
Differentially Private Weighted Sampling

Edith Cohen
Google Research
Tel Aviv University

Ofir Geri
Stanford University

Tamas Sarlos
Google Research

Uri Stemmer
Ben-Gurion University
Google Research

Abstract

Common datasets have the form of *elements* with *keys* (e.g., transactions and products) and the goal is to perform analytics on the aggregated form of *key* and *frequency* pairs. A weighted sample of keys by (a function of) frequency is a highly versatile summary that provides a sparse set of representative keys and supports approximate evaluations of query statistics. We propose *private weighted sampling* (PWS): A method that sanitizes a weighted sample as to ensure element-level differential privacy, while retaining its utility to the maximum extent possible. PWS maximizes the reporting probabilities of keys and estimation quality of a broad family of statistics. PWS improves over the state of the art even for the well-studied special case of *private histograms*, when no sampling is performed. We empirically observe significant performance gains of 20%-300% increase in key reporting for common Zipfian frequency distributions and accurate estimation with $\times 2$ -8 lower frequencies. PWS is applied as a post-processing of a non-private sample, without requiring the original data. Therefore, it can be a seamless addition to existing implementations, such as those optimized for distributed or streamed data. We believe that due to practicality and performance, PWS may become a method of choice in applications where privacy is desired.

1 Introduction

Weighted sampling schemes are often used to obtain versatile summaries of large datasets. The sample constitutes a representation of the data and also facilitates efficient estimation of many statistics. Motivated by the increasing

awareness and demand for data privacy, in this work we construct *privacy preserving* weighted sampling schemes. The privacy notion that we work with is that of *differential privacy* (Dwork et al., 2016), a strong privacy notion that is considered by many researchers to be a gold-standard for privacy preserving data analysis.

Before describing our new results, we define our setting more precisely. Consider an input dataset containing m elements, where each element contains a key x from some domain \mathcal{X} . For every key $x \in \mathcal{X}$ we write w_x to denote the multiplicity of x in the input dataset. (We also refer to w_x as the *frequency* of x in the data.) With this notation, it is convenient to represent the input dataset in its *aggregated* form $D = \{(x, w_x)\}$ containing pairs of a key and its frequency $w_x \geq 1$ in the data. Examples of such datasets are plentiful: Keys are search query strings and elements are search requests, keys are products and elements are transactions for the products, keys are locations and elements are visits by individuals, or keys are training examples and elements are activities that generate them. We aim here to protect the privacy of data elements. These example datasets tend to be very sparse, where the number of distinct keys in the data is much smaller than the size $|\mathcal{X}|$ of the domain. Yet, the number of distinct keys can be very large and samples serve as small summaries that can be efficiently stored, computed, and transmitted. We therefore aim for our private sample to retain this property and in particular only include keys that are in the dataset.

The (non-private) sampling schemes we consider are specified by a (non-decreasing) sequence $(q_i)_{i \geq 0}$ of probabilities $q_i \in [0, 1]$, where $q_0 := 0$. Such a sampling scheme takes an input dataset $D = \{(x, w_x)\}$ and returns a *sample* $S \subseteq D$, where each pair (x, w_x) is included in S independently, with probability q_{w_x} . Loosely speaking, given a (non-private) sampling scheme A , we aim in this paper to design a *privacy preserving* variant of A with the goal of preserving its “utility” to the extent possible under privacy constraints. We remark that an immediate consequence of the definition of differential privacy is that keys $x \in \mathcal{X}$ with very low frequencies cannot be included in the private sample S (except with very small probability). On the other hand, keys with high frequencies can be included

with probability (close to) 1. Private sampling schemes can therefore retain more utility when the dataset has many keys with higher frequencies or for tasks that are less sensitive to low frequency keys.

Informal Problem 1.1. *Given a (non-private) sampling scheme A , specified by a sampling function q , design a private sampling scheme that takes a dataset $D = \{(x, w_x)\}$ and outputs a “sanitized” sample $S^* = \{(x, w_x^*)\}$. Informally, the goals are:*

1. *Each pair $(x, w_x) \in D$ is sampled with probability “as close as possible” to the non-private sampling probability q_{w_x} .*
2. *The sanitized sample S^* provides utility that is “as close as possible” to that of a corresponding non-private sample S . In our constructions, the sanitized frequencies w_x^* would be random variables from which we can estimate ordinal and linear statistics with (functions of) the frequency w_x .*

Informal Problem 1.1 generalizes one of the most basic tasks in the literature of differential privacy – *privately computing histograms*. Informally, algorithms for private histograms take a dataset $D = \{(x, w_x)\}$ as input, and return, in a differentially private manner, a “sanitized” dataset $D^* = \{(x, w_x^*)\}$. It is often required that the output D^* is *sparse*, in the sense that if $w_x = 0$ then $w_x^* = 0$. Commonly, we seek to minimize the expected or maximum error of estimators applied to w^* of statistics on w . One well-studied objective is to minimize $\max_{x \in X} |w_x - w_x^*|$. The work on private histograms dates all the way back to the paper that introduced differential privacy (Dwork et al., 2016), and it has received a lot of attention since then, e.g., (Korolova et al., 2009; Hardt and Talwar, 2010; Beimel et al., 2014, 2016; Bun and Steinke, 2016; Bun et al., 2019; Balcer and Vadhan, 2018; Bun et al., 2018). Observe that the private histogram problem is a special case of Informal Problem 1.1, where $q \equiv 1$.

At first glance, one might try to solve Informal Problem 1.1 by a reduction to the private histogram problem. Specifically, we consider the baseline where the data is first “sanitized” using an algorithm for private histograms, and then a (non-private) weighted sampling algorithm is applied to the sanitized data (treating the sanitized frequencies as actual frequencies). This framework, of first sanitizing the data and then sampling it was also considered in (Cormode et al., 2012). We show that this baseline is sub-optimal, and improve upon it in several axes.

1.1 Our Contributions

Our proposed framework, *Private Weighted Sampling* (PWS), takes as input a non-private weighted sample S that is produced by a (non-private) weighted sampling scheme.

We apply a “sanitizer” to the sample S to obtain a respective privacy-preserving sample S^* . Our proposed solution has the following advantages.

Practicality. The private version is generated from the sample S as a post-processing step without the need to revisit the original dataset, which might be massive or unavailable. This means that we can augment existing implementations of non-private sampling schemes and retain their scalability and efficiency. This is particularly appealing for sampling schemes designed for massive distributed or streamed data that use small sketches and avoid a resource-heavy aggregation of the data (Gibbons and Matias, 1998; Estan and Varghese, 2002; Cohen et al., 2014; Andoni et al., 2011; Cohen et al., 2012; Cohen, 2018; Jayaram and Woodruff, 2018; Cohen and Geri, 2019; Cohen et al., 2020a). Our code is available at github.com/google-research/google-research/tree/master/private_sampling.

Benefits of end-to-end privacy analysis PWS achieves better utility compared to the baseline of first sanitizing the data and then sampling. In spirit, our gains follow from a well-known result in the literature of differential privacy stating that applying a differentially private algorithm on a random sample from the original data has the effect of boosting the privacy guarantees of the algorithm (Chaudhuri and Mishra, 2006; Kasiviswanathan et al., 2011; Bun et al., 2015). Our solution is derived from a precise *end-to-end* formulation of the privacy constraints that account for the benefits of the random sampling in our privacy analysis.

Optimal reporting probabilities. PWS is optimal in that it maximizes the probability that each key x is included in the private sample. The private reporting probability of a key x depends on the privacy parameters, frequency, and sampling rate and is at most the non-private sampling probability q_{w_x} . The derivation is provided in Section 4.

Estimation of linear statistics. Linear statistics according to a function of frequency have the form:

$$s := \sum_x L(x)g(w_x) , \quad (1)$$

where $g(w_x) \geq 0$ is a non-decreasing function of frequency with $g(0) := 0$. The most common use case is when $L(x)$ is a predicate and $g(w) := w$ and the statistics is the sum of frequencies of keys that satisfy the selection L . Our PWS sanitizer in Section 5 maintains optimal reporting probabilities and provides private information on frequencies of keys. We show that generally differential privacy does not allow for unbiased estimators for statistics without significant increase in variance. We propose biased but nonnegative and low-variance estimators.

Estimation of ordinal statistics. Ordinal statistics, such as (approximate) quantiles and top- k sets, are derived from

the order of keys that is induced by their frequencies. This order can be approximated by the order induced by PWS sanitized frequencies. We show that PWS is optimal, over all DP sanitization schemes, for a broad class of ordinal statistics. In particular, PWS maximizes the probability that *any* pair is concordant and maximizes the expected Kendall- τ rank correlation between the order induced by sanitized and true frequencies.

Improvement over prior baselines. We show analytically and empirically in Section 8 that we obtain orders of magnitude increase in reporting probability in low-frequency regimes. For estimation tasks, both PWS and prior schemes have lower error for higher frequencies but PWS obtains higher accuracy for frequencies that are $\times 2$ -8 lower than prior schemes. This is particularly helpful for datasets/selections with many mid-low frequency keys.

Improvement for private histograms. As an important special case of our results, we improve upon the state-of-the-art constructions for private (sparse) histograms (Korolova et al., 2009; Bun et al., 2019). These existing constructions obtain privacy properties by adding Laplace or Gaussian noise to the frequencies of the keys whereas we directly formulate and solve elementary constraints. Let π_i^* denote the PWS reporting probability of a key with frequency i , when applied to the special case of private histograms. Let ϕ_i denote the reporting probability of the state-of-the-art solution for private (sparse) histograms of (Korolova et al., 2009; Bun et al., 2019). Clearly π_i^* is always at least ϕ_i . We show that in low-frequency regimes we have $\pi_i^*/\phi_i \approx 2i$. Similarly for estimation tasks, PWS provides more accurate estimates in these regimes. Qualitatively, PWS and private histograms have high reporting probabilities and low estimation error for high frequencies. But PWS significantly improves on low to medium frequencies, which is important for distributions with long tails. We empirically show gains of 20%-300% in overall key reporting for Zipf-distributed frequencies. As private histograms are one of the most important building blocks in the literature of differential privacy, we believe that our improvement is significant (both in theory and in practice).

2 Related Work

The suboptimality of the Laplace mechanism for anonymization was noted by Ghosh et al. (2012). In our language, Ghosh et al. studied the non-sparse case, where all values, including 0 values, can be reported with added noise. They did not consider sampling, and studied pure differential privacy. Instead of Laplace noise, they propose the use of a symmetric Geometric distribution and establish it is optimal for certain estimation tasks. This can be viewed as a special case of what we do in that our schemes converge to that when there is no sampling, we

use pure differential privacy, and when frequencies are large (so the effect of the sparse case constraint dissipates). Ghosh et al. establish the optimality of unbiased estimators for some frequency statistics when loss is symmetric. We show that bias is necessary in the sparse case and propose estimators that control the bias and variance.

Key reporting was formulated and studied as *differentially private set union* problem (Gopi et al., 2020). They studied it without sampling, in a more general user privacy setting, and proposed a truncated Laplace noise mechanism similar to (Korolova et al., 2009; Bun et al., 2019).

Recent independent work by Desfontaines et al. (2020) derived the optimal scheme for key reporting for sparse private histograms, a special case of our solution when there is no sampling.

3 Preliminaries

We consider data in the form of a set of elements \mathcal{E} , where each element $e \in \mathcal{E}$ has a key $e.\text{key} \in \mathcal{X}$. The *frequency* of a key x , $w_x := |\{e \in \mathcal{E} \mid e.\text{key} = x\}|$, is defined as the number of elements with $e.\text{key} = x$. The *aggregated form* of the data, known in the DP literature as its *histogram*, is the set of key and frequency pairs $\{(x, w_x)\}$. We use the vector notation \mathbf{w} for the aggregated form. We will use $m := |\mathcal{E}|$ for the number of elements and n for the number of distinct keys in the data.

3.1 Weighted Sampling

We consider a very general form of without-replacement sampling schemes. Each scheme is specified by non-decreasing probabilities $(q_i)_{i \geq 1}$. The probability that a key is sampled depends on its frequency – a key with frequency i is sampled independently with probability q_i . Our proposed methods apply with any non-decreasing (q_i) .

Threshold sampling is a popular class of weighted sampling schemes. We review it for concreteness and motivation and use it in our empirical evaluation of PWS. A threshold sampling scheme (see Algorithm 1) is specified by (\mathcal{D}, f, τ) , where \mathcal{D} is a distribution, f is a function of frequency, and τ is a numeric threshold value that specifies the sampling rate. For each key we draw i.i.d. $u_x \sim \mathcal{D}$. The two common choices are $\mathcal{D} = \text{Exp}[1]$ for a probability proportional to size without replacement (ppswor) sample (Rosén, 1972) and $\mathcal{D} = U[0, 1]$ for a Poisson Probability Proportional to Size (PPS) sample (Ohlsson, 1990, 1998; Duffield et al., 2007). A key x is included in the sample if $u_x \leq \tau f(w_x)$. The probability that a key with frequency i is sampled is

$$q_i := \Pr_{u \sim \mathcal{D}}[u < f(i)\tau] . \quad (2)$$

Threshold sampling is related to bottom- k (order) sampling (Rosén, 1997; Ohlsson, 1990; Duffield et al., 2007;

Cohen and Kaplan, 2007, 2008) but instead of specifying the sample size k we specify an inclusion threshold τ . Ppswor is equivalent to drawing keys sequentially with probability proportional to $f(w_x)$. The bottom- k version stops after k keys and the threshold version has a stopping rule that corresponds to the threshold. The bottom- k version of Poisson PPS sampling is known as sequential Poisson or Priority sampling.

Algorithm 1: Threshold Sampling

```

// Threshold Sampler
Input: Dataset  $w$  of key frequency pairs  $(x, w_x)$ ; distribution  $\mathcal{D}$ , function  $f$ , threshold  $\tau$ 
Output: Sample  $S$  of key-frequency pairs from  $w$ 
begin
     $S \leftarrow \emptyset$ 
    foreach  $(x, w_x) \in w$  do
        Draw independent  $u_x \sim \mathcal{D}$ 
        if  $u_x < f(w_x)\tau$  then
             $S \leftarrow S \cup \{(x, w_x)\}$ 
    return  $S$ 
    
```

Since PWS applies a sanitizer to a sample, it inherits the efficiency of the base sampling scheme. Threshold sampling (via the respective bottom- k schemes) can be implemented efficiently using small sketches (of size expected sample size) on aggregated data that can be distributed or streamed (Duffield et al., 2007; Rosén, 1997; Ohlsson, 1998; Cohen and Kaplan, 2007). On unaggregated datasets, it can be implemented using small sketches for some functions of frequency including the moments $f(w) = w^p$ for $p \in [0, 2]$ (Cohen et al., 2012; Cohen, 2018; Cohen and Geri, 2019; Cohen et al., 2020b).

