
Unbounded Differentially Private Quantile and Maximum Estimation

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In this work we consider the problem of differentially private computation of
2 quantiles for the data, especially the highest quantiles such as maximum, but with
3 an unbounded range for the dataset. We show that this can be done efficiently
4 through a simple invocation of `AboveThreshold`, a subroutine that is iteratively
5 called in the fundamental Sparse Vector Technique, even when there is no upper
6 bound on the data. In particular, we show that this procedure can give more
7 accurate and robust estimates on the highest quantiles with applications towards
8 clipping that is essential for differentially private sum and mean estimation. In ad-
9 dition, we show how two invocations can handle the fully unbounded data setting.
10 Within our study, we show that an improved analysis of `AboveThreshold` can
11 improve the privacy guarantees for the widely used Sparse Vector Technique that
12 is of independent interest. We give a more general characterization of privacy loss
13 for `AboveThreshold` which we immediately apply to our method for improved
14 privacy guarantees. Our algorithm only requires one $O(n)$ pass through the data,
15 which can be unsorted, and each subsequent query takes $O(1)$ time. We empiri-
16 cally compare our unbounded algorithm with the state-of-the-art algorithms in
17 the bounded setting. For inner quantiles, we find that our method often performs
18 better on non-synthetic datasets. For the maximal quantiles, which we apply
19 to differentially private sum computation, we find that our method performs
20 significantly better.

21 1 Introduction

22 In statistics, quantiles are values that divide the data into specific proportions, such as median
23 that divides the data in half. Quantiles are a central statistical method for better understanding a
24 dataset. However, releasing quantile values could leak information about specific individuals within
25 a sensitive dataset. As a result, it becomes necessary to ensure that individual privacy is ensured
26 within this computation. Differential privacy offers a rigorous method for measuring the amount
27 that one individual can change the output of a computation. Due to its rigorous guarantees, differential
28 privacy has become the gold standard for measuring privacy. This measurement method then offers
29 an inherent tradeoff between accuracy and privacy with outputs of pure noise achieving perfect
30 privacy. Thus, the goal of designing algorithms for differentially private quantile computation is to
31 maximize accuracy for a given level of privacy.

32 There are a variety of previous methods for computing a given quantile of the dataset that we will
33 cover in Section 1.2, but each of these requires known bounds on the dataset. The most effective
34 and practical method invokes the exponential mechanism [Smith \(2011\)](#). For computing multiple
35 quantiles this method can be called iteratively. Follow-up work showed that it could be called
36 recursively by splitting the dataset at each call to reduce the privacy cost of composition [Kaplan
37 et al. \(2022\)](#). Further, a generalization can be called efficiently in one shot [Gillenwater et al. \(2021\)](#).

38 1.1 Our contributions

39 In this work, we offer an alternative practical and accurate approach, Unbounded Quantile Esti-
40 mation (UQE), that also invokes a well-known technique and can additionally be applied to the
41 unbounded setting. While the commonly-used technique designs a distribution to draw from that is
42 specific to the dataset, our method will simply perform a noisy guess-and-check. Initially we assume
43 there is only a lower bound on the data, as non-negative data is common in real world datasets with
44 sensitive individual information. Our method will simply iteratively increase the candidate value
45 by a small percentage and halt when the number of data points below the value exceeds the desired
46 amount dictated by the given quantile. While the relative increase will be small each iteration, the
47 exponential nature still implies that the candidate value will become massive within a reasonable
48 number of iterations. As a consequence, our algorithm can handle the unbounded setting where we
49 also show that two calls to this procedure can handle fully unbounded data. Computing multiple
50 quantiles can be achieved by applying the recursive splitting framework from Kaplan et al. (2022).

51 Performing our guess-and-check procedure with differential privacy exactly fits AboveThreshold,
52 a method that is iteratively called in the Sparse Vector Technique Dwork et al. (2009). We also
53 take a deeper look at AboveThreshold and unsurprisingly show that similar to report noisy max
54 algorithms, the noise addition can come from the Laplace, Gumbel or Exponential distributions. We
55 further push this analysis to show that for monotonic queries, a common class of queries which
56 we will also utilize in our methods, the privacy bounds for composition within the Sparse Vector
57 Technique can be further improved. Given the widespread usage of this technique,¹ we believe this
58 result is of independent interest. Furthermore, we give a more general characterization of query
59 properties that can improve the privacy bounds of AboveThreshold. We immediately utilize this
60 characterization in our unbounded quantile estimation algorithm to improve privacy guarantees.

61 While the commonly used algorithms for quantile estimation can still apply incredibly loose bounds
62 to ensure the data is contained within, this can have a substantial impact upon the accuracy for
63 estimating the highest quantiles such as maximum. This leads to an especially important application
64 for our algorithm, differentially private sum computation, which can thereby be used to compute
65 mean as well. Performing this computation practically without assumptions upon the distribution
66 often requires clipping the data and adding noise proportionally. Clipping too high adds too much
67 noise, and clipping too low changes the sum of the data too much. The highest quantiles of the data
68 are used for clipping to optimize this tradeoff. The unbounded nature of our approach fundamentally
69 allows us to estimate the highest quantiles more robustly and improve the accuracy of differentially
70 private sum computation.

71 This improvement in differentially private sum computation is further evidenced by our empirical
72 evaluation, with significant improvements in accuracy. Our empirical comparison will be upon the
73 same datasets from previous work in the bounded setting. We also compare private computation of
74 the inner quantiles on these datasets. For synthetic datasets generated from uniform or gaussian
75 distributions, we see that the more structured approach of designing a distribution for the data from
76 the exponential mechanism consistently performs better. However, for the real-world datasets, we
77 see that our unstructured approach tends to perform better even within this bounded setting. By
78 design our algorithm is less specific to the data, so our alternative approach becomes advantageous
79 when less is known about the structure and bounds of the data *a priori*. As such, for large-scale
80 privacy systems that provide statistical analysis for a wide variety of datasets, our methods will be
81 more flexible to handle greater generality accurately.

82 1.2 Background literature

83 The primary algorithm for privately computing a given quantile, by which we compare our technique,
84 applies the exponential mechanism with a utility function based upon closeness to the true quantile
85 Smith (2011). We will discuss this algorithm, which we denote as Exponential Mechanism Quantile
86 (EMQ), in greater detail in Section A. This approach was then extended to computing multiple
87 quantiles more cleverly by recursively splitting the data and establishing that only one partition of

¹For example, see Roth and Roughgarden (2010); Hardt and Rothblum (2010); Dwork et al. (2015); Nissim et al. (2016); Nissim and Stemmer (2018); Kaplan et al. (2020a,b); Hasidim et al. (2020); Bun et al. (2017); Bassily et al. (2018); Cummings et al. (2020); Ligett et al. (2017); Barthe et al. (2016a,b); Steinke and Ullman (2016); Cummings et al. (2015); Ullman (2015); Nandi and Bassily (2020); Shokri and Shmatikov (2015); Hsu et al. (2013); Sajed and Sheffet (2019); Feldman and Steinke (2017); Blum et al. (2015); Chen et al. (2016)

88 the dataset can change between neighbors thereby reducing the composition costs [Kaplan et al.](#)
 89 (2022). Additional follow-up work showed that a generalization of the utility function to multiple
 90 quantiles could be efficiently drawn upon in one shot [Gillenwater et al. \(2021\)](#). Another recent result
 91 examined this problem in the streaming data setting and gave a method that only uses strongly
 92 sub-linear space complexity [Alabi et al. \(2022\)](#).

93 Quantile computation can also be achieved through CDF estimation [Bun et al. \(2015\)](#); [Kaplan et al.](#)
 94 (2020a). However these techniques offer limited practicality as they rely upon several reductions
 95 and parameter tuning. Recursively splitting the data is also done for CDF estimation algorithms
 96 where the statistics from each split can be aggregated for quantile computation [Dwork et al. \(2010\)](#);
 97 [Chan et al. \(2011\)](#). These techniques tend to be overkill for quantile estimation and thus suffer in
 98 accuracy comparatively.

99 We will also give improved privacy analysis of the Sparse Vector Technique which was originally
 100 introduced in [Dwork et al. \(2009\)](#). A more detailed analysis of the method can be found in [Lyu et al.](#)
 101 (2017). Additional recent work has shown that more information can be output from the method at
 102 no additional privacy cost [Kaplan et al. \(2021\)](#); [Ding et al. \(2023\)](#).

103 1.3 Organization

104 We provide the requisite notation and definitions in Section 2. In Section 3, we review the
 105 AboveThreshold algorithm from the literature and show that privacy analysis can be further
 106 improved. In Section 4, we provide our unbounded quantile estimation method. In Section 5, we
 107 test our method compared to the previous techniques on synthetic and real world datasets. In
 108 Section A, we consider the estimation of the highest quantiles which has immediate application
 109 to differentially private sum and mean estimation. In Section B, we give further results on the
 110 AboveThreshold algorithm and provide the missing proofs from Section 3. In Section C, we provide
 111 further variants and extensions of our unbounded quantile estimation technique.

112 2 Preliminaries

113 We will let x, x' denote datasets in our data universe \mathcal{X} .

114 **Definition 2.1.** *Datasets $x, x' \in \mathcal{X}$ are neighboring if at most one individual's data has been changed.*

115 Note that we use the *swap* definition, but our analysis of the AboveThreshold algorithm will be
 116 agnostic to the definition of neighboring. Using this definition as opposed to the *add-subtract*
 117 definition is necessary to apply the same experimental setup as in [Gillenwater et al. \(2021\)](#). Our
 118 differentially private quantile estimation will apply to either and we will give the privacy guarantees
 119 if we instead use the *add-subtract* definition in Section C.2.

120 **Definition 2.2.** *A function $f : \mathcal{X} \rightarrow \mathbb{R}$ has sensitivity Δ if for any neighboring datasets $|f(x) -$
 121 $f(x')| \leq \Delta$*

122 **Definition 2.3.** *[Dwork et al. \(2006b,a\)](#) A mechanism $M : \mathcal{X} \rightarrow \mathcal{Y}$ is (ϵ, δ) -differentially-private (DP)
 123 if for any neighboring datasets $x, x' \in \mathcal{X}$ and $S \subseteq \mathcal{Y}$:*

$$\Pr[M(x) \in S] \leq e^\epsilon \Pr[M(x') \in S] + \delta.$$

124 We will primarily work with pure differential privacy in this work where $\delta = 0$. We will also be
 125 considering the composition properties of the Sparse Vector Technique, and the primary method for
 126 comparison will be Concentrated Differential Privacy that has become widely used in practice due to
 127 it's tighter and simpler advanced composition properties [Bun and Steinke \(2016\)](#). This definition is
 128 instead based upon Reny divergence where for probability distributions P, Q over the same domain
 129 and $\alpha > 1$

$$D_\alpha(P||Q) = \frac{1}{\alpha - 1} \ln \mathbf{E}_{z \sim P} \left[\left(\frac{P(z)}{Q(z)} \right)^{\alpha - 1} \right]$$

130 **Definition 2.4.** *[Bun and Steinke \(2016\)](#) A mechanism $M : \mathcal{X} \rightarrow \mathcal{Y}$ is ρ -zero-concentrated-
 131 differentially-private (zCDP) if for any neighboring datasets $x, x' \in \mathcal{X}$ and all $\alpha \in (1, \infty)$:*

$$D_\alpha(M(x)||M(x')) \leq \rho\alpha.$$

132 We can translate DP into zCDP in the following way.

133 **Proposition 1.** *Bun and Steinke (2016)* If M satisfies ϵ -DP then M satisfies $\frac{1}{2}\epsilon^2$ -zCDP

134 In our examination of AboveThreshold we will add different types of noise, similar to the report
 135 noisy max algorithms Ding et al. (2021). Accordingly, we will consider noise from the Laplace, Gum-
 136 bel and Exponential distributions where $\text{Lap}(b)$ has PDF $p_{\text{Lap}}(z; b)$, $\text{Gumbel}(b)$ has PDF $p_{\text{Gumbel}}(z; b)$,
 137 and $\text{Expo}(b)$ has PDF $p_{\text{Expo}}(z; b)$ where

$$p_{\text{Lap}}(z; b) = \frac{1}{2b} \exp(-|z|/b) \quad p_{\text{Gumbel}}(z; b) = \frac{1}{b} \exp\left(-\left(\frac{z}{b} + e^{-z/b}\right)\right)$$

$$p_{\text{Expo}}(z; b) = \begin{cases} \frac{1}{b} \exp(-z/b) & z \geq 0 \\ 0 & z < 0 \end{cases}$$

138 We let $\text{Noise}(b)$ denote noise addition from any of $\text{Lap}(b)$, $\text{Gumbel}(b)$, or $\text{Expo}(b)$. We will also
 139 utilize the definition of the exponential mechanism to analyze the addition of Gumbel noise.

