{mrc}}(k)}{\pi_{\mathbf{x}'}^{\text{mrc}}(k)} \leq \exp(\varepsilon) \times \frac{1 + a_0}{1 - a_0} \stackrel{(a)}{=} \exp(\varepsilon + \varepsilon_0)$$

where (a) follows from the definition of ε_0 . □

C MODIFIED MINIMAL RANDOM CODING

Let $q(\mathbf{z}|\mathbf{x})$ be an ε -LDP cap-based mechanism (see definition 3.1) for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$. Let $p(\mathbf{z})$ be the uniform distribution over \mathcal{Z} and let $\{\mathbf{z}_k\}_{k=1}^N$ be N candidates drawn from $p(\mathbf{z})$. Let θ denote the fraction of candidates inside the $\text{Cap}_{\mathbf{x}}$ associated with $q(\mathbf{z}|\mathbf{x})$. Let π^{mmrc} be the distribution over the indices $k \in [N]$ under modified minimal random coding (MMRC) obtained from Algorithm 2. Recall that $\pi^{\text{mmrc}}(k)$ is bounded by an upper threshold t_u and a lower threshold t_l (Section 3.2),

$$t_u = \frac{1}{N} \times \frac{c_1(\varepsilon, d)}{\mathbb{E}[\theta]c_1(\varepsilon, d) + (1 - \mathbb{E}[\theta])c_2(\varepsilon, d)}, \quad t_l = \frac{1}{N} \times \frac{c_2(\varepsilon, d)}{\mathbb{E}[\theta]c_1(\varepsilon, d) + (1 - \mathbb{E}[\theta])c_2(\varepsilon, d)}.$$

Similar to π^{mrc} , π^{mmrc} can be viewed as a function that maps \mathbf{x} and $(\mathbf{z}_1, \dots, \mathbf{z}_N)$ to a distribution in $[N]$. However, to reduce clutter, we will generally omit the dependence on \mathbf{x} and $(\mathbf{z}_1, \dots, \mathbf{z}_N)$. Further, since π^{mmrc} depends on $(\mathbf{z}_1, \dots, \mathbf{z}_N)$ only through θ , we will sometimes show this dependence as $\pi_{\mathbf{x}, \theta}^{\text{mmrc}}$.

Finally, let q^{mmrc} denote the distribution of \mathbf{z}_K where $K \sim \pi^{\text{mmrc}}$. That is, with δ denoting the Dirac delta function:

$$q^{\text{mmrc}}(\mathbf{z}|\mathbf{x}) := \sum_k \pi^{\text{mmrc}}(k) \delta(\mathbf{z} - \mathbf{z}_k).$$

C.1 Privacy of MMRC

In this section, we prove Theorem 3.4 i.e., we show that π^{mmrc} is a ε -LDP mechanism.

Theorem 3.4 (DP guarantee of MMRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$. Let the reference distribution $p(\cdot)$ be the uniform distribution on \mathcal{Z} . Consider any number of candidates $N \geq 1$. Then, $\pi^{\text{mmrc}}(\cdot)$ obtained from Algorithm 2 is an ε -LDP mechanism.*

Proof. For any ε -LDP cap-based $q(\cdot|\mathbf{x})$, we have the following from (1) and (6):

$$\frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \leq \exp(\varepsilon). \tag{18}$$

Further, the modification of π^{mrc} to π^{mmrc} ensures that (8) is true, that is,

$$t_l \leq \pi^{\text{mmrc}}(k) \leq t_u \quad \forall k \in [N].$$

Therefore, for any $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ and $k \in [N]$, we have

$$\frac{\pi_{\mathbf{x}}^{\text{mmrc}}(k)}{\pi_{\mathbf{x}'}^{\text{mmrc}}(k)} \leq \frac{t_u}{t_l} \stackrel{(a)}{=} \frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \stackrel{(b)}{\leq} \exp(\varepsilon),$$

where (a) follows from the definitions of t_u and t_l and (b) follows from (18). □

C.2 Supporting Lemmas to prove the utility of MMRC

To prove Theorem 3.5 (Section C.3), we prove that the expected KL divergence between π^{mrc} and π^{mmrc} can be controlled arbitrarily when the number of candidates is of the right order (Lemma C.2). To prove Lemma C.2, we first show that the KL divergence between π^{mrc} and π^{mmrc} , for a given fraction of candidates inside the $\text{Cap}_{\mathbf{x}}$, can be bounded in terms of ε (Lemma C.1).

C.2.1 The KL divergence between π^{mrc} and π^{mmrc} is small

Lemma C.1. *Let $q(\mathbf{z}|\mathbf{x})$ be an ε -LDP cap-based mechanism. Let $p(\mathbf{z})$ be the uniform distribution over \mathcal{Z} and let $\{\mathbf{z}_k\}_{k=1}^N$ be N candidates drawn from $p(\mathbf{z})$. Let θ denote the fraction of candidates inside the $\text{Cap}_{\mathbf{x}}$ associated with $q(\mathbf{z}|\mathbf{x})$. Let π^{mrc} be the distribution over the indices $k \in [N]$ under MRC obtained from Algorithm 1 and π^{mmrc} be the distribution over the indices $k \in [N]$ under MMRC obtained from Algorithm 2. Then,*

$$D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) \leq \varepsilon \log e$$

Proof. We consider three different cases depending on whether $\theta = \mathbb{E}[\theta]$, $\theta < \mathbb{E}[\theta]$ or $\theta > \mathbb{E}[\theta]$.

1. For $\theta = \mathbb{E}[\theta]$, we have $\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) = \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)$. Therefore,

$$D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) = D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot)) = 0 \leq \varepsilon \log e. \quad (19)$$

2. If $\theta < \mathbb{E}[\theta]$, then π^{mrc} violates the upper threshold t_u so that $\pi^{\text{mmrc}}(k) = t_u$ for all $k \in \text{Cap}_{\mathbf{x}}$ and we have

$$\begin{aligned} D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) &= \sum_k \pi_{\mathbf{x},\theta}^{\text{mrc}}(k) \log \frac{\pi_{\mathbf{x},\theta}^{\text{mrc}}(k)}{\pi_{\mathbf{x},\theta}^{\text{mmrc}}(k)} \\ &\stackrel{(a)}{=} \sum_{k \in \text{Cap}_{\mathbf{x}}} \frac{1}{N} \times \frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \sum_{k \notin \text{Cap}_{\mathbf{x}}} \frac{1}{N} \times \\ &\quad \frac{c_2(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \left[\log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \right. \\ &\quad \left. + \log \frac{(1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)}{(1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d) + (\mathbb{E}[\theta] - \theta)c_1(\varepsilon, d)} \right] \\ &\stackrel{(b)}{=} \frac{\theta c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \\ &\quad \frac{(1 - \theta)c_2(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \left[\log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \right. \\ &\quad \left. + \log \frac{(1 - \theta) \times c_2(\varepsilon, d)}{(1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d) + (\mathbb{E}[\theta] - \theta)c_1(\varepsilon, d)} \right] \\ &= \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \\ &\quad + \frac{(1 - \theta)c_2(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \left[\log \frac{(1 - \theta) \times c_2(\varepsilon, d)}{(1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d) + (\mathbb{E}[\theta] - \theta)c_1(\varepsilon, d)} \right] \\ &\stackrel{(c)}{\leq} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \stackrel{(d)}{\leq} \log \left(\frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d)} \right) \\ &\stackrel{(e)}{\leq} \log \frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \stackrel{(f)}{\leq} \varepsilon \log e \end{aligned} \quad (20)$$

where (a) follows from the definition of $\pi_{\mathbf{x},\theta}^{\text{mrc}}(k)$ and $\pi_{\mathbf{x},\theta}^{\text{mmrc}}(k)$, (b) follows because $|\{k : k \in \text{Cap}_{\mathbf{x}}\}| = \theta N$ and $|\{k : k \notin \text{Cap}_{\mathbf{x}}\}| = (1 - \theta)N$, (c) follows because $\log \frac{(1 - \theta) \times c_2(\varepsilon, d)}{(1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d) + (\mathbb{E}[\theta] - \theta)c_1(\varepsilon, d)} \leq 0$, (d) follows because $\theta \geq 0$, (e) follows because $\mathbb{E}[\theta] \leq 1$, and (f) follows because $c_1(\varepsilon, d)/c_2(\varepsilon, d) \leq \exp(\varepsilon)$.

3. For $\theta > \mathbb{E}[\theta]$, we have

$$\begin{aligned}
 D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) &= \sum_{z_i} \pi_{\mathbf{x},\theta}^{\text{mrc}}(k) \log \frac{\pi_{\mathbf{x},\theta}^{\text{mrc}}(k)}{\pi_{\mathbf{x},\theta}^{\text{mmrc}}(k)} \\
 &\stackrel{(a)}{=} \sum_{k \notin \text{Cap}_{\mathbf{x}}} \frac{1}{N} \times \frac{c_2(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \sum_{k \in \text{Cap}_{\mathbf{x}}} \frac{1}{N} \times \\
 &\quad \frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \left[\log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \right. \\
 &\quad \left. \log \frac{\theta c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right] \\
 &\stackrel{(b)}{=} \frac{(1 - \theta) \times c_2(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \\
 &\quad \frac{\theta c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \left[\log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} + \right. \\
 &\quad \left. \log \frac{\theta c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right] \\
 &\stackrel{(c)}{\leq} \frac{\theta c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \left(\frac{\theta c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right) \\
 &\stackrel{(d)}{\leq} \log \left(\frac{c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (1 - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right) \stackrel{(e)}{\leq} \log \frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \stackrel{(f)}{\leq} \varepsilon \log e
 \end{aligned} \tag{21}$$

where (a) follows from the definition of $\pi_{\mathbf{x},\theta}^{\text{mrc}}(k)$ and $\pi_{\mathbf{x},\theta}^{\text{mmrc}}(k)$, (b) follows because $|\{k : k \in \text{Cap}_{\mathbf{x}}\}| = \theta N$ and $|\{k : k \notin \text{Cap}_{\mathbf{x}}\}| = (1 - \theta)N$, (c) follows because $\log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \leq 0$, (d) follows because $\theta \leq 1$, (e) follows because $\mathbb{E}[\theta] \geq 0$, and (f) follows because $c_1(\varepsilon, d)/c_2(\varepsilon, d) \leq \exp(\varepsilon)$. \square

C.2.2 The expected KL divergence between the distribution of indices in MRC and MMRC can be controlled arbitrarily when N is in the right order

Lemma C.2. *Let $q(\mathbf{z}|\mathbf{x})$ be an ε -LDP cap-based mechanism. Let $p(\mathbf{z})$ be the uniform distribution over \mathcal{Z} and let $\{z_k\}_{k=1}^N$ be N candidates drawn from $p(\mathbf{z})$. Let θ denote the fraction of candidates inside the $\text{Cap}_{\mathbf{x}}$ associated with $q(\mathbf{z}|\mathbf{x})$. Let π^{mrc} be the distribution over the indices $k \in [N]$ under MRC obtained from Algorithm 1 and π^{mmrc} be the distribution over the indices $k \in [N]$ under MMRC obtained from Algorithm 2. Then,*

$$\mathbb{E}_{\theta} [D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot))] \leq \rho \times \log e \times (1 + \varepsilon)$$

where $\rho \in (0, 1)$ is a free variable that is related to N as follows:

$$N = \frac{2(\exp(\varepsilon) - 1)^2}{\rho^2} \ln \frac{2}{\rho}.$$

Proof. Let θ denote the fraction of candidates inside the cap, i.e.,

$$\theta = \frac{1}{N} \sum_{k=1}^N \mathbb{1}(z_k \in \text{Cap}_{\mathbf{x}}).$$

Therefore, we have

$$\mathbb{E}[\theta] = \mathbb{P}_{z_k \sim \text{Unif}(\mathcal{Z})}(z_k \in \text{Cap}_{\mathbf{x}}) = \mathbb{P}_{z \sim \text{Unif}(\mathcal{Z})}(z \in \text{Cap}_{\mathbf{x}}). \tag{22}$$

Now, using the Hoeffding's inequality, we have $\mathbb{P}\left\{|\theta - \mathbb{E}[\theta]| \geq \sqrt{\frac{\ln(2/\rho)}{2N}}\right\} \leq \rho$. Letting $\hat{\rho} = \sqrt{\frac{\ln(2/\rho)}{2N}}$, we have

$$\begin{aligned} \mathbb{E}_\theta [D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot))] &= \sum_{\theta} \mathbb{P}(\theta) \times D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) \\ &= \sum_{\theta: |\theta - \mathbb{E}[\theta]| \leq \hat{\rho}} \mathbb{P}(\theta) D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) + \sum_{\theta: |\theta - \mathbb{E}[\theta]| > \hat{\rho}} \mathbb{P}(\theta) D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) \end{aligned} \quad (23)$$

Now, we will upper bound $D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot))$ whenever θ is such that $|\theta - \mathbb{E}[\theta]| \leq \hat{\rho}$. As in the proof of Lemma C.1, we have three different cases depending on whether $\theta = \mathbb{E}[\theta]$, $\theta < \mathbb{E}[\theta]$ or $\theta > \mathbb{E}[\theta]$.

1. For $\theta = \mathbb{E}[\theta]$, using (19), we have $D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) = 0$.
2. For $\theta < \mathbb{E}[\theta]$, using (20), we have

$$\begin{aligned} D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) &\leq \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + \theta \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \\ &\stackrel{(a)}{=} \log \frac{c_2(\varepsilon, d) + \mathbb{E}[\theta] \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + (\mathbb{E}[\theta] - t) \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \\ &= \log \left(1 + \frac{t \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + (\mathbb{E}[\theta] - t) \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \right) \\ &\stackrel{(b)}{\leq} \frac{\log e \times t \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + (\mathbb{E}[\theta] - t) \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} \\ &\stackrel{(c)}{\leq} \log e \times t \times \left(\frac{c_1(\varepsilon, d) - c_2(\varepsilon, d)}{c_2(\varepsilon, d)} \right) \stackrel{(d)}{\leq} \log e \times \hat{\rho} \times \left(\frac{c_1(\varepsilon, d) - c_2(\varepsilon, d)}{c_2(\varepsilon, d)} \right) \end{aligned} \quad (24)$$

where (a) follows by letting $\theta = \mathbb{E}[\theta] - t$ with $t > 0$, (b) follows by using $\log(1+x) \leq x \log e$ for $x = \frac{t \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d) + (\mathbb{E}[\theta] - t) \times (c_1(\varepsilon, d) - c_2(\varepsilon, d))} > 0$, (c) follows because $\mathbb{E}[\theta] - t = \theta \geq 0$, and (d) follows because $t = \mathbb{E}[\theta] - \theta \leq \hat{\rho}$.

3. For $\theta > \mathbb{E}[\theta]$, using (21), we have

$$\begin{aligned} D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) &\leq \frac{\theta c_1(\varepsilon, d)}{c_2(\varepsilon, d) + \theta(c_1(\varepsilon, d) - c_2(\varepsilon, d))} \log \left(\frac{\theta c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right) \\ &\stackrel{(a)}{\leq} \log \left(\frac{\theta c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \times c_2(\varepsilon, d)} \right) \\ &\stackrel{(b)}{=} \log \left(\frac{(\mathbb{E}[\theta] + t) c_1(\varepsilon, d)}{\mathbb{E}[\theta] c_1(\varepsilon, d) + t c_2(\varepsilon, d)} \right) \\ &= \log \left(1 + \frac{t(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\mathbb{E}[\theta] c_1(\varepsilon, d) + t c_2(\varepsilon, d)} \right) \\ &\stackrel{(c)}{\leq} \frac{\log e \times t(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\mathbb{E}[\theta] c_1(\varepsilon, d) + t c_2(\varepsilon, d)} \\ &\stackrel{(d)}{\leq} \frac{\log e \times t(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\mathbb{E}[\theta] c_1(\varepsilon, d)} \stackrel{(e)}{\leq} \frac{\log e \times \hat{\rho}(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\mathbb{E}[\theta] c_1(\varepsilon, d)} \end{aligned} \quad (25)$$

where (a) follows because $\theta \leq 1$, (b) follows by letting $\theta = \mathbb{E}[\theta] + t$ with $t > 0$, (c) follows by using $\log(1+x) \leq x \log e$ for $x = \frac{t(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\mathbb{E}[\theta] c_1(\varepsilon, d) + t c_2(\varepsilon, d)} > 0$, (d) follows because $t > 0$, and (e) follows because $t = \theta - \mathbb{E}[\theta] \leq \hat{\rho}$.