Our methods apply with a fixed threshold τ . But the treatment extends to when the threshold is privately determined from the data. If we have a private approximation of the total count $\|f(w)\|_1 := \sum_x f(w_x)$ we can set $\tau \approx k/\|f(w)\|_1$. This provides (from the non-private sample that corresponds to the threshold) estimates with additive error $\|f(w)\|_1/\sqrt{k}$ for statistics with function of frequency $g = f$ and when L is a predicate.

3.2 Differential Privacy

The privacy requirement we consider is *element-level* differential privacy. Two datasets with aggregated forms w and w' are neighbors if $\|w - w'\|_1 = 1$, that is, the frequencies of all keys but one are the same and the difference is at most 1 for that one key. The privacy requirements are specified using two parameters $\epsilon, \delta \geq 0$.

Definition 3.1 (Dwork et al., 2016). *A mechanism M is (ϵ, δ) -differentially private if for any two neighboring inputs w, w' and set of potential outputs T ,*

$$\Pr[M(w) \in T] \leq e^\epsilon \Pr[M(w') \in T] + \delta. \quad (3)$$

Algorithm 2: Private Weighted Samples

```

// Sanitized Keys:
Input:  $(\epsilon, \delta)$ , weighted sample  $S$ , taken with non-decreasing probabilities  $(q_i)_{i \geq 1}$ 
Output: Private sample of keys  $S^*$ 
Compute  $(p_i)_{i \geq 1}$  // Reporting probabilities per freq.
begin // Sanitize using scheme
     $S^* \leftarrow \emptyset$ 
    foreach  $(x, w_x) \in S$  do
        With probability  $p_{w_x}$ ,  $S^* \leftarrow S^* \cup \{x\}$ 
    return  $S^*$ 
// Sanitized keys and frequencies:
Input:  $(\epsilon, \delta)$ , weighted sample  $S$ , taken with non-decreasing probabilities  $(q_i)_{i \geq 1}$ 
Output: Sanitized sample  $S^*$ 
Compute probability vectors  $(p_{i\bullet})_{i \geq 1}$  // Reported values
begin // Sanitize using scheme
     $S^* \leftarrow \emptyset$ 
    foreach  $(x, w_x) \in S$  do
        Draw  $j \sim p_{w_x\bullet}$  // By probability vector
        if  $j > 0$  then
             $S^* \leftarrow S^* \cup \{(x, j)\}$ 
    return  $S^*$ 
// Estimator:
Input: Sanitized sample  $S^* = \{(x, j_x)\}, \{\pi_{i,j}\}$  (where  $\pi_{ij} := p_{ij}q_i$ , functions  $g(i), L(x)$ )
Output: Estimate of the linear statistics  $\sum_x L(x)g(x)$ 
begin
    Compute  $(a_j)_{j \geq 1}$  using  $\{\pi_{ij}\}$  and  $g(i)$  // Per-key estimates for  $g()$ 
    return  $\sum_{(x, j_x) \in S^*} L(x)a_{j_x}$ 
    
```

3.3 Private Weighted Samples

Given a (non-private) weighted sample S of the data in the form of key and frequency pairs and (a representation) of the sampling probabilities $(q_i)_{i \geq 1}$ that guided the sampling, our goal is to release as much of S as we can without violating element-level differential privacy.

We consider two utility objectives. The basic objective, *sanitized keys*, is to maximize the reporting probabilities of keys in S . The private sample in this case is simply a subset of the keys in S . The refined objective is to facilitate estimates of linear frequency and order statistics. The private sample includes sanitized keys from S together with information on their frequencies. The formats of the sanitizers and estimators are provided as Algorithm 2.

4 Sanitized Keys

Algorithm 3: Compute π_i for Sanitizing Keys

```

Input:  $(\epsilon, \delta)$ , non-decreasing sampling probabilities  $(q_i)_{i \geq 1}$ , Max.Frequency
 $\pi_0 \leftarrow 0$ 
foreach  $i = 1, \dots, \text{Max.Frequency}$  do
     $\pi_i \leftarrow \min\{q_i, e^\epsilon \pi_{i-1} + \delta, 1 + e^{-\epsilon}(\pi_{i-1} + \delta - 1)\}$ 
return  $(\pi_i)_{i=1}^{\text{Max.Frequency}}$ 
    
```

A sanitizer C uses a representation of the non-decreasing $(q_i)_{i \geq 1}$ and computes respective probabilities $(p_i)_{i \geq 1}$. A non-private sample S can then be sanitized by considering

each pair $(x, w_x) \in S$ and reporting the key x independently with probability p_{w_x} .

We find it convenient to express constraints on $(p_i)_{i \geq 1}$ in terms of the *end-to-end* reporting probability of a key x with frequency i (probability that x is sampled and then reported):

$$\pi_i := p_i q_i = \Pr[x \in C(A(\mathbf{w}))].$$

Keys of frequency 0 are not sampled or reported and we have $q_0 = 0$ and $\pi_0 := 0$. The objective of maximizing p_i corresponds to maximizing π_i . We establish the following (The proof is provided in Appendix A):

Lemma 4.1. *Consider weighted sampling scheme A where keys are sampled independently according to a non-decreasing $(q_i)_{i \geq 1}$ and a key sanitizer C (Algorithm 2) is applied to the sample. Then the probabilities $p_i \leftarrow \pi_i / q_i$, where π_i are the iterates computed in Algorithm 3, are each at the maximum under the DP constraints for $C(A(\cdot))$. Moreover, $(\pi_i)_{i \geq 1}$ is non-decreasing.*

4.1 Structure and Properties of $(\pi_i)_{i \geq 1}$

The solution as computed in Algorithm 3 applies with any non-decreasing q_i . We explore properties of the solution that allow us to compute and store it more efficiently and understand the reporting loss (reduction in reporting probabilities) that is due to the privacy requirement. Proofs are provided in Appendix A.

We provide closed-form expressions of the solution π_i^* that corresponds to $q_i = 1$ for all i (aka the private histogram problem). We will use the following definition of $L(\varepsilon, \delta)$. To simplify the presentation, we assume that ε and δ are such that L is an integer (this assumption can be removed).

$$L(\varepsilon, \delta) := \frac{1}{\varepsilon} \ln \left(\frac{e^\varepsilon - 1 + 2\delta}{\delta(e^\varepsilon + 1)} \right) \approx \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{2\delta} \right) \quad (4)$$

Lemma 4.2. *When $q_i = 1$ for all i , the sequence $(\pi_i)_{i \geq 1}$ computed by Algorithm 3 has the form:*

$$\pi_i^* = \begin{cases} \delta \frac{e^{\varepsilon i} - 1}{e^\varepsilon - 1} & i \leq L + 1 \\ 1 - \delta \frac{e^{\varepsilon(2L+2-i)} - 1}{e^\varepsilon - 1} & L + 1 \leq i \leq 2L + 1 \\ 1 & i \geq 2L + 2 \end{cases} \quad (5)$$

For the general case where the q_i 's can be smaller than 1, we bound the number of frequency values for which $\pi_i < q_i$. On these frequencies, the private reporting probability is strictly lower than that of the original non-private sample, and hence there is reporting loss due to privacy.

Lemma 4.3. *There are at most $2L(\varepsilon, \delta) + 1$ values i such that $\pi_i < q_i$, where L is as defined in (4).*

We now consider the structure of the solution for threshold sampling. The solution has a particularly simple form that can be efficiently computed and represented.

Lemma 4.4. *When the sampling probabilities $(q_i)_{i \geq 1}$ are those of threshold ppswor sampling with $f(i) = i$ then the solution has the form $\pi_i = \pi_i^*$ for $i < \ell$ and $\pi_i = q_i$ for $i \geq \ell$, where $\ell = \min\{i : \pi_i^* > q_i\}$ is the lowest position with $\pi_i^* > q_i$ and π_i^* is as defined in (5).*

5 Sanitized Keys and Frequencies

Algorithm 4: Compute $(\pi_{i,j})$ for Sanitized Frequencies

Input: (ε, δ) , non-decreasing $(q_i)_{i \geq 1}$, Max.Frequency
Output: $(\pi_{i,j})$ for $0 \leq j \leq i \leq \text{Max.Frequency}$
 $\pi_{0,0} \leftarrow 1, \pi_0 = 0$
foreach $i = 1, \dots, \text{Max.Frequency}$ **do** // Iterate over rows
 $\pi_i \leftarrow \min\{q_i, e^\varepsilon \pi_{i-1} + \delta, 1 + e^{-\varepsilon}(\pi_{i-1} + \delta - 1)\}$
 // End-to-end probability to output a key with frequency i
 $\pi_{i,0} \leftarrow 1 - \pi_i$
foreach $j = 1, \dots, i - 1$ **do** // Set lower bound values;
 use $[a]_+ := \max\{a, 0\}$
 $\pi_{i,j} \leftarrow e^{-\varepsilon} \left(\sum_{h=1}^j \pi_{i-1,h} - \delta \right) - \sum_{h=1}^{j-1} \pi_{i,h}$
 $+ [e^{-\varepsilon} \pi_{i-1,0} - \pi_{i,0}]_+$
 $\pi_{i,j} \leftarrow [\pi_{i,j}]_+$
 $R \leftarrow \pi_i - \sum_{h=1}^{i-1} \pi_{i,h}$ // Remaining probability to assign
foreach $j = i, \dots, 1$ **do** // Set final values for $\pi_{i,j}$
if $R = 0$ **then Break**
 $U \leftarrow e^\varepsilon \sum_{h=j}^{i-1} \pi_{i-1,h} + \delta - \sum_{h=j+1}^i \pi_{i,h}$ // Max value allowed for $\pi_{i,j}$
if $U - \pi_{i,j} \leq R$ **then**
 $R \leftarrow R - (U - \pi_{i,j})$
 $\pi_{i,j} \leftarrow U$
else
 $\pi_{i,j} \leftarrow \pi_{i,j} + R$
 $R \leftarrow 0$
return $(\pi_{i,j})$

A frequency sanitizer C returns keys x together with sanitized information on their frequency. We use $p_{i,j}$ for the probability that C reports $j \in [t]$ for a sampled key that has frequency i , with $p_{i,0}$ being the probability that the sampled key is not reported. We have that $\sum_{j=1}^t p_{i,j}$ is the total probability that a sampled key with frequency i is reported by C . We use

$$\pi_{i,j} \leftarrow q_i p_{i,j}$$

for the end-to-end probability that a key with frequency i is sampled and reported in the private sample with sanitized value j . For notation convenience, we use $\pi_{i,0} := 1 - \sum_{j=1}^t p_{i,j}$ for the probability that a key is not reported, making $\pi_{i,\bullet}$ probability vectors. The reader can interpret the returned value j as a token from an ordered domain. The estimators we propose depend only on the order of tokens and not their values and hence are invariant to a mapping of the domain that preserves the order.

We express constraints on $(\pi_{i,j})_{i \geq 0, j \geq 0}$. For a solution to be *realizable*, we must have end-to-end reporting probabil-

ities that do not exceed the sampling probabilities:

$$\forall i, \sum_{j=1}^t \pi_{i,j} \leq q_i . \quad (6)$$

The DP constraints are provided in the sequel. Note that we must have $\sum_{j=1}^t \pi_{i,j} \leq \pi_i$, where $(\pi_i)_{i \geq 1}$ is the solution for sanitized keys (Algorithm 3), this because the sanitized frequencies DP constraints are a superset of the sanitized keys constraints – we obtain the latter in the former by considering outputs that group together all outputs with a key x with all possible values of $j > 0$. For optimality, we seek solutions that (informally) *maximally separate* the distributions of different frequencies (minimize the overlap), over all possible DP frequency reporting schemes. We will see that maximum separation can (i) always be achieved by a discrete distribution (when the maximum frequency is bounded) and (ii) can be simultaneously achieved between any pair of frequencies. In particular, we maintain optimal reporting, that is, $\sum_{j=1}^t \pi_{i,j} = \pi_i$ and $\pi_{i,0} = \bar{\pi}_i$. The solutions we express are such that for $i_1 > i_2$, $\pi_{i_1, \bullet}$ (first-order) stochastically dominates $\pi_{i_2, \bullet}$: That is, for any h , the probability of a token $j \geq h$ is non-decreasing with frequency.

We present two algorithms that express $(\pi_{i,j})$. Algorithm 4 provides a simplified construction, where t is equal to the maximum frequency and we always report $j \leq i$ for a key with frequency i . The sanitizer satisfies realizability and DP and has optimal key reporting but attains maximum separation only under some restrictions. The values $\pi_{i,j}$ are specified in order of increasing i , where the row $\pi_{i, \bullet}$ is set so that the probability mass of π_i is pushed to the extent possible to higher j values.

Algorithm 5 specifies PDFs (f_i) for a frequency sanitizer. The PDFs have a discrete point mass at 0 (that corresponds to the probability of not reporting) and are piecewise constant elsewhere. The scheme is a refinement of the scheme of Algorithm 4 and, as we shall see, for any (q_i) and (ε, δ) , it maximally separates sanitized values for different frequencies. The construction introduces at most $3m$ distinct breakpoints for frequencies up to m and can be discretized to have an equivalent $(\pi_{i,j})$ form with $j \in [3m]$. (More details and proofs are provided in Appendix C.)

Theorem 5.1. *The sanitizer with $(\pi_{i,j})$ expressed in Algorithm 5 satisfies:*

1. $\forall i, \sum_{j=1}^i \pi_{i,j} = \pi_i$, and in particular, (6) holds and the sanitizer is realizable.
2. (ε, δ) -DP
3. *Maximum separation: For each i , there is an index c_i so that subject to the above and to row $\pi_{i-1, \bullet}$, for all $j' \leq c_i$, the sum $\sum_{j=1}^{j'} \pi_{i,j}$ is at a minimum and for all $j' \geq c_i$, the sum $\sum_{j=j'}^i \pi_{i,j}$ is at a maximum.*

The $(\pi_{i,j})$ expressed by Algorithm 4 satisfy maximum separation (Property 3) under the particular restrictions on the reported values (that only i different outputs are possible for frequencies up to i). (The proofs are provided in Appendix B)

The $(\pi_{i,j})$ expressed by Algorithm 5 and then discretized satisfy maximum separation (Property 3) unconditionally, over all DP frequency sanitization schemes.

For the special case where $q_i = 1$ for all i , Algorithm 4 provides maximum separation. We provide a closed-form expression for the solution $\pi_{i,j}^*$.