140 **Definition 2.5.** *McSherry and Talwar (2007)* The Exponential Mechanism is a randomized mapping
 141 $M : \mathcal{X} \rightarrow \mathcal{Y}$ such that

$$\Pr[M(x) = y] \propto \exp\left(\frac{\epsilon \cdot q(x, y)}{2\Delta}\right)$$

142 where $q : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ has sensitivity Δ .

143 3 Improved Analysis for Sparse Vector Technique

144 In this section, we review the AboveThreshold algorithm from the literature. To our knowledge,
 145 this technique has only been used with Laplace noise in the literature. Unsurprisingly, we show
 146 that Gumbel and Exponential noise can also be applied, with the former allowing for a closed
 147 form expression of each output probability. We further show that for monotonic queries the
 148 privacy analysis of the Sparse Vector Technique, which iteratively applies AboveThreshold, can
 149 be improved. Most proofs will be pushed to the appendix where we also give a more general
 150 characterization of query properties that can improve the privacy bounds of AboveThreshold.

151 3.1 Above Threshold Algorithm

152 We first provide the algorithm for AboveThreshold where noise can be applied from any of the
 153 Laplace, Gumbel or Exponential distributions.

Algorithm 1 AboveThreshold

Require: Input dataset x , a stream of queries $\{f_i : \mathcal{X} \rightarrow \mathbb{R}\}$ with sensitivity Δ , and a threshold T

- 1: Set $\hat{T} = T + \text{Noise}(\Delta/\epsilon_1)$
 - 2: **for** each query i **do**
 - 3: Set $v_i = \text{Noise}(\Delta/\epsilon_2)$
 - 4: **if** $f_i(x) + v_i \geq \hat{T}$ **then**
 - 5: Output \top and halt
 - 6: **else**
 - 7: Output \perp
 - 8: **end if**
 - 9: **end for**
-

154 We will also define a common class of queries within the literature that is often seen to achieve a
 155 factor of 2 improvement in privacy bounds.

156 **Definition 3.1.** We say that stream of queries $\{f_i : \mathcal{X} \rightarrow \mathbb{R}\}$ with sensitivity Δ is monotonic if for any
 157 neighboring $x, x' \in \mathcal{X}$ we have either $f_i(x) \leq f_i(x')$ for all i or $f_i(x) \geq f_i(x')$ for all i .

158 To our knowledge all previous derivations of AboveThreshold in the literature apply Lap noise
159 which gives the following privacy guarantees.

160 **Lemma 3.1** (Theorem 4 and 5 of [Lyu et al. \(2017\)](#)). *If the noise addition is Lap then Algorithm 1 is*
161 *$(\epsilon_1 + 2\epsilon_2)$ -DP for general queries and is $(\epsilon_1 + \epsilon_2)$ -DP for monotonic queries.*

162 Given that Expo noise is one-sided Lap noise, it can often be applied for comparative algorithms such
163 as this one and report noisy max as well. We will show this extension in Section B for completeness.

164 **Corollary 3.1.** *If the noise addition is Expo then Algorithm 1 is $(\epsilon_1 + 2\epsilon_2)$ -DP for general queries and*
165 *$(\epsilon_1 + \epsilon_2)$ -DP for monotonic queries.*

166 While the proofs for Expo noise generally follow from the Lap noise proofs, it will require different
167 techniques to show that Gumbel noise can be applied as well. In particular, we utilize the known
168 connection between adding Gumbel noise and the exponential mechanism.

169 **Lemma 3.2.** *If the noise addition is Gumbel and $\epsilon_1 = \epsilon_2$ then Algorithm 1 is $(\epsilon_1 + 2\epsilon_2)$ -DP for general*
170 *queries and $(\epsilon_1 + \epsilon_2)$ -DP for monotonic queries.*

171 We defer the proof of this to the appendix. In all of our empirical evaluations we will use Expo
172 noise in our calls to AboveThreshold because it has the lowest variance for the same parameter.
173 While we strongly believe that Expo noise will be most accurate under the same noise parameters,
174 we leave a more rigorous examination to future work. We also note that this examination was
175 implicitly done for report noisy max between Gumbel and Expo noise in [McKenna and Sheldon](#)
176 [\(2020\)](#), where their algorithm is equivalent to adding Expo noise [Ding et al. \(2021\)](#), and Expo noise
177 was shown to be clearly superior.

178 3.2 Improved privacy analysis for Sparse Vector Technique

179 In this section, we further consider the iterative application of AboveThreshold which is known
180 as the sparse vector technique. We show that for monotonic queries, we can improve the privacy
181 analysis of sparse vector technique to obtain better utility for the same level of privacy. Our primary
182 metric for measuring privacy through composition will be zCDP which we defined in Section 2 and
183 has become commonly used particularly due to the composition properties. We further show in
184 Section B that our analysis also enjoys improvement under the standard definition of differential
185 privacy. These improved properties immediately apply to our unbounded quantile estimation
186 algorithm as our queries will be monotonic.

187 **Theorem 1.** *If the queries are monotonic, then for any noise addition of Lap, Gumbel, or Expo we*
188 *have that Algorithm 1 is $\frac{1}{2}\epsilon^2$ -zCDP where $\epsilon = \frac{\epsilon_1}{2} + \epsilon_2$. If the noise addition is Gumbel then we further*
189 *require $\epsilon_1 = \epsilon_2$*

190 Note that applying Proposition 1 will instead give $\epsilon = \epsilon_1 + \epsilon_2$. It will require further techniques to
191 reduce this by $\epsilon_1/2$ which will immediately allow for better utility with the same privacy guarantees.
192 This bound also follows the intuitive factor of 2 improvement that is often expected for monotonic
193 queries. The analysis will be achieved through providing a range-bounded property, a definition
194 that was introduced in [Durfee and Rogers \(2019\)](#). This definition is ideally suited to characterizing
195 the privacy loss of selection algorithms, by which we can view AboveThreshold. As such it will also
196 enjoy the improved composition bounds upon the standard differential privacy definition shown
197 in [Dong et al. \(2020\)](#). This range bounded property was then unified with zCDP with improved
198 privacy guarantees in [Cesar and Rogers \(2021\)](#).

199 We give a proof of this theorem along with further discussion in Section B.3. We will also give
200 a generalized characterization of when we can take advantage of properties of the queries to
201 tighten the privacy bounds in AboveThreshold. We will immediately utilize this characterization
202 to improve the privacy guarantees for our method in Section C.

203 4 Unbounded Quantile Estimation

204 In this section we give our method for unbounded quantile estimation. We focus upon the lower
205 bounded setting which we view as most applicable to real-world problems, where non-negative
206 data is incredibly common, particularly for datasets that contain information about individuals.

207 This method can be symmetrically apply to upper bounded data, and in Section C.1 we will show
 208 how this approach can be extended to the fully unbounded setting.

209 For the quantile problem we will assume our data $x \in \mathbb{R}^n$. Given quantile $q \in [0, 1]$ and dataset
 210 $x \in \mathbb{R}^n$ the goal is to find $t \in \mathbb{R}$ such that $|\{x_j \in x | x_j < t\}|$ is as close to qn as possible.

211 4.1 Unbounded quantile mechanism

212 The idea behind unbounded quantile estimation will be a simple guess-and-check method that will
 213 just invoke AboveThreshold. In particular, we will guess candidate values t such that $|\{x_j \in x | x_j < t\}|$
 214 is close to qn . We begin with the smallest candidate, recall that we assume lower bounded data
 215 here and generalize later, and iteratively increase by a small percentage. At each iteration we
 216 check if $|\{x_j \in x | x_j < t\}|$ has exceeded qn , and terminate when it does, outputting the most recent
 217 candidate. In order to achieve this procedure privately, we will simply invoke AboveThreshold.
 218 We also discuss how this procedure can be achieved efficiently in Section 4.2.

219 Thus we give our algorithm here where the lower bound of the data, ℓ , is our starting candidate and
 220 β is the scale at which the threshold increases. Given that we want our candidate value to increase
 221 by a small percentage it must start as a positive value. As such we will essentially just shift the data
 222 such that the lower bound is instead 1. In all our experiments we set $\beta = 1.001$, so the increase is by
 223 0.1% each iteration.

Algorithm 2 Unbounded quantile mechanism

Require: Input dataset x , a quantile q , a lower bound ℓ , and parameter $\beta > 1$

1: Run AboveThreshold with x , $T = qn$ and $f_i(x) = |\{x_j \in x | x_j - \ell + 1 < \beta^i\}|$

2: Output $\beta^k + \ell - 1$ where k is the query that AboveThreshold halted at

224 Given that our method simply calls AboveThreshold it will enjoy all the privacy guarantees from
 225 Section 3.1. Furthermore we will show that our queries are monotonic.

226 **Lemma 4.1.** For any sequence of thresholds $\{t_i \in \mathbb{R}\}$ let $f_i(x) = |\{x_j \in x | x_j < t_i\}|$ for all i . For any
 227 neighboring dataset under Definition 2.1, we have that $\{f_i\}$ are monotonic queries with sensitivity 1.

228 *Proof.* Let x_j be the value that differs between neighbors x, x' . Define $S_{x,t} = \{x_j \in x | x_j < t\}$. We
 229 consider the case $x'_j > x_j$ and the other will follow symmetrically. For all thresholds $t_i \in (-\infty, x_j]$ we
 230 have $x_j \notin S_{x,t_i}$ and $x'_j \notin S_{x',t_i}$, so $f_i(x) = f_i(x')$. For all thresholds $t_i \in (x'_j, \infty)$ we have $x_j \in S_{x,t_i}$ and
 231 $x'_j \in S_{x',t_i}$, so $f_i(x) = f_i(x')$. Finally, for all thresholds $t_i \in (x_j, x'_j]$ we have $x_j \in S_{x,t_i}$ and $x'_j \notin S_{x',t_i}$, so
 232 $f_i(x) = f_i(x') + 1$. Therefore, $f_i(x) \geq f_i(x')$ for all i , and the sensitivity is 1.