Therefore, for θ such that $|\theta - \mathbb{E}[\theta]| \leq \hat{\rho}$, we have the following from (24) and (25):

$$D_{\text{KL}}(\pi_{\mathbf{x},\theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x},\theta}^{\text{mmrc}}(\cdot)) \leq \frac{\log e \times \hat{\rho}(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\min\{c_2(\varepsilon, d), \mathbb{E}[\theta] c_1(\varepsilon, d)\}} \stackrel{(a)}{=} \frac{\log e \times \hat{\rho}(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{\min\{c_2(\varepsilon, d), c_1(\varepsilon, d) \mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}})\}}$$

$$\begin{aligned}
 &\stackrel{(b)}{\leq} \frac{2 \log e \times \hat{\rho}(c_1(\varepsilon, d) - c_2(\varepsilon, d))}{c_2(\varepsilon, d)} \\
 &\stackrel{(c)}{\leq} 2 \log e \times \hat{\rho}(\exp(\varepsilon) - 1)
 \end{aligned} \tag{26}$$

where (a) follows from (22), (b) follows because $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}}) \geq c_2(\varepsilon, d)/2c_1(\varepsilon, d)$ from the definition of cap-based mechanisms, and (c) follows because $c_1(\varepsilon, d)/c_2(\varepsilon, d) \leq \exp(\varepsilon)$.

Using (26) and Lemma C.1 in (23), we have

$$\begin{aligned}
 \mathbb{E}_{\theta} [D_{\text{KL}}(\pi_{\mathbf{x}, \theta}^{\text{mrc}}(\cdot) \parallel \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(\cdot))] &\leq \sum_{\theta: |\theta - \mathbb{E}[\theta]| \leq \hat{\rho}} \mathbb{P}(\theta) \times 2 \log e \times \hat{\rho}(\exp(\varepsilon) - 1) + \sum_{\theta: |\theta - \mathbb{E}[\theta]| > \hat{\rho}} \mathbb{P}(\theta) \times \varepsilon \log e \\
 &\stackrel{(a)}{\leq} 2 \log e \times \hat{\rho}(\exp(\varepsilon) - 1) + \rho \varepsilon \log e \\
 &\stackrel{(b)}{\leq} 2 \log e \times \sqrt{\frac{\ln(2/\rho)}{2N}} (\exp(\varepsilon) - 1) + \rho \varepsilon \log e \\
 &\stackrel{(c)}{\leq} \log e \times \rho(1 + \varepsilon)
 \end{aligned}$$

where (a) follows because $\mathbb{P}(|\theta - \mathbb{E}[\theta]| \leq \hat{\rho}) \leq 1$ and $\mathbb{P}(|\theta - \mathbb{E}[\theta]| \geq \hat{\rho}) \leq \rho$, (b) follows by plugging in $\hat{\rho} = \sqrt{\frac{\ln(2/\rho)}{2N}}$, and (c) follows by plugging in N . \square

C.3 Utility of MMRC

In this section, we first prove Theorem 3.5 i.e., we show that, with number of candidates exponential in ε , samples drawn from q^{mmrc} will be similar to the samples drawn from q^{mrc} in terms of ℓ_2 error.

Then, in Theorem C.1, we show that MMRC can simulate any ε -LDP cap-based mechanism in a nearly lossless fashion with about ε bits of communication.

C.3.1 Utility of MMRC with respect to q^{mrc}

Theorem 3.5 (Utility of MMRC). *Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$. Let the reference distribution $p(\cdot)$ be the uniform distribution on \mathcal{Z} . Let N denote the number of candidates. Then, q^{mmrc} is such that*

$$\begin{aligned}
 \mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] &\leq \mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] \\
 &\quad + \sqrt{\frac{\rho(1 + \varepsilon)}{2}} \max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2
 \end{aligned} \tag{9}$$

where $\rho \in (0, 1)$ is such that $N = \frac{2(\exp(\varepsilon) - 1)^2}{\rho^2} \ln \frac{2}{\rho}$.

Proof. We will first upper bound the difference between $\mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2]$ and $\mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2]$ in terms of the total variation distance between q^{mrc} and q^{mmrc} . Due to a property of the total variation distance (e.g., Song et al., 2016), we have

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] - \mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] \leq \max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2 \times \|q^{\text{mrc}}(\mathbf{z}|\mathbf{x}) - q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})\|_{\text{TV}}. \tag{27}$$

Next, we will upper bound the total variation distance between q^{mrc} and q^{mmrc} using Pinsker's inequality as follows:

$$\|q^{\text{mrc}}(\mathbf{z}|\mathbf{x}) - q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})\|_{\text{TV}} \leq \sqrt{\frac{1}{2 \log e} D_{\text{KL}}(q^{\text{mrc}}(\mathbf{z}|\mathbf{x}) \parallel q^{\text{mmrc}}(\mathbf{z}|\mathbf{x}))}. \tag{28}$$

Next, we will upper bound the KL divergence between $q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ and $q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})$. To that end, for every $\mathbf{x} \in \mathcal{X}$, let $p^{\text{mrc}}(z_1, \dots, z_N, K, z_K|\mathbf{x})$ denote the joint distribution of the candidates z_1, \dots, z_N drawn from $p(\mathbf{z})$, the transmitted index K under MRC, and the sample z_K corresponding to K . We have

$$p^{\text{mrc}}(z_1, \dots, z_N, K, z_K|\mathbf{x}) = p(z_1, \dots, z_N|\mathbf{x}) \times p^{\text{mrc}}(K|z_1, \dots, z_N, \mathbf{x}) \times p^{\text{mrc}}(z_K|z_1, \dots, z_N, K, \mathbf{x})$$

$$\begin{aligned}
 &\stackrel{(a)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times p^{\text{mrc}}(K|\mathbf{z}_1, \dots, \mathbf{z}_N, \mathbf{x}) \times p^{\text{mrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) \\
 &\stackrel{(b)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mrc}}(k) \times p^{\text{mrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) \\
 &\stackrel{(c)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mrc}}(k)
 \end{aligned} \tag{29}$$

where (a) follows because $\mathbf{z}_1, \dots, \mathbf{z}_N$ are independent of \mathbf{x} , (b) follows because $p^{\text{mrc}}(K|\mathbf{z}_1, \dots, \mathbf{z}_N, \mathbf{x}) = \pi_{\mathbf{x}, \theta}^{\text{mrc}}(k)$, and (c) follows because $p^{\text{mrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) = 1$ (note that \mathbf{z}_K can be viewed as a function of $(\mathbf{z}_1, \dots, \mathbf{z}_N, K)$).

Similarly, for every $\mathbf{x} \in \mathcal{X}$, let $p^{\text{mmrc}}(\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{z}_K|\mathbf{x})$ denote the joint distribution of the candidates $\mathbf{z}_1, \dots, \mathbf{z}_N$ drawn from $p(\mathbf{z})$, the transmitted index K under MMRC, and the sample \mathbf{z}_K corresponding to K . We have

$$\begin{aligned}
 p^{\text{mmrc}}(\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{z}_K|\mathbf{x}) &= p(\mathbf{z}_1, \dots, \mathbf{z}_N|\mathbf{x}) \times p^{\text{mmrc}}(K|\mathbf{z}_1, \dots, \mathbf{z}_N, \mathbf{x}) \times p^{\text{mmrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) \\
 &\stackrel{(a)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times p^{\text{mmrc}}(K|\mathbf{z}_1, \dots, \mathbf{z}_N, \mathbf{x}) \times p^{\text{mmrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) \\
 &\stackrel{(b)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(k) \times p^{\text{mmrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) \\
 &\stackrel{(c)}{=} p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(k)
 \end{aligned} \tag{30}$$

where (a) follows because $\mathbf{z}_1, \dots, \mathbf{z}_N$ are independent of \mathbf{x} , (b) follows because $p^{\text{mmrc}}(K|\mathbf{z}_1, \dots, \mathbf{z}_N, \mathbf{x}) = \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(k)$, and (c) follows because $p^{\text{mmrc}}(\mathbf{z}_K|\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{x}) = 1$.

We are now in a position to upper bound the KL divergence between $q^{\text{mrc}}(\mathbf{z}_K|\mathbf{x})$ and $q^{\text{mmrc}}(\mathbf{z}_K|\mathbf{x})$:

$$\begin{aligned}
 D_{\text{KL}}(q^{\text{mrc}}(\mathbf{z}|\mathbf{x}) \| q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})) &\stackrel{(a)}{\leq} D_{\text{KL}}(p^{\text{mrc}}(\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{z}_K|\mathbf{x}) \| p^{\text{mmrc}}(\mathbf{z}_1, \dots, \mathbf{z}_N, K, \mathbf{z}_K|\mathbf{x})) \\
 &\stackrel{(b)}{=} D_{\text{KL}}(p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mrc}}(k) \| p(\mathbf{z}_1, \dots, \mathbf{z}_N) \times \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(k)) \\
 &\stackrel{(c)}{=} \mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} [D_{\text{KL}}(\pi_{\mathbf{x}, \theta}^{\text{mrc}}(k) \| \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(k))] \\
 &\stackrel{(d)}{=} \mathbb{E}_{\theta} [D_{\text{KL}}(\pi_{\mathbf{x}, \theta}^{\text{mrc}}(\cdot) \| \pi_{\mathbf{x}, \theta}^{\text{mmrc}}(\cdot))] \stackrel{(e)}{\leq} \rho \times \log e \times (1 + \varepsilon)
 \end{aligned} \tag{31}$$

where (a) follows because by the chain rule for KL-divergence, (b) follows from (29) and (30), (c) follows by the definition of KL-divergence, (d) follows because π^{mrc} and π^{mmrc} depend on $\mathbf{z}_1, \dots, \mathbf{z}_N$ only via θ for cap-based mechanisms, and (e) follows from Lemma C.2 because $N = \frac{2(\exp(\varepsilon)-1)^2}{\rho^2} \ln \frac{2}{\rho}$. Combining (27), (28), and (31), we have

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] \leq \mathbb{E}_{q^{\text{mrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] + \sqrt{\frac{\rho(1+\varepsilon)}{2}} \times \max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2.$$

□

Remark C.1. For bounded ε -LDP mechanisms such as `PrivUnit2` and `Subset Selection`, the term $\max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2$ in (9) is of the same order as $\mathbb{E}_q[\|\mathbf{z} - \mathbf{x}\|_2^2]$. Therefore, by picking a large N in Theorem 3.5 (i.e. $\log N \geq C\varepsilon$ for a sufficiently large C), ρ can be made arbitrarily small and the estimation error of MMRC can be arbitrarily close to the estimation error of MRC.

C.3.2 Utility of MMRC with respect to q

Theorem C.1. Consider any input alphabet \mathcal{X} , output alphabet \mathcal{Z} , data $\mathbf{x} \in \mathcal{X}$, and ε -LDP cap-based mechanism $q(\cdot|\mathbf{x})$. Let the reference distribution $p(\cdot)$ be the uniform distribution on \mathcal{Z} . Let N denote the number of candidates. Then, q^{mmrc} is such that

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\mathbf{z} - \mathbf{x}\|_2^2] \leq \mathbb{E}_q [\|\mathbf{z} - \mathbf{x}\|_2^2] + \sqrt{\frac{\rho(1+\varepsilon)}{2}} \times \max_{\mathbf{x}, \mathbf{z}} \|\mathbf{z} - \mathbf{x}\|_2^2 + \frac{2\alpha}{1-\alpha} \times \sqrt{\mathbb{E}_q[\|\mathbf{z} - \mathbf{x}\|_4^4]}$$

holds with probability at least $1 - 2\alpha$ where

$$\alpha = \sqrt{2^{-c\varepsilon} + 2^{-c^2/\log e} + 1}.$$

and c and $\rho \in (0, 1)$ are free variables such that

$$N = \max \left\{ 2^{(\log e + 4c)\varepsilon}, \frac{2(\exp(\varepsilon) - 1)^2}{\rho^2} \ln \frac{2}{\rho} \right\}$$

Proof. The proof follows from Theorem 3.1 and Theorem 3.5. \square

D PRELIMINARY ON PrivUnit_2

First, we briefly recap the PrivUnit_2 mechanism (q^{pu}) proposed in Bhowmick et al. (2018). PrivUnit_2 is a private sampling scheme when the input alphabet \mathcal{X} is the d -dimensional unit ℓ_2 sphere \mathbb{S}^{d-1} . More formally, given a vector $\mathbf{x} \in \mathbb{S}^{d-1}$, PrivUnit_2 (see Algorithm 3) draws a vector \mathbf{z} from a spherical cap $\{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\}$ with probability $p_0 \geq 1/2$ or from its complement $\{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle < \gamma\}$ with probability $1 - p_0$, where $\gamma \in [0, 1]$ and p_0 are constants that trade accuracy and privacy. In other words, the conditional density $q^{\text{pu}}(\mathbf{z}|\mathbf{x})$ is:

$$q^{\text{pu}}(\mathbf{z}|\mathbf{x}) = \begin{cases} p_0 \times \frac{2}{A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})} & \text{if } \langle \mathbf{x}, \mathbf{z} \rangle \geq \gamma \\ (1 - p_0) \times \frac{2}{2A(1, d) - A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})} & \text{otherwise} \end{cases} \quad (32)$$

where $A(1, d)$ denotes the area of \mathbb{S}^{d-1} and $I_x(a, b)$ denotes the regularized incomplete beta function.

Algorithm 3: Privatized Unit Vector: PrivUnit_2

Require: $\mathbf{x} \in \mathbb{S}^{d-1}$, $\gamma \in [0, 1]$, $p_0 \geq 1/2$.

Draw random vector \mathbf{z} according to the distribution

$$\mathbf{z} = \begin{cases} \text{uniform on } \{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\} & \text{with probability } p_0 \\ \text{uniform on } \{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle < \gamma\} & \text{otherwise.} \end{cases}$$

Set $\alpha = \frac{d-1}{2}$, $\tau = \frac{1+\gamma}{2}$, and

$$m_{\text{pu}} = \frac{(1 - \gamma^2)^\alpha}{2^{d-2}(d-1)} \left[\frac{p_0}{B(\alpha, \alpha) - B(\tau; \alpha, \alpha)} - \frac{1 - p_0}{B(\tau; \alpha, \alpha)} \right] \quad (33)$$

return $\hat{\mathbf{x}}^{\text{pu}} = \frac{\mathbf{z}}{m_{\text{pu}}}$

Given its inputs \mathbf{x} , γ , and p_0 , Algorithm 3 returns an estimator $\hat{\mathbf{x}}^{\text{pu}} := \mathbf{z}/m_{\text{pu}}$ which is differentially private and unbiased where m_{pu} is a scaling factor. The choice of γ described in Theorem D.1 ensures differential privacy and the choice of the scaling factor m described in (33) ensures unbiasedness where

$$B(x; \alpha, \beta) := \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt \quad \text{where } B(\alpha, \beta) := B(1; \alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha + \beta)}$$

denotes the incomplete beta function.

D.1 PrivUnit_2 is a differentially private mechanism

The following theorem borrowed from Bhowmick et al. (2018) describes the choice of γ and provides the precise associated differential privacy guarantee of the PrivUnit_2 mechanism.

Theorem D.1 (Bhowmick et al. (2018)). *Let $\gamma \in [0, 1]$ and $p_0 = \frac{e^{\varepsilon_0}}{1+e^{\varepsilon_0}}$. Then algorithm $\text{PrivUnit}_2(\cdot, \gamma, p_0)$ is*

$\varepsilon = (\bar{\varepsilon} + \varepsilon_0)$ -differentially private whenever $\gamma \geq 0$ is such that

$$\bar{\varepsilon} \geq \log \frac{1 + \gamma \cdot \sqrt{2(d-1)/\pi}}{\left(1 - \gamma \cdot \sqrt{2(d-1)/\pi}\right)_+}, \quad \text{i.e.} \quad \gamma \leq \frac{e^{\bar{\varepsilon}} - 1}{e^{\bar{\varepsilon}} + 1} \sqrt{\frac{\pi}{2(d-1)}},$$

or

$$\bar{\varepsilon} \geq 1/2 \log(d) + \log 6 - \frac{d-1}{2} \log(1 - \gamma^2) + \log \gamma \quad \text{and} \quad \gamma \geq \sqrt{\frac{2}{d}}.$$

Here, ε can be viewed as the total privacy budget. Typically, μ fraction of this budget is allocated for the spherical cap threshold γ and $1 - \mu$ fraction is allocated to the probability parameter p_0 with which a particular spherical cap is chosen i.e., $\bar{\varepsilon} = \mu\varepsilon$ and $\varepsilon_0 = (1 - \mu)\varepsilon$ for some $\mu \in [0, 1]$. While the parameter μ can be optimized over as described in [Feldman and Talwar \(2021\)](#), we will view it as a constant for convenience. Our results on MRC and MMRC simulating PrivUnit_2 can be easily extended to the setup where μ needs to be optimized over.

D.2 PrivUnit_2 is unbiased and order-optimal

The following lemma borrowed from [Bhowmick et al. \(2018\)](#) shows that the output of the PrivUnit_2 mechanism (a) is unbiased, (b) has a bounded norm, and (c) has order-optimal utility.