Lemma 5.2. *Let the DP parameters (ε, δ) be such that $L(\varepsilon, \delta)$ as in (4) is integral. Let $(\pi_{i,j}^*)$ be the solution computed in Algorithm 4 for $q_i = 1$ for all i . Then the matrix with entries $\pi_{i,j}^*$ for $i, j \geq 1$ has a lower triangular form, with the non-zero entries as follows:*

$$\text{For } j \in \{\max\{1, i - 2L\}, \dots, i\}, \pi_{i,j}^* = \pi_{i-j+1}^* - \pi_{i-j}^* .$$

$$\text{Equivalently, } \pi_{i,j}^* = \begin{cases} \delta e^{(i-j)\varepsilon} & \text{if } 0 \leq i - j \leq L \\ \delta e^{(2L-(i-j))\varepsilon} & \text{if } L + 1 \leq i - j \leq 2L . \end{cases}$$

Algorithm 5: Compute (f_i) for Sanitizing Frequencies

Input: (ε, δ) , non-decreasing sampling probabilities $(q_i)_{i \geq 1}$,
 Max.Frequency
Output: $(f_i)_{i=0}^{\text{Max.Frequency}}$, where $f_i : [0, i]$ // PDF of sanitized
 frequency for frequency i : discrete mass at
 $f_i(0)$ (probability of not reporting) and
 density on $(0, i]$
 $f_0(0) \leftarrow 1; \pi_i \leftarrow 0$ // Keys with frequency 0 are never
 reported
for $i \leftarrow 1$ **to** Max.Frequency **do** // Specify $f_i : [0, i]$
 $\pi_i \leftarrow \min\{q_i, e^\varepsilon \pi_{i-1} + \delta, 1 + e^{-\varepsilon}(\pi_{i-1} + \delta - 1)\}$;
 $f_i(0) \leftarrow 1 - \pi_i$ // Reporting probability for i
 $f_i(i-1, i] \leftarrow \min\{\pi_i, \delta\}$
 // Represent a function $f_L : (0, i-1]$ that "lower
 bounds" f_i
if $\max\{0, e^{-\varepsilon} f_{i-1}(0) - f_i(0)\} + \int_{0^+}^{i-1} f_{i-1}(x) dx \leq \delta$ **then**
 $f_L(0, i-1] \leftarrow 0$
else
 $b_i \leftarrow z$ that solves
 $\max\{0, e^{-\varepsilon} f_{i-1}(0) - f_i(0)\} + \int_{0^+}^z f_{i-1}(x) dx = \delta$
 // Well defined, as from DP, we always have
 $\max\{0, e^{-\varepsilon} f_{i-1}(0) - f_i(0)\} \leq \delta$
 $f_L(0, b_i] \leftarrow 0$
for $x \in (b_i, i-1]$ **do** $f_L(x) \leftarrow e^{-\varepsilon} f_{i-1}(x)$
 // Point where $f_i(x) - f_{i-1}(x)$ switches sign
 $c_i \leftarrow z$ that solves
 $\int_0^z f_L(x) dx + e^\varepsilon \int_z^{i-1} f_{i-1}(x) dx = \pi_i - \min\{\pi_i, \delta\}$
 // Any solution $z \in (0, i-1]$
for $x \in (c_i, i-1]$ **do** $f_i(x) = e^\varepsilon f_{i-1}(x)$
for $x \in (0, c_i]$ **do** $f_i(x) \leftarrow f_L(x)$
return $(f_i)_{i=0}^{\text{Max.Frequency}}$

6 Estimation of Ordinal Statistics

The sanitized frequencies can be used for estimation of statistics specified with respect to the actual frequencies. In this section we consider *ordinal* statistics, that only depend on the order of frequencies but not on their nominal values. Ordinal statistics include (approximate) top- k

set, quantiles, rank of a key, set of keys with a higher (or lower) rank than a specified key, and more. We approximate ordinal statistics from the ordering of keys that is induced by sanitized frequencies. The quality of estimated ordinal statistics is determined by the match between the order induced by exact frequencies and the order induced by sanitized frequencies. We say that the two orders are *concordant* on a subset of keys $\{x_i\}$, when they match on that subset. Since the output of our sanitizer is stochastic, we consider the probability of a subset being concordant. When sanitized values are discrete and two keys have the same sanitized value, we use probability of 0.5 that two keys are concordant.

We define (see Appendix D) a measure of separation between distributions f_{i_1} and f_{i_2} at a certain quantile value α and show that the (f_i) constructed by Algorithm 5 maximize it pointwise for any i_1, i_2, α . This measure generalizes and follows from Property 3 stated in Theorem 5.1. As a corollary we show (The proof is provided in Appendix D):

Corollary 6.1. *The sanitizing scheme specified by the (f_i) computed by Algorithm 5 maximizes the following: The probability that a subset of keys is concordant, the probability that a key is correctly ordered with respect to all other keys, and the expected Kendall- τ rank correlation.*

Note that we get optimality in a strong sense – there is no Pareto front where concordant probability on some pairs of frequencies needs to be reduced in order to get a higher value for other pairs.

7 Estimation of Linear Frequency Statistics

The objective is to estimate statistics of the form

$$s := \sum_x L(x)g(w_x) . \quad (7)$$

We briefly review estimators for the non-private setting where the sample consists of pairs (x, w_x) of keys and their frequency. We use the per-key inverse-probability estimators (Horvitz and Thompson, 1952) (also known as importance sampling). The estimate $\widehat{g(w_x)}$ of $g(w_x)$ is 0 if key x is not included in the sample and otherwise the estimate is

$$a_{w_x} := \frac{g(w_x)}{q_{w_x}} . \quad (8)$$

These estimates are nonnegative, a desired property for nonnegative values, and are also unbiased when $q_{w_x} > 0$. The estimate for the query statistics (7) is

$$\hat{s} := \sum_{(x, w_x)} L(x)\widehat{g(w_x)} = \sum_{(x, w_x) \in S} L(x)a_{w_x} . \quad (9)$$

Since the estimate is 0 for keys not represented in the sample, it can be computed from the sample. The variance of a

per-key estimate for a key with frequency i is $g(i)^2(\frac{1}{q_i} - 1)$ and the variance of the sum estimator (9) is

$$\text{Var}[\hat{s}] = \sum_x L(x)^2 g(w_x)^2 \left(\frac{1}{q_{w_x}} - 1 \right) .$$

These inverse-probability estimates are optimal for the sampling scheme in that they minimize the sum of per-key variance under unbiasedness and non-negativity constraints. We note that the quality of the estimates depends on the match between $g(i)$ and q_i : Probability Proportional to Size (PPS), where $q_i \propto g(i)$ is most effective and minimizes the sum of per-key variance for the sample size. Our aim here is to optimize what we can do privately when q and g are given.

7.1 Estimation with Sanitized Samples

We now consider estimation from sanitized samples S^* . We specify our estimators $(a_j)_{j \geq 1}$ in terms of the reported sanitized frequencies j . The estimate is 0 for keys that are not reported and are a_j when reported with value j . The estimate of the statistics is

$$\hat{s} := \sum_{(x, j) \in S^*} L(x)a_j . \quad (10)$$

As for choosing $(a_j)_{j \geq 1}$, a first attempt is the unique unbiased estimator: The unbiasedness constraints

$$\forall i, \sum_{j=1}^i \pi_{ij} a_j = g(i)$$

form a triangular system with a unique solution $(a_j)_{j \geq 1}$:

$$a_i \leftarrow \frac{g(i) - \sum_{j=1}^{i-1} \pi_{i,j} a_j}{\pi_{i,i}} .$$

However, $(a_j)_{j \geq 1}$ may include negative values and estimates have high variance. We argue that bias is unavoidable with privacy: First, the inclusion probability of keys with frequency $w_x = 1$ can not exceed δ . Therefore, the variance contribution of the key to any unbiased estimate is at least $1/\delta$. Typically, δ is chosen so that $1/\delta \gg nk$, where n is the support size and k the sample size, so this error can not be mitigated. Second, we show in Appendix F that even for the special case of $q = 1$, any unbiased estimator applied to the output of any sanitized keys and frequencies scheme with optimal reporting probabilities must assume negative values. That is, DP schemes do not admit unbiased nonnegative estimators without compromising reporting probabilities. We therefore seek estimators that are biased but balance bias and variance and are nonnegative. In our evaluation we use the following Maximum Likelihood estimator (MLE):

$$a_j \leftarrow \frac{g(i)}{\pi_i}, \text{ where } i = \arg \max_h \pi_{hj} . \quad (11)$$

This estimate is "right" for the frequency i for which the probability of reporting j is maximized. The estimate can be biased up or down. Another estimator with desirable properties is proposed in Appendix E.

We express the expected value, bias, Mean Squared Error (MSE), and variance of the per-key estimate for a key with frequency i :

$$\begin{aligned} E_i &:= \sum_{j=1}^i \pi_{i,j} a_j \\ \text{Bias}_i &:= E_i - g(i) \\ \text{MSE}_i &:= \bar{\pi}_i g(i)^2 + \sum_{j=1}^i \pi_{i,j} (a_j - g(i))^2 \\ \text{Var}_i &:= \text{MSE}_i - \text{Bias}_i^2. \end{aligned}$$

For the sum estimate (10) we get:

$$\begin{aligned} \text{Bias}[\hat{s}] &= \sum_x L(x) \text{Bias}_{w_x} \\ \text{Var}[\hat{s}] &= \sum_x L(x)^2 \text{Var}_{w_x} \\ \text{MSE}[\hat{s}] &= \text{Var}[\hat{s}] + \text{Bias}[\hat{s}]^2 \\ \text{NRMSE}[\hat{s}] &= \frac{\sqrt{\text{MSE}[\hat{s}]}}{s}. \end{aligned} \quad (12)$$

Note that the variance component of the normalized squared error $\text{MSE}[\hat{s}]/s^2$ decreases linearly with support size whereas the bias component may not. We therefore consider both the variance and bias of the per-key estimators and qualitatively seek low bias and "bounded" variance. We measure quality of statistics estimators using the Normalized Root Mean Squared Error (NRMSE).

8 Performance Analysis

We study the performance of PWS on the key reporting and estimation objectives and compare with a baseline method that provides the same privacy guarantees. We use precise expressions (not simulations) to compute probabilities, bias, variance, and MSE of the different methods.

8.1 Private Histograms Baseline

We review the *Stability-based Histograms* (SbH) method of (Bun et al., 2019; Korolova et al., 2009; Vadhan, 2017), which we use as a baseline. SbH, provided as Algorithm 6, is designed for the special case when $q_i = 1$ for all frequencies. The input S is the full data of pairs of keys and positive frequencies (x, w_x) . The private output S^* is a subset of the keys in the data with positive sanitized frequencies (x, w_x^*) .

The SbH method is considered the state of the art for sparse histograms (only keys with $w_x > 0$ can be reported). The method returns non-negative $w_x^* > 0$ sanitized frequencies.

Algorithm 6: Stability-based Histograms (SbH)

Input: (ε, δ) , $S = \{(x, w_x)\}$ where $w_x > 0$
Output: Key value pairs S^*

```

 $S^* \leftarrow \emptyset$  // Initialize private histogram
 $T \leftarrow (1/\varepsilon) \ln(1/\delta) + 1$  // Threshold
foreach  $(x, w_x) \in S$  do
     $w_x^* \leftarrow w_x + \text{Lap}[\frac{1}{\varepsilon}]$  // Add Laplace random variable
    if  $w_x^* \geq T$  then
         $S^* \leftarrow S^* \cup (x, w_x^*)$ 
    
```

For the case of no sampling, we compare PWS (with $q \equiv 1$) with SbH. We use the SbH sanitized frequencies directly for estimation. For sampling, our baseline is *sampled-SbH*: The data is first sanitized using SbH and then sampled, using a weighted sampling algorithm with q , while treating the sanitized frequencies as actual frequencies. For estimation, we apply the estimator (9) (which in this context is biased). To facilitate comparison with SbH and sampled-SbH we express the reporting probabilities, bias, and variance in Appendix G.

8.2 Reporting Probabilities: No Sampling

We start with the case of no sampling and the objective of maximizing the number of privately reported keys. We compare the PWS (optimal) probabilities π^* (5) to the baseline SbH (Bun et al., 2019; Korolova et al., 2009; Vadhan, 2017) reporting probabilities ϕ (49). Figure 1 shows reporting probability per frequency for selected DP parameters. We can see that with both private methods the reporting probability reaches 1 for high frequencies but PWS (Opt) reaches the maximum earlier and is significantly higher than ϕ along the way. Analytically from the expressions we can see that for $i \leq L(\varepsilon, \delta)$, $\pi_i^*/\phi_i \in [2, 2/\varepsilon]$ and for lower i we have $\pi_i^*/\phi_i \approx 2i$. We can also see that π_{2L+1}^* reaches 1 at $2L \approx \frac{2}{\varepsilon} \ln(\varepsilon/\delta)$ whereas $\phi_i > 1 - \delta$ for $i \approx \frac{2}{\varepsilon} \ln(1/\delta)$. The ratio between the frequency values when maximum reporting is reached is $\approx \ln(1/\delta) / \ln(\varepsilon/\delta)$.

Figure 2 shows the expected numbers of reported keys with PWS (Opt) and SbH for frequency distributions that are Zipf $[\alpha]$ with $\alpha = 1, 2$ as we sweep the privacy parameter δ . Overall we see that PWS gains 20%-300% in the number of keys reported over baseline. Note that as expected, the optimal PWS reports *all* keys when $\delta = 1$ (i.e., no privacy guarantees) but SbH incurs reporting loss.

We additionally evaluate PWS and compare it to SbH on two real-world datasets:

1. ABC: The words of news headlines from the Australian Broadcasting Corporation. The keys are words and the frequency is the respective number of occurrences (Kulkarni, 2017).

2. SO: The multi-graph of Stack Overflow where keys are nodes in the graph and frequencies are undirected degrees (Paranjape et al., 2017).

Figure 3 shows the expected numbers of reported keys on these datasets. We note that the results are similar to what was observed on the synthetic Zipf datasets.




Figure 1: Key reporting probability for frequency. No sampling ($q = 1$) with PWS (Opt) and SbH for $(\epsilon, \delta) = (0.1, 0.01), (0.01, 10^{-6})$




Figure 2: Expected fraction of keys that are privately reported with PWS (Opt) and SbH for Zipf $[\alpha]$ frequency distributions. For $\alpha = 1, 2$, privacy parameters $\epsilon = 0.1$ and sweeping δ between 1 and 10^{-8} . Left: The respective ratio of PWS to SbH.




Figure 3: Evaluation of PWS on real-world datasets (without sampling). Left: The ratio of reported keys with PWS to SbH. Center and Right: Fraction of total keys reported with sampled-SbH and PWS as we sweep the parameter δ .

8.3 Reporting Probabilities with Sampling

Figure 4 shows reporting probabilities with PWS (optimal reporting probabilities), sampled-SbH, and non-private sampling, for representative sampling rates and privacy parameters. As expected, for sufficiently large frequencies both private methods have reporting probabilities that match the sampling probabilities q of the non-private scheme. But PWS reaches q at a lower frequency than sampled-SbH and has significantly higher reporting probabilities for lower frequencies. Figure 5 shows the fraction of keys reported for Zipf distributions as we sweep the sampling rate (threshold τ). PWS reports more keys than sampled-SbH and the gain persists also with low sampling




Figure 4: Reporting probability as a function of frequency. For ppswor sampling with threshold τ (q), PWS private samples (Opt), and sampled-SbH private samples.

rates. We can see that with PWS, thanks to end-to-end privacy analysis, the reporting loss due to sampling mitigates the reporting loss needed for privacy – reporting approaches that of the non-private sampling when the sampling rate τ approaches δ . Sampled-SbH, on the other hand, incurs reporting loss due to privacy on top of the reporting loss due to sampling. Figures 6 and 7 show the expected fraction of reported keys on the real-world datasets ABC and SO.