233 □

234 Note that for the *swap* definition of neighboring the threshold remains constant. We will discuss
 235 how to extend our algorithm and further improve the privacy bounds for the *add-subtract* definition
 236 of neighboring in Section C.2.

237 4.2 Simple and scalable implementation

238 In this section we show how our call to AboveThreshold can be done with a simple linear time
 239 pass through the data, and each subsequent query takes $O(1)$ time. While the running time could
 240 potentially be infinite, if we set $\beta = 1.001$, then after 50,000 iterations our threshold is already over
 241 10^{21} and thus highly likely to have halted. Unless the scale of the data is absurdly high or the β
 242 value chosen converges to 1, our guess-and-check process will finish reasonably quickly.²

243 In our initial pass through the data, for each data point x_j we will find the index i such that
 244 $\beta^i \leq x_j - \ell + 1 < \beta^{i+1}$, which can be done by simply computing $\lfloor \log_\beta(x_j - \ell + 1) \rfloor$ as our lower
 245 bound ensures $x_j - \ell + 1 \geq 1$. Using a dictionary or similar data structure we can efficiently store
 246 $|\{x_j \in x | \beta^i \leq x_j - \ell + 1 < \beta^{i+1}\}|$ for each i with the default being 0. This preprocessing does not

²Note that the process can also be terminated at any time without affecting the privacy guarantees.

247 require sorted data and takes $O(n)$ arithmetic time, where we note that the previous algorithms
248 also measure runtime arithmetically.

249 Finally, for each query if we already have $\{x_j \in x | x_j - \ell + 1 < \beta^i\}$, then we can add $\{x_j \in x | \beta^i \leq$
250 $x_j - \ell + 1 < \beta^{i+1}\}$ in $O(1)$ time to get $\{x_j \in x | x_j - \ell + 1 < \beta^{i+1}\}$. Inductively, each query will take
251 $O(1)$ time. We provide the code in Section D.3 for easier reproducibility.

252 4.3 Extension to multiple quantiles

253 The framework for computing multiple quantiles set up in Kaplan et al. (2022) is agnostic to the
254 technique used for computing a single quantile. Their method will first compute the middle quantile
255 and split the data according to the result. Through recursive application the number of levels of
256 computation will be logarithmic. Furthermore, at each level we can see that at most one partition of
257 the data will differ between neighbors, allowing for instead a logarithmic number of compositions.
258 As such our approach can easily be applied to this recursive splitting framework to achieve the
259 same improvements in composition. This will require some minor updating of their proofs to the
260 *swap* definition that we will do in Section C.

261 5 Empirical Evaluation

262 In this section we empirically evaluate our approach compared to the previous approaches. We
263 give further detail of the previous approaches, particularly EMQ, in Section A along with strong
264 intuition upon why our approach will better handle maximal quantile estimation for data clipping.
265 We first go over the datasets and settings used for our experiments which will follow recent related
266 work Gillenwater et al. (2021); Kaplan et al. (2022). Next we evaluate how accurately our method
267 estimates quantiles for the different datasets in the bounded setting. Finally, we will consider the
268 application of computing differentially private sum, which also gives mean computation, and show
269 how our algorithm allows for a significantly more robust and accurate method when tight bounds
270 are not known for the dataset.

271 5.1 Datasets

272 We borrow the same setup and datasets as Gillenwater et al. (2021); Kaplan et al. (2022). We test our
273 algorithm compared to the state-of-the-art on six different datasets. Two datasets will be synthetic.
274 One draws 10,000 data points from the uniform distribution in the range $[-5, 5]$ and the other draws
275 10,000 data points from the normal distribution with mean zero and standard deviation of five. Two
276 datasets will come from Soumik (2019) with 11,123 data points, where one has book ratings and the
277 other has book page counts. Two datasets will come from Dua and Graf (2019) with 48,842 data
278 points, where one has the number of hours worked per week and the other has the age for different
279 people. We provide histograms of these datasets for better understanding in Section D.

280 5.2 Quantile estimation experiments

281 For our quantile estimation experiments, for a given quantile $q \in [0, 1]$ we consider the error of
282 outcome o_q from one of the private methods to be $|o_q - t_q|$ where t_q is the true quantile value. We use
283 the in-built quantile function in the numpy library with the default settings to get the true quantile
284 value. As in previous related works, we randomly sample 1000 datapoints from each dataset and
285 run the quantile computation on each method. This process is then iterated upon 100 times and the
286 error is averaged. We set $\varepsilon = 1$ as in the previous works, which will require setting $\varepsilon_1 = \varepsilon_2 = 1/2$
287 for the call to AboveThreshold in our method.

288 We will also tighten the ranges to the following, $[-5, 5]$ for the uniform dataset, $[-25, 25]$ for
289 the normal dataset, $[0, 10]$ for the ratings dataset, $[0, 10000]$ for the pages dataset, $[0, 100]$ for the
290 hours dataset, and $[0, 100]$ for the ages dataset. Given that EMQ suffers performance when many
291 datapoints are equal we add small independent noise to our non-sythetic datasets. This noise will
292 be from the normal distribution with standard deviation 0.001 for the ratings dataset and 0.1 for the
293 other three that have integer values. Our method does not require the noise addition but we will
294 use the perturbed dataset for fair comparison. True quantiles are still computed upon the original
295 data. For the datasets with integer values we rounded each output to the nearest integer. For our
296 method we set $\beta = 1.001$ for all datasets.

297 For these experiments we only compare our method UQE and the previous method EMQ, using
 298 the implementations from [Gillenwater et al. \(2021\)](#). The other procedures, discussed in Section A.1
 299 are more generalized and thus for this specific setting do not perform nearly as well which can be
 300 seen in the previous experiments [Gillenwater et al. \(2021\)](#); [Kaplan et al. \(2022\)](#), so we omit them
 301 from our results. For this experiment we consider estimating each quantile from 5% to 95% at a 1%
 302 interval. In Figure 1 we plot the mean absolute error of each normalized by the mean absolute error
 303 of UQE to make for an easier visualization.

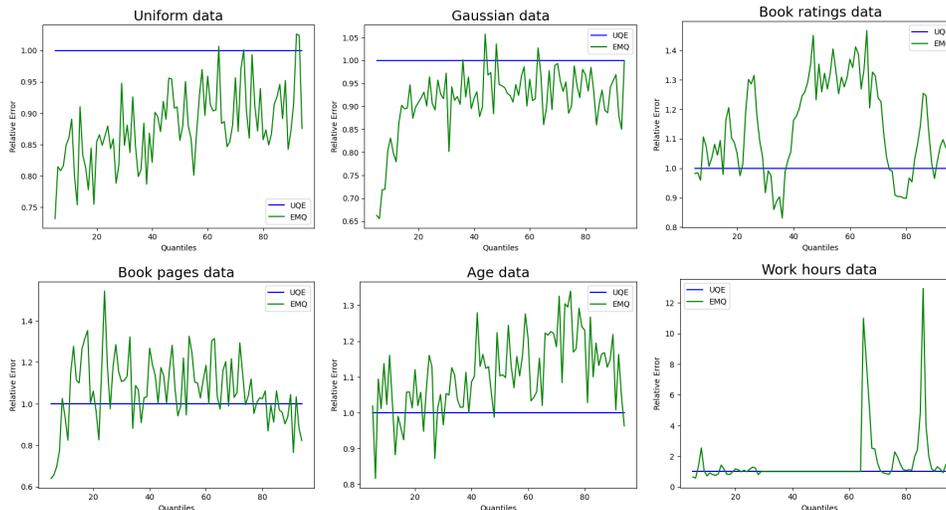


Figure 1: Plots of $UQE = (\text{mean absolute error UQE}) / (\text{mean absolute error UQE})$ and $EMQ = (\text{mean absolute error EMQ}) / (\text{mean absolute error UQE})$. Normalizing in this way will make for an easier visualization. When EMQ is below UQE then it's error is lower, and when EMQ is above UQE then it's error is higher

304 As we can see in Figure 1, EMQ consistently performs better on synthetic data, and UQE tends
 305 to performs better on the non-synthetic data. This fits with our intuition that UQE will be best
 306 suited to situations where the data is unstructured and less is known about the dataset beforehand
 307 because our guess-and-check methodology is designed to better handle ill-behaving datasets.

308 5.3 Sum estimation experiments

309 As the primary application of our method we will also be considering differentially private sum
 310 computation, which can thereby compute mean as well. We will be using the following 2ϵ -DP
 311 general procedure for computing the sum of non-negative data:

- 312 1. Let $\Delta = \text{PrivateQuantile}(x, q, \epsilon)$ where PrivateQuantile is any differentially private
 313 computation algorithm and $q \approx 1$.
- 314 2. Output $\text{Lap}(\Delta/\epsilon) + \sum_{j=1}^n \min(x_j, \Delta)$

315 We further test this upon the non-synthetic datasets. For large-scale privacy systems that provide
 316 statistical analysis for a wide variety of datasets, if we use a PrivateQuantile that requires an
 317 upper bound then we ideally want this bound to be agnostic to the dataset. This is particularly true
 318 for sum computations upon groupby queries as the range and size can differ substantially amongst
 319 groups. As such, we fix the range at $[0, 10000]$ to encompass all the datasets. We will otherwise use
 320 the same general setup as in Section 5.2.

321 We will measure the error of this procedure as the absolute error of the output and the true sum of
 322 the dataset. For each of the 100 iterations of choosing 1000 samples randomly from the full dataset,
 323 we also add Lap noise 100 times. Averaging over all these iterations gives our mean absolute error.

324 We will run this procedure with $\epsilon \in \{0.1, 0.5, 1\}$. Further we will use $q = 0.99$ always for our method,
 325 but give the absolute error for the best performing $q \in \{0.95, 0.96, 0.97, 0.98, 0.99\}$ for the other

Privacy	Method	Ratings data	Pages data	Ages data	Work hours data
$\epsilon = 1$	UQE	4.78 _{0.21}	4385.23 _{2077.16}	103.05 _{16.04}	180.48 _{44.92}
	EMQ	5.73 _{0.54}	4324.38 _{2343.25}	187.06 _{33.55}	339.06 _{75.82}
	AT	8.75 _{0.31}	4377.45 _{2340.93}	293.11 _{117.32}	471.98 _{174.07}
$\epsilon = 0.5$	UQE	9.22 _{0.31}	7102.34 _{3093.13}	180.61 _{27.03}	277.89 _{77.60}
	EMQ	6906.13 _{7123.36}	7601.47 _{3963.52}	678.22 _{1931.11}	2131.80 _{4669.11}
	AT	29.01 _{115.04}	7491.97 _{4367.31}	473.19 _{238.19}	582.58 _{369.29}
$\epsilon = 0.1$	UQE	44.59 _{1.79}	21916.37 _{6423.75}	821.77 _{157.68}	981.10 _{219.66}
	EMQ	45861.79 _{28366.46}	46552.09 _{27944.18}	50558.32 _{28843.72}	47185.58 _{29437.22}
	AT	11351.77 _{21423.40}	30830.49 _{20422.05}	14490.06 _{25268.71}	9928.18 _{20159.17}

Table 1: Mean absolute error for differentially private sum estimation. The standard deviation over the 100 iterations is also provided for each in the subscript. UQE = Our unbounded quantile estimation method. EMQ = The exponential mechanism based quantile estimation method. AT = The aggregate tree method for quantile estimation. For our method we only use $q = 0.99$. For the others we use the best performance for $q \in \{0.95, 0.96, 0.97, 0.98, 0.99\}$.