Proposition D.1 ([Bhowmick et al. \(2018\)](#)). *Let $\hat{\mathbf{x}}^{p_u} = \text{PrivUnit}_2(\mathbf{x}, \gamma, p_0)$ for some $\mathbf{x} \in \mathbb{S}^{d-1}$, $\gamma \in [0, 1]$, and $p_0 \in [1/2, 1]$. Then, $\mathbb{E}[\hat{\mathbf{x}}^{p_u}] = \mathbf{x}$. Further, assume that $0 \leq \varepsilon \leq d$. Then, there exists a numerical constant $c < \infty$ such that if γ saturates either of the two inequalities (34), then $\gamma \gtrsim \min\{\varepsilon/\sqrt{d}, \sqrt{\varepsilon/d}\}$, and*

$$\|\hat{\mathbf{x}}^{p_u}\|_2 \leq c \cdot \sqrt{\frac{d}{\varepsilon} \vee \frac{d}{(e^\varepsilon - 1)^2}}.$$

Additionally, $\mathbb{E}[\|\hat{\mathbf{x}}^{p_u} - \mathbf{x}\|_2^2] \lesssim \frac{d}{\varepsilon} \vee \frac{d}{(e^\varepsilon - 1)^2}$.

D.3 PrivUnit_2 is a cap-based mechanism

The randomness in the estimator $\hat{\mathbf{x}}^{p_u}$ obtained from the $\text{PrivUnit}_2(\mathbf{x}, \gamma, p_0)$ mechanism comes from \mathbf{z} . Therefore, we obtain a convenient expression for the conditional distribution of \mathbf{z} conditioned on \mathbf{x} i.e., $q^{p_u}(\mathbf{z}|\mathbf{x})$. Define $\text{Cap}_{\mathbf{x}} := \{\mathbf{z} | \langle \mathbf{x}, \mathbf{z} \rangle \geq \gamma\}$. Recall from (34) that γ is a function of ε and d . Further, as described in Section D.1, when the budget split parameter μ is known, p_0 can be viewed as a function of ε . Then, the conditional distribution $q^{p_u}(\mathbf{z}|\mathbf{x})$ in (32) can be written as follows:

$$q^{p_u}(\mathbf{z}|\mathbf{x}) = \begin{cases} c_1(\varepsilon, d) & \text{if } \mathbf{z} \in \text{Cap}_{\mathbf{x}} \\ c_2(\varepsilon, d) & \text{if } \mathbf{z} \notin \text{Cap}_{\mathbf{x}} \end{cases} \quad (35)$$

where $c_1(\varepsilon, d) = p_0 \times \frac{2}{A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})}$ and $c_2(\varepsilon, d) = (1 - p_0) \times \frac{2}{2A(1, d) - A(1, d)I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})}$ are functions of ε and d .

Further, $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \mathcal{Z}_{\mathbf{x}}) = \frac{I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2})}{2}$. Therefore,

$$\frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \times \mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \mathcal{Z}_{\mathbf{x}}) = \frac{p_0 \times (2 - I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2}))}{2(1 - p_0)} \stackrel{(a)}{\geq} \frac{p_0}{2(1 - p_0)} \stackrel{(b)}{\geq} \frac{1}{2}$$

where (a) follows because $I_{1-\gamma^2}(\frac{d-1}{2}, \frac{1}{2}) \leq 1$ and (b) follows because $p_0 \geq 1/2$.

E SIMULATING PrivUnit_2 USING MINIMAL RANDOM CODING

In this section, we simulate PrivUnit_2 using MRC analogous to how we simulate PrivUnit_2 using MMRC in Section 4. First, in Appendix E.1, we provide an unbiased estimator for MRC simulating PrivUnit_2 . Next, in Appendix

E.2 we provide the utility guarantee associated with MRC simulating PrivUnit_2 . To do that, first, in Appendix E.2.1, we show that when the number of candidates N is exponential in ε , the scaling factor m_{mrc} is close to the scaling parameter associated with PrivUnit_2 (i.e., m_{pu}). Next, in Appendix E.2.2, we provide the relationship between the mean squared error associated with MRC simulating PrivUnit_2 and the mean squared error associated with PrivUnit_2 . In Appendix E.2.3, we combine everything and show that, for mean estimation, MRC can simulate PrivUnit_2 in a near-lossless manner while only using on the order of ε -bits of communication. Finally, in Appendix E.3, we provide some empirical comparisons.

E.1 Unbiased Minimal Random Coding simulating PrivUnit_2

Consider the PrivUnit_2 ε -LDP mechanism q^{pu} described in Section 2 with parameters p_0 and γ . PrivUnit_2 is a cap-based mechanism with $\text{Cap}_{\mathbf{x}} = \{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\}$ as discussed in Appendix D. Let π^{mrc} be the distribution and $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ be the candidates obtained from Algorithm 1 when the reference distribution is $\text{Unif}(\mathbb{S}^{d-1})$. Let $K \sim \pi^{\text{mrc}}(\cdot)$. Define $p_{\text{mrc}} := \mathbb{P}(\mathbf{z}_K \in \text{Cap}_{\mathbf{x}})$ to be the probability with which the sampled candidate \mathbf{z}_K belongs to the spherical cap associated with PrivUnit_2 . Define m_{mrc} as the scaling factor in (2) when p_0 in (2) is replaced by p_{mrc} . Define $\hat{\mathbf{x}}^{\text{mrc}} := \mathbf{z}_K / m_{\text{mrc}}$ as the estimator of the MRC mechanism simulating PrivUnit_2 . The following Lemma shows that $\hat{\mathbf{x}}^{\text{mrc}}$ is an unbiased estimator.

Lemma E.1. *Let $\hat{\mathbf{x}}^{\text{mrc}}$ be the estimator of the MRC mechanism simulating PrivUnit_2 as defined above. Then, $\mathbb{E}_{q^{\text{mrc}}}[\hat{\mathbf{x}}^{\text{mrc}}] = \mathbf{x}$.*

Proof. For $k \in [N]$, let $A_k := \mathbf{1}(\mathbf{z}_k \in \text{Cap}_{\mathbf{x}})$. Then, $p_{\text{mrc}} = \mathbb{P}(A_K = 1)$. Using the definition of $\hat{\mathbf{x}}^{\text{mrc}}$, we have

$$\mathbb{E}_{q^{\text{mrc}}}[\hat{\mathbf{x}}^{\text{mrc}}] = \frac{1}{m_{\text{mrc}}} \mathbb{E}_{q^{\text{mrc}}}[\mathbf{z}_K].$$

Let us evaluate $\mathbb{E}_{q^{\text{mrc}}}[\mathbf{z}_K]$. We have

$$\begin{aligned} \mathbb{E}_{q^{\text{mrc}}}[\mathbf{z}_K] &\stackrel{(a)}{=} \mathbb{E}_{K, \mathbf{z}_1, \dots, \mathbf{z}_N}[\mathbf{z}_K] \\ &\stackrel{(b)}{=} \mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} \left[\sum_{k=1}^N \pi_{\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_N}^{\text{mrc}}(k) \times \mathbf{z}_k \right] \\ &\stackrel{(c)}{=} \mathbb{E}_{A_1, \dots, A_N} \left[\mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} \left[\sum_{k=1}^N \pi_{\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_N}^{\text{mrc}}(k) \times \mathbf{z}_k \mid A_1, \dots, A_N \right] \right] \\ &\stackrel{(d)}{=} \sum_{k=1}^N \mathbb{E}_{A_1, \dots, A_N} \left[\mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} \left[\pi_{\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_N}^{\text{mrc}}(k) \times \mathbf{z}_k \mid A_1, \dots, A_N \right] \right] \\ &\stackrel{(e)}{=} \sum_{k=1}^N \mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} \left[\mathbf{z}_k \mid A_1, \dots, A_N \right] \right] \\ &\stackrel{(f)}{=} \sum_{k=1}^N \mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_1, \dots, \mathbf{z}_N} \left[\mathbf{z}_k \mid A_k \right] \right] \\ &\stackrel{(g)}{=} \sum_{k=1}^N \mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} \left[\mathbf{z}_k \mid A_k \right] \right] \\ &\stackrel{(h)}{=} \sum_{k=1}^N \mathbb{E}_{A_k} \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} \left[\mathbf{z}_k \mid A_k \right] \mid A_k \right] \right] \\ &\stackrel{(i)}{=} \sum_{k=1}^N \mathbb{P}(A_k = 1) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} \left[\mathbf{z}_k \mid A_k \right] \mid A_k = 1 \right] \right] \\ &\quad + \sum_{k=1}^N \mathbb{P}(A_k = 0) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} \left[\mathbf{z}_k \mid A_k \right] \mid A_k = 0 \right] \right] \\ &= \sum_{k=1}^N \mathbb{P}(A_k = 1) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_k=1, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} \left[\mathbf{z}_k \mid A_k = 1 \right] \right] \right] \end{aligned}$$

$$\begin{aligned}
 & + \sum_{k=1}^N \mathbb{P}(A_k = 0) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_k=0, \dots, A_N}^{\text{mrc}}(k) \mathbb{E}_{\mathbf{z}_k} [\mathbf{z}_k | A_k = 0] \right] \right] \\
 & \stackrel{(j)}{=} \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 1] \sum_{k=1}^N \mathbb{P}(A_k = 1) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_k=1, \dots, A_N}^{\text{mrc}}(k) \right] \right] \\
 & + \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 0] \sum_{k=1}^N \mathbb{P}(A_k = 0) \left[\mathbb{E}_{A_1, \dots, A_N} \left[\pi_{\mathbf{x}, A_1, \dots, A_k=0, \dots, A_N}^{\text{mrc}}(k) \right] \right] \\
 & \stackrel{(k)}{=} \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 1] \sum_{k=1}^N \mathbb{P}(A_k = 1) \pi_{\mathbf{x}, A_k=1}^{\text{mrc}}(k) \\
 & + \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 0] \sum_{k=1}^N \mathbb{P}(A_k = 0) \pi_{\mathbf{x}, A_k=0}^{\text{mrc}}(k) \\
 & \stackrel{(l)}{=} \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 1] \mathbb{P}(A_K = 1) + \mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 0] \mathbb{P}(A_K = 0) \\
 & \stackrel{(m)}{=} m_{\text{mrc}} \mathbf{x}
 \end{aligned} \tag{36}$$

where (a) follows because the randomness in q^{mrc} comes from the randomness in $K, \mathbf{z}_1, \dots, \mathbf{z}_N$, (b) follows by calculating the expectation over K and showing the dependence of π^{mrc} on $\mathbf{z}_1, \dots, \mathbf{z}_N$ explicitly, (c) follows by the tower property of expectation, (d) follows by linearity of expectation, (e) follows because $\pi_{\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_N}^{\text{mrc}}(k) = \pi_{\mathbf{x}, A_1, \dots, A_N}^{\text{mrc}}(k)$ since π^{mrc} depends on $\mathbf{z}_1, \dots, \mathbf{z}_N$ via A_1, \dots, A_N , (f) follows because \mathbf{z}_k is independent of $A_1, \dots, A_{k-1}, A_{k+1}, \dots, A_N$ given A_k , (g) follows by marginalizing $\mathbf{z}_1, \dots, \mathbf{z}_{k-1}, \mathbf{z}_{k+1}, \dots, \mathbf{z}_N$, (h) follows by the tower property of expectation, (i) follows by evaluating the expectation over A_k , (j) follows because $\mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 1] := \mathbb{E}_{\mathbf{z}_k} [\mathbf{z}_k | A_k = 1]$ and $\mathbb{E}_{\mathbf{z}} [\mathbf{z} | A = 0] := \mathbb{E}_{\mathbf{z}_k} [\mathbf{z}_k | A_k = 0]$ are constants for every $k \in [N]$, (k) follows by marginalizing A_1, \dots, A_N , (l) follows from the definitions of $\mathbb{P}(A_K = 1)$ and $\mathbb{P}(A_K = 0)$, and (m) follows from rotational symmetry (see the proof of Lemma 4.1 in [Bhowmick et al. \(2018\)](#) for details). Therefore, we can write

$$\mathbb{E}_{q^{\text{mrc}}} [\hat{\mathbf{x}}^{\text{mrc}}] = \frac{1}{m_{\text{mrc}}} \mathbb{E}_{q^{\text{mrc}}} [\mathbf{z}_K] \stackrel{(a)}{=} \mathbf{x}$$

where (a) follows from (36). \square

E.2 Utility of Minimal Random Coding simulating PrivUnit_2

E.2.1 The scaling factors of PrivUnit_2 and MRC are close when N is of the right order

In the following Lemma, we show that when the number of candidates N is exponential in ε , then the scaling parameters associated with PrivUnit_2 and the MRC scheme simulating PrivUnit_2 are close.

Lemma E.2. *Let N denote the number of candidates used in the MRC scheme. Let $K \sim \pi^{\text{mrc}}$ where π^{mrc} is the distribution over the indices $[N]$ associated the MRC scheme simulating $\text{PrivUnit}_2(\mathbf{x}, \gamma, p_0)$. Consider any $\lambda > 0$. Then, the scaling factor m_{p_u} associated with PrivUnit_2 and the scaling factor m_{mrc} associated with the MRC scheme simulating PrivUnit_2 are such that*

$$m_{p_u} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$$

as long as

$$N \geq 2e^{2\varepsilon} \left(\frac{2(1+\lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1+\lambda)}{\lambda(p_0 - 1/2)} \right).$$

Proof. Following the proofs of Lemma 4.1 and Proposition 4 in [Bhowmick et al. \(2018\)](#), we can write $m_{p_u} = \gamma_+ p_0 + \gamma_- (1 - p_0)$ and $m_{\text{mrc}} = \gamma_+ p_{\text{mrc}} + \gamma_- (1 - p_{\text{mrc}})$ where

$$\gamma_+ \triangleq \frac{(1 - \gamma^2)^\alpha}{2^{d-2}(d-1)(B(\alpha, \alpha) - B(\tau; \alpha, \alpha))}, \quad \text{and} \quad \gamma_- \triangleq \frac{(1 - \gamma^2)^\alpha}{2^{d-2}(d-1)(B(\tau; \alpha, \alpha))}.$$

Therefore, we have

$$\frac{1}{m_{\text{mrc}}} - \frac{1}{m_{\text{pu}}} = \frac{m_{\text{pu}} - m_{\text{mrc}}}{m_{\text{pu}} \cdot m_{\text{mrc}}} = \frac{1}{m_{\text{pu}}} \left(\frac{(\gamma_+ - \gamma_-) \cdot (p_0 - p_{\text{mrc}})}{((\gamma_+ - \gamma_-)p_{\text{mrc}} + \gamma_-)} \right) = \frac{1}{m_{\text{pu}}} \left(\frac{p_0 - p_{\text{mrc}}}{p_{\text{mrc}} + \frac{\gamma_-}{\gamma_+ - \gamma_-}} \right) \quad (37)$$

From [Bhowmick et al. \(2018\)](#), we have $\gamma_- \leq 0 \leq \gamma_+$ and $|\gamma_+| \geq |\gamma_-|$. These inequalities imply $\frac{\gamma_-}{\gamma_+ - \gamma_-} \geq -\frac{1}{2}$. Plugging this in (37), we have

$$\frac{1}{m_{\text{mrc}}} - \frac{1}{m_{\text{pu}}} \leq \frac{1}{m_{\text{pu}}} \left(\frac{p_0 - p_{\text{mrc}}}{p_{\text{mrc}} - 1/2} \right) = \frac{1}{m_{\text{pu}}} \left(\frac{1}{\frac{p_0 - 1/2}{p_0 - p_{\text{mrc}}} - 1} \right) \quad (38)$$

We will now upper bound $\frac{p_0 - p_{\text{mrc}}}{p_0 - 1/2}$. We start by obtaining convenient expressions for p_{mrc} and p_0 . To compute $p_{\text{mrc}} = \mathbb{P}(\mathbf{z}_K \in \text{Cap}_{\mathbf{x}})$, recall that θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$. Let $c_1(\varepsilon, d)$ and $c_2(\varepsilon, d)$ be as defined in (35). Let $\bar{c}_1(\varepsilon, d) = c_1(\varepsilon, d) \times A(1, d)$ and $\bar{c}_2(\varepsilon, d) = c_2(\varepsilon, d) \times A(1, d)$. It is easy to see from Algorithm 3 and (35) that $\mathbb{P}(\mathbf{z}_k \in \text{Cap}_{\mathbf{x}}) = \bar{c}_1(\varepsilon, d)/p_0$. Further, since \mathbf{z}_k are generated uniformly at random,

$$\theta \sim \frac{1}{N} \text{Binom} \left(N, \frac{\bar{c}_1(\varepsilon, d)}{p_0} \right),$$

so we have

$$\begin{aligned} p_{\text{mrc}} &= \mathbb{P} \{ \mathbf{z}_K \in \text{Cap}_{\mathbf{x}} \} = \mathbb{E} [\mathbb{P} \{ \mathbf{z}_K \in \text{Cap}_{\mathbf{x}} | \theta \}] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\frac{\bar{c}_1(\varepsilon, d)\theta}{\bar{c}_1(\varepsilon, d)\theta + \bar{c}_2(\varepsilon, d)(1-\theta)} \right] \\ &= \frac{\bar{c}_1(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)} \mathbb{E} \left[\frac{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))\theta}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))\theta + \bar{c}_2(\varepsilon, d)} \right] \\ &\stackrel{(b)}{=} \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \mathbb{E} \left[\frac{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}{\bar{c}_2(\varepsilon, d)} - \frac{1}{\theta + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}} \right] \end{aligned} \quad (39)$$

where (a) follows from (7) because q^{pu} is a cap-based mechanism and (b) follows by simple manipulations.