Figure 5: Fraction of total keys reported with threshold-ppswor, sampled-SbH, and PWS (Opt), as we sweep the sampling rate τ . For Zipf $[\alpha]$, $\epsilon = 0.1$ and $\delta = 0.001$ the gains in reporting of PWS over Sampled-SbH are at least 230% ($\alpha = 0.5$), 97% ($\alpha = 1$) and 37% ($\alpha = 2$).




Figure 6: Evaluation on real-world datasets (with sampling): fraction of total keys reported with threshold-ppswor (non-private), sampled-SbH, and PWS, as we sweep the sampling rate τ .

8.4 Estimation of Linear Statistics

We evaluate estimation quality for linear statistics (7) when $g(w) = w$ and $L(x)$ is a selection predicate. The statistics is simply the sum of frequencies of selected keys. We compare performance of PWS with the MLE estimator (11), the baseline sampled-SbH, and for reference, the estimator of the respective non-private sample (9). Figure 8 (top) shows normalized bias Bias_i/i as a function of the frequency i for the two private methods (the non-private estimator is unbiased and not shown). With both




Figure 7: Ratio of expected number of keys that are privately reported with PWS to SbH for the ABC and SO datasets, as we sweep the sampling rate τ .

methods, the bias decreases with frequency and diminishes for $i \gg 2\epsilon^{-1} \ln(1/\delta)$. PWS has lower bias at lower frequencies than SbH, allowing for more accurate estimates on a broader range. We can see that with PWS, the bias decreases when the sampling rate (τ) decreases and diminishes when τ approaches δ . This is a benefit of the end-to-end privacy analysis. The bias of the baseline method does not change with sampling rate.

Figure 9 shows the normalized variance Var_i/i^2 per frequency i for representative parameter settings. The private methods PWS and sampled-SbH maintain low variance across frequencies: The value is fractional with no sampling and is of the order of that of the non-private unbiased estimator with sampling. In particular this means that the bias is a good proxy for performance and that the improvement in bias of PWS with respect to baseline does not come with a hidden cost in variance. For high frequencies (not shown), keys with all methods are included with probability (close to) 1. The non-private method that reports exact frequencies have 0 variance whereas the private methods maintain a low variance, but the normalized variance diminishes for all methods.

For statistics estimation, the per-key performance suggest that when the selection has many high frequency keys, the private methods perform well and are similar to non-private sampling. When the selection is dominated by very low frequencies, the private methods perform poorly and well below the respective non-private sample. But for low to medium frequencies, PWS can provide drastic improvements over SbH and the gain increases with lower sampling rates. Figure 8 (bottom) shows the NRMSE as a function of sampling rate for estimating the sum of frequencies on a selection of 2×10^5 keys with frequencies uniformly drawn between 1 and 200. We can see that the error of non-private sampling and of sampled-SbH decreases with higher sampling rate. Note the perhaps counter-intuitive phenomenon that PWS (MLE) hits its sweet spot midway: This is due to a balance of the two components of the error, the variance which increases and the bias that decreases when the sampling rate decreases. Also note that PWS significantly improves over SbH also with no sampling ($\tau = 1$).

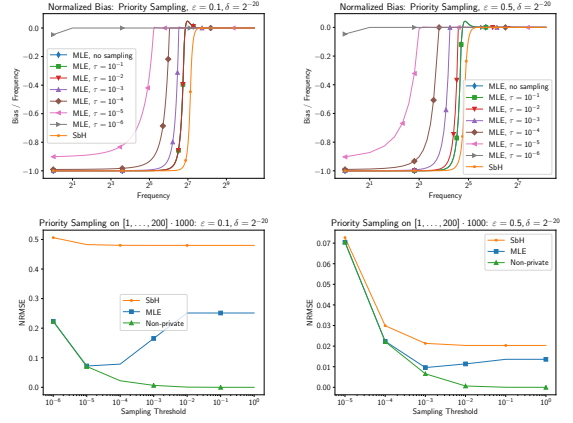


Figure 8: Top: Normalized bias for PWS (MLE) and sampled-SbH as a function of frequency, for different sampling rates. The bias of the sampled-SbH estimates (shown once) does not change with sampling rate. Bottom: NRMSE as a function of sampling rate for a selection of 2×10^5 keys with frequencies drawn uniformly $[1, 200]$.

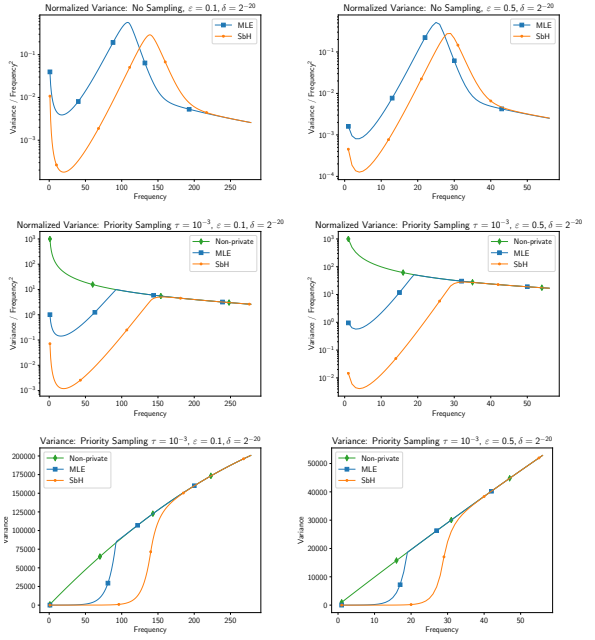


Figure 9: Normalized variance Var_i/i^2 and variance Var_i for PWS (MLE) and sampled-SbH as a function of the frequency i .

Conclusion

We presented Private Weighted Sampling (PWS), a method to post-process a weighted sample and produce a version that is differentially private. Our private samples maximize the number of reported keys subject to the privacy constraints and support estimation of linear and ordinal statistics. We demonstrate significant improvement over prior

methods for both reporting and estimation tasks, even for the well studied special case of private histograms (when there is no sampling).

An appealing direction for future work is to explore the use of PWS to design *composable* private sketches, e.g., in the context of coordinated samples. Threshold and bottom- k samples of different datasets are *coordinated* when using consistent $\{u_x\}$. Coordinated samples generalize MinHash sketches and support estimation of similarity measures and statistics over multiple datasets (Brewer et al., 1972; Saavedra, 1995; Cohen, 1997; Broder, 2000; Cohen, 2014a,b).

Acknowledgments

Part of this work was done while Ofir Geri was an intern at Google Research. This work was partially supported by Moses Charikar’s Google Faculty Research Award.

References

- A. Andoni, R. Krauthgamer, and K. Onak. Streaming algorithms via precision sampling. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, 2011. URL <https://doi.org/10.1109/FOCS.2011.82>.
- V. Balcer and S. P. Vadhan. Differential privacy on finite computers. In A. R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPICs*, pages 43:1–43:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPICs.ITCS.2018.43. URL <https://doi.org/10.4230/LIPICs.ITCS.2018.43>.
- A. Beimel, H. Brenner, S. P. Kasiviswanathan, and K. Nissim. Bounds on the sample complexity for private learning and private data release. *Mach. Learn.*, 94(3):401–437, 2014. doi: 10.1007/s10994-013-5404-1. URL <https://doi.org/10.1007/s10994-013-5404-1>.
- A. Beimel, K. Nissim, and U. Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. *Theory Comput.*, 12(1):1–61, 2016. doi: 10.4086/toc.2016.v012a001. URL <https://doi.org/10.4086/toc.2016.v012a001>.
- K. R. W. Brewer, L. J. Early, and S. F. Joyce. Selecting several samples from a single population. *Australian Journal of Statistics*, 14(3):231–239, 1972.
- A. Z. Broder. Identifying and filtering near-duplicate documents. In *Proc. of the 11th Annual Symposium on Combinatorial Pattern Matching*, volume 1848 of *LNCS*, pages 1–10. Springer, 2000.
- M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In M. Hirt and A. D. Smith, editors, *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658, 2016. doi: 10.1007/978-3-662-53641-4_24. URL https://doi.org/10.1007/978-3-662-53641-4_24.
- M. Bun, K. Nissim, U. Stemmer, and S. P. Vadhan. Differentially private release and learning of threshold functions. In V. Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 634–649. IEEE Computer Society, 2015. doi: 10.1109/FOCS.2015.45. URL <https://doi.org/10.1109/FOCS.2015.45>.
- M. Bun, C. Dwork, G. N. Rothblum, and T. Steinke. Composable and versatile privacy via truncated CDP. In I. Diakonikolas, D. Kempe, and M. Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 74–86. ACM, 2018. doi: 10.1145/3188745.3188946. URL <https://doi.org/10.1145/3188745.3188946>.
- M. Bun, K. Nissim, and U. Stemmer. Simultaneous private learning of multiple concepts. *J. Mach. Learn. Res.*, 20:94:1–94:34, 2019. URL <http://jmlr.org/papers/v20/18-549.html>.
- K. Chaudhuri and N. Mishra. When random sampling preserves privacy. In *Proceedings of the 26th Annual International Conference on Advances in Cryptology, CRYPTO’06*, page 198–213, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3540374329. doi: 10.1007/11818175_12. URL https://doi.org/10.1007/11818175_12.
- E. Cohen. Size-estimation framework with applications to transitive closure and reachability. *J. Comput. System Sci.*, 55:441–453, 1997.
- E. Cohen. Distance queries from sampled data: Accurate and efficient. In *KDD*. ACM, 2014a. full version: <http://arxiv.org/abs/1203.4903>.
- E. Cohen. Estimation for monotone sampling: Competitiveness and customization. In *PODC*. ACM, 2014b. URL <http://arxiv.org/abs/1212.0243>. full version <http://arxiv.org/abs/1212.0243>.
- E. Cohen. Stream sampling framework and application for frequency cap statistics. *ACM Trans. Algorithms*, 14(4):52:1–52:40, 2018. URL <https://doi.org/10.1145/3234338>. preliminary version published in *KDD 2015*. arXiv:<http://arxiv.org/abs/1502.05955>.
- E. Cohen and O. Geri. Sampling sketches for concave sub-linear functions of frequencies. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

- E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *ACM PODC*, 2007.
- E. Cohen and H. Kaplan. Tighter estimation using bottom-k sketches. In *Proceedings of the 34th VLDB Conference*, 2008. URL <http://arxiv.org/abs/0802.3448>.
- E. Cohen, G. Cormode, and N. Duffield. Don't let the negatives bring you down: Sampling from streams of signed updates. In *Proc. ACM SIGMETRICS/Performance*, 2012.
- E. Cohen, N. Duffield, H. Kaplan, C. Lund, and M. Thorup. Algorithms and estimators for accurate summarization of unaggregated data streams. *J. Comput. System Sci.*, 80, 2014.
- E. Cohen, O. Geri, and R. Pagh. Composable sketches for functions of frequencies: Beyond the worst case. In *ICML*, 2020a. URL <https://arxiv.org/abs/2004.04772>.
- E. Cohen, R. Pagh, and D. P. Woodruff. Wor and p 's: Sketches for ℓ_p -sampling without replacement. In *NeurIPS*, 2020b. URL <https://arxiv.org/abs/2007.06744>.
- G. Cormode, C. Procopiuc, D. Srivastava, and T. T. L. Tran. Differentially private summaries for sparse data. In *Proceedings of the 15th International Conference on Database Theory, ICDT '12*, page 299–311, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450307918. doi: 10.1145/2274576.2274608. URL <https://doi.org/10.1145/2274576.2274608>.
- D. Desfontaines, J. Voss, B. Gipson, and C. Mandayam. Differentially private partition selection, 2020. URL <https://arxiv.org/abs/2006.03684>.
- N. Duffield, M. Thorup, and C. Lund. Priority sampling for estimating arbitrary subset sums. *J. Assoc. Comput. Mach.*, 54(6), 2007.
- C. Dwork, F. McSherry, K. Nissim, and A. D. Smith. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016. doi: 10.29012/jpc.v7i3.405. URL <https://doi.org/10.29012/jpc.v7i3.405>.
- C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *SIGCOMM*. ACM, 2002.
- A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. *SIAM J. Comput.*, 41(6):1673–1693, 2012. URL <https://doi.org/10.1137/09076828X>.
- P. Gibbons and Y. Matias. New sampling-based summary statistics for improving approximate query answers. In *SIGMOD*. ACM, 1998.
- S. Gopi, P. Gulhane, J. Kulkarni, J. H. Shen, M. Shokouhi, and S. Yekhanin. Differentially private set union. In *ICML*, 2020. URL <https://arxiv.org/abs/2002.09745>.
- M. Hardt and K. Talwar. On the geometry of differential privacy. In L. J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 705–714. ACM, 2010. doi: 10.1145/1806689.1806786. URL <https://doi.org/10.1145/1806689.1806786>.
- D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- R. Jayaram and D. P. Woodruff. Perfect ℓ_p sampling in a data stream. In *FOCS*, 2018. URL <https://doi.org/10.1109/FOCS.2018.00058>.
- S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. D. Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, 2011. doi: 10.1137/090756090. URL <https://doi.org/10.1137/090756090>.
- A. Korolova, K. Kenthapadi, N. Mishra, and A. Ntoulas. Releasing search queries and clicks privately. In J. Quezada, G. León, Y. S. Maarek, and W. Nejdl, editors, *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, pages 171–180. ACM, 2009. doi: 10.1145/1526709.1526733. URL <https://doi.org/10.1145/1526709.1526733>.
- R. Kulkarni. A million news headlines [csv data file]. <https://www.kaggle.com/therohk/million-headlines/home>, 2017.
- E. Ohlsson. Sequential poisson sampling from a business register and its application to the swedish consumer price index. Technical Report 6, Statistics Sweden, 1990.
- E. Ohlsson. Sequential poisson sampling. *J. Official Statistics*, 14(2):149–162, 1998.
- A. Paranjape, A. R. Benson, and J. Leskovec. Motifs in temporal networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM '17*, pages 601–610, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450346757. doi: 10.1145/3018661.3018731. URL <https://doi.org/10.1145/3018661.3018731>.
- B. Rosén. Asymptotic theory for successive sampling with varying probabilities without replacement, I. *The Annals of Mathematical Statistics*, 43(2):373–397, 1972. URL <http://www.jstor.org/stable/2239977>.
- B. Rosén. Asymptotic theory for order sampling. *J. Statistical Planning and Inference*, 62(2):135–158, 1997.

P. J. Saavedra. Fixed sample size pps approximations with a permanent random number. In *Proc. of the Section on Survey Research Methods*, pages 697–700, Alexandria, VA, 1995. American Statistical Association.

S. Vadhan. *The Complexity of Differential Privacy*. 04 2017. ISBN 978-3-319-57047-1. doi: 10.1007/978-3-319-57048-8_7.

A Proofs: Sanitized Keys

We establish that $(\pi_i)_{i \geq 1}$ as computed by Algorithm 3 are maximum under the DP constraints and are non-decreasing.