326 methods. It is important to note that this value would have to be chosen ahead of time, which
327 would add more error to the other methods. The previous methods we consider here are again the
328 EMQ, but also the aggregate tree (AT) methods, were we use both the implementation along with
329 generally best performing height (3) and branching factor (10) from [Gillenwater et al. \(2021\)](#). We
330 also implemented the bounding technique using inner quartile range within Algorithm 1 of [Smith](#)
331 [\(2011\)](#), but this performed notably worse than the others so we omitted the results from our table.

332 As we can see in Table 1, our method is far more robust and accurate. Furthermore, for our method
333 the choice of q remained constant and we can see that our results still stayed consistently accurate
334 when ϵ changed. Note that the noise added to the clipped sum is also scaled proportional to ϵ so the
335 amount the error increased as ϵ decreased for our method is what would be expected proportionally.
336 Once again these findings are consistent with our intuition. Our technique is more robust to
337 differing datasets and privacy parameters, and especially better performing for this important use
338 case.

339 Recall that the sampled data had size 1000 so dividing accordingly can give the error on mean
340 estimates. There is a long line of literature on differentially private mean estimation.³ To our knowl-
341 edge, all of these more complex algorithms either require assumptions upon the data distribution,
342 such as sub-Gaussian or bounded moments, or bounds upon the data range or related parameters,
343 and most often require both. These results also focus upon proving strong theoretical guarantees
344 of accuracy with respect to asymptotic sample complexity. We first note that our approach will
345 provide better initial bounds upon the data as seen in our experiments, which directly improve
346 the theoretical guarantees in the results that require a data range. But also our focus here is upon
347 practical methods that are agnostic to data distributions and more widely applicable to real-world
348 data. Consequently, a rigorous comparison among all of these methods would be untenable and
349 outside the scope of this work.

350 We also discuss how further tuning our parameters can give additional improvement for our method
351 in Section D.2. While these parameters must be set in advance, and we kept ours fixed within the
352 experiments for this reason, there could be improved default settings that we leave to future work.

353 References

- 354 Alabi, D., Ben-Eliezer, O., and Chaturvedi, A. (2022). Bounded space differentially private quantiles.
355 *arXiv preprint arXiv:2201.03380*.
- 356 Barthe, G., Fong, N., Gaboardi, M., Grégoire, B., Hsu, J., and Strub, P.-Y. (2016a). Advanced proba-
357 bilistic couplings for differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on*
358 *Computer and Communications Security*, pages 55–67.

³See [Kamath et al. \(2022\)](#) for an extensive review of this literature

- 359 Barthe, G., Gaboardi, M., Grégoire, B., Hsu, J., and Strub, P.-Y. (2016b). Proving differential privacy
360 via probabilistic couplings. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in*
361 *Computer Science*, pages 749–758.
- 362 Bassily, R., Thakkar, O., and Guha Thakurta, A. (2018). Model-agnostic private learning. *Advances*
363 *in Neural Information Processing Systems*, 31.
- 364 Blum, A., Morgenstern, J., Sharma, A., and Smith, A. (2015). Privacy-preserving public information
365 for sequential games. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer*
366 *Science*, pages 173–180.
- 367 Bun, M., Nissim, K., Stemmer, U., and Vadhan, S. (2015). Differentially private release and learning
368 of threshold functions. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*,
369 pages 634–649. IEEE.
- 370 Bun, M. and Steinke, T. (2016). Concentrated differential privacy: Simplifications, extensions, and
371 lower bounds. In *Theory of Cryptography: 14th International Conference, TCC 2016-B, Beijing,*
372 *China, October 31–November 3, 2016, Proceedings, Part I*, pages 635–658. Springer.
- 373 Bun, M., Steinke, T., and Ullman, J. (2017). Make up your mind: The price of online queries in
374 differential privacy. In *Proceedings of the twenty-eighth annual ACM-SIAM symposium on discrete*
375 *algorithms*, pages 1306–1325. SIAM.
- 376 Cesar, M. and Rogers, R. (2021). Bounding, concentrating, and truncating: Unifying privacy loss
377 composition for data analytics. In *Algorithmic Learning Theory*, pages 421–457. PMLR.
- 378 Chan, T.-H. H., Shi, E., and Song, D. (2011). Private and continual release of statistics. *ACM*
379 *Transactions on Information and System Security (TISSEC)*, 14(3):1–24.
- 380 Chen, Y., Machanavajjhala, A., Reiter, J. P., and Barrientos, A. F. (2016). Differentially private
381 regression diagnostics. In *ICDM*, pages 81–90.
- 382 Cummings, R., Kearns, M., Roth, A., and Wu, Z. S. (2015). Privacy and truthful equilibrium selection
383 for aggregative games. In *Web and Internet Economics: 11th International Conference, WINE 2015,*
384 *Amsterdam, The Netherlands, December 9–12, 2015, Proceedings 11*, pages 286–299. Springer.
- 385 Cummings, R., Krehbiel, S., Lut, Y., and Zhang, W. (2020). Privately detecting changes in unknown
386 distributions. In *International Conference on Machine Learning*, pages 2227–2237. PMLR.
- 387 Ding, Z., Kifer, D., Steinke, T., Wang, Y., Xiao, Y., Zhang, D., et al. (2021). The permute-and-flip mech-
388 anism is identical to report-noisy-max with exponential noise. *arXiv preprint arXiv:2105.07260*.
- 389 Ding, Z., Wang, Y., Xiao, Y., Wang, G., Zhang, D., and Kifer, D. (2023). Free gap estimates from the
390 exponential mechanism, sparse vector, noisy max and related algorithms. *The VLDB Journal*,
391 32(1):23–48.
- 392 Dong, J., Durfee, D., and Rogers, R. (2020). Optimal differential privacy composition for exponential
393 mechanisms. In *International Conference on Machine Learning*, pages 2597–2606. PMLR.
- 394 Dua, D. and Graf, C. (2019). Uci machine learning repository. In <http://archive.ics.uci.edu/ml>.
- 395 Durfee, D. and Rogers, R. M. (2019). Practical differentially private top-k selection with pay-what-
396 you-get composition. *Advances in Neural Information Processing Systems*, 32.
- 397 Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. L. (2015). Preserving
398 statistical validity in adaptive data analysis. In *Proceedings of the forty-seventh annual ACM*
399 *symposium on Theory of computing*, pages 117–126.
- 400 Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. (2006a). Our data, ourselves:
401 Privacy via distributed noise generation. In *Advances in Cryptology-EUROCRYPT 2006: 24th*
402 *Annual International Conference on the Theory and Applications of Cryptographic Techniques, St.*
403 *Petersburg, Russia, May 28–June 1, 2006. Proceedings 25*, pages 486–503. Springer.

- 404 Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006b). Calibrating noise to sensitivity in private
405 data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006,*
406 *New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer.
- 407 Dwork, C., Naor, M., Pitassi, T., and Rothblum, G. N. (2010). Differential privacy under continual
408 observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages
409 715–724.
- 410 Dwork, C., Naor, M., Reingold, O., Rothblum, G. N., and Vadhan, S. (2009). On the complexity of
411 differentially private data release: efficient algorithms and hardness results. In *Proceedings of the*
412 *forty-first annual ACM symposium on Theory of computing*, pages 381–390.
- 413 Feldman, V. and Steinke, T. (2017). Generalization for adaptively-chosen estimators via stable
414 median. In *Conference on learning theory*, pages 728–757. PMLR.
- 415 Gillenwater, J., Joseph, M., and Kulesza, A. (2021). Differentially private quantiles. In *International*
416 *Conference on Machine Learning*, pages 3713–3722. PMLR.
- 417 Hardt, M. and Rothblum, G. N. (2010). A multiplicative weights mechanism for privacy-preserving
418 data analysis. In *2010 IEEE 51st annual symposium on foundations of computer science*, pages
419 61–70. IEEE.
- 420 Hasidim, A., Kaplan, H., Mansour, Y., Matias, Y., and Stemmer, U. (2020). Adversarially robust
421 streaming algorithms via differential privacy. *Advances in Neural Information Processing Systems*,
422 33:147–158.
- 423 Hsu, J., Roth, A., and Ullman, J. (2013). Differential privacy for the analyst via private equilibrium
424 computation. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*,
425 pages 341–350.
- 426 Kamath, G., Mouzakis, A., Singhal, V., Steinke, T., and Ullman, J. (2022). A private and
427 computationally-efficient estimator for unbounded gaussians. In *Conference on Learning Theory*,
428 pages 544–572. PMLR.
- 429 Kaplan, H., Ligett, K., Mansour, Y., Naor, M., and Stemmer, U. (2020a). Privately learning thresholds:
430 Closing the exponential gap. In *Conference on Learning Theory*, pages 2263–2285. PMLR.
- 431 Kaplan, H., Mansour, Y., and Stemmer, U. (2021). The sparse vector technique, revisited. In *Conference*
432 *on Learning Theory*, pages 2747–2776. PMLR.
- 433 Kaplan, H., Schnapp, S., and Stemmer, U. (2022). Differentially private approximate quantiles. In
434 *International Conference on Machine Learning*, pages 10751–10761. PMLR.
- 435 Kaplan, H., Sharir, M., and Stemmer, U. (2020b). How to find a point in the convex hull privately.
436 *SoCG*.
- 437 Ligett, K., Neel, S., Roth, A., Waggoner, B., and Wu, S. Z. (2017). Accuracy first: Selecting a differential
438 privacy level for accuracy constrained erm. *Advances in Neural Information Processing Systems*,
439 30.
- 440 Lyu, M., Su, D., and Li, N. (2017). Understanding the sparse vector technique for differential privacy.
441 *Proceedings of the VLDB Endowment*, 10(6).
- 442 McKenna, R. and Sheldon, D. R. (2020). Permute-and-flip: A new mechanism for differentially
443 private selection. *Advances in Neural Information Processing Systems*, 33:193–203.
- 444 McSherry, F. and Talwar, K. (2007). Mechanism design via differential privacy. In *48th Annual IEEE*
445 *Symposium on Foundations of Computer Science (FOCS’07)*, pages 94–103. IEEE.
- 446 Nandi, A. and Bassily, R. (2020). Privately answering classification queries in the agnostic pac model.
447 In *Algorithmic Learning Theory*, pages 687–703. PMLR.
- 448 Nissim, K., Raskhodnikova, S., and Smith, A. (2007). Smooth sensitivity and sampling in private
449 data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*,
450 pages 75–84.