To compute p_0 , observe that we have the following relationship between $\bar{c}_1(\varepsilon, d)$, $\bar{c}_2(\varepsilon, d)$, and p_0 from (35):

$$\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{1 - p_0}{\bar{c}_2(\varepsilon, d)} = 1$$

Using this and with some simple manipulations, we have

$$p_0 = \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \left(\frac{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}{\bar{c}_2(\varepsilon, d)} - \frac{1}{\mathbb{E}[\theta] + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}} \right) \quad (40)$$

From (39) and (40), we have

$$\begin{aligned} p_0 - p_{\text{mrc}} &= \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \left(\mathbb{E} \left[\frac{1}{\theta + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}} - \frac{1}{\mathbb{E}[\theta] + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}} \right] \right) \\ &= \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \left(\mathbb{E} \left[\frac{\mathbb{E}[\theta] - \theta}{\left(\theta + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)} \right) \left(\mathbb{E}[\theta] + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)} \right)} \right] \right). \end{aligned}$$

Now, using the Hoeffding's inequality, we have $\mathbb{P}\left\{|\theta - \mathbb{E}[\theta]| \geq \sqrt{\frac{\ln(2/\beta)}{2N}}\right\} \leq \beta$. Conditioned on the event $\left\{|\theta - \mathbb{E}[\theta]| \leq \sqrt{\frac{\ln(2/\beta)}{2N}}\right\}$ and using the fact that $|p_0 - p_{\text{mrc}}| \leq 1$, we have

$$p_0 - p_{\text{mrc}} \leq \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \left(\frac{\sqrt{\frac{\ln(2/\beta)}{2N}}}{\left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} - \sqrt{\frac{\ln(2/\beta)}{2N}} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right) \left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right)} \right) + \beta \quad (41)$$

where we have also plugged in $\mathbb{E}[\theta] = \frac{p_0}{\bar{c}_1(\varepsilon, d)}$. Now, we can lower bound $\left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right)$ as follows:

$$\left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right) \stackrel{(a)}{\geq} \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)} \stackrel{(b)}{\geq} \frac{1}{\exp(\varepsilon) - 1}$$

where (a) follows by lower bounding $p_0/\bar{c}_1(\varepsilon, d)$ by 0 and (b) follows because we have $\bar{c}_1(\varepsilon, d)/\bar{c}_2(\varepsilon, d) \leq \exp(\varepsilon)$. Further, if we pick $N \geq 2 \ln(2/\beta) (\exp(\varepsilon) - 1)^2$, then

$$\sqrt{\frac{\ln(2/\beta)}{2N}} \leq \frac{1}{2} \times \frac{1}{\exp(\varepsilon) - 1} \leq \frac{1}{2} \left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right). \quad (42)$$

Using (42) in (41), we have

$$\begin{aligned} p_0 - p_{\text{mrc}} &\leq \frac{\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)}{(\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))^2} \left(\frac{2\sqrt{\frac{\ln(2/\beta)}{2N}}}{\left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right) \left(\frac{p_0}{\bar{c}_1(\varepsilon, d)} + \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d)}\right)} \right) + \beta \\ &= \left(\frac{2\bar{c}_1(\varepsilon, d)\bar{c}_2(\varepsilon, d)\sqrt{\frac{\ln(2/\beta)}{2N}}}{\left(p_0 \left(1 - \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d)}\right) + \bar{c}_2(\varepsilon, d)\right)^2} \right) + \beta \\ &\stackrel{(a)}{\leq} \left(\frac{2\bar{c}_1(\varepsilon, d)}{\bar{c}_2(\varepsilon, d)} \sqrt{\frac{\ln(2/\beta)}{2N}} \right) + \beta \stackrel{(b)}{\leq} \left(2\exp(\varepsilon) \sqrt{\frac{\ln(2/\beta)}{2N}} \right) + \beta \stackrel{(c)}{\leq} \frac{\lambda(p_0 - 1/2)}{1 + \lambda}. \end{aligned} \quad (43)$$

where (a) follows because $p_0 \left(1 - \frac{\bar{c}_2(\varepsilon, d)}{\bar{c}_1(\varepsilon, d)}\right) \geq 0$, (b) follows because we have $\bar{c}_1(\varepsilon, d)/\bar{c}_2(\varepsilon, d) \leq \exp(\varepsilon)$ and (c) follows if we pick

$$\beta \leq \frac{\lambda(p_0 - 1/2)}{2(1 + \lambda)} \quad \text{and} \quad N \geq \frac{2 \exp(2\varepsilon) \ln(2/\beta)}{\left(\frac{\lambda(p_0 - 1/2)}{1 + \lambda} - \beta\right)^2} = 2 \exp(2\varepsilon) \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)}\right)^2 \ln\left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)}\right) \quad (44)$$

Further, it is easy to verify that (42) holds since the choice of N in (44) is such that $N \geq \frac{1}{2} \ln(2/\beta) (\exp(\varepsilon) - 1)^2$.

Now, rearranging (43) gives us an upper bound on $\frac{p_0 - p_{\text{mrc}}}{p_0 - 1/2}$, i.e.,

$$\frac{p_0 - p_{\text{mrc}}}{p_0 - 1/2} \leq \frac{\lambda}{1 + \lambda}. \quad (45)$$

Using (45) in (38), we have

$$\frac{1}{m_{\text{mrc}}} - \frac{1}{m_{\text{pu}}} \leq \frac{\lambda}{m_{\text{pu}}}. \quad (46)$$

Rearranging (46) completes the proof. \square

E.2.2 Relationship between mean squared errors associated with PrivUnit_2 and MRC simulating PrivUnit_2

In the following Proposition, we show that if the scaling factor m_{mrc} is close to the scaling parameter m_{pu} , then the mean squared error associated with MRC simulating PrivUnit_2 (i.e., $\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with PrivUnit_2 (i.e., $\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]$).

Proposition E.1. *Let $q^{\text{pu}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP PrivUnit_2 mechanism with parameters p_0 and γ and estimator $\hat{\mathbf{x}}^{\text{pu}}$. Let $q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ denote the MRC privatization mechanism simulating PrivUnit_2 with N candidates and estimator $\hat{\mathbf{x}}^{\text{mrc}}$. Let m_{pu} denote the scaling factor associated with PrivUnit_2 and m_{mrc} denote the scaling factor associated with the MRC scheme simulating PrivUnit_2 . Consider any $\lambda > 0$. If $m_{\text{pu}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$, then*

$$\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \leq (1 + \lambda)^2 \mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] + 2(1 + \lambda)(2 + \lambda) \sqrt{\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + (2 + \lambda)^2.$$

Proof. We will start by upper bounding $1/m_{\text{pu}}$ in terms of $\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]$. First, observe that

$$\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\| \stackrel{(a)}{\geq} \|\hat{\mathbf{x}}^{\text{pu}}\| - \|\mathbf{x}\| \stackrel{(b)}{\geq} \frac{1}{m_{\text{pu}}} - 1 \quad (47)$$

where (a) follows from the triangle inequality and (b) follows because $\|\hat{\mathbf{x}}^{\text{pu}}\| = 1/m_{\text{pu}}$ and $\|\mathbf{x}\| \leq 1$. Next, we have

$$\frac{1}{m_{\text{pu}}} = \frac{1}{m_{\text{pu}}} - 1 + 1 \stackrel{(a)}{\leq} \sqrt{\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + 1 \quad (48)$$

where (a) follows from (47). We will now upper bound $\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2]$. We have

$$\begin{aligned} \mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] &= \mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}}\|_2^2] + \|\mathbf{x}\|_2^2 - 2\langle \mathbb{E}_{q^{\text{mrc}}} [\hat{\mathbf{x}}^{\text{mrc}}], \mathbf{x} \rangle \\ &\stackrel{(a)}{\leq} \mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}}\|_2^2] + \|\mathbf{x}\|_2^2 + 2\sqrt{\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}}\|_2^2] \cdot \|\mathbf{x}\|_2^2} \\ &\stackrel{(b)}{\leq} \left(\frac{1}{m_{\text{mrc}}}\right)^2 + 1 + \frac{2}{m_{\text{mrc}}} \\ &\stackrel{(c)}{\leq} \left(\frac{1 + \lambda}{m_{\text{pu}}}\right)^2 + 1 + \frac{2(1 + \lambda)}{m_{\text{pu}}} \\ &\stackrel{(d)}{\leq} (1 + \lambda)^2 \mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] + 2(1 + \lambda)(2 + \lambda) \sqrt{\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + (2 + \lambda)^2 \end{aligned}$$

where (a) follows from Cauchy–Schwarz inequality, (b) follows because $\|\hat{\mathbf{x}}^{\text{mrc}}\| = 1/m_{\text{mrc}}$ and $\|\mathbf{x}\| \leq 1$, (c) follows from Lemma E.2 (which shows $m_{\text{pu}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$), and (d) follows using (48) and some simple manipulations. \square

In the following Lemma, we show that with on the order of ε -bits of communication, the mean squared error associated with MRC simulating PrivUnit_2 (i.e., $\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with PrivUnit_2 (i.e., $\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]$).

Lemma E.3. *Let $q^{\text{pu}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP PrivUnit_2 mechanism with parameters p_0 and γ and estimator $\hat{\mathbf{x}}^{\text{pu}}$. Let $q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ denote the MRC privatization mechanism simulating PrivUnit_2 with N candidates and estimator $\hat{\mathbf{x}}^{\text{mrc}}$. Consider any $\lambda > 0$. Then,*

$$\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \leq (1 + \lambda)^2 \mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] + 2(1 + \lambda)(2 + \lambda) \sqrt{\mathbb{E}_{q^{\text{pu}}} [\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + (2 + \lambda)^2$$

as long as

$$N \geq 2e^{2\varepsilon} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)}\right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)}\right).$$

Proof. The proof follows from Proposition E.1 and Lemma E.2. \square

E.2.3 Simulating PrivUnit_2 using Minimal Random Coding

The following Theorem shows that, for mean estimation, MRC can simulate PrivUnit_2 in a near-lossless manner (when n is large and λ is small) while only using on the order of ε bits of communication.

Theorem E.1. *Let $r_{\text{ME}}(\hat{\mu}^{p^u}, q^{p^u})$ and $r_{\text{ME}}(\hat{\mu}^{mrc}, q^{mrc})$ be the empirical mean estimation error for PrivUnit_2 with parameter p_0 and MRC simulating PrivUnit_2 with N candidates respectively. Consider any $\lambda > 0$. Then,*

$$r_{\text{ME}}(\hat{\mu}^{mrc}, q^{mrc}) \leq (1 + \lambda)^2 r_{\text{ME}}(\hat{\mu}^{p^u}, q^{p^u}) + 2(1 + \lambda)(2 + \lambda) \sqrt{\frac{r_{\text{ME}}(\hat{\mu}^{p^u}, q^{p^u})}{n}} + \frac{(2 + \lambda)^2}{n}.$$

as long as

$$N \geq 2e^{2\varepsilon} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)} \right).$$

Proof. The proof follows directly from Lemma E.3 since for all $i \in [n]$, \hat{x}_i^{mrc} are independent of each other as well as unbiased. \square

E.3 Empirical Comparisons

In this section, we compare MRC simulating PrivUnit_2 (using its approximate DP guarantee) against PrivUnit_2 and SQKR for mean estimation with $d = 500$ and $n = 5000$. We use the same data generation scheme described in Section 4.3 and set $\delta = 10^{-6}$. As before, SQKR uses $\#$ -bits = ε because it leads to a poor performance if $\#$ -bits $> \varepsilon$. We show the privacy-accuracy tradeoffs for these three methods in Figure 3. We see that MRC simulating PrivUnit_2 can attain the accuracy of the uncompressed PrivUnit_2 for the range of ε 's typically considered by LDP mechanisms while only using $(3\varepsilon/\ln 2) + 6$ bits. In comparison with the results from Section 4.3, the results in this section come with an approximate guarantee ($\delta = 10^{-6}$) and with a higher number of bits of communication. In other words, along with the obvious gains of pure privacy instead of approximate privacy, MMRC results in a lower communication cost (and therefore a lower computation cost) compared to MRC.




Figure 3: Comparing PrivUnit_2 , MRC simulating PrivUnit_2 and SQKR for mean estimation in terms of ℓ_2 error vs ε with $d = 500$, $n = 5000$, and $\#$ bits = $(3\varepsilon/\ln 2) + 6$.

F MODIFIED MINIMAL RANDOM CODING SIMULATING PrivUnit_2

In this section, we prove Lemma 4.1 (in Appendix F.1) and Theorem 4.1 (in Appendix F.2.3). To prove Theorem 4.1, first, in Appendix F.2.1, we show that when the number of candidates N is exponential in ε , the scaling factor m_{mrc} is close to the scaling parameter associated with PrivUnit_2 (i.e., m_{pu}). Next, in Appendix F.2.2, we provide the relationship between the mean squared error associated with MMRC simulating PrivUnit_2 and the mean squared error associated with PrivUnit_2 . Finally, in Appendix F.3, we provide some empirical comparisons in addition to the ones in Section 4.3 between MMRC simulating PrivUnit_2 and PrivUnit_2 .

F.1 Unbiased Modified Minimal Random Coding simulating PrivUnit_2

Consider the PrivUnit_2 ε -LDP mechanism q^{pu} described in Section 2 with parameters p_0 and γ . PrivUnit_2 is a cap-based mechanism with $\text{Cap}_{\mathbf{x}} = \{\mathbf{z} \in \mathbb{S}^{d-1} \mid \langle \mathbf{z}, \mathbf{x} \rangle \geq \gamma\}$ as discussed in Appendix D. Let π^{mmrc} be the distribution and $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ be the candidates obtained from Algorithm 2 when the reference distribution is $\text{Unif}(\mathbb{S}^{d-1})$. Let $K \sim \pi^{\text{mmrc}}(\cdot)$. Define $p_{\text{mmrc}} := \mathbb{P}(\mathbf{z}_K \in \text{Cap}_{\mathbf{x}})$ to be the probability with which the sampled candidate \mathbf{z}_K belongs to the spherical cap associated with PrivUnit_2 . Define m_{mmrc} as the scaling factor in (2) when p_0 in (2) is replaced by p_{mmrc} . Define $\hat{\mathbf{x}}^{\text{mmrc}} := \mathbf{z}_K/m_{\text{mmrc}}$ as the estimator of the MMRC mechanism simulating PrivUnit_2 .

Lemma 4.1. *Let $\hat{\mathbf{x}}^{\text{mmrc}}$ be the estimator of the MMRC mechanism simulating PrivUnit_2 as defined above. Then, $\mathbb{E}_{q^{\text{mmrc}}}[\hat{\mathbf{x}}^{\text{mmrc}}] = \mathbf{x}$.*

Proof. The proof is similar to the proof of Lemma E.1. \square

F.2 Utility of Modified Minimal Random Coding simulating PrivUnit_2

F.2.1 The scaling factors of PrivUnit_2 and MMRC are close when N is of the right order

In the following Lemma, we show that when the number of candidates N is exponential in ε , then the scaling parameters associated with PrivUnit_2 and the MMRC scheme simulating PrivUnit_2 are close.

Lemma F.1. *Let N denote the number of candidates used in the MMRC scheme. Let $K \sim \pi^{\text{mmrc}}$ where π^{mmrc} is the distribution over the indices $[N]$ associated the MMRC scheme simulating $\text{PrivUnit}_2(\mathbf{x}, \gamma, p_0)$. Consider any $\lambda > 0$. Then, the scaling factor m_{pu} associated with PrivUnit_2 and the scaling factor m_{mmrc} associated with the MMRC scheme simulating PrivUnit_2 are such that*

$$m_{\text{pu}} - m_{\text{mmrc}} \leq \lambda \cdot m_{\text{mmrc}}$$

as long as

$$N \geq \frac{e^{2\varepsilon}}{2} \left(\frac{2(1+\lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1+\lambda)}{\lambda(p_0 - 1/2)} \right).$$

Proof. The proof follows a structure similar to the proof of Lemma E.2. As in the proof of Lemma E.2, we have

$$\frac{1}{m_{\text{mmrc}}} - \frac{1}{m_{\text{pu}}} \leq \frac{1}{m_{\text{pu}}} \left(\frac{1}{\frac{p_0 - 1/2}{p_0 - p_{\text{mmrc}}} - 1} \right)$$

We will now upper bound $\frac{p_0 - p_{\text{mmrc}}}{p_0 - 1/2}$. We start by obtaining expressions for p_{mmrc} and p_0 .