Proof of Lemma 4.1. We use the notation $\bar{\pi}_i := 1 - \pi_i$ for the probability of a key not being included in $C(A())$. Since key inclusions in the (original or sanitized) sample are independent, the probability of a particular sanitized sample Z (set of keys, possibly empty) has the product form

$$P_{\mathbf{w}}(Z) := \prod_{x \in Z} \pi_{w_x} \prod_{x \notin Z} \bar{\pi}_{w_x} .$$

Since $\pi_i = q_i p_i$ is the product of two probabilities, one that is given (q_i) and one that we set (p_i) then our solution for $(\pi_i)_{i \geq 1}$ is realizable if and only if it satisfies the constraints for all i :

$$\pi_i \leq q_i . \quad (13)$$

We now consider the DP constraints. Consider two neighboring datasets \mathbf{w} and \mathbf{w}' and the two cases (i) For some $i \geq 1$ there is a key x such that $w_x = i$ and $w'_x = i - 1$ (ii) For some $i \geq 0$ there is a key x such that $w_x = i$ and $w'_x = i + 1$.

We consider sets T of possible outputs that we partition to outputs T^+ that include the key x and outputs T^- that do not include x . We use the notation

$$Q^+ = \sum_{Z \in T^+} P_{\mathbf{w}_{-x}}(Z \setminus \{x\})$$

$$Q^- = \sum_{Z \in T^-} P_{\mathbf{w}_{-x}}(Z)$$

for the respective sums over these sets of outputs of the probability projected on keys other than x .

The general DP constraints on a set of outputs T have one of the following form corresponding to our two cases:

$$Q^+ \pi_i + Q^- \bar{\pi}_i \leq e^\varepsilon (Q^+ \pi_{i-1} + Q^- \bar{\pi}_{i-1}) + \delta \quad (14)$$

$$Q^+ \pi_i + Q^- \bar{\pi}_i \leq e^\varepsilon (Q^+ \pi_{i+1} + Q^- \bar{\pi}_{i+1}) + \delta . \quad (15)$$

We observe, assuming monotonicity of π_i , that constraints (14) are strictest when Q^- is as small as possible. This

because $\bar{\pi}_i \leq \bar{\pi}_{i-1}$ implies $Q^- \bar{\pi}_i \leq e^\varepsilon Q^- \bar{\pi}_{i-1}$ (using $e^\varepsilon \geq 1$), so the larger Q^- is, the less strict the inequality becomes. This is achieved when $T^- = \emptyset$ and thus $Q^- = 0$, that is, T only includes outputs that include x . Similarly, assuming $\pi_i \geq \pi_{i-1}$, the constraints are strictest when Q^+ is as large as possible. That is T^+ includes all possible outcomes on keys other than x and thus $Q^+ = 1$. A similar argument shows that constraints (15) are strictest when $Q^+ = 0$ and $Q^- = 1$.

We obtain that the DP constraints simplify to

$$i \geq 1: \pi_i \leq e^\varepsilon \pi_{i-1} + \delta \quad (16)$$

$$i \geq 0: \bar{\pi}_i \leq e^\varepsilon \bar{\pi}_{i+1} + \delta . \quad (17)$$

We show below that the solution $(\pi_i)_{i \geq 1}$ of the simplified constraints (as subset of all constraints) is non-decreasing and hence any solution to the full set of constraints must also be non-decreasing and thus the assumption that led to the simplification is valid.

We observe that the constraints (16) and (17) are upper bounds on π_i in terms of π_{i-1} and the feasibility constraint (13) is also an upper bound on π_i . Therefore, each iterate π_i computed in Algorithm 3 attains the maximum possible value by the constraints, provided that π_{i-1} is at its maximum value. The claim that each π_i is maximized follows by induction.

Finally, to establish that $(\pi_i)_{i \geq 1}$ are non-decreasing we show that each term in the minimum that determines π_{i+1} is at least π_i :

- (13): $q_{i+1} \geq q_i \geq \pi_i$
- (16): $e^\varepsilon \pi_i + \delta \geq \pi_i$
- (17): $1 + e^{-\varepsilon} (\pi_i + \delta - 1) = (1 - e^{-\varepsilon}) + e^{-\varepsilon} (\pi_i + \delta) \geq (1 - e^{-\varepsilon}) \pi_i + e^{-\varepsilon} (\pi_i + \delta) = \pi_i + e^{-\varepsilon} \delta \geq \pi_i$

□

Consider the iterates π_i and the corresponding constraints sequence of the minimum among the three constraints: (13), (16), and (17). We will slightly abuse notation and use these references to constraints in expressions. We first consider the relation between (16) and (17):

Lemma A.1. *The constraints sequence has all positions with (16) preceding all positions with (17). In the typical settings of $\varepsilon \ll 1$, the highest position i before the transition has $\pi_i \leq (1 - \delta)/2$.*

Proof. The ratio of constraints as a function of $x = \pi_i$ is:

$$\frac{(16)}{(17)} = \frac{e^\varepsilon x + \delta}{1 + e^{-\varepsilon} (x + \delta - 1)} .$$

This is an increasing function for $x \in (0, 1]$. Therefore the iterates π_i are such that initially (16) is smaller (ratio is

lower than 1) and then (17) is smaller (ratio is larger than 1). Solving for the crossing point (ratio equal to one) we get

$$x = \frac{1 - \delta}{1 + e^\varepsilon} \approx \frac{1 - \delta}{2 + \varepsilon},$$

using the first order approximation $e^z \approx 1 + z$ which holds when $z \ll 1$. Also note that $x \leq (1 - \delta)/2$. \square

A subsequence of (π_i) where all constraints are (16) or all (17) has a compact form:

Lemma A.2. *The iterates $\pi_{i+1} = e^\varepsilon \pi_i + \delta$ on a subsequence with only (16) that starts at i_0 can be compactly expressed for $i > i_0$:*

$$\pi_i = \pi_{i_0} e^{(i-i_0)\varepsilon} + \delta \frac{e^{(i-i_0)\varepsilon} - 1}{e^\varepsilon - 1}. \quad (18)$$

Similarly, for a subsequence with only (17) constraints where $\bar{\pi}_{i+1} = e^{-\varepsilon}(\bar{\pi}_i - \delta)$ we get:

$$\bar{\pi}_i = \bar{\pi}_{i_0} e^{-(i-i_0)\varepsilon} - \delta e^{-\varepsilon} \frac{1 - e^{-(i-i_0)\varepsilon}}{1 - e^{-\varepsilon}}. \quad (19)$$

Proof. The iterates form a geometric series. \square

We now establish the closed-form expressions of the solution π_i^* that corresponds to $q_i = 1$ for all i .

Proof of Lemma 4.2. Since $q_i = 1$ for all i , the constraint sequence includes only (16) and (17) constraints (until the minimum exceeds 1, in which $\pi_i = 1$ at this position and all subsequent positions). From Lemma A.1 we know it has the form (16)*(17)*. We have $\pi_1^* = \delta$.

We have $\pi_1^* = \delta$ and hence while the constraint (16) holds. Using (18) (Lemma A.2) we have $\pi_i^* = \delta \frac{e^{i\varepsilon} - 1}{e^\varepsilon - 1}$. From the proof of Lemma A.1, we have (16) in the constraint sequence until $\pi_i^* > \frac{1-\delta}{1+e^\varepsilon}$. From our choice of L , we have

$$\pi_L^* = \delta \frac{e^{L\varepsilon} - 1}{e^\varepsilon - 1} = \frac{1 - \delta}{1 + e^\varepsilon}.$$

Therefore, both (16) and (17) hold at position L . We have $\pi_{L+1}^* = e^\varepsilon \pi_L^* + \delta$. From our choice of L we have $\bar{\pi}_{L+1}^* = \pi_L^*$ and $\bar{\pi}_L^* = \pi_{L+1}^*$. We apply (19) (Lemma A.2) with $i_0 = L$ to obtain the claim. Note the symmetry of the solution where for $1 \leq i \leq L$, $\bar{\pi}_{2L+1-i}^* = \pi_i^*$. \square

We are now ready to bound the number of positions i where $\pi_i < q_i$:

Proof of Lemma 4.3. We have $\pi_i < q_i$ if and only if the i th position in the constraint sequence has (16) or (17).

The sequence of π_i is non-decreasing with at most one transition from (16) to (17). For i such that $q_i \leq \delta$ we have the

constraint (13). Hence $\pi_i \geq \delta = \pi_1^*$ at the first position with $\pi_i < q_i$. Each application of the minimum of (16) and (17) increases π_i . The total increase is larger when the initial value is larger. Let h be the i th position with (16) or (17). Because of the monotone increase we can show by induction that $\pi_h \geq \pi_i^*$. Since there are $2L + 1$ positions in $(\pi_i^*)_{i \geq 1}$ with value $\pi_i^* < 1$, there can be at most $2L + 1$ positions in $(\pi_i)_{i \geq 1}$ with $\pi_i < q_i$. \square

With threshold sampling (say by moments of frequency) we have a closed form for q_i . Using this and Lemma A.2 we can express the solution π_i with computation that depends on the number of transitions in the constraint sequence. The following Lemma bounds the number of such transitions. The proof of Lemma 4.4 follows as a special case:

Lemma A.3. *For threshold ppswor with $p = 1$ the constraints sequence has the regular-expression form*

$$(16)^*(17)^*(13)^*.$$

For priority sampling with $p \leq 1$ and for priority with $p \geq 1$ when $\tau \geq \delta$, all (13) constraints must follow all (16) constraints in the constraint sequence.

Proof. From Lemma A.1 there is at most one transition from (16) to (17).

We now consider the relation between (13) and (16). When $\pi_i = q_i$, we will have (13) \leq (16) if

$$\rho(i, \delta) := \frac{q_{i+1} - \delta}{q_i} \leq e^\varepsilon. \quad (20)$$

We need to establish that once (20) holds for $i = i_0$, it continues to hold for $i \geq i_0$. Equivalently, establishing the claim for all $\varepsilon > 0$ is equivalent to establishing that $\rho(i, \delta)$ is non-increasing with i when $\rho(i, \delta) > 1$. That is, that $\rho(i, \delta)$ is non-increasing, equivalently, that the partial derivative satisfies

$$\frac{\partial \rho(i, \delta)}{\partial i} \leq 0 \quad (21)$$

when

$$q_{i+1} - q_i \geq \delta. \quad (22)$$

In some of the derivations i will be convenient to work with the continuous form of (22):

$$\frac{\partial q_i}{\partial i} \geq \delta. \quad (23)$$

Consider ppswor threshold sampling with $p = 1$. Recall that $q_i = 1 - e^{-\tau i}$. Therefore, the condition (22) is $e^{-\tau i}(1 - e^{-\tau}) \geq \delta$. It suffices to check (21) when $\delta + e^{-\tau} \leq 1$. Substituting and solving (21) we obtain that the derivative is negative when $\delta + e^{-t} \leq 1$. Therefore,

there can be at most one transition from (16) to (13) for ppswor with $p = 1$.

We next consider priority threshold sampling $q_i = \min\{1, \tau i^p\}$. For $i^p \tau \geq 1$, $q_i = 1$ and (22) does not hold. Therefore it suffices to consider $q_i = \tau i^p < 1$ and

$$\rho(i, \delta) = \frac{(i+1)^p - \frac{\delta}{\tau}}{i^p} .$$

When $p \leq 1$, $(i+1)^p - i^p \leq 1$ and thus condition (22) does not hold when $\tau \leq \delta$. Hence it suffices to consider $p \geq 1$ or $\delta < \tau$. Using the continuous form (23) with $q_i = \tau i^p$ we get that it is satisfied when

$$i^{p-1} \geq \delta / (p\tau) . \quad (24)$$

By solving (21) we get

$$(i+1)^{p-1} > \frac{\delta}{\tau} .$$

This holds for all $i \geq 1$ and $p \geq 1$ when $\delta \leq \tau$. Since it suffices to for the solution to hold for (24), we obtain the claim for $p \leq 1$. Combining, we obtain that there can be at most one transition from (16) to (13) for priority sampling with $p \leq 1$ and for $p \geq 1$ when $\tau \geq \delta$.

We next consider the relation between (13) and (17). We define

$$\bar{\rho}(i, \delta) := \frac{\bar{q}_i - \delta}{\bar{q}_{i+1}} .$$

When $\pi_i = q_i$, we will have (13) \leq (17) if

$$\bar{\rho}(i, \delta) \leq e^\varepsilon .$$

To establish the claim for all $\varepsilon > 0$ it is equivalent to establish that $\bar{\rho}(i, \delta)$ is non-increasing when it is greater than 1. Equivalently, that $\bar{q}_i - \bar{q}_{i+1} = q_{i+1} - q_i > \delta$ (same as (22) and (23)) implies that $\bar{\rho}(i, \delta)$ is non-increasing. That is,

$$\frac{\partial \bar{\rho}(i, \delta)}{\partial i} \leq 0 . \quad (25)$$

For ppswor with $p = 1$ we get

$$\bar{\rho}(i, \delta) = \frac{e^{-\tau i} - \delta}{e^{-\tau(i+1)}}$$

and that (25) holds for all $\delta \geq 0$ and $\tau > 0$. \square

B Proofs: Sanitized Keys and Frequencies

We establish properties of the values $\{\pi_{i,j}\}$ computed by Algorithm 4.

Proof of Theorem 5.1. We write the DP constraints in terms of sets T of potential outputs on pairs of neighboring datasets \mathbf{w} and \mathbf{w}' . We follow the proof of Lemma 4.1. Consider two neighboring datasets \mathbf{w} and \mathbf{w}' and the two cases: (i) For some $i \geq 1$ there is a key x such that $w_x = i$ and $w'_x = i - 1$. (ii) For some $i \geq 0$ there is a key x such that $w_x = i$ and $w'_x = i + 1$.

We consider a set of outputs T . A potential output $Z \in T$ is a set of key value pairs (y, j) where key y is reported with value j . For purposes of this proof we partition T to sets T_j according to the output on key x . If $(x, j) \in Z$ we place $Z \in T_j$ and if key x is not in Z we place Z in T_0 .

We denote by Q_j the respective combined probability of outputs T_j when projected on all keys other than x (equivalently, the probability of T_j when key x is removed from the dataset). The general DP constraints have the form

$$\Pr[C(A(\mathbf{w})) \in T] \leq e^\varepsilon \Pr[C(A(\mathbf{w}')) \in T] + \delta .$$

We have

$$\Pr[C(A(\mathbf{w})) \in T] = \sum_{j=0}^i Q_j \pi_{i,j} .$$

From the two choices of the neighboring dataset \mathbf{w}' we have one of:

$$\Pr[C(A(\mathbf{w}')) \in T] = \sum_{j=0}^{i-1} Q_j \pi_{i-1,j}$$

$$\Pr[C(A(\mathbf{w}')) \in T] = \sum_{j=0}^{i+1} Q_j \pi_{i+1,j} .$$

The corresponding DP constraints are:

$$i \geq 1 : \sum_{j=0}^i Q_j \pi_{i,j} \leq e^\varepsilon \sum_{j=0}^{i-1} Q_j \pi_{i-1,j} + \delta \quad (26)$$

$$i \geq 1 : \sum_{j=0}^i Q_j \pi_{i,j} \leq e^\varepsilon \sum_{j=0}^{i+1} Q_j \pi_{i+1,j} + \delta . \quad (27)$$

Considering any particular set of values $\pi_{i,j}$, the strictest constraints of the form (26) would have $Q_j = 1$ for j where $\pi_{i,j} > \pi_{i-1,j}$ and $Q_j = 0$ otherwise. Similarly for constraints of the form (27), the strictest would have $Q_j = 1$ for j where $\pi_{i,j} > \pi_{i+1,j}$ and $Q_j = 0$ otherwise. By "strictest" constraints we mean that if the values $\{\pi_{i,j}\}$ satisfy these selected constraints, they satisfy all DP constraints.