- 451 Nissim, K. and Stemmer, U. (2018). Clustering algorithms for the centralized and local models. In
 452 *Algorithmic Learning Theory*, pages 619–653. PMLR.
- 453 Nissim, K., Stemmer, U., and Vadhan, S. (2016). Locating a small cluster privately. In *Proceedings*
 454 *of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages
 455 413–427.
- 456 Roth, A. and Roughgarden, T. (2010). Interactive privacy via the median mechanism. In *Proceedings*
 457 *of the forty-second ACM symposium on Theory of computing*, pages 765–774.
- 458 Sajed, T. and Sheffet, O. (2019). An optimal private stochastic-mab algorithm based on optimal
 459 private stopping rule. In *International Conference on Machine Learning*, pages 5579–5588. PMLR.
- 460 Shokri, R. and Shmatikov, V. (2015). Privacy-preserving deep learning. In *Proceedings of the 22nd*
 461 *ACM SIGSAC conference on computer and communications security*, pages 1310–1321.
- 462 Smith, A. (2011). Privacy-preserving statistical estimation with optimal convergence rates. In
 463 *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822.
- 464 Soumik (2019). Goodreads-books dataset. In <https://www.kaggle.com/jealousleopard/goodreadsbooks>.
- 465 Steinke, T. and Ullman, J. (2016). Between pure and approximate differential privacy. *Journal of*
 466 *Privacy and Confidentiality*.
- 467 Ullman, J. (2015). Private multiplicative weights beyond linear queries. In *Proceedings of the 34th*
 468 *ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 303–312.

469 A Maximal Quantiles Estimation

470 In this section, we specifically consider applying our unbounded quantile estimation to computing
 471 the highest quantiles. This is commonly needed for clipping the dataset to privately compute the
 472 sum and mean, which we empirically test in Section 5.3. Our primary goal here is to build intuition
 473 upon why our approach fundamentally gives more accurate and robust estimations of the highest
 474 quantiles when only loose upper bounds on the data are known. We first give more details upon the
 475 previous methods, particularly the EMQ method. Then we give a more detailed look at how these
 476 methods are negatively affected by loose upper bounds and why ours performs well in comparison.

477 A.1 Previous techniques

478 The most effective and practical previous method for quantile estimation, EMQ, builds a distribu-
 479 tion over an assumed bounded range $[a, b]$ through an invocation of the exponential mechanism.
 480 Assuming the data is sorted, it will partition the range based upon the data and select an interval
 481 $[x_j, x_{j+1}]$ with probability proportional to

$$\exp\left(-\frac{\epsilon|j - qn|}{2}\right) (x_{j+1} - x_j)$$

482 where the intervals $[a, x_1]$ and $[x_n, b]$ are also considered. A uniformly random point is then drawn
 483 from within the selected interval. Note that this utility function is not monotonic and will not enjoy
 484 the same improved privacy bounds.

485 There are other approaches that recursively partition the data range while computing statistics
 486 on each partition, then aggregate these statistics across logarithmic levels to reduce the noise for
 487 quantile computation Dwork et al. (2010); Chan et al. (2011). Another technique is to instead utilize
 488 the local sensitivity while maintaining privacy by using a smoothing procedure that can also be
 489 used for computing quantiles Nissim et al. (2007). There is also a similar approach to ours within
 490 Chen et al. (2016) for bounding an unbounded dataset that increases the bounds by a factor of 2
 491 each iteration but we will discuss in Section A.2 why this variant performs substantially worse.

492 **A.2 Effect of bounds upon maximum quantiles**

493 As mentioned, EMQ will construct a probability distribution over the range $[a, b]$. The corresponding
 494 PDF is a unimodal step function of the intervals defined by the data with the peak interval $[x_j, x_{j+1}]$
 495 being such that j closest to qn . Note then that if $b \gg x_n$ then the probability of selecting $[x_n, b]$
 496 increases dramatically. The exponential decay as the PDF moves away from the peak interval
 497 diminishes the impact of $[x_n, b]$ significantly if n is far away from qn . However, for computing
 498 the highest quantiles, we want q close to 1 by definition, and the looseness of the upper bound
 499 drastically effects the accuracy.

500 In contrast, as soon as the candidate value for our method exceeds the true quantile value, each
 501 successive query will have an output of at least qn which is the threshold. The probability of
 502 continuing for λ more queries then exponentially decreases in λ .⁴ For ease of comparison, we can
 503 modify our Algorithm 2 such that in the last step, the output is drawn uniformly from $[\beta^{k-1}, \beta^k]$,
 504 assuming $\ell = 1$ for simplicity. This would then create a continuous probability distribution that
 505 would similarly be a step function with each interval $[\beta^{k-1}, \beta^k]$ being a step.

506 For better intuition we plot the approximate PDF of each in Figure 2. We use the Gumbel noise
 507 in the call to AboveThreshold to take advantage of the closed form expression in Lemma B.2 for
 508 easier PDF computation.

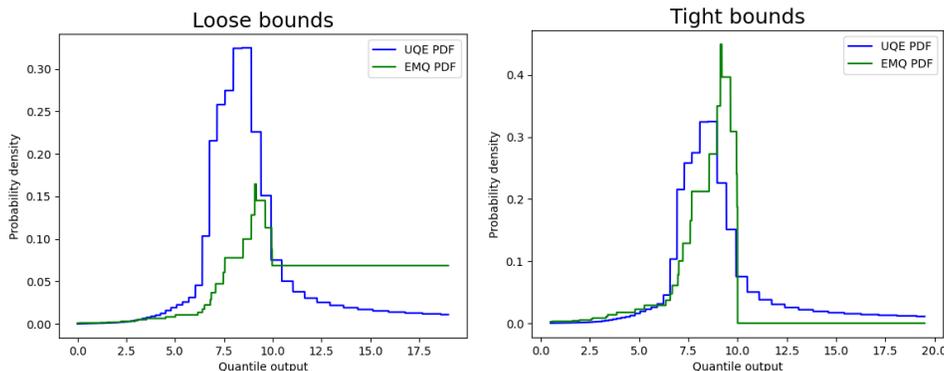


Figure 2: Illustrative example of how the approximate PDF of the previous method, EMQ is affected by looser upper bounds compared to our unbounded method UQE. A small amount of data was drawn uniformly from $[0, 10]$ and we set $q = 0.9$, so accurate output would be about 9. The left and right side assumed a range of $[0, 20]$ and $[0, 10]$ respectively for EMQ.

509 As we can see in Figure 2, the upper bound, b , increasing from 10 to 20 dramatically increases the
 510 probability of selecting a point in $[x_n, b]$ which changes the normalization for the other intervals,
 511 significantly altering the PDF of EMQ. Given that our method is unbounded, the PDF for UQE stays
 512 the same in both figures and sees the expected exponential decay once the candidate quantile passes
 513 the true quantile.

514 There are also alternative methods for bounding the data range by computing the interquartile
 515 range more accurately and scaling up. However these methods make strong assumptions upon the
 516 data distribution being close to the normal distribution as well as bounds upon the third moment
 517 Smith (2011). For real-world datasets, the tails of the data can vary significantly, and scaling up
 518 in a data-agnostic manner can often be similarly inaccurate. We also found this to be true in our
 519 experiments.

520 The smooth sensitivity framework will run into similar issues when q is closer to 1 because fewer
 521 data points need to be changed to push the quantile value to the upper bound. If this upper bound is
 522 large, then the smooth local sensitivity will still be substantial. Aggregate tree methods are slightly
 523 more robust to loose upper bounds as they aren't partitioning specific to the data. However, for

⁴We expect λ to most likely be small and this only adds a factor of β^λ additional error. Note that setting β too high, such as $\beta = 2$ which gives a similar algorithm to Chen et al. (2016), implies that even taking five additional queries will lead to a value that is over 32 times too large.

524 accurate estimates the partitioning of the range still requires the data be well-distributed across the
 525 evenly split partitions, which is significantly affected by loose upper bounds.

526 B Improved analysis for Sparse Vector Technique

527 In this section we complete the analysis for improving the privacy bounds for sparse vector technique
 528 with monotonic queries. We will first give a more generalized characterization of query properties
 529 by which we can improve the privacy bounds for AboveThreshold. While this characterization is
 530 more complex, we will immediately utilize it in Section C to improve the privacy bounds of our
 531 methods under an alternate common definition of neighboring. Additionally, the class of queries
 532 that we apply it to are not monotonic, so improvements can be made in a multitude of settings.
 533 Next we will review the definition of range-bounded, a property of privacy mechanisms that can be
 534 used to improve composition bounds, and apply it to our setting. Finally we will utilize these results
 535 to show that the privacy analysis of the sparse vector technique can be improved more generally,
 536 but also specifically for monotonic queries.

537 B.1 Generalized characterization

538 We provide a more general characterization for the privacy loss of AboveThreshold that is specific
 539 to the input stream of queries. This will be achieved by providing a one-sided privacy parameter
 540 for each pair of neighboring datasets where order matters. For a given pair x, x' , our goal will be to
 541 upper bound the output distribution of x by the output distribution of x' up to an exponential factor
 542 of the one-sided privacy parameter. The specificity of the parameter to the neighboring datasets
 543 will reduce conciseness but it is for this reason that we will be able to further improve privacy
 544 analysis. This precise characterization will also help improve bounds for our method in Section C
 545 without requiring onerous analysis.

546 **Definition B.1.** For a stream of queries $\{f_i\}$ with sensitivity Δ , we define our one-sided privacy loss
 547 for neighboring data sets x, x' as

$$\varepsilon(x, x') = \max_k \left(\frac{\varepsilon_1}{\Delta} \Delta_k(x, x') + \frac{\varepsilon_2}{\Delta} \max\{0, \Delta_k(x, x') - (f_k(x') - f_k(x))\} \right)$$

548 where $\Delta_k(x, x') = \max_{i < k} \max\{0, f_i(x') - f_i(x)\}$.

549 Given that our definition is meant to encompass the one-sided privacy loss for AboveThreshold
 550 we will now prove that fact for Expo and Gumbel noise (and Lap follows equivalently to Expo). We
 551 first prove this conjecture for Expo noise.

552 **Lemma B.1.** For a stream of queries $\{f_i\}$ with sensitivity Δ , threshold T and neighboring data sets
 553 x, x' , if we run Algorithm 1 with Expo noise, then for any given outcome $\{\perp^{k-1}, \top\}$ we have

$$\Pr[\text{AboveThreshold}(x, \{f_i\}, T)] = \{\perp^{k-1}, \top\} \leq \exp(\varepsilon(x, x')) \Pr[\text{AboveThreshold}(x', \{f_i\}, T)] = \{\perp^{k-1}, \top\}$$

554 *Proof.* Let $v_i \sim \text{Expo}(\Delta/\varepsilon_2)$ denote the noise drawn for query f_i , and let $v \sim \text{Expo}(\Delta/\varepsilon_1)$ denote the
 555 noise drawn for the threshold. We will fix the randomness of $v_{i < k}$ for datasets x and x' , and let
 556 $\tau = \max_{i < k} f_i(x) + v_i$ and $\tau' = \max_{i < k} f_i(x') + v_i$.