To compute $p_{\text{mmrc}} := \mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}}\}$, recall that θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$. Let $c_1(\varepsilon, d)$ and $c_2(\varepsilon, d)$ be as defined in (35). Let $\bar{c}_1(\varepsilon, d) = c_1(\varepsilon, d) \times A(1, d)$ and $\bar{c}_2(\varepsilon, d) = c_2(\varepsilon, d) \times A(1, d)$. It is easy to see from Algorithm 3 and (35) that $\mathbb{P}(\mathbf{z}_k \in \text{Cap}_{\mathbf{x}}) = \bar{c}_1(\varepsilon, d)/p_0$. Further, since \mathbf{z}_k are generated uniformly at random,

$$\theta \sim \frac{1}{N} \text{Binom} \left(N, \frac{\bar{c}_1(\varepsilon, d)}{p_0} \right),$$

so we have

$$\begin{aligned} p_{\text{mmrc}} &= \mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}}\} = \mathbb{E}[\mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}}|\theta\}] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\frac{\theta \bar{c}_1(\varepsilon, d)}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} \times \mathbf{1}(\theta \leq \mathbb{E}[\theta]) \right. \\ &\quad \left. + \frac{\mathbb{E}[\theta] \bar{c}_1(\varepsilon, d) + (\theta - \mathbb{E}[\theta]) \bar{c}_2(\varepsilon, d)}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} \times \mathbf{1}(\theta > \mathbb{E}[\theta]) \right] \end{aligned} \quad (49)$$

where (a) follows from Algorithm 2.

Similarly, with some simple manipulations on the definition of p_0 , we have

$$p_0 = \frac{\mathbb{E}[\theta] \bar{c}_1(\varepsilon, d)}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} \quad (50)$$

From (49) and (50), we have

$$\begin{aligned} p_0 - p_{\text{mmrc}} &= \frac{\mathbb{E}[\bar{c}_1(\varepsilon, d)(\mathbb{E}[\theta] - \theta) \times \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + \bar{c}_2(\varepsilon, d)(\mathbb{E}[\theta] - \theta) \times \mathbf{1}(\theta > \mathbb{E}[\theta])]}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} \\ &\stackrel{(a)}{\leq} \frac{\mathbb{E}[\bar{c}_1(\varepsilon, d)(\mathbb{E}[\theta] - \theta) \times \mathbf{1}(\theta \leq \mathbb{E}[\theta])]}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} \end{aligned}$$

where (a) follows because $(\mathbb{E}[\theta] - \theta) \times \mathbf{1}(\theta > \mathbb{E}[\theta]) \leq 0$. Now, using the Hoeffding's inequality, we have $\mathbb{P}\left\{|\theta - \mathbb{E}[\theta]| \geq \sqrt{\frac{\ln(2/\beta)}{2N}}\right\} \leq \beta$. Conditioned on the event $\left\{|\theta - \mathbb{E}[\theta]| \leq \sqrt{\frac{\ln(2/\beta)}{2N}}\right\}$ and using the fact that $|p_0 - p_{\text{mmrc}}| \leq 1$, we have

$$\begin{aligned} p_0 - p_{\text{mmrc}} &\leq \frac{\bar{c}_1(\varepsilon, d) \sqrt{\frac{\ln(2/\beta)}{2N}}}{\bar{c}_2(\varepsilon, d) + \mathbb{E}[\theta] (\bar{c}_1(\varepsilon, d) - \bar{c}_2(\varepsilon, d))} + \beta \\ &\stackrel{(a)}{\leq} \frac{\bar{c}_1(\varepsilon, d)}{\bar{c}_2(\varepsilon, d)} \sqrt{\frac{\ln(2/\beta)}{2N}} + \beta \stackrel{(b)}{\leq} \exp(\varepsilon) \sqrt{\frac{\ln(2/\beta)}{2N}} + \beta \stackrel{(c)}{\leq} \frac{\lambda(p_0 - 1/2)}{1 + \lambda}. \end{aligned}$$

where (a) follows because $\mathbb{E}[\theta] \geq 0$, (b) follows because $\bar{c}_1(\varepsilon, d)/\bar{c}_2(\varepsilon, d) \leq e^\varepsilon$, and (c) follows if we pick

$$\beta \leq \frac{\lambda(p_0 - 1/2)}{2(1 + \lambda)} \quad \text{and} \quad N \geq \frac{\exp(2\varepsilon) \ln(2/\beta)}{2 \left(\frac{\lambda(p_0 - 1/2)}{1 + \lambda} - \beta \right)^2} = \frac{\exp(2\varepsilon)}{2} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)} \right).$$

The rest of the proof is similar to the proof of Lemma E.2. \square

F.2.2 Relationship between the mean squared errors associated with PrivUnit₂ and MMRC simulating PrivUnit₂

In the following Proposition, we show that if the scaling factor m_{mmrc} is close to the scaling parameter m_{pu} , then the mean squared error associated with MMRC simulating PrivUnit₂ (i.e., $\mathbb{E}_{q^{\text{mmrc}}}[\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with PrivUnit₂ (i.e., $\mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]$).

Proposition F.1. *Let $q^{\text{pu}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP PrivUnit₂ mechanism with parameters p_0 and γ and estimator $\hat{\mathbf{x}}^{\text{pu}}$. Let $q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})$ denote the MMRC privatization mechanism simulating PrivUnit₂ with N candidates and estimator $\hat{\mathbf{x}}^{\text{mmrc}}$. Let m_{pu} denote the scaling factor associated with PrivUnit₂ and m_{mmrc} denote the scaling factor associated with the MMRC scheme simulating PrivUnit₂. Consider any $\lambda > 0$. If $m_{\text{pu}} - m_{\text{mmrc}} \leq \lambda \cdot m_{\text{mmrc}}$, then*

$$\mathbb{E}_{q^{\text{mmrc}}}[\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2] \leq (1 + \lambda)^2 \mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] + 2(1 + \lambda)(2 + \lambda) \sqrt{\mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + (2 + \lambda)^2.$$

Proof. The proof is similar to the proof of Proposition E.1. \square

In the following Lemma, we show that with on the order of ε -bits of communication, the mean squared error associated with MMRC simulating PrivUnit₂ (i.e., $\mathbb{E}_{q^{\text{mmrc}}}[\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with PrivUnit₂ (i.e., $\mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]$).

Lemma F.2. *Let $q^{\text{pu}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP PrivUnit₂ mechanism with parameters p_0 and γ and estimator $\hat{\mathbf{x}}^{\text{pu}}$. Let $q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})$ denote the MMRC privatization mechanism simulating PrivUnit₂ with N candidates and estimator $\hat{\mathbf{x}}^{\text{mmrc}}$ as defined above. Consider any $\lambda > 0$. Then,*

$$\mathbb{E}_{q^{\text{mmrc}}}[\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2] \leq (1 + \lambda)^2 \mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2] + 2(1 + \lambda)(2 + \lambda) \sqrt{\mathbb{E}_{q^{\text{pu}}}[\|\hat{\mathbf{x}}^{\text{pu}} - \mathbf{x}\|_2^2]} + (2 + \lambda)^2$$

as long as

$$N \geq \frac{e^{2\varepsilon}}{2} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)} \right).$$

Proof. The proof follows from Proposition F.1 and Lemma F.1. \square

F.2.3 Simulating PrivUnit_2 using Modified Minimal Random Coding

Now, we provide a proof of Theorem 4.1.

Theorem 4.1. *Let $r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{p_u}, q^{p_u})$ and $r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{\text{mmrc}}, q^{\text{mmrc}})$ be the empirical mean estimation error for PrivUnit_2 with parameter p_0 and MMRC simulating PrivUnit_2 with N candidates respectively. Consider any $\lambda > 0$. Then,*

$$r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{\text{mmrc}}, q^{\text{mmrc}}) \leq (1 + \lambda)^2 r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{p_u}, q^{p_u}) + 2(1 + \lambda)(2 + \lambda) \sqrt{\frac{r_{\text{ME}}(\hat{\boldsymbol{\mu}}^{p_u}, q^{p_u})}{n}} + \frac{(2 + \lambda)^2}{n}.$$

as long as

$$N \geq \frac{e^{2\varepsilon}}{2} \left(\frac{2(1 + \lambda)}{\lambda(p_0 - 1/2)} \right)^2 \ln \left(\frac{4(1 + \lambda)}{\lambda(p_0 - 1/2)} \right). \quad (11)$$

Proof. The proof follows directly from Lemma F.2 since for all $i \in [n]$, $\hat{\boldsymbol{x}}_i^{\text{mmrc}}$ are independent of each other as well as unbiased. \square

F.3 Additional Empirical Comparisons

In Section 4.3, we empirically demonstrated the privacy-accuracy-communication tradeoffs of MMRC simulating PrivUnit_2 against PrivUnit_2 and SQKR in terms of ℓ_2 error vs #bits and ℓ_2 error vs ε (see Figure 1). In this section, we provide comparisons between these methods in terms of ℓ_2 error vs d (see Figure 4 (left)) and ℓ_2 error vs n (see Figure 4 (right)) for a fixed ε ($=6$) and a fixed #bits ($=11$). As before, SQKR uses #bits $= \varepsilon$ for both because it leads to a poor performance if #bits $> \varepsilon$.




Figure 4: Comparing PrivUnit_2 , MMRC simulating PrivUnit_2 and SQKR for mean estimation with $\varepsilon = 6$ and #bits $= 11$. **Left:** ℓ_2 error vs d for $n = 5000$. **Right:** ℓ_2 error vs n for $d = 500$.

G PRELIMINARY ON Subset Selection

In this section, we briefly recap the **Subset Selection** (SS) mechanism proposed in Ye and Barg (2018). Let $\boldsymbol{x} = (x_1, x_2, \dots, x_d) \in \{0, 1\}^d$ be the one-hot representation of an input symbol in $\mathcal{X} = [d] = \{1, \dots, d\}$ ¹¹. Let $q^{\text{ss}}(\boldsymbol{z}|\boldsymbol{x})$ be the **Subset Selection** mechanism defined in Ye and Barg (2018) where the output alphabet is the set of all d -bit binary strings with Hamming weight $s \in [d]$, i.e.,

$$\mathcal{Z} = \left\{ \boldsymbol{z} = (z_1, z_2, \dots, z_d) \in \{0, 1\}^d : \sum_{i=1}^d z_i = s \right\}. \quad (51)$$

¹¹With a slight abuse of notation, when context is clear, we sometime use $\boldsymbol{x} = i$ for some $i \in [d]$ to indicate the one-hot representation of symbol i

Given $\mathbf{x} \in \mathcal{X}$, **Subset Selection** maps it to $\mathbf{z} \in \mathcal{Z}$ with the following conditional probability:

$$q^{\text{ss}}(\mathbf{z}|\mathbf{x} = i) := \begin{cases} \frac{e^\varepsilon}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}} & \text{if } \mathbf{z} \in \mathcal{Z}_i \\ \frac{1}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}} & \text{if } \mathbf{z} \in \mathcal{Z} \setminus \mathcal{Z}_i \end{cases} \quad (52)$$

where $\mathcal{Z}_i = \{\mathbf{z} = (z^{(1)}, \dots, z^{(d)}) \in \mathcal{Z} : z^{(i)} = 1\}$ is the set of elements in \mathcal{Z} with 1 in the i^{th} location.

[Ye and Barg \(2018\)](#) show that the marginal distribution of \mathbf{z} is a linear function of that of \mathbf{x} . In particular, if we define $p_i := \mathbb{P}\{x_i = 1\}$ for all $i \in [d]$ and let $\mathbf{z} \sim q^{\text{ss}}(\cdot|\mathbf{x})$, then (53) is due to (5) in [Ye and Barg \(2018\)](#),

$$q_i^{\text{ss}} := \mathbb{P}\{z_i = 1\} = \frac{\binom{d-1}{s-1}e^\varepsilon p_i + \left(\binom{d-2}{s-2}e^\varepsilon + \binom{d-2}{s-1}\right)(1-p_i)}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}} \quad (53)$$

$$\begin{aligned} &= \frac{s(d-s)(e^\varepsilon - 1)}{(d-1)(s(e^\varepsilon - 1) + d)} p_i + \frac{s((s-1)e^\varepsilon + (d-s))}{(d-1)(s(e^\varepsilon - 1) + d)} \\ &= m_{\text{ss}} \cdot p_i + b_{\text{ss}}, \end{aligned} \quad (54)$$

where

$$m_{\text{ss}} := \frac{s(d-s)(e^\varepsilon - 1)}{(d-1)(s(e^\varepsilon - 1) + d)}, \quad b_{\text{ss}} := \frac{s((s-1)e^\varepsilon + (d-s))}{(d-1)(s(e^\varepsilon - 1) + d)}. \quad (55)$$

The final estimator of \mathbf{x} is denoted by $\hat{\mathbf{x}}^{\text{ss}}$ and is defined as $\frac{1}{m_{\text{ss}}} \cdot (\mathbf{z} - b_{\text{ss}} \cdot \mathbf{1}_d)$, where $\mathbf{1}_d \triangleq [1, \dots, 1]^\top \in \mathbb{R}^d$. In other words, m_{ss} and b_{ss} are used to de-bias the outcome \mathbf{z} . The scheme is summarized in [Algorithm 4](#).

Algorithm 4: Subset Selection

Require: $\mathbf{x} \in [d]$, $s \in [d]$.

Draw a s -hot random vector \mathbf{z} according to the distribution $q^{\text{ss}}(\mathbf{z}|\mathbf{x})$ in (52).

return $\hat{\mathbf{x}}^{\text{ss}} = \frac{1}{m_{\text{ss}}} \cdot (\mathbf{z} - b_{\text{ss}} \cdot \mathbf{1}_d)$

G.1 Subset Selection is unbiased and order-optimal

The following proposition borrowed from [Ye and Barg \(2018\)](#) shows that the output of the **Subset Selection** mechanism (a) is unbiased and (b) has order-optimal utility.

Proposition G.1. *Let $\hat{\mathbf{x}}^{\text{ss}} = \text{Subset Selection}(\mathbf{x}, s)$ for some $\mathbf{x} \in \mathcal{X}$ and $s \in [d]$. Then, $\mathbb{E}[\hat{\mathbf{x}}^{\text{ss}}] = \mathbf{x}$. Further, the ℓ_2 estimation error is*

$$\mathbb{E} \left[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2 \right] = \left(\frac{(s(d-2)+1)e^{2\varepsilon}}{(d-s)(e^\varepsilon - 1)^2} + \frac{2(d-2)}{(e^\varepsilon - 1)^2} + \frac{(d-2)(d-s)+1}{s(e^\varepsilon - 1)^2} - \sum_i p_i^2 \right).$$

Moreover, if we pick $s := \lceil \frac{d}{1+e^\varepsilon} \rceil$, then

$$\mathbb{E} \left[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2 \right] = \frac{d}{\min(e^\varepsilon, (e^\varepsilon - 1)^2, d)},$$

which is order-optimal.

G.2 Subset Selection is a cap-based mechanism

As discussed in [Section 3](#), q^{ss} defined in (3) is a cap-based mechanism with $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$, $c_1(\varepsilon, d) = \frac{e^\varepsilon}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}}$,

and $c_2(\varepsilon, d) = \frac{1}{\binom{d-1}{s-1}e^\varepsilon + \binom{d-1}{s}}$.

Further, $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \mathcal{Z}_{\mathbf{x}}) = \frac{\binom{d-1}{s-1}}{\binom{d}{s}} = \frac{s}{d}$. Therefore,

$$\frac{c_1(\varepsilon, d)}{c_2(\varepsilon, d)} \times \mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \mathcal{Z}_{\mathbf{x}}) = e^\varepsilon \times \frac{s}{d} \stackrel{(a)}{=} \frac{e^\varepsilon}{d} \times \lceil \frac{d}{1+e^\varepsilon} \rceil \geq \frac{e^\varepsilon}{d} \times \frac{d}{1+e^\varepsilon} \stackrel{(b)}{\geq} \frac{1}{2}$$

where (a) follows by plugging in $s = \lceil \frac{d}{1+e^\varepsilon} \rceil$ and (b) follows because $\varepsilon \geq 0$.