Therefore, taking the union of all these "strictest" sets of constraints over all $\{\pi_{i,j}\}$ we obtain that without loss of generality it suffices to solve for constraints of the follow-

ing form: For $i \geq 1$ and all $J \subset \{0, \dots, i\}$:

$$\sum_{j \in J} \pi_{i,j} \leq e^\epsilon \sum_{j \in J} \pi_{i-1,j} + \delta \quad (28)$$

$$\sum_{j \in J} \pi_{i,j} \leq e^\epsilon \sum_{j \in J} \pi_{i+1,j} + \delta . \quad (29)$$

We re-arrange the set of constraints (29) and get the equivalent set for all i and $J \subset \{0, \dots, i\}$:

$$\sum_{j \in J} \pi_{i,j} \geq e^{-\epsilon} \left(\sum_{j \in J} \pi_{i-1,j} - \delta \right) . \quad (30)$$

For sets J , the constraints (28) and (6) determine upper bounds on $\sum_{j \in J} \pi_{i,j}$ as determined by $\sum_{j \in J} \pi_{i-1,j}$ for $j \in J$. The constraints (30) determine lower bounds.

The solution is constructed by increasing i , so that row i is set after rows $h < i$ are set. We first set $\pi_{i,0} = 1 - \pi_i$. We set other entries so that

$$\sum_{j=1}^i \pi_{i,j} = \pi_i \quad (31)$$

(we will establish inductively that we can always satisfy (31)). Note that such settings imply that (6) are satisfied. At the high level, we set the entries $\pi_{i,\bullet}$ using two passes: The first pass is performed in increasing order by $j \geq 1$ and we greedily compute the minimum values $D_{i,j}$ we can have for $\pi_{i,j}$ so that constraints of type (30) are satisfied. The second pass is performed in decreasing order by $j \leq i$ and greedily sets $\pi_{i,j}$ to the maximum value we can while satisfying the constraints (28) and (6). The second pass finishes at an entry h so that $\sum_{j=1}^{h-1} D_{i,j} + \sum_{j=h}^i \pi_{i,j} = \pi_i$. We set all entries $1 \leq j < h$ to $\pi_{i,j} \leftarrow D_{i,j}$. A subtlety is that during the passes we only use the respective constraints (30) and (28) for J of prefix or suffix forms, but we establish that the constructed solution satisfies these constraints for all J .

We now elaborate on the first pass. The value $D_{i,j}$ for $j \geq 1$ is set to the minimum needed for $\pi_{i,j}$ so that a setting of $\pi_{i,h} \leftarrow D_{i,h}$ for $h < j$ satisfies (30) for $J = \{0, \dots, j\}$ and $J = \{1, \dots, j\}$ (We refer to such J as being in *prefix form*). Specifically, for $j = 1, \dots, i-1$ in order we get:

$$D_{i,j} \leftarrow \max\left\{0, e^{-\epsilon} \left(\sum_{h=1}^j \pi_{i-1,h} - \delta \right) - \sum_{h=1}^{j-1} D_{i,h} + \max\{0, e^{-\epsilon} \pi_{i-1,0} - \pi_{i,0}\} \right\} .$$

Let the breakpoint b_1 be maximum such that

$$\sum_{j=1}^{b_1} \pi_{i-1,j} + \max\{0, e^{-\epsilon} \pi_{i-1,0} - \pi_{i,0}\} \leq \delta .$$

Then $D_{i,\bullet}$ has $D_{i,j} = 0$ for $j \leq b_1$ and $D_{i,j} = e^{-\epsilon} \pi_{i-1,j}$ for $j > b_1 + 1$. It follows that

$$\forall j \geq 1, D_{i,j} \leq e^{-\epsilon} \pi_{i-1,j} . \quad (32)$$

Also note (from the form of the constraints) that for all $j \geq 1$, $\sum_{h=1}^j D_{i,j}$ is a lower bound on $\sum_{h=1}^j \pi_{i,j}$ for all $\pi_{i,\bullet}$ that satisfy (30) on all J with prefix form. That is, we could not get a lower prefix with a non-greedy setting of $D_{i,j}$.

We now provide details on the second pass. We visit entries $j \leq i$ in decreasing order and set $\pi_{i,j} > D_{i,j}$ as follows. Let

$$D_i \leftarrow \sum_{h=1}^{i-1} D_{i,h} .$$

From (32) we get

$$D_i \leq \sum_{h=1}^{i-1} \pi_{i-1,h} = \pi_{i-1} \leq \pi_i . \quad (33)$$

We compute the initial probability mass to allocate $R \leftarrow \pi_i - D_i$. From (33) we get that $R \geq 0$. While allocating we maintain that $\sum_{h=j}^i (\pi_{i,h} - D_{i,h}) \leq R$. For j , we first compute the maximum value U we can have for $\pi_{i,j}$ given values we already set for $\pi_{i,h}$ for $h > j$ so that the constraint (28) with $J = \{j, \dots, i\}$ is satisfied:

$$U \leftarrow e^\epsilon \sum_{h=j}^{i-1} \pi_{i-1,h} + \delta - \sum_{h=j+1}^i \pi_{i,h} .$$

We then compute the increase $\Delta \leftarrow U - D_{i,j}$. If $\Delta \leq R$ we set $\pi_{i,j} \leftarrow U$ and $R \leftarrow R - \Delta$. Otherwise, we set $\pi_{i,j} \leftarrow D_{i,j} + R$ and $R \leftarrow 0$. The pass terminates when $R = 0$ and we let b_2 be the j value when the pass terminates. ($\pi_{i,j}$ for $1 \leq j < b_2$ are set to $D_{i,j}$). The solution has the following structure. If initially $R \leq \delta$, the second pass stops at $b_2 = i$. Otherwise, $\pi_{i,i} = \delta$ and for $b_2 < j \leq i-1$, $\pi_{i,j} = e^\epsilon \pi_{i-1,j}$. Note that (due to the form of the constraints) is that the construction maximizes $\sum_{h=j}^i \pi_{i,j}$ for all $\pi_{i,\bullet}$ that satisfy (28) on all J with suffix form.

We next establish that the constructed $\pi_{i,\bullet}$ satisfy (31). For that, we need to show that we can always exhaust R . Note that if the situation is that current values have $\sum_{j=1}^i \pi_{i,j} < \pi_i$ we must have slack and $U > \pi_{i,j}$ in the constraints. This since $\sum_{j=1}^i \pi_{i,j} < \pi_i \leq e^\epsilon \pi_{i-1}$. We obtain $U = \pi_i - \sum_{j=2}^i \pi_{i,j} > \pi_{i,1}$ when processing $j = 1$.

Finally, we need to establish that the constructed $\pi_{i,\bullet}$ satisfies constraints (30) and (28) for all J .

From the construction, we have that any solution with $\pi_{i,j} \geq D_{i,j}$ for all j , and in particular, the one constructed, satisfy (30) for all J that has the form $\{0, \dots, h\}$ or $\{1, \dots, h\}$ for some h . Consider now an arbitrary J . All indices with $j > b_1$ have $\pi_{i,j} \geq e^{-\epsilon} \pi_{i-1,j}$. Therefore

$J' \leftarrow J \setminus \{b_1 + 1, \dots, i\}$ is such that if (30) holds for J' , it must hold for J . The entries $1 \leq j < b_1$ have $\pi_{i,j} = 0$. Thus $J'' \leftarrow J' \cup \{1, \dots, b_1 - 1\}$ is such that if (30) holds for J'' , it must hold for J' . Now note that J'' has a prefix form starting at 0 or 1 and ending at $b_1 - 1$ or b_1 . Therefore, (30) holds for J'' and hence, also holds for J .

From the construction, we know that (28) is satisfied for any J that has a suffix form (h, \dots, i) for $h \geq b_2$. Consider now an arbitrary J . Now recall that for all $j < i$, $\pi_{i,j} \leq e^\varepsilon \pi_{i-1,j}$. Therefore, if (28) holds for $J' = J \setminus \{0, \dots, i-1\}$, it must hold for J . But either $J' = \{i\}$ or J' is empty and in both cases (28) holds. \square

Proof of Lemma 5.2. We need to show that $(\pi_{i,j}^*)_{j=\max\{1, i-2L\}}$ is a suffix of the following sequence of length $2L + 1$:

$$\delta \cdot (1, e^\varepsilon, \dots, e^{\varepsilon(L-1)}, e^{\varepsilon L}, e^{\varepsilon(L-1)}, \dots, e^\varepsilon, 1) \quad (34)$$

We inductively verify that the construction in Algorithm 4 is such that if form (34) holds for $\pi_{i,\bullet}^*$, it also holds for $\pi_{i+1,\bullet}^*$. We use properties of the solution outlined in the proof of Theorem 5.1.

Note that for $i \leq 2L + 1$, π_i^* is equal to the sum of the length- i suffix of the sequence and for $i > 2L + 1$ we have $\pi_i^* = 1$.

For $i \leq L$: the algorithm first sets lower bounds on $(\pi_{i+1,j}^*)_{j=1}^i$ that are pointwise below the respective $\pi_{i,j}^*$. The setting of final values will set $\pi_{i+1,i+1}^* \leftarrow \delta$ and for $h = 1, \dots, \min\{i, L\}$:

$$\pi_{i+1,i+1-h}^* \leftarrow e^\varepsilon \pi_{i,i+1-h}^* .$$

For $L + 1 \leq i \leq 2L$: We have that $\overline{\pi_{i+1}^*} = \pi_{i+1,0}^* = e^{-\varepsilon}(\pi_{i,0}^* - \delta)$. Therefore, the lower bounds the algorithm computes for $(\pi_{i+1,j}^*)_{j=1}^i$ are respectively $e^{-\varepsilon} \pi_{i,j}^*$. This is consistent with the prefix for $j \leq i - L$. The setting of final values will be according to the suffix until the sum is equal to π_{i+1}^* , which will be the case when $j = i - L$.

For $i \geq 2L + 1$: In this case $\pi_i^* = 1$ and $\pi_{i+1}^* = 1$ and therefore $\pi_{i,0}^* = 0$. We have from our assumption $\pi_{i,j}^* = 0$ for $j < i - (2L)$. Therefore, the respective lower bounds on $\pi_{i+1,j}^*$ are 0. We have $\pi_{i,i-2L}^* = \delta$ and thus the lower bound $\pi_{i+1,i-2L}^*$ is 0 and we obtain lower bounds. The lower bounds on $\pi_{i+1,i+1-h}^*$ for $1 \leq h < 2L$ will respectively be $e^{-\varepsilon} \pi_{i,i+1-h}^*$. The setting of final values will set $\pi_{i+1,i+1}^* \leftarrow \delta$ and set for $h = 1, \dots, L$:

$$\pi_{i+1,i+1-h}^* \leftarrow e^\varepsilon \pi_{i,i+1-h}^*$$

At that point we allocated all of the probability mass of 1 and have the claimed form (34) for row $i + 1$. \square

C Frequency Sanitizer with Maximum Separation

We express a general form of frequency sanitizers and then establish that the sanitizer of Algorithm 5 can be discretized and has the stated properties.

C.1 General Form of Frequency Sanitizers

We express a general form of *frequency sanitizers* (frequency reporting schemes), in order to facilitate a discussion of when a scheme is optimal. We consider the sparse case, where a scheme never reports keys with frequency 0. A scheme is designed for a given non-decreasing sampling probabilities $(q_i)_{i \geq 1}$ and privacy parameters (ε, δ) as in Algorithm 2 and is specified by a sequence $(f_i)_{i \geq 0}$ of probability density functions (PDF) over the support $\mathfrak{R}_{\geq 0}$. The PDF f_i is the distribution of end-to-end reported values for a key with frequency $i \geq 0$. Each f_i has discrete probability mass at the point $x = 0$, which corresponds to the event that the key is not reported. To simplify presentation, we denote the value of the discrete point mass by $f_i(0)$ and assume that there is continuous mass density on $(0, \infty)$. Note that for realizability we must have $f_i(0) \geq 1 - q_i$. In particular, the function f_0 , which specifies the reporting function for a key with frequency 0, has a discrete mass of 1 at $x = 0$ and is 0 elsewhere (in our sparse case, keys with frequency 0 are not sampled or reported). The end-to-end sampling and sanitization is equivalent to drawing independently, for each key x with frequency i in the dataset, $j_x \sim f_i$, and reporting (x, j_x) if $j_x > 0$.

The sanitization of a set of *sampled* key and frequency pairs is performed as in Algorithm 2, with the following distributions f_i^- playing the role of the discrete $p_{i,\bullet}$: $f_i^-(0) := (f_i(0) - (1 - q_i))/q_i$ and $f_i^-(x) = f_i(x)/q_i$ otherwise. The sanitization processes each (sampled) key x with frequency i by drawing independently a sanitized frequency $j_x \sim f_i^-$ and if $j_x > 0$, reporting the pair (x, j_x) . The reader can verify that the end-to-end distributions (f_i) are equivalent to sampling with (q_i) and sanitizing with (f_i^-) . For each f_i , we denote the respective Cumulative Distribution Function (CDF) by F_i and the inverse CDF by the sets $F_i^{-1}(\alpha) := \{x \mid \{F_i(x) = \alpha\}\}$, defined for $\alpha \in [f_i(0), 1)$.

Lemma C.1. *A frequency sanitizer specified by $(f_i)_{i \geq 0}$ is (ε, δ) -DP if and only if for any measurable set $J \subset \mathfrak{R}_{\geq 0}$*

$$\forall i \geq 0, \forall J \subset \mathfrak{R}_{\geq 0}, \int_J f_i(x) dx \leq e^\varepsilon \int_J f_{i-1}(x) dx + \delta \quad (35)$$

$$\forall i \geq 0, \forall J \subset \mathfrak{R}_{\geq 0}, \int_J f_i(x) dx \geq e^{-\varepsilon} \left(\int_J f_{i-1}(x) dx - \delta \right) \quad (36)$$

(Note that we allow J to include or exclude the discrete point mass at $x = 0$)

Proof. Variant on the proof of Theorem 5.1 \square

C.2 A Refined Frequency Sanitizer

Algorithm 5 specifies PDFs (f_i) for a frequency sanitizer. The scheme is a refinement of the scheme of Algorithm 4 and, as we shall see, for any (q_i) and (ε, δ) , it maximally separates sanitized values for different frequencies. At a high level, the sub-optimality in the $(\pi_{i,j})$ of Algorithm 4 stems from frequencies i for which (i) R is not exactly exhausted at U for some output value or (ii) there is no prefix of discretized outputs $\pi_{i,\bullet}$ where exactly $\delta = \sum_{j=1}^{b_i} \pi_{i,j} + \max\{0, e^{-\varepsilon} \pi_{i-1,0} - \pi_{i,0}\}$. The scheme of Algorithm 5 (which can also be discretized) introduces additional outputs so that the above breakpoints align wholly with outputs. Note that for the special case of $(\pi_{i,j}^*)$ with integral L , Algorithm 4 (Lemma 5.2) yields the same scheme as Algorithm 5. This because b_i and c_i are always integral and no new breakpoints are introduced by Algorithm 5.