557 It suffices then to show that for any fixed randomness $v_{i < k}$ we have

$$\Pr_{v, v_k} [\tau < T + v < f_k(x) + v_k] \leq \exp(\varepsilon(x, x')) \Pr_{v, v_k} [\tau' < T + v < f_k(x') + v_k]$$

558 We can then prove this by showing that for every pair of draws v, v_k that satisfies $\tau < T + v < f_k(x) + v_k$,
 559 there is a unique pair v', v'_k that satisfies $\tau' < T + v' < f_k(x') + v'_k$ and that

$$p_{\text{Expo}}(v; \Delta/\varepsilon_1) \cdot p_{\text{Expo}}(v_k; \Delta/\varepsilon_2) \leq \exp(\varepsilon(x, x')) p_{\text{Expo}}(v'; \Delta/\varepsilon_1) \cdot p_{\text{Expo}}(v'_k; \Delta/\varepsilon_2)$$

560 By definition, draws from the exponential distribution must take non-negative values so we must
 561 also ensure that $v' \geq v$ and $v'_k \geq v_k$. As such, we let $\Delta_1 = \max\{0, \tau' - \tau\}$ and $\Delta_2 = \max\{0, \Delta_1 - (f_k(x') -$

562 $f_k(x)\}$, and our injective mapping will be $v' = v + \Delta_1$ and $v'_k = v_k + \Delta_2$. It is then straightforward to
 563 see that if $\tau < T + v$ we must have $\tau' < T + v + \Delta_1$. Similarly, if $T + v < f_k(x) + v_k$ then $T + v + \Delta_1 <$
 564 $f_k(x') + v_k + \Delta_2$ because $\Delta_2 \geq \Delta_1 - (f_k(x') - f_k(x))$. The PDF of the exponential distribution then
 565 gives us $p_{\text{Expo}}(v; \Delta/\varepsilon_1) \leq \exp(\frac{\varepsilon_1 \Delta_1}{\Delta}) p_{\text{Expo}}(v'; \Delta/\varepsilon_1)$ and $p_{\text{Expo}}(v_k; \Delta/\varepsilon_2) \leq \exp(\frac{\varepsilon_2 \Delta_2}{\Delta}) p_{\text{Expo}}(v'_k; \Delta/\varepsilon_2)$

566 Due to the fixing of randomness, we have that $\tau' - \tau \leq \max_{i < k} (f_i(x') - f_i(x))$ so $\Delta_1 \leq \Delta_k(x, x')$,
 567 which implies our desired inequality above.

568 □

569 This proof then easily applies to Lap as well. The proof for Gumbel will differ though as we no
 570 longer have similar closeness of PDF properties, but can instead utilize our closed form expressions.

571 **Lemma B.2.** *Given a dataset x , a stream of queries $\{f_i\}$ and threshold T , for any outcome $\{\perp^{k-1}, \top\}$
 572 from running AboveThreshold with Gumbel noise and $\varepsilon = \varepsilon_1 = \varepsilon_2$, then*

$$\Pr [\text{AboveThreshold}(x, \{f_i\}, T) = \{\perp^{k-1}, \top\}] = \frac{\exp(\frac{\varepsilon}{\Delta} f_k(x))}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^k \exp(\frac{\varepsilon}{\Delta} f_i(x))} \cdot \frac{\exp(\frac{\varepsilon}{\Delta} T)}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^{k-1} \exp(\frac{\varepsilon}{\Delta} f_i(x))}$$

573 *Proof.* We add noise $\text{Gumbel}(\Delta/\varepsilon)$ to all $f_i(x)$ and the threshold T , and in order to output $\{\perp^{k-1}, \top\}$
 574 we must have that for the first k queries, the noisy value of $f_k(x)$ is the largest and the noisy
 575 threshold is the second largest. It is folklore in the literature that adding Gumbel noise (with the
 576 same noise parameter) and choosing the largest index is equivalent to the exponential mechanism.
 577 This can also be extended to showing that adding Gumbel noise and choosing the top- k indices in
 578 order is equivalent to the peeling exponential mechanism. The peeling exponential mechanism
 579 first selects the top index using the exponential mechanism and removes it from the candidate set,
 580 repeating iteratively until accumulating the top- k indices. A formal proof of this folklore result is
 581 also provided in Lemma 4.2 of [Durfee and Rogers \(2019\)](#). Applying this result and the definition of
 582 the exponential mechanism gives the desired equality.

583 □

584 Interestingly, we can similarly show that AboveThreshold with Gumbel noise is equivalent to an
 585 iterative exponential mechanism that we provide in Section B.4. The closed form expression will
 586 actually allow for a slightly improved characterization that we also utilize in Section C.2.

587 **Lemma B.3.** *For a stream of queries $\{f_i\}$ with sensitivity Δ , threshold T and neighboring data sets
 588 x, x' , if we run Algorithm 1 with Gumbel noise with $\varepsilon_1 = \varepsilon_2$, then for any outcome $\{\perp^{k-1}, \top\}$ we have*

$$\Pr [\text{AboveThreshold}(x, \{f_i\}, T) = \{\perp^{k-1}, \top\}] \leq \exp(\varepsilon(x, x')) \Pr [\text{AboveThreshold}(x', \{f_i\}, T) = \{\perp^{k-1}, \top\}]$$

589 and further we can relax $\Delta_k(x, x')$ to $\max_{i < k} (f_i(x') - f_i(x))$.

590 *Proof.* Let $\varepsilon = \varepsilon_1 = \varepsilon_2$. From Lemma B.2 we know that

$$\Pr [\text{AboveThreshold}(x, \{f_i\}, T) = \{\perp^{k-1}, \top\}] = \frac{\exp(\frac{\varepsilon}{\Delta} f_k(x))}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^k \exp(\frac{\varepsilon}{\Delta} f_i(x))} \cdot \frac{\exp(\frac{\varepsilon}{\Delta} T)}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^{k-1} \exp(\frac{\varepsilon}{\Delta} f_i(x))}$$

591 Similar to the proof of Lemma B.1, we let $\Delta_1 = \max_{i < k} (f_i(x') - f_i(x))$ and $\Delta_2 = \max\{0, \Delta_1 + (f_k(x) -$
 592 $f_k(x'))\}$. We first want to show that

$$\frac{\exp(\frac{\varepsilon}{\Delta} T)}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^{k-1} \exp(\frac{\varepsilon}{\Delta} f_i(x))} \leq e^{\frac{\varepsilon \Delta_1}{\Delta}} \frac{\exp(\frac{\varepsilon}{\Delta} T)}{\exp(\frac{\varepsilon}{\Delta} T) + \sum_{i=1}^{k-1} \exp(\frac{\varepsilon}{\Delta} f_i(x'))}$$

593 This inequality reduces to $\sum_{i=1}^{k-1} \exp(\frac{\epsilon}{\Delta} f_i(x')) \leq \exp(\epsilon \Delta_1 / \Delta) \sum_{i=1}^{k-1} \exp(\frac{\epsilon}{\Delta} f_i(x))$ which follows from
 594 our definition of Δ_1 . Next we want to show that

$$\frac{\exp(\frac{\epsilon}{\Delta} f_k(x))}{\exp(\frac{\epsilon}{\Delta} T) + \sum_{i=1}^k \exp(\frac{\epsilon}{\Delta} f_i(x))} \leq e^{\frac{\epsilon \Delta_2}{\Delta}} \frac{\exp(\frac{\epsilon}{\Delta} f_k(x'))}{\exp(\frac{\epsilon}{\Delta} T) + \sum_{i=1}^k \exp(\frac{\epsilon}{\Delta} f_i(x'))}$$

595 First we let $\Delta'_1 = \max_{t \leq k} (f_t(x') - f_t(x))$ and $\Delta'_2 = f_k(x) - f_k(x')$. We see that we equivalently have
 596 $\Delta_2 = \Delta'_1 + \Delta'_2$ because if $\Delta'_1 > \Delta_1$ then $\Delta'_1 = f_k(x') - f_k(x) = -\Delta'_2$, and $\Delta_1 - \Delta'_1 < 0$ which implies
 597 $\Delta'_1 + \Delta'_2 = \max\{0, \Delta_1 + \Delta'_2\} = 0$. Furthermore, we see that $\exp(\frac{\epsilon}{\Delta} f_k(x)) = \exp(\epsilon \Delta'_2 / \Delta) \exp(\frac{\epsilon}{\Delta} f_k(x'))$
 598 and $\sum_{i=1}^k \exp(\frac{\epsilon}{\Delta} f_i(x')) \leq \exp(\epsilon \Delta'_1 / \Delta) \sum_{i=1}^k \exp(\frac{\epsilon}{\Delta} f_i(x))$ which implies the inequality above.

599 Combining our two inequalities along with the fact that $\Delta_1 \leq \Delta_k(x, x')$ completes the proof.

600

□

601 Now that we've bounded the one-sided privacy loss for neighboring datasets for each type of noise,
 602 this will immediately imply an overall privacy bound over all neighbors that utilizes this general
 603 characterization.

604 **Corollary B.1.** *For a stream of queries $\{f_i\}$ with sensitivity Δ and threshold T , Algorithm 1 with Lap,
 605 Gumbel, or Expo noise is ϵ -DP where*

$$\epsilon = \max_{x \sim x'} \epsilon(x, x')$$

606 and for Gumbel noise we must have $\epsilon_1 = \epsilon_2$.

607 B.2 Range-bounded definition and properties

608 The definition of *range-bounded* was originally introduced in Durfee and Rogers (2019) to tighten
 609 the privacy bounds on the composition of exponential mechanisms. The analysis was further
 610 extended in Dong et al. (2020) to give the optimal composition of range-bounded mechanisms. This
 611 definition was then unified with other definitions in Cesar and Rogers (2021).

612 **Definition B.2** (Durfee and Rogers (2019)). *A mechanism $M : \mathcal{X} \rightarrow \mathcal{Y}$ is ϵ -range-bounded if for
 613 any neighboring datasets $x, x' \in \mathcal{X}$ and outcomes $y, y' \in \mathcal{Y}$:*

$$\frac{\Pr[M(x) = y]}{\Pr[M(x') = y]} \leq e^\epsilon \frac{\Pr[M(x) = y']}{\Pr[M(x') = y']}$$

614 If we have bounds upon this property that are stronger than those immediately implied by DP, then
 615 it was shown in Dong et al. (2020) that substantial improvements can be made in bounding the
 616 overall DP over the composition of these mechanisms. While these improvements can be applied
 617 to our results here as well, we focus upon the privacy bounds with respect to zCDP as it is much
 618 cleaner to work with in providing strong composition bounds. From the proof of Lemma 3.4 in
 619 Cesar and Rogers (2021) we see that the range-bounded property can give a significant improvement
 620 in zCDP properties compared to similar DP guarantees.

621 **Proposition 2** (Cesar and Rogers (2021)). *If a mechanism M is ϵ -range-bounded then it is $\frac{1}{8}\epsilon^2$ -zCDP.*

622 Our general characterization is already set up to provide range-bounded properties of
 623 AboveThreshold. This will allow us to give tighter guarantees on the range-boundedness which
 624 can be translated to better privacy composition bounds.

625 **Corollary B.2.** *For a stream of queries $\{f_i\}$ and threshold T , Algorithm 1 with Lap, Gumbel, or Expo
 626 noise is ϵ -range-bounded where*

$$\epsilon = \max_{x \sim x'} (\epsilon(x, x') + \epsilon(x', x))$$

627 and for Gumbel noise we must have $\epsilon_1 = \epsilon_2$.