H SIMULATING Subset Selection USING MINIMAL RANDOM CODING

In this section, we simulate `Subset Selection` using MRC analogous to how we simulate `Subset Selection` using MMRC in Section 5. First, in Appendix H.1, we provide an unbiased estimator for MRC simulating `Subset Selection`. Next, in Appendix H.2 we provide the utility guarantee associated with MRC simulating `Subset Selection`. To do that, first, in Appendix H.2.1, we show that when the number of candidates N is exponential in ε , the scaling factor m_{mrc} is close to the scaling parameter associated with `Subset Selection` (i.e., m_{ss}). Next, in Appendix H.2.2, we provide the relationship between the mean squared error associated with MRC simulating `Subset Selection` and the mean squared error associated with `Subset Selection`. In Appendix H.2.3, we combine everything and show that, for frequency estimation, MRC can simulate `Subset Selection` in a near-lossless manner while only using on the order of ε -bits of communication. Finally, in Appendix H.3, we provide some empirical comparisons.

H.1 Unbiased Minimal Random Coding simulating Subset Selection

Consider the `Subset Selection` ε -LDP mechanism q^{ss} with parameter s as described in Section 2 and Appendix G. Let π^{mrc} be the distribution and $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ be the candidates obtained from Algorithm 1 when the reference distribution is $\text{Unif}(\mathcal{Z})$ where \mathcal{Z} is as defined in (51). Let θ denote the fraction of candidates inside $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$ where $\mathcal{Z}_{\mathbf{x}}$ is the set of elements in \mathcal{Z} with 1 in the same location as \mathbf{x} . It is easy to see that $\theta \sim \frac{1}{N} \text{Binom}(N, \frac{s}{d})$. Let $q_i^{\text{mrc}} = \mathbb{P}(z_i = 1)$ where $\mathbf{z} \sim q^{\text{mrc}}(\cdot | \mathbf{x})$ i.e., $q_i^{\text{mrc}} = \mathbb{P}\{(\mathbf{z}_K)_i = 1\}$ where $K \sim \pi^{\text{mrc}}(\cdot)$.

The following lemma shows that the marginal distribution of q_i^{mrc} can be written as a linear function of p_i similar to q_i^{ss} in (54). This allows us to provide an unbiased estimator for MRC simulating `Subset Selection`.

Lemma H.1. *Let $K \sim \pi^{\text{mrc}}(\cdot)$ and $q_i^{\text{mrc}} = \mathbb{P}\{(\mathbf{z}_K)_i = 1\}$ for $i \in [d]$. Then,*

$$q_i^{\text{mrc}} = p_i m_{\text{mrc}} + b_{\text{mrc}}$$

where

$$m_{\text{mrc}} := \mathbb{E} \left[\frac{\theta e^\varepsilon}{e^\varepsilon \theta + (1 - \theta)} \right] - \frac{1}{d-1} \mathbb{E} \left[s - \frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1 - \theta)} \right],$$

$$b_{\text{mrc}} := \frac{1}{d-1} \mathbb{E} \left[s - \frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1 - \theta)} \right].$$

Further, $\hat{\mathbf{x}}_{\text{mrc}} := (\mathbf{z}_K - b_{\text{mrc}} \cdot \mathbf{1}_d) / m_{\text{mrc}}$ is an unbiased estimator of \mathbf{x} , i.e., $\mathbb{E}[\hat{\mathbf{x}}_{\text{mrc}}] = \mathbf{x}$.

Proof. We have

$$\mathbb{P}\{(\mathbf{z}_K)_i = 1\} = \sum_j p_j \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\} \stackrel{(a)}{=} p_i \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = i\} + (1 - p_i) \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\}. \quad (56)$$

where (a) follows by symmetry. Next, we compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = i\}$ and $\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\}$ separately.

To compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = i\}$, recall that θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} . From Appendix G.2, recall that $c_1(\varepsilon, d) := \frac{e^\varepsilon}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$, $c_2(\varepsilon, d) :=$

$\frac{1}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$. Further, since \mathbf{z}_k are generated uniformly at random,

$$\theta \sim \frac{1}{N} \text{Binom} \left(N, \frac{\binom{d-1}{s-1}}{\binom{d-1}{s}} \right) = \frac{1}{N} \text{Binom} \left(N, \frac{s}{d} \right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = i\} &= \mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}} | \mathbf{x} = i\} \stackrel{(a)}{=} \mathbb{E}[\mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}} | \mathbf{x} = i, \theta\}] = \mathbb{E}\left[\frac{c_1(\varepsilon, d)\theta}{c_1(\varepsilon, d)\theta + (1-\theta)c_2(\varepsilon, d)}\right] \\ &\stackrel{(b)}{=} \mathbb{E}\left[\frac{e^\varepsilon\theta}{e^\varepsilon\theta + (1-\theta)}\right], \end{aligned} \quad (57)$$

where (a) follows by the law of total probability and (b) is due to $c_1(\varepsilon, d)/c_2(\varepsilon, d) = e^\varepsilon$.

To compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\}$, we decompose it into

$$\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\} = \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1 | \mathbf{x} = j\} + \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\}, \quad (58)$$

for any $j \neq i$ and calculate each of the terms separately.

As before, let θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} . Further, let $\bar{\theta}$ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} as well as have 1 in the j^{th} location. Since \mathbf{z}_k are generated uniformly at random,

$$\bar{\theta} \sim \frac{1}{N} \text{Binom}\left(N\theta, \frac{\binom{d-2}{s-2}}{\binom{d-1}{s-1}}\right) = \frac{1}{N} \text{Binom}\left(N\theta, \frac{s-1}{d-1}\right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1 | \mathbf{x} = j\} &\stackrel{(a)}{=} \mathbb{E}_\theta [\mathbb{E}_{\bar{\theta}} [\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1 | \mathbf{x} = j, \bar{\theta}, \theta\}]] \\ &= \mathbb{E}_\theta \left[\mathbb{E}_{\bar{\theta}} \left[\frac{c_1(\varepsilon, d)\bar{\theta}}{c_1(\varepsilon, d)\theta + (1-\theta)c_2(\varepsilon, d)} \right] \right] \\ &\stackrel{(b)}{=} \frac{s-1}{d-1} \mathbb{E}_\theta \left[\frac{c_1(\varepsilon, d)\theta}{c_1(\varepsilon, d)\theta + (1-\theta)c_2(\varepsilon, d)} \right] \stackrel{(c)}{=} \frac{s-1}{d-1} \mathbb{E} \left[\frac{e^\varepsilon\theta}{e^\varepsilon\theta + (1-\theta)} \right] \end{aligned} \quad (59)$$

where (a) follows by the law of total probability, (b) follows because $\mathbb{E}[\bar{\theta}] = \frac{s-1}{d-1} \times \theta$, and (c) is due to $c_1(\varepsilon, d)/c_2(\varepsilon, d) = e^\varepsilon$.

Similarly, to compute the term $\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\}$, let $\bar{\theta}$ denote the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} as well as have 0 in the j^{th} location. Since \mathbf{z}_k are generated uniformly at random,

$$\bar{\theta} \sim \frac{1}{N} \text{Binom}\left(N(1-\theta), \frac{\binom{d-2}{s-1}}{\binom{d-1}{s}}\right) = \frac{1}{N} \text{Binom}\left(N(1-\theta), \frac{s}{d-1}\right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\} &\stackrel{(a)}{=} \mathbb{E}_\theta [\mathbb{E}_{\bar{\theta}} [\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j, \bar{\theta}, \theta\}]] \\ &= \mathbb{E}_\theta \left[\mathbb{E}_{\bar{\theta}} \left[\frac{c_2(\varepsilon, d)\bar{\theta}}{c_1(\varepsilon, d)\theta + (1-\theta)c_2(\varepsilon, d)} \right] \right] \\ &\stackrel{(b)}{=} \frac{s}{d-1} \mathbb{E}_\theta \left[\frac{c_2(\varepsilon, d)(1-\theta)}{c_1(\varepsilon, d)\theta + (1-\theta)c_2(\varepsilon, d)} \right] \\ &\stackrel{(c)}{=} \frac{s}{d-1} \mathbb{E} \left[\frac{(1-\theta)}{e^\varepsilon\theta + (1-\theta)} \right], \end{aligned} \quad (60)$$

where (a) follows by the law of total probability, (b) follows because $\mathbb{E}[\bar{\theta}] = \frac{s}{d-1} \times (1-\theta)$, and (c) is due to $c_1(\varepsilon, d)/c_2(\varepsilon, d) = e^\varepsilon$. Using (59) and (60) in (58), we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\} &= \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1 | \mathbf{x} = j\} + \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\} \\ &= \frac{s-1}{d-1} \mathbb{E} \left[\frac{e^\varepsilon\theta}{e^\varepsilon\theta + (1-\theta)} \right] + \frac{s}{d-1} \mathbb{E} \left[\frac{(1-\theta)}{e^\varepsilon\theta + (1-\theta)} \right] \end{aligned}$$

$$= \frac{1}{d-1} \left(s - \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1-\theta)} \right] \right) \quad (61)$$

Combining everything, we have

$$\begin{aligned} q_i^{\text{mrc}} &= \mathbb{P} \{ (\mathbf{z}_K)_i = 1 \} \\ &\stackrel{(a)}{=} p_i \times [\mathbb{P} \{ (\mathbf{z}_K)_i = 1 | \mathbf{x} = i \} - \mathbb{P} \{ (\mathbf{z}_K)_i = 1 | \mathbf{x} = j \}] + \mathbb{P} \{ (\mathbf{z}_K)_i = 1 | \mathbf{x} = j \}. \\ &\stackrel{(b)}{=} p_i \times \left[\mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1-\theta)} \right] - \frac{1}{d-1} \left(s - \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1-\theta)} \right] \right) \right] + \frac{1}{d-1} \left(s - \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \theta + (1-\theta)} \right] \right) \\ &\stackrel{(c)}{=} p_i m_{\text{mrc}} + b_{\text{mrc}} \end{aligned}$$

where (a) follows from (56), (b) follows from (57) and (61), and (c) follows from the definitions of m_{mrc} and b_{mrc} .

Note that the above conclusion holds for all prior distribution $\mathbf{p} = (p_1, \dots, p_d)$ such that $\mathbf{x} \sim \mathbf{p}$. Thus by setting $\mathbf{p} = \mathbf{x}$ (here \mathbf{x} is viewed as a one-hot vector), i.e., letting \mathbf{p} be the point mass distribution at \mathbf{x} , we have

$$\mathbb{E}[\hat{\mathbf{x}}_{\text{mrc}}] = (\mathbb{E}[\mathbf{z}_K] - b_{\text{mrc}} \cdot \mathbf{1}_d) / m_{\text{mrc}} = (q^{\text{mrc}} - b_{\text{mrc}} \cdot \mathbf{1}_d) / m_{\text{mrc}} = ((m_{\text{mrc}} \cdot \mathbf{p} + b_{\text{mrc}} \cdot \mathbf{1}_d) - b_{\text{mrc}} \cdot \mathbf{1}_d) / m_{\text{mrc}} = \mathbf{p} \stackrel{(a)}{=} \mathbf{x},$$

where (a) is due to our construction of \mathbf{p} . □

H.2 Utility of Minimal Random Coding simulating Subset Selection

H.2.1 The scaling factors of Subset Selection and MRC are close when N is of the right order

In the following Lemma, we show that when the number of candidates N is exponential in ε , then the scaling parameters associated with **Subset Selection** and the MRC scheme simulating **Subset Selection** are close.

Lemma H.2. *Let N denote the number of candidates used in the MRC scheme. Let $K \sim \pi^{\text{mrc}}$ where π^{mrc} is the distribution over the indices $[N]$ associated the MRC scheme simulating **Subset Selection**. Consider any $\lambda > 0$. Then, the scaling factors m_{ss} and b_{ss} associated with **Subset Selection** and the scaling factors m_{mrc} and b_{mrc} associated with the MRC scheme simulating **Subset Selection** are such that*

$$m_{\text{ss}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$$

and $b_{\text{ss}} \leq b_{\text{mrc}}$ as long as

$$N \geq \frac{2(e^\varepsilon + 3)^2(1 + \lambda)^2}{0.24^2 \lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right).$$

Proof. First, we will obtain convenient expressions for m_{ss} and b_{ss} defined in (55). We can write

$$m_{\text{ss}} := \left(\frac{\mathbb{E}[\theta] e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \right) - \frac{1}{d-1} \left(s - \frac{e^\varepsilon \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \right) \quad (62)$$

$$b_{\text{ss}} := \frac{1}{d-1} \left(s - \frac{e^\varepsilon \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \right). \quad (63)$$

To verify these, we simply plug $\mathbb{E}[\theta] = \frac{s}{d}$ into (62) resulting in:

$$m_{\text{ss}} = \frac{d}{d-1} \frac{se^\varepsilon}{se^\varepsilon + (d-s)} - \frac{s}{d-1} = \frac{dse^\varepsilon - s^2e^\varepsilon - s(d-s)}{(d-1)(se^\varepsilon + d-s)} = \frac{s(d-s)(e^\varepsilon - 1)}{(d-1)(se^\varepsilon + d-s)}.$$

and into (63) resulting in:

$$b_{\text{ss}} = \frac{1}{d-1} \left(s - \frac{se^\varepsilon}{se^\varepsilon + d-s} \right) = \frac{1}{d-1} \left(\frac{s^2e^\varepsilon + s(d-s) - se^\varepsilon}{se^\varepsilon + d-s} \right) = \frac{1}{d-1} \left(\frac{s(s-1)e^\varepsilon + s(d-s)}{se^\varepsilon + d-s} \right).$$

Recall the definitions of b_{ss} and m_{ss} from Lemma H.1. Applying Jensen's inequality on the concave function $x \mapsto \frac{x}{x+c}$ for some $c > 0$ yields $m_{\text{mrc}} \leq m_{\text{ss}}$ and $b_{\text{mrc}} \geq b_{\text{ss}}$.

Now, we will bound $|m_{\text{mrc}} - m_{\text{ss}}|$:

$$\begin{aligned} |m_{\text{mrc}} - m_{\text{ss}}| &= \left(\frac{d}{d-1} \right) \left(\frac{\mathbb{E}[\theta]e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \mathbb{E} \left[\frac{\theta e^\varepsilon}{e^\varepsilon \theta + (1 - \theta)} \right] \right) \\ &\stackrel{(a)}{\leq} 2 \left(\frac{\mathbb{E}[\theta]e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \mathbb{E} \left[\frac{\theta e^\varepsilon}{e^\varepsilon \theta + (1 - \theta)} \right] \right) \\ &= 2 \left(\mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)\theta + 1)} \right] \right), \end{aligned} \quad (64)$$

where (a) holds since $d \geq 2$. Next, we condition on the event $\mathcal{E} := \left\{ |\mathbb{E}[\theta] - \theta| \leq \sqrt{\frac{\ln(2/\beta)}{2N}} \right\}$, which has probability $\mathbb{P}_\theta \{\mathcal{E}\} \geq 1 - \beta$ by Hoeffding's inequality. We continue to upper bound (64):

$$\begin{aligned} |m_{\text{mrc}} - m_{\text{ss}}| &= 2 \left(\mathbb{P}\{\mathcal{E}\} \mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)\theta + 1)} \middle| \mathcal{E} \right] + \mathbb{P}\{\mathcal{E}^c\} \mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)\theta + 1)} \middle| \mathcal{E}^c \right] \right) \\ &\stackrel{(a)}{\leq} 2 \left(\mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)\theta + 1)} \middle| \mathcal{E} \right] + \beta \right) \\ &\stackrel{(b)}{\leq} 2 \left(\mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)E[\theta]/2 + 1)} \middle| \mathcal{E} \right] + \beta \right) \\ &\leq 4\mathbb{E} \left[\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)^2} \middle| \mathcal{E} \right] + 2\beta \stackrel{(c)}{\leq} 4\sqrt{\frac{\ln(2/\beta)}{2N}} \frac{e^\varepsilon(1 + e^\varepsilon)^2}{4e^{2\varepsilon}} + 2\beta \\ &= \sqrt{\frac{\ln(2/\beta)}{2N}} \left(e^\varepsilon + 2 + \frac{1}{e^\varepsilon} \right) + 2\beta \leq \sqrt{\frac{\ln(2/\beta)}{2N}} (e^\varepsilon + 3) + 2\beta, \end{aligned} \quad (65)$$

where (a) holds since

$$\frac{(\mathbb{E}[\theta] - \theta) e^\varepsilon}{((e^\varepsilon - 1)\mathbb{E}[\theta] + 1)((e^\varepsilon - 1)\theta + 1)} = \frac{\mathbb{E}[\theta]e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \frac{\theta e^\varepsilon}{e^\varepsilon \theta + (1 - \theta)} \leq 1,$$

(b) holds if we pick N large enough so that $|\theta - \mathbb{E}[\theta]| \leq \frac{\mathbb{E}[\theta]}{2}$ for which a sufficient condition is $\sqrt{\frac{\ln(2/\beta)}{2N}} \leq \frac{\mathbb{E}[\theta]}{2}$ i.e., $N \geq \frac{2\ln(2/\beta)}{\mathbb{E}[\theta]^2} = 2(d/s)^2 \ln(2/\beta)$, and (c) holds since $\mathbb{E}[\theta] = s/d \geq 1/(1 + e^\varepsilon)$. Notice that the constraint $N \geq 2(d/s)^2 \ln(2/\beta)$ in inequality (b) can be further satisfied as long as $N \geq 2\ln(2/\beta)(1 + e^\varepsilon)^2$ since $s/d \geq 1/(1 + e^\varepsilon)$.