Theorem C.2. *The (f_i) computed by Algorithm 5 and the sanitizer they specify satisfy:*

1. For all i , $f_i(0) = 1 - \pi_i$.
2. The sanitizer is (ε, δ) -DP.
3. For each i , there is a $c_i \geq 0$ such that subject to the above and f_{i-1} , $\int_J f_i(x) dx$ is maximized for all $J = (z, i-1]$ for $z \geq c_i$ and is minimized for all $J = [0, z]$ for $z \leq c_i$.

Proof. We first establish a more limited claim, that for each i there is c_i so that the DP constraints are satisfied for all intervals of the form $(x, i]$ for $x > c_i$ and $[0, x]$ or $(0, x]$ for all $x \leq c_i$. We will also show that under the conditions above, $\int_J f_i(x) dx$ is maximized for all $J = (z, i-1]$ for $z \geq c_i$ and is minimized for all J of the form $[0, z]$ for $z \leq c_i$.

Each f_i is constructed from f_{i-1} as follows.

From constraint 17, the maximum probability mass that can be placed on J for which $\int_J f_{i-1}(x) dx = 0$ is $\min\{\delta, \pi_i\}$. The construction places this maximum amount on $(i-1, i]$.

We construct a function $f_L : (0, i-1]$ that are the minimum values needed to satisfy (36) for all intervals $[0, z]$ or $(0, z]$ for $z \in (0, i-1]$. We show that indeed f_L has this property: The respective constraints are

$$\begin{aligned} \bar{\pi}_i + \int_{0+}^z f_i(x) dx &\geq e^{-\varepsilon} (\bar{\pi}_{i-1} + \int_{0+}^z f_{i-1}(x) dx - \delta) \\ \int_{0+}^z f_i(x) dx &\geq e^{-\varepsilon} (\int_{0+}^z f_{i-1}(x) dx - \delta) , \end{aligned}$$

(where with abuse of notation we use \int_{0+}^z to exclude the discrete mass point at $x = 0$). Combining, we get

$$\int_{0+}^z f_i(x) dx \geq \max\{e^{-\varepsilon} \bar{\pi}_{i-1} - \bar{\pi}_i, 0\} + e^{-\varepsilon} (\int_{0+}^z f_{i-1}(x) dx - \delta) . \quad (37)$$

Note that the right hand side could be negative but note it is non-decreasing. We compute a point b_i so that it is positive for all $z > b_i$ and at most 0 for $z < b_i$. The function $f_L(x)$ specified in the algorithm is equal to the maximum of the two constraints

$$\int_0^z f_L(x) dx = \max\{e^{-\varepsilon} \bar{\pi}_{i-1} - \bar{\pi}_i, 0\} + e^{-\varepsilon} (\int_{0+}^z f_{i-1}(x) dx - \delta) .$$

Finally, we note that given the placement on $(i-1, i]$, constraints (36) on intervals $(z, i]$ are satisfied if and only if $\int_z^{i-1} f_i(x) dx \leq e^\varepsilon (\int_z^{i-1} f_{i-1}(x) dx)$. The solution sets them at maximum value (equality) for all $z \geq c_i$.

We next establish that the frequency sanitizer specified by the (f_i) is (ε, δ) -DP (satisfies the DP constraints (Lemma C.1).

From the construction, the functions (f_i) satisfy that for all i , there is $c_i \geq 0$, so that

$$\begin{aligned} \forall x \in (c_i, \infty], f_i(x) &\geq e^\varepsilon f_{i-1}(x) \\ \forall x \in (0, c_i], f_i(x) &\leq e^{-\varepsilon} f_{i-1}(x) . \end{aligned}$$

We already established that (f_i) satisfies the constraints (35) on sets J that are intervals of the form (z, ∞) and satisfies the constraints (36) on intervals of the form $[0, z]$ (included the discrete point mass at $x = 0$) or $(0, z)$ (do not include the discrete point mass).

We will show that the constraints are satisfied for any J . For each set J , define the partition $J \setminus \{0\} = J^+ \cup J^-$ where $J^+ = J \cap (c_i, \infty)$ and $J^- = J \cap (0, c_i)$. Consider the constraint (35) for J . Noting that we always have $\bar{\pi}_i \leq \bar{\pi}_{i-1}$, the constraint will hold for J if it holds for J^+ . In turn, the constraint holds for J^+ if it holds for (c_i, ∞) . Consider the constraint (36) for J . If $\bar{\pi}_i > e^{-\varepsilon} \bar{\pi}_{i-1}$ or if $0 \notin J$, the constraints holds for J if it holds for J^- . In turn, it holds for J^- if it holds for $(0, c_i]$. Otherwise, the constraint holds for J if it holds for $\{0\} \cup J^-$. In turn, it will hold for $\{0\} \cup J^-$ if it holds for $[0, c_i]$. \square

Lemma C.3. *The functions $(f_i)_{i=1}^m$ are piecewise constant on $(0, \infty)$ and have at most $3m$ distinct breakpoints in total.*

Proof. The specification of f_i on $(0, i]$ is constructed from $f_{i-1} : (0, i-1]$. A constant value is assigned on $(i-1, i]$. The construction computes two points b_i and c_i so that $f_{i-1}(x) = 0$ on $(0, b_i]$, is a constant ($e^{-\varepsilon}$) times f_{i-1} on $x \in (b_i, c_i]$ and is a constant (e^ε) times f_{i-1} on $x \in (c_i, i-1]$. Therefore, f_i has at most 3 more breakpoints than f_{i-1} . The breakpoints are $\{i\}_{i=1}^m \cup \{b_i\}_{i=1}^m \cup \{c_i\}_{i=1}^m$. \square

The sanitizer can be discretized by collapsing intervals between consecutive breakpoints to discrete output points with the respective probability mass. The discretization does not impact estimation or privacy: We can always map back from the discrete sanitized frequencies to "simulate" the respective continuous ones (and vice versa) by drawing uniformly at random from a corresponding interval.

D Maximum Separation and Ordinal Statistics

D.1 Maximum Separation

We define a measure of separation between distributions f_{i_1} and f_{i_2} at a certain quantile value α and show that the (f_i) constructed by Algorithm 5 maximize it pointwise.

For (f_i) we define for all i_1, i_2 and $\alpha \in (f_{i_2}(0), 1)$ the respective functions

$$T_{i_1, i_2}(\alpha) := F_{i_1}(\inf F_{i_2}^{-1}(\alpha)). \quad (38)$$

Intuitively, lower values of $T_{i_1, i_2}(\alpha)$ for $i_1 > i_2$ and higher values for $i_1 < i_2$ mean we can separate better the data frequencies i_1, i_2 from their respective sanitized values and have higher probability of a pair being concordant. We show that the (f_i) constructed by Algorithm 5 yield a scheme that simultaneously optimizes all $T_{i_1, i_2}(\alpha)$ values:

Theorem D.1. *The (f_i) constructed by Algorithm 5, for all $\alpha \in (f_{i_2}(0), 1)$, minimize $T_{i_1, i_2}(\alpha)$ for all $i_1 > i_2$ and maximize $T_{i_1, i_2}(\alpha)$ for all $i_1 < i_2$, α , over all DP frequency sanitizers.*

Proof. The construction of $f_i : (0, i-1]$ from $f_{i-1} : (0, i-1)$ results in $T_{i, i-1}$ that is invariant to the actual distribution of f_{i-1} . From Theorem C.2, $f_i(0)$ is at a minimum (that depends only on $f_{i-1}(0)$ and the mass of $1-f_i(0)$ is pushed as high as possible for any prefix. We now consider T_{i_1, i_2} for $i_1 > i_2$ and note that it can be expressed in terms of $T_{i_1, i_1+1}, \dots, T_{i_2-1, i_2}$, using (repeatedly if needed):

$$\begin{aligned} T_{i+2, i}(\alpha) &= F_{i+2}(\inf F_i^{-1}(\alpha)) \\ &= F_{i+2}(F_{i+1}^{-1}(F_{i+1}(F_i^{-1}(\alpha)))) \\ &= F_{i+2}(F_{i+1}^{-1}(T_{i+1, i}(\alpha))) \\ &= T_{i+2, i+1}(T_{i+1, i}(\alpha)). \end{aligned}$$

Finally, note from the expressions that when $T_{i+1, i}(\alpha)$ are at a maximum for all i and α then so is T_{i_1, i_2} for all $i_1 > i_2$ and α . \square

D.2 Ordinal Statistics

Proof of Corollary 6.1. Conveniently, we can express properties of the order induced by sanitized frequencies using only T and $(f_i(0))_{i \geq 1}$:

The probability that a pair of keys with frequencies $i_1 > i_2$ is concordant is

$$f_{i_2}(0)(1-f_{i_1}(0)) + \frac{1}{2}f_{i_2}(0)f_{i_1}(0) + \int_{f_{i_2}(0)}^1 (1-T_{i_1, i_2}(\alpha))d\alpha. \quad (39)$$

For a key x , keys Y with $w_y > w_x$ for $y \in Y$ and keys Z with $w_z < w_x$ for $z \in Z$, the probability that all $\{(x, y)\}_{y \in Y}$ and $\{(x, z)\}_{z \in Z}$ pairs are concordant is

$$\begin{aligned} &f_{w_x}(0) \prod_{y \in Y} (1 - f_{w_y}(0)) \frac{1}{2^{|Z|}} \prod_{z \in Z} f_{w_z}(0) + \\ &\int_{f_{w_x}(0)}^1 f_{w_x}(\alpha) \prod_{y \in Y} (1 - T_{w_x, w_y}(\alpha)) \prod_{z \in Z} T_{w_x, w_z}(\alpha) d\alpha. \end{aligned} \quad (40)$$

The probability that a set of keys with frequencies $i_1 < i_2 < \dots < i_k$ is concordant is a sum over expressions that have a form of a constant that depends on $\{f_{i_j}(0)\}_{j \in [k]}$ times an expression of the form $\int_{f_{i_1}(0)}^1 \int_{T_{i_2, i_1}(\alpha_1)}^1 \dots d\alpha_k \dots d\alpha_1$.

Note that all these expressions are non-decreasing in $T_{i_1, i_2}(\alpha)$ for $i_1 > i_2$ and non-increasing for $i_1 < i_2$. Therefore, these expressions are maximized for the (f_i) of Algorithm 5.

The expected Kendall- τ rank correlation increases with the expected number of concordant pairs. Since the scheme maximizes the expected number of concordant pairs, it also maximizes the expected Kendall- τ rank correlation. \square

E The Biased-Down Estimator

$$a_j \leftarrow \min_{i | \pi_{i,j} > 0} \frac{g(i) - \sum_{h=1}^{j-1} a_h \pi_{i,h}}{\pi_i - \sum_{h=1}^{j-1} \pi_{i,h}}. \quad (41)$$

The sequence (a_j) is non-decreasing and is guaranteed not to over-estimate (see details in Appendix E).

Note (from the structure of $\pi_{i,j}$) that the minimum is over at most $2L((\varepsilon, \delta))$ values of i and each sum can be over at most $2L(\varepsilon, \delta)$ positive $\pi_{i,h}$ entries (between $i \pm L(\varepsilon, \delta)$).

Lemma E.1. *The estimator expressed by the sequence (41)*

$$a_j \leftarrow \min_{i | \pi_{i,j} > 0} \frac{g(i) - \sum_{h=1}^{j-1} a_h \pi_{i,h}}{\pi_i - \sum_{h=1}^{j-1} \pi_{i,h}}. \quad (42)$$

is biased-down and non-decreasing.

Proof. The estimate a_j is always set to be at most the value needed to have an unbiased estimate for $g(i)$ when $a_i = a_j$ for $i \geq j$. Therefore, the estimate can only be biased down.

Let $r_{i,j} = \frac{g(i) - \sum_{h=1}^{j-1} a_h \pi_{i,h}}{\pi_i - \sum_{h=1}^{j-1} \pi_{i,h}}$ and recall that a_j is set to the minimum over applicable i of $r_{i,j}$. Now note that $r_{i,j}$ is

non-decreasing with j because a_j is always set to be at most $r_{i,j}$. Since for each j we take a minimum over a set of values that can only be larger, a_j may only increase. \square

F Limits on Private Non-negative Unbiased Estimation

We show that private weighted sampling schemes with optimal key reporting generally do not admit non-negative and unbiased estimation of frequencies. The lemma below considers the case when there is no sampling, but the argument extends to sampling schemes where $\pi_i = \pi_i^*$ for an appropriate prefix of the sequence.

Lemma F.1. *Consider $q \equiv 1$ and (any) keys and frequencies sanitizer with optimal π_i^* reporting of keys. Then there is no unbiased and nonnegative estimator for frequencies.*

Proof. Consider a sanitized keys and frequencies scheme for $q = 1$. The scheme reports a key with frequency i with (optimal) probability π_i^* . When a key is reported, the scheme reports a token as a sanitized frequency. Let $\mathcal{T}[i]$ be the distribution on output tokens for a key with frequency i . To make this a distribution we use the special output \perp for the (probability π_i^*) event that the key is not reported. Using our notation we have

$$\Pr_{t \sim \mathcal{T}[i]} [t \neq \perp] = \pi_i^* .$$

Consider a token t that has positive probability to be reported with $i = 1$, that is, $\Pr_{z \sim \mathcal{T}[1]} [z = t] > 0$. We argue that for any $h \leq L(\varepsilon, \delta)$ (where L is as defined in (4))

$$\Pr_{z \sim \mathcal{T}[h]} [z = t] = e^{(h-1)\varepsilon} \Pr_{z \sim \mathcal{T}[1]} [z = t] . \quad (43)$$

The argument follows from the privacy constraints for maintaining the maximum key reporting probabilities of π_i^* . The maximum probability with frequency h for tokens that are not reported for frequency $h - 1$ is δ . Therefore, to have $\pi_h^* = \pi_{h-1}^* e^\varepsilon + \delta$ the reporting probability of each token reported for $h - 1$ must increase by a factor of at least e^ε .