628 *Proof.* We can rewrite the inequality from Definition B.2 as equivalently

$$\Pr [M(x) = y] \cdot \Pr [M(x') = y'] \leq e^\epsilon \Pr [M(x') = y] \cdot \Pr [M(x) = y']$$

629 Our claim then follows immediately by applying Lemma B.1 and Lemma B.3 □

630 B.3 Improved composition analysis of SVT

631 In this section we complete our analysis for improving the privacy bounds for sparse vector
632 technique. While our generalized characterization was in a more complex form, we simplify it here
633 for general queries and monotonic queries. This will then match up to the privacy guarantees that
634 we gave in Section 3 for AboveThreshold.

635 **Lemma B.4.** *Given a stream of queries $\{f_i\}$ with sensitivity Δ then for any neighboring datasets*
636 *$\epsilon(x, x') \leq \epsilon_1 + 2\epsilon_2$. If the queries are monotonic then for any neighboring datasets $\epsilon(x, x') \leq \epsilon_1 + \epsilon_2$,*
637 *and furthermore, $\epsilon(x, x') + \epsilon(x', x) \leq \epsilon_1 + 2\epsilon_2$.*

638 *Proof.* By Definition 2.2 we must always have $\Delta_k(x, x') \leq \Delta$ for any neighboring datasets. Fur-
639 thermore, we must also have $f_k(x') - f_k(x) \geq -\Delta$ for any neighboring datasets. This implies
640 $\epsilon(x, x') \leq \epsilon_1 + 2\epsilon_2$ for any neighboring datasets.

641 Now assume the queries are monotonic. Without loss of generality assume $f_k(x) \geq f_k(x')$ for all
642 k . Therefore $\Delta_k(x, x') = 0$ and $f_k(x') - f_k(x) \geq -\Delta$, which implies $\epsilon(x, x') \leq \epsilon_2$. Similarly we have
643 $\Delta_k(x', x) \leq \Delta$ but also $f_k(x) - f_k(x') \geq 0$, which implies $\epsilon(x', x) \leq \epsilon_1 + \epsilon_2$. Combining these implies
644 $\epsilon(x', x) \leq \epsilon_1 + \epsilon_2$ and $\epsilon(x', x) + \epsilon(x, x') \leq \epsilon_1 + 2\epsilon_2$ for any neighboring datasets. □

645 This lemma along with Corollary B.1 immediately imply Corollary 3.1 and Lemma 3.2

646 *Proof of Theorem 1.* From Lemma B.4 and Corollary B.2 we have that Algorithm 1 is $\epsilon_1 + 2\epsilon_2$ -range-
647 bounded. Applying Proposition 2 then gives the zCDP guarantees. □

648 The sparse vector technique is simply an iterative call to AboveThreshold and as such the privacy
649 bounds will come from the composition. By analyzing the composition through zCDP which
650 has become commonplace in the privacy community, we can immediately improve the privacy
651 guarantees for monotonic queries in the sparse vector technique. Furthermore we can also improve
652 the privacy guarantees under the standard definition using the improved composition bounds for
653 range-bounded mechanisms from Dong et al. (2020).

654 We also note here that the generalized Sparse Vector Technique given in Lyu et al. (2017) only adds
655 noise once to the threshold and only has to scale the noise to the number of calls to AboveThreshold
656 for the queries and not the threshold. Our analysis can extend to give the same properties for Lap
657 and Expo noise, but utilizing advanced composition properties will give far more accuracy. More
658 specifically, the threshold and queries having noise proportional to \sqrt{c}/ϵ , is preferable to all the
659 queries having noise proportional to c/ϵ , and only the threshold proportional to $1/\epsilon$, where c is the
660 number of calls to AboveThreshold.

661 B.4 Iterative exponential mechanism

662 It is folklore that the peeling exponential mechanism is equivalent to taking the top- k after adding
663 Gumbel noise, as formally shown in Lemma 4.2 of Durfee and Rogers (2019). We show a similar
664 property here for AboveThreshold with Gumbel noise that it is equivalent to iteratively running
665 the exponential mechanism.

666 **Lemma B.5.** *If $\epsilon_1 = \epsilon_2 = \epsilon/2$ then Algorithm 1 with Gumbel noise gives the equivalent output*
667 *distribution to Algorithm 3.*

668 *Proof.* We first define

$$p_k = \frac{\exp(\frac{\epsilon}{2\Delta} f_k(x))}{\exp(\frac{\epsilon}{2\Delta} T) + \sum_{i=1}^k \exp(\frac{\epsilon}{2\Delta} f_i(x))}$$

Algorithm 3 Iterative Exponential Mechanism

Require: Input dataset x , a stream of queries $\{f_i : \mathcal{X} \rightarrow \mathbb{R}\}$ with sensitivity Δ , and a threshold T

- 1: **for** each query i **do**
- 2: Run exponential mechanism with $\mathcal{Y} = \{0, \dots, i\}$ where $q(x, 0) = T$ and $q(x, j) = f_j(x)$ for all $j > 0$
- 3: **if** i is selected **then**
- 4: Output \top and **halt**
- 5: **else**
- 6: Output \perp
- 7: **end if**
- 8: **end for**

669 By construction, the probability of Algorithm 3 outputting $\{\perp^{k-1}, \top\}$ is equal to $p_k \prod_{i=1}^{k-1} (1 - p_i)$.
670 Furthermore, by our definition of p_k we have

$$1 - p_k = \frac{\exp(\frac{\epsilon}{2\Delta}T) + \sum_{i=1}^{k-1} \exp(\frac{\epsilon}{2\Delta}f_i(x))}{\exp(\frac{\epsilon}{2\Delta}T) + \sum_{i=1}^k \exp(\frac{\epsilon}{2\Delta}f_i(x))}$$

671 Through telescoping cancellation

$$\prod_{i=1}^{k-1} (1 - p_i) = \frac{\exp(\frac{\epsilon}{2\Delta}T)}{\exp(\frac{\epsilon}{2\Delta}T) + \sum_{i=1}^{k-1} \exp(\frac{\epsilon}{2\Delta}f_i(x))}$$

672 From Lemma B.2 we then see that the output probabilities are equivalent.

673

□

674 C Additional unbounded quantile estimation results

675 In this section we first extend our method to the fully unbounded setting by simply making two
676 calls to AboveThreshold and essentially searching through both positive and negative numbers in
677 each respective call. Next we show how our methods can also extend to the *add-subtract* definition
678 of neighboring. Further, we'll apply our approach to the framework set up in Kaplan et al. (2022)
679 for computing multiple quantiles and extend it to the *swap* definition.

680 C.1 Fully unbounded quantile estimation

681 Recall that our unbounded quantile estimation algorithm assumed that the data was lower bounded.
682 It then slowly increased that bound by a small percentage until the appropriate amount of data fell
683 below the threshold for the given quantile. In order to extend to the fully unbounded setting, we
684 will simply first apply this guess and check method to the positive numbers and then apply it to the
685 negative numbers.

Algorithm 4 Fully unbounded quantile mechanism

Require: Input dataset x , a quantile q , and parameter $\beta > 1$

- 1: Run AboveThreshold with x , $T = qn$ and $f_i(x) = |\{x_j \in x | x_j + 1 < \beta^i\}|$
- 2: Run AboveThreshold with x , $T = (1 - q)n$ and $f_i(x) = |\{x_j \in x | x_j - 1 > -\beta^i\}|$
- 3: If the first AboveThreshold halts at $k > 0$ then output $\beta^k - 1$
- 4: If the second AboveThreshold halts at $k > 0$ then output $-\beta^k + 1$
- 5: Otherwise return 0

686 Note that the second call could equivalently be achieved by flipping the sign of all the datapoints
687 and again applying queries $f_i(x) = |\{x_j \in x | x_j + 1 < \beta^i\}|$. Therefore all our privacy guarantees from
688 Section 4 will still apply, but composing over two calls to AboveThreshold, which is the Sparse
689 Vector Technique.

690 In the first call to AboveThreshold we are assuming a lower bound of 0, and searching the positive
691 numbers. If it terminates immediately then it is likely that more than a q th fraction of the data
692 is below 0. We then symmetrically search through the negative numbers by assuming an upper
693 bound of 0. If this halts immediately then it is likely that the quantile is already near 0. We could
694 also apply other variants of this, such as three total calls to get the maximum and minimum of the
695 data if we wanted full data bounds.

696 Further note that computing the smallest quantiles will be challenging for our algorithm because it
697 may take many queries to get to the appropriate threshold, but each query will have a reasonable
698 chance of terminating if q is very small. To account for these queries in the lower-bounded setting,
699 we can invert all the datapoints and instead search for quantile $1 - q$ on this transformed data, then
700 invert our resulting estimate. We would also need to reduce our parameter β a reasonable amount
701 in this setting.

702 C.2 Extension to *add-subtract* neighbors

703 We also extend our results to the *add-subtract* definition of neighboring datasets.

704 **Definition C.1.** *Datasets $x, x' \in \mathcal{X}$ are neighboring if one of them can be obtained from the other by*
705 *adding or removing one individual's data.*

706 Under this definition we see that our threshold $T = qn$ is no longer fixed, so we will instead set
707 each query to be $f_i(x) = |\{x_j \in x \mid x_j - \ell + 1 < \beta^i\}| - qn$. Unfortunately this query will no longer be
708 monotonic, so we will instead take advantage of our more general characterization in Section B.1 to
709 further tighten the privacy bounds in this setting.

710 **Lemma C.1.** *Given the stream of queries $f_i(x) = |\{x_j \in x \mid x_j - \ell + 1 < \beta^i\}| - qn$ with sensitivity*
711 *$\Delta = 1$, for any neighboring datasets x, x' under Definition C.1 and quantile $q \in [0, 1]$ we have*
712 *$\epsilon(x, x') \leq \max\{(1 - q)\epsilon_1, q\epsilon_1 + \epsilon_2\}$.*

713 *Proof.* Without loss of generality, assume that x has one more individual's data, so $x \in \mathbb{R}^n$ and
714 $x' \in \mathbb{R}^{n-1}$. Let $g_i(x) = |\{x_j \in x \mid x_j - \ell + 1 < \beta^i\}|$. By construction we must have $g_i(x) - g_i(x') \in \{0, 1\}$
715 because x has an additional datapoint. Furthermore, because the thresholds are increasing, if
716 $g_k(x) - g_k(x') = 1$ then $g_i(x) - g_i(x') = 1$ for all $i > k$. Similarly if $g_k(x) - g_k(x') = 0$ then
717 $g_i(x) - g_i(x') = 0$ for all $i < k$. We further see that $f_i(x) - f_i(x') = g_i(x) - g_i(x') - q$ for all i .

718 First consider the case when $g_k(x) - g_k(x') = 0$. We therefore have $g_i(x) - g_i(x') = 0$ for all $i < k$
719 so $\Delta_k(x, x') = q$, and also $\Delta_k(x, x') - (f_k(x') - f_k(x)) = 0$, so $\epsilon(x, x') \leq q\epsilon_1$. Similarly, we have
720 $\Delta_i(x', x) = 0$ and $\Delta_k(x, x') - (f_k(x') - f_k(x)) = q$, so $\epsilon(x', x) \leq q\epsilon_2$.