Next, we lower bound m_{ss} in (62):

$$\begin{aligned} m_{\text{ss}} &= \left(\frac{d}{d-1} \right) \left(\frac{\mathbb{E}[\theta]e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \frac{s}{d} \right) \geq \frac{\mathbb{E}[\theta]e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \frac{s}{d} \\ &\stackrel{(a)}{=} \frac{s}{d} \left[\frac{(e^\varepsilon - 1)(d - s)}{(e^\varepsilon - 1) \cdot s + d} \right] = \frac{s}{d} \left[\frac{(e^\varepsilon - 1)(1 - s/d)}{(e^\varepsilon - 1) \cdot s/d + 1} \right] \\ &\stackrel{(b)}{\geq} \frac{1}{1 + e^\varepsilon} \left[\frac{(e^\varepsilon - 1) \left(\frac{e^\varepsilon}{1 + e^\varepsilon} - \frac{1}{d} \right)}{(e^\varepsilon - 1) \left(\frac{1}{1 + e^\varepsilon} + \frac{1}{d} \right) + 1} \right] \\ &\stackrel{(c)}{\geq} \frac{1}{1 + e^\varepsilon} \left[\frac{(e^\varepsilon - 1) \frac{e^\varepsilon - 1}{1 + e^\varepsilon}}{(e^\varepsilon - 1) \left(\frac{2}{1 + e^\varepsilon} \right) + 1} \right] = \frac{(e^\varepsilon - 1)^2}{(3e^\varepsilon - 1)(e^\varepsilon + 1)} \\ &\stackrel{(d)}{\geq} \frac{(e - 1)^2}{(3e - 1)(e + 1)} \geq 0.24, \end{aligned} \quad (66)$$

where (a) holds by plugging in $\mathbb{E}[\theta] = s/d$, (b) holds since $s = \lceil d/(1 + e^\varepsilon) \rceil$ (so $\frac{1}{1 + e^\varepsilon} \leq \frac{s}{d} \leq \frac{1}{1 + e^\varepsilon} + \frac{1}{d}$), (c) holds since we only focus on the regime where $\varepsilon \leq d - 1$ (so $\frac{1}{d} \leq \frac{1}{1 + e^\varepsilon}$), and (d) holds by observing that $f(x) := \frac{(x-1)^2}{(3x-1)(x+1)}$

is an increasing function for $x \geq 1$ and we have $\varepsilon \geq 1$. Putting things together, we obtain

$$\frac{m_{ss} - m_{\text{mrc}}}{m_{\text{mrc}}} = \frac{m_{ss} - m_{\text{mrc}}}{m_{ss} - (m_{ss} - m_{\text{mrc}})} \stackrel{(a)}{\leq} \frac{\sqrt{\frac{\ln(2/\beta)}{2N}}(e^\varepsilon + 3) + 2\beta}{0.24 - \left(\sqrt{\frac{\ln(2/\beta)}{2N}}(e^\varepsilon + 3) + 2\beta\right)} \stackrel{(b)}{\leq} \lambda,$$

where (a) follows from (65) and (66) and (b) follows as long as

$$\sqrt{\frac{\ln(2/\beta)}{2N}}(e^\varepsilon + 3) + 2\beta \leq \frac{0.24\lambda}{1 + \lambda}. \quad (67)$$

To ensure (67), we let

$$\beta \leq \frac{0.24\lambda}{4(1 + \lambda)} \quad \text{and} \quad N \geq \frac{1}{2} \left(\frac{(e^\varepsilon + 3)}{\frac{0.24\lambda}{(1 + \lambda)} - 2\beta} \right)^2 \ln(2/\beta) = \frac{2(e^\varepsilon + 3)^2(1 + \lambda)^2}{0.24^2\lambda^2} \ln\left(\frac{8(1 + \lambda)}{0.24\lambda}\right).$$

It is easy to verify that this choice of N satisfies $N \geq 2 \ln(2/\beta)(1 + e^\varepsilon)^2$. \square

H.2.2 Relationship between mean squared errors associated with Subset Selection and MRC simulating Subset Selection

In the following Proposition, we show that if m_{mrc} is close to m_{ss} and $b_{\text{mrc}} \geq b_{ss}$, then the mean squared error associated with MRC simulating Subset Selection (i.e., $\mathbb{E}_{q^{\text{mrc}}}[\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with Subset Selection (i.e., $\mathbb{E}_{q^{\text{ss}}}[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$).

Proposition H.1. *Let $q^{\text{ss}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP Subset Selection mechanism with estimator $\hat{\mathbf{x}}^{\text{ss}}$. Let $q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ denote the MRC privatization mechanism simulating Subset Selection with N candidates and estimator $\hat{\mathbf{x}}^{\text{mrc}}$. Let m_{ss} and b_{ss} denote the scaling factors associated with Subset Selection and m_{mrc} and b_{mrc} denote the scaling factors associated with the MRC scheme simulating Subset Selection. Consider any $\lambda > 0$. If $m_{\text{pu}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$ and $b_{\text{mrc}} \geq b_{ss}$, then*

$$\mathbb{E}_{q^{\text{mrc}}}[\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) \mathbb{E}_{q^{\text{ss}}}[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$$

Proof. We have

$$\mathbb{E}_{q^{\text{mrc}}}[\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \stackrel{(a)}{=} \sum_{i=1}^d \text{Var}(\hat{x}_i^{\text{mrc}}) \stackrel{(b)}{=} \left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i \text{Var}((z_K)_i) \stackrel{(c)}{=} \left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}).$$

where (a) follows because \mathbf{x} is a constant, (b) follows because $\hat{\mathbf{x}}^{\text{mrc}} = (\mathbf{z}_K - b_{\text{mrc}})/m_{\text{mrc}}$, and (c) follows because $(z_K)_i \sim \text{Ber}(q_i^{\text{mrc}})$. Similarly, we have We have

$$\mathbb{E}_{q^{\text{ss}}}[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2] \stackrel{(a)}{=} \sum_{i=1}^d \text{Var}(\hat{x}_i^{\text{ss}}) \stackrel{(b)}{=} \left(\frac{1}{m_{ss}}\right)^2 \sum_i \text{Var}(z_i) \stackrel{(c)}{=} \left(\frac{1}{m_{ss}}\right)^2 \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}).$$

where (a) follows because \mathbf{x} is a constant, (b) follows because $\hat{\mathbf{x}}^{\text{ss}} = (\mathbf{z} - b_{ss})/m_{ss}$, and (c) follows because $z_i \sim \text{Ber}(q_i^{\text{ss}})$.

Now, let us look at the difference i.e.,

$$\begin{aligned} & \mathbb{E}_{q^{\text{mrc}}}[\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] - \mathbb{E}_{q^{\text{ss}}}[\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2] \\ &= \left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - \left(\frac{1}{m_{ss}}\right)^2 \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}) \\ &\leq \left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i (q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - q_i^{\text{ss}}(1 - q_i^{\text{ss}})) + \left[\frac{1}{m_{\text{mrc}}^2} - \frac{1}{m_{ss}^2}\right] \left(\sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}})\right). \end{aligned}$$

Now, first, we will bound $\left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i (q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - q_i^{\text{ss}}(1 - q_i^{\text{ss}}))$. To that end, observe that $m_{\text{pu}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$ implies

$$\frac{1}{m_{\text{mrc}}} \leq (1 + \lambda) \frac{1}{m_{\text{ss}}}. \quad (68)$$

Further, we have

$$\begin{aligned} q_i^{\text{mrc}} &\stackrel{(a)}{=} m_{\text{mrc}} p_i + b_{\text{mrc}} \stackrel{(b)}{=} q_i^{\text{ss}} + (m_{\text{mrc}} - m_{\text{ss}}) p_i + (b_{\text{mrc}} - b_{\text{ss}}) \stackrel{(c)}{\geq} q_i^{\text{ss}} - \lambda \cdot m_{\text{mrc}} \cdot p_i + (b_{\text{mrc}} - b_{\text{ss}}) \\ &\stackrel{(d)}{\geq} q_i^{\text{ss}} - \lambda \cdot m_{\text{mrc}} \cdot p_i \stackrel{(e)}{\geq} q_i^{\text{ss}} - \lambda \cdot m_{\text{ss}} \cdot p_i \stackrel{(f)}{\geq} (1 - \lambda) q_i^{\text{ss}}, \end{aligned} \quad (69)$$

where (a) follows from Lemma H.1, (b) follows from (54), (c) follows because $m_{\text{pu}} - m_{\text{mrc}} \leq \lambda \cdot m_{\text{mrc}}$, (d) follows because $b_{\text{mrc}} \geq b_{\text{ss}}$, (e) follows because $m_{\text{ss}} \geq m_{\text{mrc}}$ as seen in Lemma H.2, and (f) follows because $b_{\text{ss}} \geq 0$. Next, we have

$$\frac{q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - q_i^{\text{ss}}(1 - q_i^{\text{ss}})}{q_i^{\text{ss}}(1 - q_i^{\text{ss}})} = \frac{(q_i^{\text{ss}} - q_i^{\text{mrc}})(q_i^{\text{ss}} + q_i^{\text{mrc}} - 1)}{q_i^{\text{ss}}(1 - q_i^{\text{ss}})} \stackrel{(a)}{\leq} \frac{\lambda q_i^{\text{ss}}(q_i^{\text{ss}} + q_i^{\text{mrc}} - 1)}{q_i^{\text{ss}}(1 - q_i^{\text{ss}})} \stackrel{(b)}{\leq} \frac{\lambda}{1 - q_i^{\text{ss}}} \quad (70)$$

where (a) follows from (69) and (b) follows since $q_i^{\text{ss}} \leq 1$ and $q_i^{\text{mrc}} \leq 1$.

Let us now upper bound q_i^{ss} . We have

$$q_i^{\text{ss}} = m_{\text{ss}} \cdot p_i + b_{\text{ss}} \stackrel{(a)}{\leq} m_{\text{ss}} + b_{\text{ss}} \stackrel{(b)}{=} \left(\frac{\mathbb{E}[\theta] e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \right) \stackrel{(c)}{\leq} \frac{1}{2} \quad (71)$$

where (a) follows because $p_i \leq 1$, (b) follows from (62) and (63), and (c) follows because $\mathbb{E}[\theta] = \frac{s}{d} \geq \frac{1}{e^\varepsilon + 1}$. Combining (70) and (71), and then re-arranging results in

$$\sum_i q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}) \leq 2\lambda \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}).$$

Together with (68), we obtain

$$\left(\frac{1}{m_{\text{mrc}}}\right)^2 \sum_i (q_i^{\text{mrc}}(1 - q_i^{\text{mrc}}) - q_i^{\text{ss}}(1 - q_i^{\text{ss}})) \leq \frac{2\lambda(1 + \lambda)^2}{m_{\text{ss}}^2} \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}).$$

To bound $\left[\frac{1}{m_{\text{mrc}}^2} - \frac{1}{m_{\text{ss}}^2}\right] (\sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}))$, simply note that (68) implies $\frac{1}{m_{\text{mrc}}^2} \leq (1 + \lambda)^2 \frac{1}{m_{\text{ss}}^2}$ resulting in

$$\left[\frac{1}{m_{\text{mrc}}^2} - \frac{1}{m_{\text{ss}}^2}\right] \left(\sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}})\right) \leq \frac{2\lambda + \lambda^2}{m_{\text{ss}}^2} \left(\sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}})\right).$$

Combining everything, we have

$$\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \leq (1 + 2\lambda(1 + \lambda)^2 + 2\lambda + \lambda^2) \frac{1}{m_{\text{mrc}}^2} \sum_i q_i^{\text{ss}}(1 - q_i^{\text{ss}}) = (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) \mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2] \quad \square$$

In the following Lemma, we show that with on the order of ε -bits of communication, the mean squared error associated with MRC simulating Subset Selection (i.e., $\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with Subset Selection (i.e., $\mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$).

Lemma H.3. *Let $q^{\text{ss}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP Subset Selection mechanism with estimator $\hat{\mathbf{x}}^{\text{ss}}$. Let $q^{\text{mrc}}(\mathbf{z}|\mathbf{x})$ denote the MRC privatization mechanism simulating Subset Selection with N candidates and estimator $\hat{\mathbf{x}}^{\text{mrc}}$. Consider any $\lambda > 0$. Then,*

$$\mathbb{E}_{q^{\text{mrc}}} [\|\hat{\mathbf{x}}^{\text{mrc}} - \mathbf{x}\|_2^2] \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) \mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$$

as long as

$$N \geq \frac{2(e^\varepsilon + 3)^2(1 + \lambda)^2}{0.24^2 \lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right).$$

Proof. The proof follows from Proposition H.1 and Lemma H.2. □

H.2.3 Simulating Subset Selection using Minimal Random Coding

The following Theorem shows that, for frequency estimation, MRC can simulate `Subset Selection` in a near-lossless manner (when λ is small) while only using on the order of ε bits of communication.

Theorem H.1. *Let $r_{\text{FE}}(\hat{\Pi}^{ss}, q^{ss})$ and $r_{\text{FE}}(\hat{\Pi}^{mrc}, q^{mrc})$ be the empirical frequency estimation error for `Subset Selection` and MRC simulating `Subset Selection` with N candidates respectively. Consider any $\lambda > 0$. Then*

$$r_{\text{FE}}(\hat{\Pi}^{mrc}, q^{mrc}) \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) r_{\text{FE}}(\hat{\Pi}^{ss}, q^{ss}),$$

as long as

$$N \geq \frac{2(e^\varepsilon + 3)^2(1 + \lambda)^2}{0.24^2\lambda^2} \ln\left(\frac{8(1 + \lambda)}{0.24\lambda}\right).$$

Proof. The proof follows directly from Lemma H.3 since for all $i \in [n]$, \hat{x}_i^{mrc} are independent of each other as well as unbiased. \square

H.3 Empirical Comparisons

In this section, we compare MRC simulating `Subset Selection` (using its approximate DP guarantee) against `Subset Selection` and RHR for frequency estimation with $d = 500$ and $n = 5000$. We use the same data generation scheme described in Section 5.3 and set $\delta = 10^{-6}$. As before, RHR uses $\#$ -bits = ε because it leads to a poor performance if $\#$ -bits $> \varepsilon$. We show the privacy-accuracy tradeoffs for these three methods in Figure 5. We see that MRC simulating `Subset Selection` can attain the accuracy of the uncompressed `Subset Selection` for the range of ε 's typically considered by LDP mechanisms while only using $(3\varepsilon/\ln 2) + 6$ bits. In comparison with the results from Section 5.3, the results in this section come with an approximate guarantee ($\delta = 10^{-6}$) and with a higher number of bits of communication. In other words, along with the obvious gains of pure privacy instead of approximate privacy, MMRC results in a lower communication cost (and therefore a lower computation cost) compared to MRC.




Figure 5: Comparing `Subset Selection`, MRC simulating `Subset Selection` and SQKR for frequency estimation in terms of ℓ_2 error vs ε with $d = 500$, $n = 5000$, and $\#$ bits = $(3\varepsilon/\ln 2) + 6$.

I MODIFIED MINIMAL RANDOM CODING SIMULATING `Subset Selection`

In this section, we prove Lemma 5.1 (in Appendix I.1) and Theorem 5.1 (in Appendix I.2.3). To prove Theorem 5.1, first, in Appendix I.2.1, we show that when the number of candidates N is exponential in ε , the scaling factor m_{mrc} is close to the scaling parameter associated with `Subset Selection` (i.e., m_{ss}). Next, in Appendix I.2.2, we provide the relationship between the mean squared error associated with MMRC simulating `Subset Selection` and the mean squared error associated with `Subset Selection`. Finally, in Appendix I.3, we provide some empirical comparisons in addition to the ones in Section 5.3 between MMRC simulating `Subset Selection` and `Subset Selection`.