We now consider estimation. A general estimator for this scheme returns an estimate with expected value a_t for output token t . Note that any unbiased estimator must be 0 when a key is not reported ($a_\perp = 0$) and for all h we have:

$$h = \mathbf{E}_{z \sim \mathcal{T}[h]} a_z \quad (44)$$

Let T_1 be the set of possible output tokens $t \neq \perp$ such that $\Pr_{z \sim \mathcal{T}[1]} [z = t] > 0$. We have $\Pr_{z \sim \mathcal{T}[1]} [z \in T_1] = \pi_1^* = \delta$ and from unbiasedness (44) with $h = 1$:

$$\mathbf{E}_{z \sim \mathcal{T}[1]} a_z = \frac{1}{\delta} .$$

Consider now the estimates for a key with frequency $h \leq L(\varepsilon, \delta)$. We use (44) to obtain:

$$h = \mathbf{E}_{z \sim \mathcal{T}[h]} a_z = \mathbf{E}_{z \sim \mathcal{T}[h]} I_{z \notin T_1} a_z + \mathbf{E}_{z \sim \mathcal{T}[h]} I_{z \in T_1} a_z$$

We will show that we can have that the second term is larger than h which will mean that the first term is negative. We use (43):

$$\begin{aligned} \mathbf{E}_{z \sim \mathcal{T}[h]} I_{z \in T_1} a_z &= \sum_{t \in T_1} \Pr_{z \sim \mathcal{T}[h]} [z = t] a_t \\ &= \sum_{t \in T_1} \Pr_{z \sim \mathcal{T}[1]} [z = t] e^{\varepsilon(h-1)} a_t \\ &= e^{\varepsilon(h-1)} \sum_{t \in T_1} \Pr_{z \sim \mathcal{T}[1]} [z = t] a_t \\ &= e^{\varepsilon(h-1)} \mathbf{E}_{z \sim \mathcal{T}[1]} a_z = e^{\varepsilon(h-1)} . \end{aligned}$$

We obtain that when $h < e^{\varepsilon(h-1)}$ holds, which is the case for example when $\varepsilon = 1$ and $h = 2$, we have $a_t < 0$ on some tokens t . This because the contribution to the expectation of the estimate of frequency h that is only due to outputs T_1 already exceed the value h . Therefore, we must have negative values $a_t < 0$ on at least some tokens $t \notin T_1$. \square

G SbH Baseline: Expressions

In this section we derive expressions for inclusion probabilities, bias, and error for the baseline method of Stability-based Histograms (Bun et al., 2019) (SbH) (see Section 8.1). We use these expressions in our empirical and analytical evaluation.

In this section we treat the weighted sampling probabilities q_i as a continuous function for $i \geq 0$ and use estimators that are continuous functions of a reported j . For consistency with other parts of the paper we maintain the discrete indices notation i, j . For a key with frequency i , we express the probability density $\phi_{i,j}$ that the key is sampled and reported with frequency $j \geq T$. The distribution $\text{Lap}[1/\varepsilon]$ is a combination of $(1/2)\text{Exp}[\varepsilon]$ and $(-1/2)\text{Exp}[\varepsilon]$.

$$\begin{aligned} j \geq i : \phi_{i,j} &= \frac{1}{2} q_j \varepsilon e^{-\varepsilon(j-i)} \\ j \leq i : \phi_{i,j} &= \frac{1}{2} q_j \varepsilon e^{-\varepsilon(i-j)} . \end{aligned}$$

The respective overall reporting probability for a key with frequency i is

$$\phi_i = \int_T^\infty \phi_{i,j} dj .$$

For estimation, we follow (8). For reported frequency j to estimate $g(i)$ for a key with frequency i we use:

$$a_j := \frac{g(j)}{q_j} .$$

Note that since this is applied after the privacy transform, the estimator is biased. But for keys with frequencies where $g(j)$ is likely to be close to $g(i)$ and q_j close to q_i this estimate would be closer to a direct use of (8) on the original data. The expected value and MSE of the estimate for a key with frequency i are:

$$\begin{aligned} \mathbf{E}_i &= \int_T^\infty a_j \phi_{i,j} dj \\ \text{MSE}_i &= (g(i))^2 \cdot (1 - \phi_i) \\ &\quad + \int_T^\infty (a_j - g(i))^2 \phi_{i,j} dj . \end{aligned}$$

Using these per-frequency expressions, we can express the MSE and bias of sum estimators for linear statistics as in Section 7.1.

G.1 Explicit expressions

Substituting $\phi_{i,j}$ and a_j we obtain explicit expressions in terms of $(q_i)_{i \geq 1}$ and $g()$:

$$\phi_i = \begin{cases} i > T : & \frac{1}{2}\varepsilon \left(\int_i^\infty q_j e^{-\varepsilon(j-i)} dj + \int_T^i q_j e^{-\varepsilon(i-j)} dj \right) \\ & = \frac{1}{2}\varepsilon \left(e^{\varepsilon i} \int_i^\infty q_j e^{-\varepsilon j} dj + e^{-\varepsilon i} \int_T^i q_j e^{\varepsilon j} dj \right) \\ i \leq T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \int_T^\infty q_j e^{-\varepsilon j} dj . \end{cases} \quad (45)$$

The expected value \mathbf{E}_i of the estimate of $g(i)$ is:

$$\mathbf{E}_i = \begin{cases} i > T : & \frac{1}{2}\varepsilon \left(e^{\varepsilon i} \int_i^\infty q_j a_j e^{-\varepsilon j} dj + e^{-\varepsilon i} \int_T^i q_j a_j e^{\varepsilon j} dj \right) \\ & = \frac{1}{2}\varepsilon \left(e^{\varepsilon i} \int_i^\infty g(j) e^{-\varepsilon j} dj + e^{-\varepsilon i} \int_T^i g(j) e^{\varepsilon j} dj \right) \\ i \leq T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \int_T^\infty g(j) e^{-\varepsilon j} dj . \end{cases} \quad (46)$$

Note that \mathbf{E}_i (and hence the bias $\mathbf{E}_i - g(i)$) does not depend on the sampling q . We express the expected value for the special case when $g(i) = i$:

$$\mathbf{E}_i = \begin{cases} i > T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \varepsilon^{-2} e^{-\varepsilon i} (\varepsilon i + 1) \\ & \quad + \frac{1}{2}\varepsilon \varepsilon^{-2} e^{-\varepsilon i} (e^{\varepsilon i} (\varepsilon i - 1) - e^{\varepsilon T} (\varepsilon T - 1)) \\ & \quad = i - \frac{1}{2} e^{-\varepsilon(i-T)} (T - \frac{1}{\varepsilon}) \\ i \leq T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \varepsilon^{-2} e^{-\varepsilon T} (\varepsilon T + 1) \\ & \quad = \frac{1}{2} e^{-\varepsilon(T-i)} (T + \frac{1}{\varepsilon}) . \end{cases} \quad (47)$$

The MSE (general q and g) is:

$$\begin{aligned} \text{MSE}_i &= (1 - \phi_i) g(i)^2 + \\ &\quad \int_T^\infty q_j (a_j - g(i))^2 \text{PDF}_{\text{Lap}[1/\varepsilon]}(j - i) dj \\ &= -2g(i) \mathbf{E}_i + (g(i))^2 \phi_i + (g(i))^2 (1 - \phi_i) + \\ &\quad \int_T^\infty \frac{(g(j))^2}{q_j} \text{PDF}_{\text{Lap}[1/\varepsilon]}(j - i) dj \\ &= (g(i))^2 - 2g(i) \mathbf{E}_i + \\ &\quad \int_T^\infty \frac{(g(j))^2}{q_j} \text{PDF}_{\text{Lap}[1/\varepsilon]}(j - i) dj \\ &= (g(i))^2 - 2g(i) \mathbf{E}_i + \\ &\quad \begin{cases} i > T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \int_i^\infty \frac{g(j)^2}{q_j} e^{-\varepsilon j} dj + \\ & \frac{1}{2}\varepsilon e^{-\varepsilon i} \int_T^i \frac{g(j)^2}{q_j} e^{\varepsilon j} dj \\ i \leq T : & \frac{1}{2}\varepsilon e^{\varepsilon i} \int_T^\infty \frac{g(j)^2}{q_j} e^{-\varepsilon j} dj . \end{cases} \end{aligned} \quad (48)$$

G.2 Expressions for Private histograms ($q \equiv 1$)

We now express the reporting probabilities ϕ_i for the case where no sampling is subsequently performed ($q \equiv 1$). From (45) we obtain:

$$\phi_i = 1 - \text{CDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(T - i) = \begin{cases} i \geq T : & 1 - \frac{1}{2\delta} e^{-(i-1)\varepsilon} \\ i < T : & \frac{1}{2} \delta e^{\varepsilon(i-1)} \end{cases} \quad (49)$$

By substituting $q \equiv 1$ and $g(i) = i$ in (48) we get

$$\text{MSE}_i = \begin{cases} i > T : & \frac{2}{\varepsilon^2} - e^{-\varepsilon(i-T)} \left(\frac{1}{2} T^2 - iT + \frac{i-T}{\varepsilon} + \frac{1}{\varepsilon^2} \right) \\ i \leq T : & i^2 + e^{-\varepsilon(T-i)} \left(\frac{1}{2} T^2 - iT - \frac{i-T}{\varepsilon} + \frac{1}{\varepsilon^2} \right) \end{cases} \quad (50)$$

G.3 Expressions with sampling

The sampling schemes we consider are parameterized by $\tau > 0$. For threshold ppswor sampling $q_j = 1 - e^{-\tau f(j)}$ or threshold Poisson $q_j = \min\{1, \tau f(j)\}$.

We express ϕ_i for ppswor threshold sampling and function of frequency $f(w) = w$:

$$\phi_i = \begin{cases} i \geq T : & 1 - \frac{1}{2} e^{-\varepsilon(i-T)} - \frac{\varepsilon}{2(\varepsilon+\tau)} e^{-i\tau} \\ & \quad + I_{\varepsilon \neq \tau} \frac{\varepsilon}{2(\varepsilon-\tau)} (e^{-\varepsilon(i-T)-\tau T} - e^{-\tau i}) \\ & \quad - I_{\varepsilon = \tau} \left(\frac{1}{2} \varepsilon (i - T) e^{-\varepsilon i} \right) \\ i \leq T : & \frac{1}{2} e^{-\varepsilon(T-i)} \left(1 - \frac{\varepsilon}{\varepsilon+\tau} e^{-T\tau} \right) \end{cases} \quad (51)$$

We now consider priority sampling with threshold τ and $f(w) = w$. We start from expressing the inclusion probability ϕ_i . Recall that in the non-private case, the inclusion probability of i is $q_i = \min\{\tau i, 1\}$.

If $T \geq \frac{1}{\tau}$,

$$\begin{aligned}\phi_i &= \int_T^\infty q_j \text{PDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(j-i) dj \\ &= \int_T^\infty \text{PDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(j-i) dj \\ &= 1 - \text{CDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(T-i) \\ &= \begin{cases} i \geq T: & 1 - \frac{1}{2\delta} e^{-(i-1)\varepsilon} \\ i < T: & \frac{1}{2} \delta e^{\varepsilon(i-1)} \end{cases}\end{aligned}$$

Otherwise, $T < \frac{1}{\tau}$.

$$\begin{aligned}\phi_i &= \int_T^\infty q_j \text{PDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(j-i) dj \\ &= \int_T^{\frac{1}{\tau}} \tau j \text{PDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(j-i) dj \\ &\quad + \int_{\frac{1}{\tau}}^\infty \text{PDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(j-i) dj \\ &= \tau \int_T^{\frac{1}{\tau}} j \cdot \frac{1}{2} \varepsilon e^{-\varepsilon|j-i|} dj + 1 - \text{CDF}_{\text{Lap}[\frac{1}{\varepsilon}]}(\frac{1}{\tau} - i)\end{aligned}$$

To compute the inclusion probability, we consider three cases:

1. $i \leq T$. In that case,

$$\phi_i = \frac{\tau}{2} \left(\left(T + \frac{1}{\varepsilon} \right) e^{(i-T)\varepsilon} - \frac{1}{\varepsilon} e^{(i-\frac{1}{\tau})\varepsilon} \right).$$

2. $T < i < 1/\tau$. In that case,

$$\phi_i = \tau \left(i - \frac{1}{2\varepsilon} e^{\varepsilon(i-\frac{1}{\tau})} - \frac{1}{2} \left(T - \frac{1}{\varepsilon} \right) e^{\varepsilon(T-i)} \right).$$

3. $i \geq 1/\tau$. In that case,

$$\phi_i = 1 - \frac{\tau}{2\varepsilon} e^{\varepsilon(\frac{1}{\tau}-i)} - \frac{\tau}{2} \left(T - \frac{1}{\varepsilon} \right) e^{\varepsilon(T-i)}.$$

To compute the MSE, we use Eq. (48), and need to compute $\int_T^\infty \frac{(g(j))^2}{q_j} \text{PDF}_{\text{Lap}[1/\varepsilon]}(j-i) dj$. In our implementation, we wrote functions that evaluate the integrals:

$$\begin{aligned}\int x e^{\varepsilon x} dx &= \frac{1}{\varepsilon^2} e^{\varepsilon x} (\varepsilon x - 1) + C \\ \int x e^{-\varepsilon x} dx &= -\frac{1}{\varepsilon^2} e^{-\varepsilon x} (\varepsilon x + 1) + C \\ \int x^2 e^{\varepsilon x} dx &= \frac{1}{\varepsilon^3} e^{\varepsilon x} (\varepsilon^2 x^2 - 2\varepsilon x + 2) + C \\ \int x^2 e^{-\varepsilon x} dx &= -\frac{1}{\varepsilon^3} e^{-\varepsilon x} (\varepsilon^2 x^2 + 2\varepsilon x + 2) + C.\end{aligned}$$

Then we considered the following cases in order to compute the MSE. If $i \leq T$, we need to compute $\frac{1}{2} \varepsilon e^{\varepsilon i} \int_T^\infty \frac{j^2}{\min\{\tau j, 1\}} e^{-\varepsilon j} dj$, and we have two cases:

1. $1/\tau < T$. In that case, the integral $\int_T^\infty \frac{j^2}{\min\{\tau j, 1\}} e^{-\varepsilon j} dj$ becomes $\int_T^\infty j^2 e^{-\varepsilon j} dj$.
2. $1/\tau \geq T$. In that case, the integral $\int_T^\infty \frac{j^2}{\min\{\tau j, 1\}} e^{-\varepsilon j} dj$ becomes $\frac{1}{\tau} \int_T^{\frac{1}{\tau}} j e^{-\varepsilon j} dj + \int_{\frac{1}{\tau}}^\infty j^2 e^{-\varepsilon j} dj$.

Similarly, if $i > T$, we need to compute

$$\frac{1}{2} \varepsilon e^{\varepsilon i} \int_i^\infty \frac{j^2}{\min\{\tau j, 1\}} e^{-\varepsilon j} dj + \frac{1}{2} \varepsilon e^{-\varepsilon i} \int_T^i \frac{j^2}{\min\{\tau j, 1\}} e^{\varepsilon j} dj.$$

and consider the three cases: (i) $1/\tau < T$, (ii) $T \leq 1/\tau < i$, and (iii) $i \leq 1/\tau$.