721 Next consider the case when $g_k(x) - g_k(x') = 1$. Therefore we have $\Delta_k(x, x') \leq q$ and also
722 $f_k(x') - f_k(x) = q - 1$ which implies $\Delta_k(x, x') - (f_k(x') - f_k(x)) \leq 1$, so $\epsilon(x, x') \leq q\epsilon_1 + \epsilon_2$. Similarly,
723 we have that $\Delta_k(x', x) \leq 1 - q$, and thus $\Delta_k(x', x) - (f_k(x) - f_k(x')) = 0$, so $\epsilon(x', x) \leq (1 - q)\epsilon_1$.

724 Combining these bounds gives our desired result.

725 □

726 With these improved bounds we can then show that our methods also extend to the *add-subtract*
727 definition of neighboring with tighter privacy guarantees.

728 **Corollary C.1.** *If we run Algorithm 2 with $\epsilon_1 = \epsilon_2$ for a quantile q under Definition C.1 then it is*
729 *$(q\epsilon_1 + \epsilon_2)$ -DP. Further if the noise in AboveThreshold is Gumbel then it achieves $\frac{1}{8}(\epsilon_1 + \epsilon_2)^2$ -zCDP.*

730 *Proof.* Combining Lemma C.1 and Corollary B.1 immediately imply the first claim. For the second
731 claim, we further examine the proof Lemma C.1 and reconsider the case when $g_k(x) - g_k(x') = 0$.
732 Also note the relaxation for Gumbel noise in Lemma B.3. We then have $\Delta_k(x', x) = -q$ and thus
733 $\Delta_k(x, x') - (f_k(x') - f_k(x)) = 0$, which instead implies $\epsilon(x', x) \leq 0$ for this case. Therefore under
734 these conditions, for neighbors x, x' we have without loss of generality that $\epsilon(x, x') \leq q\epsilon_1 + \epsilon_2$ and
735 $\epsilon(x', x) \leq (1 - q)\epsilon_1$. Therefore it must be $(\epsilon_1 + \epsilon_2)$ -range-bounded and applying Proposition 2 gives
736 the zCDP guarantees.

737 □

738 C.3 Extension to multiple quantile estimation

739 As previously established, the framework in Kaplan et al. (2022) is agnostic to the single quantile
740 estimation method. However their proof is for the *add-subtract* neighbors, although they note that
741 it can be extended easily to the *swap* neighbors. We also discuss here how it can be extended for
742 completeness.

743 For the *swap* neighboring definition, at each level of the recursive partitioning scheme either two
744 partitions differ under the *add-subtract* definition, or one partition differs under the *swap* definition.
745 Applying their framework to our algorithm, we will instead compute all the thresholds in advance
746 of the splitting. Accordingly these thresholds will remain fixed. If the thresholds are fixed then we
747 can actually further improve our privacy bounds for the *add-subtract* definition. Once again we
748 will use our more general characterization from Section B.1.

749 **Lemma C.2.** *Given the stream of queries $f_i(x) = |\{x_j \in x \mid x_j - \ell + 1 < \beta^i\}|$ with sensitivity $\Delta = 1$, for
750 any neighboring datasets x, x' under Definition C.1 we have $\varepsilon(x, x') \leq \max\{\varepsilon_1, \varepsilon_2\}$*

751 *Proof.* Without loss of generality, assume that x has one more individuals data, so $x \in \mathbb{R}^n$ and
752 $x' \in \mathbb{R}^{n-1}$. By construction we must have $f_i(x) - f_i(x') \in \{0, 1\}$ because x has an additional datapoint.
753 Furthermore, because the thresholds are increasing, if $f_k(x) - f_k(x') = 0$ then $f_i(x) - f_i(x') = 0$ for
754 all $i < k$. Thus the case of $f_k(x) - f_k(x') = 0$ is easy.

755 Instead consider $f_k(x) - f_k(x') = 1$. We know $\Delta_k(x, x') = 0$ so $\Delta_k(x, x') - (f_k(x') - f_k(x)) = 1$, so
756 $\varepsilon(x, x') \leq \varepsilon_2$. Further we must have $\Delta_k(x', x) \leq 1$ and $\Delta_k(x', x) - (f_k(x) - f_k(x')) = 0$, so $\varepsilon(x', x) \leq \varepsilon_1$.
757 Combining these implies $\varepsilon(x', x) \leq \max\{\varepsilon_1, \varepsilon_2\}$ for any neighboring datasets.

758 □

759 Applying these bounds we see that applying the framework of Kaplan et al. (2022) to the *swap*
760 definition either leads to one composition of $(\varepsilon_1 + \varepsilon_2)$ -DP for the *swap* definition or two compositions
761 of $\max\{\varepsilon_1, \varepsilon_2\}$ for the *add-subtract* definition at each level. Setting $\varepsilon_1 = \varepsilon_2$ we then have that the
762 privacy cost of computing m quantiles with this framework will be equivalent to $\log(m + 1)$
763 compositions of $(\varepsilon_1 + \varepsilon_2)$ -DP.

764 C.4 Applying permute-and-flip framework to EMQ

765 It would be an interesting future direction to see if the EMQ method could also instead effectively
766 utilize the permute-and-flip framework from McKenna and Sheldon (2020) that was shown to
767 improve accuracy. While this framework is equivalent to adding Expo noise for report noisy max,
768 the quantile problem requires considering an infinite domain, which makes the Expo noise addition
769 procedure impossible. Using this alternate, but equivalent Ding et al. (2021), approach can make this
770 possible. If we look at Algorithm 3 from McKenna and Sheldon (2020) we could also extend this to
771 drawing a single point uniformly in the interval, then drawing the Bernoulli and removing it from
772 the interval. However due to the continuous nature of the interval, it's unlikely that performing
773 this sampling without replacement will have any noticeable improvement over sampling with
774 replacement which is equivalent to the exponential mechanism. Furthermore, we cannot run
775 this framework as efficiently because there are no corresponding nice closed form expressions
776 compared to the exponential mechanism. Ideally, we would modify this approach to draw each
777 interval proportional to it's length, then draw the Bernoulli and remove the interval. However it is
778 critical to note that this does not give an identical distribution because the sampling is done without
779 replacement. Accordingly this simple modification would not work, but there could potentially be
780 other effective changes to fit this framework that we leave to future work.

781 Furthermore, we note that adding Expo noise for report noisy max does not achieve the range-
782 bounded property that the exponential mechanism enjoys Durfee and Rogers (2019). So the compo-
783 sition improvements for zCDP from Proposition 2 could not be applied.

784 D Further Experiment Details

785 In this section we provide some follow-up details from our empirical evaluation.

786 **D.1 Data histograms**

787 We provide histograms of our datasets for better understanding in Figure 3.

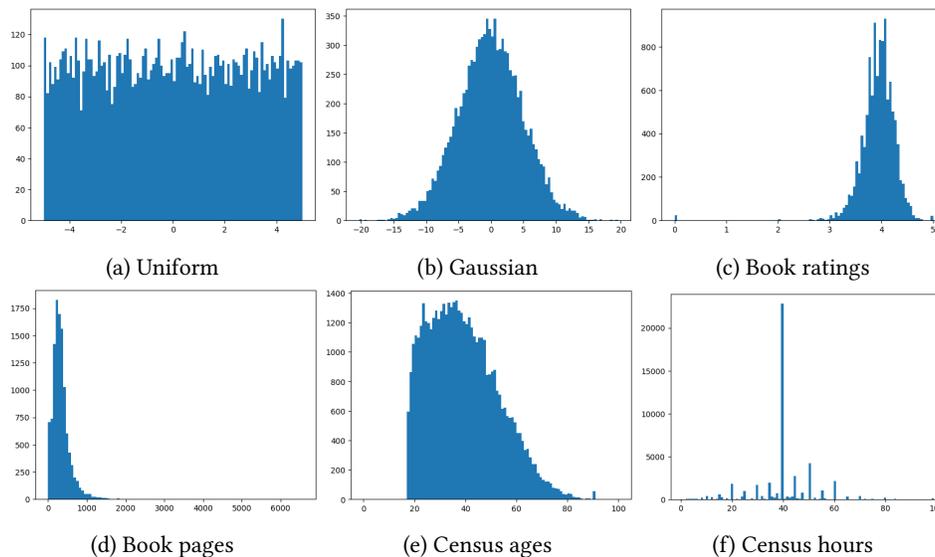


Figure 3: Histograms for each of our datasets.

788 **D.2 Parameter tuning**

789 We kept our β parameter fixed in all experiments for consistency but also to make our method
790 data agnostic. However, our choice was aggressively small in order to achieve higher precision in
791 the inner quantile estimation comparison in Section 5.2. This choice was still highly resilient to
792 changes in ε for our sum experiments as we see our error only scaled proportional to the increase
793 in noise. But for more significant decreases in ε or in the data size, i.e. the conditions under which
794 all private algorithms suffer substantial accuracy loss, this choice of β could be too small. Those
795 settings imply that the noise added is larger and the distance between queries and threshold shrinks,
796 so our method is more likely to terminate earlier than desired. Smaller β values will then intensify
797 this issue as the candidate values increase more slowly. For the clipping application, we generally
798 think using a value of $\beta = 1.01$ would be a more stable choice. In fact, some preliminary testing
799 shows that this setting actually improves our results in Table 1. Furthermore, increasing β will also
800 reduce the number of queries and thus the computational cost. We leave to future work a more
801 thorough analysis of this parameter to determine a good default value that performs well agnostic
802 to the input data. Additionally, all our methods and proofs only require an increasing sequence
803 of candidate values and it's possible that other potential sequences would be even more effective.
804 For example, if tight upper and lower bounds are known on the data, such as within the recursive
805 computation of multiple quantiles, then it likely makes more sense to simply uniformly partition
806 the interval and check in increasing order. But we leave more consideration upon this to future
807 work as well.

808 Our choice of $q = 0.99$ in the differentially private sum experiments was also to maintain consistency
809 with the choices for the previous method but also to keep variance of the estimates lower. This
810 will create some negative bias as we then expect to clip the data. As we can see in our illustrative
811 example Figure 2, the PDF will exponentially decrease once we pass the true quantile value, but
812 will do so less sharply once all the queries have value n . Accordingly, setting $q = 1.0$ would add
813 slightly more variance to the estimation but initial testing showed improvement in error. However,
814 if the user would prefer slightly higher variance to avoid negative bias, then setting the threshold
815 at n or even $n + 1/\varepsilon$, would make it far more likely that the process terminates with a value slightly
816 above the maximum. This is particularly useful for heavy-tailed data, where clipping at the 99th
817 percentile can have an out-sized impact on the bias. We leave a more rigorous examination of these
818 bias-variance tradeoffs for good default settings to future work.

819 D.3 Implementation code

820 For ease of implementation, we provide some simple python code to run our method with access
821 to the respective Noise generator of the users choosing. In our experiments we used exponential
822 noise from the numpy library.

```
823 def unboundedQuantile(data, l, b, q, eps_1, eps_2):  
824     d = defaultdict(int)  
825     for x in data:  
826         i = math.log(x-l+1, b) // 1  
827         d[i] += 1  
828  
829  
830     t = q * len(data) + noise(1/eps_1)  
831     cur, i = 0, 0  
832     while True:  
833         cur += d[i]  
834         i += 1  
835         if cur + noise(1/eps_2) > t:  
836             break  
837     return b**i - l + 1  
838
```