I.1 Unbiased Modified Minimal Random Coding simulating Subset Selection

Consider the `Subset Selection` ε -LDP mechanism q^{ss} described in Section 2 with $s := \lceil \frac{d}{1+\varepsilon} \rceil$. `Subset Selection` is cap-based mechanism as discussed in Section 3 and Appendix G with $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{z} \sim \text{Unif}(\mathcal{Z})}(\mathbf{z} \in \text{Cap}_{\mathbf{x}}) = s/d$. Let π^{mmrc} be the distribution and $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ be the candidates obtained from Algorithm 2 when the reference distribution is $\text{Unif}(\mathcal{Z})$ where \mathcal{Z} is as defined in (51). Let θ denote the fraction of candidates inside $\text{Cap}_{\mathbf{x}} = \mathcal{Z}_{\mathbf{x}}$ where $\mathcal{Z}_{\mathbf{x}}$ is the set of elements in \mathcal{Z} with 1 in the same location as \mathbf{x} . It is easy to see that $\theta \sim \frac{1}{N} \text{Binom}(N, \frac{s}{d})$. Let $q_i^{\text{mmrc}} = \mathbb{P}(z_i = 1)$ where $\mathbf{z} \sim q^{\text{mmrc}}(\cdot|\mathbf{x})$ i.e., $q_i^{\text{mmrc}} = \mathbb{P}\{(\mathbf{z}_K)_i = 1\}$ where $K \sim \pi^{\text{mmrc}}(\cdot)$.

Lemma I.1. *Let $K \sim \pi^{\text{mmrc}}(\cdot)$ and $q_i^{\text{mmrc}} = \mathbb{P}\{(\mathbf{z}_K)_i = 1\}$ for $i \in [d]$. Then,*

$$q_i^{\text{mmrc}} = p_i m_{\text{mmrc}} + b_{\text{mmrc}}$$

where

$$m_{\text{mmrc}} := \frac{d}{d-1} \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + \frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta > \mathbb{E}[\theta]) \right] - \frac{s}{d-1} \quad (72)$$

$$b_{\text{mmrc}} := \frac{1}{d-1} \left(s - \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + \frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta > \mathbb{E}[\theta]) \right] \right). \quad (73)$$

Proof. Following the proof of Lemma H.1, we compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = i\}$ and $\mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = j\}$ separately. To compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = i\}$, recall that θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} . From Appendix G.2, recall that $c_1(\varepsilon, d) := \frac{e^\varepsilon}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$, $c_2(\varepsilon, d) :=$

$\frac{1}{\binom{d-1}{s-1} e^\varepsilon + \binom{d-1}{s}}$. Further, since \mathbf{z}_k are generated uniformly at random,

$$\theta \sim \frac{1}{N} \text{Binom} \left(N, \frac{\binom{d-1}{s-1}}{\binom{d-1}{s}} \right) = \frac{1}{N} \text{Binom} \left(N, \frac{s}{d} \right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = i\} &= \mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}}|\mathbf{x} = i\} \stackrel{(a)}{=} \mathbb{E}[\mathbb{P}\{\mathbf{z}_K \in \text{Cap}_{\mathbf{x}}|\mathbf{x} = i, \theta\}] \\ &\stackrel{(b)}{=} \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + \frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbf{1}(\theta > \mathbb{E}[\theta]) \right] \end{aligned} \quad (74)$$

where (a) follows by the law of total probability and (b) is due to Algorithm 2 and $c_1(\varepsilon, d)/c_2(\varepsilon, d) = e^\varepsilon$.

To compute $\mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = j\}$, we decompose it into

$$\mathbb{P}\{(\mathbf{z}_K)_i = 1|\mathbf{x} = j\} = \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1|\mathbf{x} = j\} + \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0|\mathbf{x} = j\}, \quad (75)$$

for any $j \neq i$ and calculate each of the terms separately.

As before, let θ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} . Further, let $\bar{\theta}$ denotes the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} as well as have 1 in the j^{th} location. Since \mathbf{z}_k are generated uniformly at random,

$$\bar{\theta} \sim \frac{1}{N} \text{Binom} \left(N\theta, \frac{\binom{d-2}{s-2}}{\binom{d-1}{s-1}} \right) = \frac{1}{N} \text{Binom} \left(N\theta, \frac{s-1}{d-1} \right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1|\mathbf{x} = j\} &\stackrel{(a)}{=} \mathbb{E}_\theta [\mathbb{E}_{\bar{\theta}} [\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1|\mathbf{x} = j, \bar{\theta}, \theta\}]] \\ &\stackrel{(b)}{=} \mathbb{E}_\theta \left[\mathbb{E}_{\bar{\theta}} \left[\frac{e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \times \bar{\theta} \right] \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + \mathbb{E}_{\bar{\theta}} \left[\frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{\theta (e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta]))} \times \bar{\theta} \right] \mathbf{1}(\theta > \mathbb{E}[\theta]) \right] \end{aligned}$$

$$\stackrel{(c)}{=} \frac{s-1}{d-1} \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta \leq \mathbb{E}[\theta]) + \frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta > \mathbb{E}[\theta]) \right]. \quad (76)$$

where (a) follows by the law of total probability, (b) follows from Algorithm 2, and (c) is because $\mathbb{E}[\bar{\theta}] = \frac{s-1}{d-1} \times \theta$.

Similarly, to compute the term $\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\}$, let $\bar{\theta}$ denote the fraction of candidates that belong inside the $\text{Cap}_{\mathbf{x}}$ i.e., have 1 in the same location as \mathbf{x} as well as have 0 in the j^{th} location. Since \mathbf{z}_k are generated uniformly at random,

$$\bar{\theta} \sim \frac{1}{N} \text{Binom} \left(N(1 - \theta), \frac{\binom{d-2}{s-1}}{\binom{d-1}{s}} \right) = \frac{1}{N} \text{Binom} \left(N(1 - \theta), \frac{s}{d-1} \right),$$

so we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\} &\stackrel{(a)}{=} \mathbb{E}_\theta [\mathbb{E}_{\bar{\theta}} [\mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j, \bar{\theta}, \theta\}]] \\ &\stackrel{(b)}{=} \mathbb{E}_\theta \left[\mathbb{E}_{\bar{\theta}} \left[\frac{\bar{\theta}}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta > \mathbb{E}[\theta]) + \frac{(1 - \mathbb{E}[\theta]) + (\mathbb{E}[\theta] - \theta) e^\varepsilon}{(1 - \theta)(e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta]))} \bar{\theta} \cdot \mathbb{1}(\theta \leq \mathbb{E}[\theta]) \right] \right] \\ &\stackrel{(c)}{=} \frac{s}{d-1} \mathbb{E} \left[\frac{(1 - \theta)}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta > \mathbb{E}[\theta]) + \frac{(1 - \mathbb{E}[\theta]) + (\mathbb{E}[\theta] - \theta) e^\varepsilon}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta \leq \mathbb{E}[\theta]) \right]. \end{aligned} \quad (77)$$

where (a) follows by the law of total probability, (b) follows from Algorithm 2, and (c) is because $\mathbb{E}[\bar{\theta}] = \frac{s}{d-1} \times \theta$. Using (76) and (77) in (75), we have

$$\begin{aligned} \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\} &= \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 1 | \mathbf{x} = j\} + \mathbb{P}\{(\mathbf{z}_K)_i = 1, (\mathbf{z}_K)_j = 0 | \mathbf{x} = j\} \\ &= \frac{1}{d-1} \left(s - \mathbb{E} \left[\frac{e^\varepsilon \theta}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta \leq \mathbb{E}[\theta]) + \frac{e^\varepsilon \mathbb{E}[\theta] + \theta - \mathbb{E}[\theta]}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{1}(\theta > \mathbb{E}[\theta]) \right] \right). \end{aligned} \quad (78)$$

Combining everything, we have

$$q_i^{\text{mmrc}} = \mathbb{P}\{(\mathbf{z}_K)_i = 1\} = p_i \times [\mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = i\} - \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\}] + \mathbb{P}\{(\mathbf{z}_K)_i = 1 | \mathbf{x} = j\} \stackrel{(a)}{=} p_i m_{\text{mmrc}} + b_{\text{mmrc}}$$

where (a) follows from (74) and (78), and the definitions of m_{mmrc} and b_{mmrc} . \square

Lemma 5.1. *Let $\hat{\mathbf{x}}^{\text{mmrc}}$ be the estimator of the MMRC mechanism simulating Subset Selection as defined above. Then, $\mathbb{E}[\hat{\mathbf{x}}^{\text{mmrc}}] = \mathbf{x}$.*

Proof. Given Lemma I.1, the proof follows from the proof of Lemma H.1. \square

I.2 Utility of Modified Minimal Random Coding simulating Subset Selection

I.2.1 The scaling factors of Subset Selection and MMRC are close when N is of the right order

In the following Lemma, we show that when the number of candidates N is exponential in ε , then the scaling parameters associated with Subset Selection and the MMRC scheme simulating Subset Selection are close.

Lemma I.2. *Let N denote the number of candidates used in the MMRC scheme. Let $K \sim \pi^{\text{mmrc}}$ where π^{mmrc} is the distribution over the indices $[N]$ associated the MMRC scheme simulating Subset Selection. Consider any $\lambda > 0$. Then, the scaling factors m_{ss} and b_{ss} associated with Subset Selection and the scaling factors m_{mmrc} and b_{mmrc} associated with the MMRC scheme simulating Subset Selection are such that*

$$m_{ss} - m_{\text{mmrc}} \leq \lambda \cdot m_{\text{mmrc}}$$

and $b_{ss} \leq b_{\text{mmrc}}$ as long as

$$N \geq \frac{2(e^\varepsilon + 1)^2(1 + \lambda)^2}{0.24^2 \lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right).$$

Proof. The proof is similar to the proof of Lemma H.2. We only show the key steps here.

From (73) and (63), we have

$$\begin{aligned} b_{\text{mmrc}} - b_{\text{ss}} &= \frac{1}{d-1} \cdot \frac{1}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta > \mathbb{E}[\theta])] \\ &\stackrel{(a)}{\geq} \frac{1}{d-1} \cdot \frac{1}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [(\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta > \mathbb{E}[\theta])] \\ &= \frac{1}{d-1} \cdot \frac{1}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [(\mathbb{E}[\theta] - \theta)] = 0. \end{aligned}$$

where (a) follows because $e^\varepsilon \geq 1$. From (72) and (62), we have

$$\begin{aligned} m_{\text{ss}} - m_{\text{mmrc}} &= \frac{d}{d-1} \cdot \frac{1}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta]) + (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta > \mathbb{E}[\theta])] \\ &\leq \frac{d}{d-1} \cdot \frac{1}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])] \\ &\stackrel{(a)}{\leq} \frac{2}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \cdot \mathbb{E} [e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])] \end{aligned} \quad (79)$$

where (a) holds since $d \geq 2$. Next, we condition on the event $\mathcal{E} := \left\{ |\mathbb{E}[\theta] - \theta| \leq \sqrt{\frac{\ln(2/\beta)}{2N}} \right\}$, which has probability $\mathbb{P}_\theta \{\mathcal{E}\} \geq 1 - \beta$ by Hoeffding's inequality. We continue to upper bound (79):

$$\begin{aligned} m_{\text{ss}} - m_{\text{mmrc}} &= 2 \left(\mathbb{P}\{\mathcal{E}\} \mathbb{E} \left[\frac{e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \middle| \mathcal{E} \right] + \mathbb{P}\{\mathcal{E}^c\} \mathbb{E} \left[\frac{e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \middle| \mathcal{E}^c \right] \right) \\ &\stackrel{(a)}{\leq} 2 \left(\mathbb{E} \left[\frac{e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \middle| \mathcal{E} \right] + \beta \right) \\ &\stackrel{(b)}{\leq} (1 + e^\varepsilon) \sqrt{\frac{\ln(2/\beta)}{2N}} + 2\beta \end{aligned}$$

where (a) holds since

$$\frac{e^\varepsilon (\mathbb{E}[\theta] - \theta) \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} = \frac{e^\varepsilon \mathbb{E}[\theta] \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} - \frac{e^\varepsilon \theta \cdot \mathbf{1}(\theta \leq \mathbb{E}[\theta])}{e^\varepsilon \mathbb{E}[\theta] + (1 - \mathbb{E}[\theta])} \leq 1,$$

and (b) holds since $\mathbb{E}[\theta] = s/d \geq 1/(1 + e^\varepsilon)$.

The rest of the proof is similar to the proof of Lemma H.2. \square

I.2.2 Relationship between the mean squared errors associated with Subset Selection and MMRC simulating Subset Selection

In the following Proposition, we show that if m_{mmrc} is close to m_{ss} and $b_{\text{mmrc}} \geq b_{\text{ss}}$, then the mean squared error associated with MMRC simulating Subset Selection (i.e., $\mathbb{E}_{q^{\text{mmrc}}} [\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with Subset Selection (i.e., $\mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$).

Proposition I.1. *Let $q^{\text{ss}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP Subset Selection mechanism with estimator $\hat{\mathbf{x}}^{\text{ss}}$. Let $q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})$ denote the MMRC privatization mechanism simulating Subset Selection with N candidates and estimator $\hat{\mathbf{x}}^{\text{mmrc}}$. Let m_{ss} and b_{ss} denote the scaling factors associated with Subset Selection and m_{mmrc} and b_{mmrc} denote the scaling factors associated with the MMRC scheme simulating Subset Selection. Consider any $\lambda > 0$. If $m_{\text{pu}} - m_{\text{mmrc}} \leq \lambda \cdot m_{\text{mmrc}}$ and $b_{\text{mmrc}} \geq b_{\text{ss}}$, then*

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2] \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) \mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$$

Proof. The proof is similar to the proof of Proposition H.1. \square

In the following Lemma, we show that with on the order of ε -bits of communication, the mean squared error associated with MMRC simulating Subset Selection (i.e., $\mathbb{E}_{q^{\text{mmrc}}} [\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2]$) is close to the mean squared error associated with Subset Selection (i.e., $\mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2]$).

Lemma I.3. *Let $q^{\text{ss}}(\mathbf{z}|\mathbf{x})$ be the ε -LDP Subset Selection mechanism with parameters d and $s = \lceil \frac{d}{1+\varepsilon} \rceil$ and estimator $\hat{\mathbf{x}}^{\text{ss}}$. Let $q^{\text{mmrc}}(\mathbf{z}|\mathbf{x})$ denote the MMRC privatization mechanism simulating Subset Selection with N candidates and estimator $\hat{\mathbf{x}}^{\text{mmrc}}$ as defined above. Consider any $\lambda > 0$. Then,*

$$\mathbb{E}_{q^{\text{mmrc}}} [\|\hat{\mathbf{x}}^{\text{mmrc}} - \mathbf{x}\|_2^2] \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) \mathbb{E}_{q^{\text{ss}}} [\|\hat{\mathbf{x}}^{\text{ss}} - \mathbf{x}\|_2^2],$$

as long as

$$N \geq \frac{2(e^\varepsilon + 1)^2(1 + \lambda)^2}{0.24^2\lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right).$$

Proof. The proof follows from Proposition I.1 and Lemma I.2. \square

I.2.3 Simulating Subset Selection using Modified Minimal Random Coding

Now, we provide a proof of Theorem 5.1.

Theorem 5.1. *Let $r_{\text{FE}}(\hat{\Pi}^{\text{ss}}, q^{\text{ss}})$ and $r_{\text{FE}}(\hat{\Pi}^{\text{mmrc}}, q^{\text{mmrc}})$ be the empirical mean estimation error for Subset Selection and MMRC simulating Subset Selection with N candidates respectively. Consider any $\lambda > 0$. Then*

$$r_{\text{FE}}(\hat{\Pi}^{\text{mmrc}}, q^{\text{mmrc}}) \leq (1 + 4\lambda + 5\lambda^2 + 2\lambda^3) r_{\text{FE}}(\hat{\Pi}^{\text{ss}}, q^{\text{ss}}),$$

as long as

$$N \geq \frac{2(e^\varepsilon + 1)^2(1 + \lambda)^2}{0.24^2\lambda^2} \ln \left(\frac{8(1 + \lambda)}{0.24\lambda} \right). \quad (13)$$

Proof. The proof follows directly from Lemma I.3 since for all $i \in [n]$, $\hat{\mathbf{x}}_i^{\text{mmrc}}$ are independent of each other as well as unbiased. \square

I.3 Additional Empirical Comparisons

In Section 5.3, we empirically demonstrated the privacy-accuracy-communication tradeoffs of MMRC simulating Subset Selection against Subset Selection and RHR in terms of ℓ_2 error vs #bits and ℓ_2 error vs ε (see Figure 2). In this section, we provide comparisons between these methods in terms of ℓ_2 error vs d (see Figure 6 (left)) and ℓ_2 error vs n (see Figure 6 (right)) for a fixed $\varepsilon (=6)$ and a fixed #bits (=14). As before, RHR uses #bits = ε for both because it leads to a poor performance if #bits $> \varepsilon$.




Figure 6: Comparing Subset Selection, MMRC simulating Subset Selection and RHR for frequency estimation with $\varepsilon = 6$ and #bits = 14. **Left:** ℓ_2 error vs d for $n = 5000$. **Right:** ℓ_2 error vs n for $d = 500$.